7 Fictions, Representations, and Reality

Margaret Morrison

1. INTRODUCTION

There are many ways in which we use unrealistic representations for modeling the physical world. In some cases we construct models that we know to be false—but not false in the sense that they involve idealization or abstraction from real properties or situations; most models do this. Instead, they are considered false because they describe a situation that cannot, no matter how many corrections are added, be physically true of the phenomenon in question. Maxwell's ether models are a case in point. No one understood or believed that the structure of the ether consisted of idle wheels and rotating vortices, yet those types of models were the foundation of Maxwell's first derivation of the electromagnetic field equations.

Other instances of model building involve mathematical abstractions that are also not accurate representations of physical phenomena. For example, in his work on population genetics, R. A. Fisher (1918, 1922) assumed an analogy between populations of genes and the way that statistical mechanics models populations of molecules in a gas. These populations contained an infinite number of genes that act independently of each other. His inspiration was the velocity distribution law, which gave results about pressure (among other things) from highly idealized assumptions about the molecular structure of a gas. This kind of abstraction (assuming infinite populations of genes) was crucial in enabling Fisher to show that selection indeed operated in Mendelian populations.[1] In situations like this where we have mathematical abstractions that are *necessary* for arriving at a certain result there is no question of relaxing or correcting the assumptions in the way we de-idealize cases like frictionless planes and so on; the abstractions are what make the model work.

Another type of model or modeling assumption(s) that also resists corrections is the type used to treat a specific kind of problem. Deviations from those particular situations typically do not involve corrections to the model's assumptions but the introduction of a new model that describes the situation somewhat differently. For example, the typical Fisher–Wright model in modern population genetics assumes that generations are discrete—they do not overlap. However, if we want to examine changes in allele frequencies when we have overlapping generations, we don't simply add corrections or parameters to the discrete generation models; instead a different model is used (Moran model) for which parallels to the Fisher–Wright models can be drawn. A similar situation occurs in the case of nuclear models. The kinds of assumptions about nuclear structure required by the liquid drop model for explaining and predicting fission are very different from those contained in the shell model for explaining magic numbers.[2] These kinds of models may be fictional in a way that is similar to Maxwell's ether models, but the important point is that they are not open to the addition of correction factors or de-idealization.

Finally, we have the relatively straightforward cases of idealization where a law or a model idealizes or leaves out a particular property but allows for the addition of correction factors that bring the model system closer (in representational terms) to the physical system being modelled or described. The Hardy–Weinberg law is a good example of a case where violations of certain conditions, like random mating, may mean that the law fails to apply, but in other situations, depending on the type and degree of deviation from idealized conditions (e.g., no mutation), the law may continue to furnish reasonably accurate predictions. In other words, the population will continue to have Hardy-Weinberg proportions in each generation but the allele frequencies will change with that condition. The simple pendulum is a familiar example where we know how to add correction factors that bring the model closer to concrete phenomena. The key in each of these cases is that we know how to manipulate the idealizations to get the outcomes we want.

Because each of the cases I have just mentioned has a different structure, the important question is whether they exemplify anything *different* about the way unrealistic representations yield reliable knowledge.[3] I see this not as a logical problem of deriving true conclusions from false premises but rather an epistemic one that deals with the way false representations transmit information about concrete cases.[4] The latter is a problem that in some sense pervades many cases of knowledge acquisition to a greater or lesser extent—think of the use of metaphors in transmitting information. It is tempting to classify all of the examples just given as instances of fictional representation and then ask how fictions give us knowledge of real-world situations. I want to argue that this would be a mistake. The language of fictions is at once too broad and too narrow. Although it encompasses the fact that none of these representations is realistic, it fails to capture specifics of the relation that certain kinds of model-representations have to real systems. I claim that this is because the processes involved in both abstraction and idealization are not typically the same as those involved in constructing fictional models/representations. Introducing a mathematical abstraction that is necessary for obtaining certain results involves a different type of activity than constructing a model you know to be false in order to see whether certain analogies or similarities can be established. To simply classify all forms of nonaccurate description as fictions is to ignore the different ways scientific representation is linked with explanation and understanding.

I want to address the problem of unrealistic representation by introducing a finer grained distinction that uses the notion of fictional representation to refer only to the kind of models that fall into the category occupied by Maxwell's ether models. My reasons for doing so reflect the practice behind the construction of these types of models. Fictional models are deliberately intended as imaginary accounts whose physical similarity to the target system is not immediately obvious. Instead, one needs to examine the specific details of the model in order to establish the appropriate kinds of relations. Contrast this with the use of idealization where we have conditions that have been deliberately omitted or idealized (frictionless planes) in order to facilitate calculation or to illustrate a general principle for a simple case. Here we usually know immediately what the purported relation to the target system actually is.

In keeping with my rather narrow account of fictional representation, I want to also suggest a way of thinking about abstraction that differentiates it from idealization in the following way: Where idealization distorts or omits properties that are often not necessary for the problem at hand, abstraction (typically mathematical in nature) *introduces* a specific type of representation that is not amenable to correction and is necessary for explanation/prediction of the target system. What is crucial about abstraction, characterized in this way, is that it highlights the fact that the process is not simply one of adding back and taking away as characterized in the literature; instead it shows how certain kinds of mathematical representations are essential for explaining/predicting concrete phenomena.[5] Once again, my aim in distinguishing these different kinds of representation is to call attention to the fact that the notion of a 'fiction' is not sufficiently rich to capture the various ways that mathematical abstraction and idealization function in explanation and prediction.

So, what is it about the *structure* of each of these types of representation that makes them successful in a particular context? As I mentioned earlier, this isn't a logical problem in that we aren't so much concerned with the informational content or argument structure (that we can get true information from false premises) but rather with the features of the *representation* that produces knowledge. For example, what aspects of populations characterized by the Hardy–Weinberg law account for the latter's success in predicting genotype frequencies?[6] Similarly, what was the essential feature in Maxwell's ether model that led to the derivation of the field equations? So, there are really two interrelated issues here: (1) How do fictional representations provide reliable information, and (2) what is essentially different about the way that fictional models as opposed to abstraction and idealization accomplish this? My intuition is that although I can draw some general conclusions about the differences between fictions, abstractions, and idealizations, the answer to (1) will be a highly context-specific affair. That is, the way a fictional model produces information will depend largely on the nature of the model itself and what we want it to do. That said, it is important to recognize that there are stories to tell about how knowledge is produced in these cases that go beyond the simple appeal to heuristic power.

Let me begin by discussing some general issues related to fictions and falsity and then go on to examine a case of a fictional model, an idealized law, and a mathematical abstraction in order to illustrate some of the differences among them and attempt to answer some of the questions just raised.

2. FABLES, FICTIONS, AND FACTS

One of the things that is especially puzzling about fictional scientific representation is the relationship it bears to, say, literary fictions. Although the latter describe worlds that we know are not real, the intention, at least in more meritorious works of fiction, is often to shed light on various aspects of our life in the real world. To that extent some kind of parallel relationship exists between the two worlds that makes the fictional one capable of "touching", in some sense, the real one. But what does this "touching" consist of? There are aspects of the fictional world that we take to be representative of the real world but only because we can draw certain parallels or assume certain similarity relations hold. For example, many of the relationships described in the novels of Simone de Beauvoir can be easily assimilated to her own experiences and life with Sartre. In other words, even though the characters are not real, the dynamic that exists between them may be an accurate depiction of the dynamic between real individuals.

But what about scientific fictions? There too we have a relationship between the real world and the world described by our models. We also want to understand certain features of those models as making some type of *realistic* claim about the world. So, although we sometimes trade in analogies or metaphors, the goal is to represent the world in as realistic a way as we are able or in as realistic a way that will facilitate calculation or understanding of some aspect of the target system. Sometimes the relation between the fictional and the real is understood in terms of the abstract and the concrete, where abstract entities or concepts are understood as fictional versions of concrete realistic entities. And, unlike the literary case, we don't always understand the model, as a whole, to be a fictional entity; sometimes there are aspects of the model that are intended as realistic representations of the system we are interested in. The question, of course, is how exactly these models transmit reliable information about physical systems.

Nancy Cartwright (1999a), in a paper called "Fables and Models," draws on the ideas of Lessing about the relationship between fables and morals as a way of shedding light on the relation between the abstract and the concrete. Lessing sees a fable as a way of providing a graspable, intuitive content for abstract symbolic judgements (the moral). Fables are like fictions—they are stories that allegedly tell us something about the world we inhabit. The interesting thing about Lessing's fables is that they aren't allegories. Why is that? Because allegories function in terms of similarity relations; they don't have literal meaning and say something similar to what they seem to say. But, as Cartwright notes, for Lessing similarity is not the right relation to focus on. The relation between the moral and the fable is that of the general to the more specific; indeed, it is a misusage to say that the special bears a similarity to the general or that the individual has a similarity with its type. In more concrete terms, Lessing's account of the relation between the fable and the moral is between the abstract symbolic claim and its more concrete manifestation.

Cartwright claims that this mirrors what is going on in physics. We use abstract concepts that need "fitting out" in a particular way using more concrete models; the laws of physics are like the morals and the models like the fables. It is because force is an abstract concept that it can only "exist in particular mechanical models" (1999a, p. 46). What she concludes is that the laws are true only of objects in the model in the way that the morals are true of their fables; so, continuing on with Lessing's analysis, we would say that the model is an instance of the law. In that sense the model is less abstract than the law; but what about the relation of the model to the world? Cartwright says she is inclined to think that even when models fit "they do not fit very exactly" (1999a, p. 48). This provides a context in which to understand Cartwright's claim about how laws can be both false and have broad applicability. They are literally false about the world, yet the concrete models that instantiate them are what constrain their application, and it is because the models also don't have an exact fit that they have such broad applicability.

But what exactly does this view entail when we move from the model to the world? If, as it seems to suggest, we can only talk about laws in the context of the fictional world described by the model, then how do we connect the fictional model with the real world that we are interested in explaining/predicting/describing? Should we understand reality as an instance of the model in the way that the model is an instance of the law, or do we need to invoke the similarity relation as a way of understanding why the model works in the way that it does? It isn't clear to me how talking in terms of instances or appealing to the general/specific relation tells us much here. When we need to know the features of the model that are instantiated in the world, we are essentially asking how the model is similar to the world. In other words, when we want to know some *facts* about the world, we need to move beyond the relation of law to model to one of model and world. Understanding the model as a concrete instantiation of a law doesn't guarantee that the model bears any similarity to the physical world. Maxwell's initial ether model was an instantiation of laws of hydrodynamics, yet it was a highly fictional representation of the ether/electromagnetic field. So, if our laws only say something about the world in virtue of the relation that the model bears to the world, the question becomes one of determining what, exactly, fictional models say about the world and the way in which they do it. The problem, however, is that if all models are fictions then we seem forced to conclude that science provides information about the world in the same way that novels do.

This characterization seems unhelpful, primarily because it fails to do justice to the way models are used in unearthing aspects of the world we want to understand. Put differently, we need to know the variety of ways models can represent the world if we are to have faith in those representations as sources of knowledge. To say force is an abstract concept that exists only in models leaves us with no insight about how to deal with physical forces that we encounter in the world. And, as I mentioned earlier, to characterize models generally as fictions doesn't tell us much either. We need a finer grained distinction that will capture the various types of unrealistic representations that are used in model construction and how those representations function in an explanatory or predictive way. Fictional representation is just one type. In order to see how we might make sense of the idea that fictional models can provide us with information about the world, and how we can retrieve information directly from idealized laws/models and mathematical abstractions, let us look at some examples of how this happens.

3. FICTIONAL MECHANISMS YIELD ACCURATE PREDICTIONS: HOW THE MODEL PROVIDES INFORMATION

In the various stages of development of the electromagnetic theory, Maxwell used a variety of tools that included different forms of a fictional ether model as well as physical analogies. Each of these played an important role in developing both mathematical and physical ideas that were crucial to the formulation and conceptual understanding of field theory. In order to appreciate exactly how a field theory emerged from these fictional representations, we need to start with Maxwell's 1856 representation of Faraday's electromagnetic theory in what he called a "mathematically precise yet visualizable form" (Maxwell, 1856; hereafter *FL*).[7] The main idea in Faraday's account was that the seat of electromagnetic phenomena was in the spaces surrounding wires and magnets, not in the objects themselves. He used iron filings to visualize the patterns of these forces in space, referring to the spatial distribution as lines of force that constituted a kind of field. Electrical charges were conceived as epiphenomena that were manifestations of the termination points of the lines of force, and as such, they had no independent existence. On this picture the field was primary with charges and currents emerging from it.

The method Maxwell employed involved both mathematical and physical analogies between stationary fields and the motion of an incompressible fluid that flowed through tubes (where the lines of force are represented by the tubes). Using the formal equivalence between the equations of heat flow and action at a distance, Maxwell substituted the flow of the ideal fluid for the distant action. Although the pressure in the tubes varied inversely as the distance from the source, the crucial difference was that the energy of the system was in the tubes rather than being transmitted at a distance. The direction of the tubes indicated the direction of the fluid in the way that the lines of force indicated the direction and intensity of a current. Both the tubes and the lines of force satisfied the same partial differential equations. The purpose of the analogy was to illustrate the mathematical similarity of the laws, and although the fluid was a purely fictional entity it provided a visual representation of this new field theoretic approach to electromagnetism.

What Maxwell's analogy did was furnish what he termed a physical "conception" for Faraday's lines of force; a conception that involved a fictional representation, yet provided a mathematical account of electromagnetic phenomena as envisioned on this field theoretic picture. The method of physical analogy, as Maxwell referred to it, marked the beginning of what he saw as progressive stages of development in theory construction. Physical analogy was intended as a middle ground between a purely mathematical formula and a physical hypothesis. It was important as a visual representation because it enabled one to see electromagnetic phenomena in a new way. Although the analogy did provide a model (in some sense), it was merely a descriptive account of the distribution of the lines in space with no mechanism for understanding the forces of attraction and repulsion between magnetic poles.

A physical account of how the behavior of magnetic lines could give rise to magnetic forces was further developed in a paper entitled "On Physical Lines of Force" (Maxwell, 1861–1862; hereafter *PL*). The paper marked the beginnings of his famous ether model, which described the magnetic field in terms of the rotation of the ether around the lines of force. The idea of a rotating ether was first put forward by Kelvin, who explained the Faraday effect (the rotation of the plane of polarized light by magnets) as a result of the rotation of molecular vortices in a fluid ether. In order to allow for the rotation of the adjacent vortices, the forces that caused the motion of the medium and the occurrence of electric currents, the model consisted of layers of rolling particles between the vortices. The forces exerted by the vortices on these particles were the cause of electromotive force, with changes in their motion corresponding to electromagnetic induction. That Maxwell thought of this as a purely fictional representation is obvious from the following quotation:

> The conception of a particle having its motion connected with that of a vortex by perfect rolling contact may appear somewhat awkward. I do not bring it forward as a connexion existing in nature, or even as that which one would willingly assent to as an electrical hypothesis. It is,

> however, a mode of connexion which is mechanically conceivable, and easily investigated. (Maxwell, 1965, vol. 1, p. 486)

One of the problems with the fluid vortex model was how to mechanically explain the transmission of rotation from the exterior to the interior parts of each cell. How could a fluid surface exert tangential forces on the particles? In order to remedy this and extend the model to electrostatics Maxwell developed an elastic solid model made up of spherical cells endowed with elasticity which made the medium capable of sustaining elastic waves. In addition, the wave theory of light was based on the notion of an elastic medium that could account for transverse vibrations; hence it was quite possible that the electromagnetic medium might possess the same property. The fluid vortices had uniform angular velocity and rotated as a rigid sphere, but elastic vortices would produce the deformations and displacements in the medium that needed to be incorporated. With this elasticized medium Maxwell now needed to explain the condition of a body with respect to the surrounding medium when it is said to be charged with electricity and to account for the forces acting between electrical bodies.

According to the Faraday picture, electric lines of force were primary and electric charge was simply a manifestation of the terminating points on the lines of force. If charge was to be identified with the accumulation of particles in some portion of the medium, then it was necessary to have some way of representing that. In other words, how is it possible to represent charge as existing in a field? In the case of a charged capacitor with dielectric material between the plates, the dielectric material itself was

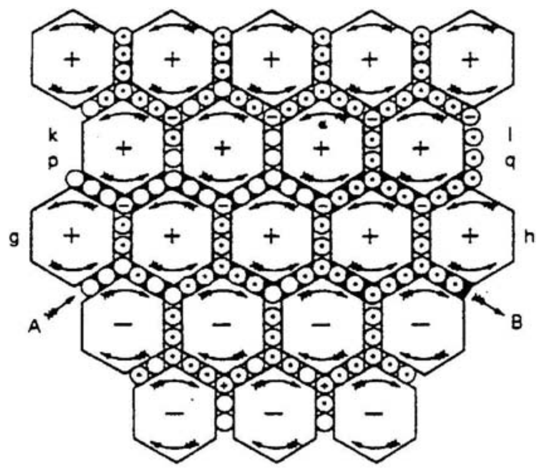

Figure 7.1 Maxwell's vortex ether model. AB is a current of electricity, with the large spaces representing the vortices and the smaller circles the idle wheels.

the primary seat of the "inductive" state and the plates served merely as bounding surfaces where the chain of polarized particles was terminated. So, what is taken to be the charge on a conductor is nothing but the apparent surface charge of the adjacent dielectric medium.

The elastic vortices that constituted the ether or medium were separated by electric particles whose action on the vortex cells resulted in a type of distortion. In other words, the effect of an electromotive force (the tangential force with which the particles are pressed by the matter of the cells) is represented as a distortion of the cells caused by a change in position of the electric particles. That, in turn, gave rise to an elastic force that set off a chain reaction. Maxwell saw the cell distortion as a displacement of electricity within each molecule, with the total effect over the entire medium producing a "general displacement of electricity in a given direction" (1965, vol. 1, p. 491). Understood literally, the notion of displacement meant that the elements of the dielectric had changed positions. And, because changes in displacement involved a motion of electricity Maxwell argued that they should be "treated as" currents in the positive or negative direction according to whether displacement was increasing or diminishing. Displacement also served as a model for dielectric polarization—electromotive force was responsible for distorting the cells and its action on the dielectric produced a state of polarization. From this we can get some sense of just how complex the model really was and how, despite its fictional status, its various intricacies provided a representation of important features of the electromagnetic field. But most important was how this account of the displacement current(s) furnished the appropriate mathematical representation that would give rise to the field equations.

In the original version of the ether model displacement was calculated only in terms of the rotation of the vortices without any distortion, but in the elastic model displacement made an additional contribution to the electric current. In order to show how the transmission of electricity was possible in a medium, a modification of Ampere's law was required in order to generalize it to the case of open circuits. What is important here, however, is not the modification of Ampere's law per se but rather the way in which the model informed that modification. If we think for a moment about what Maxwell was trying to achieve, namely, a field theoretic representation of electromagnetism, then it becomes obvious that some way of treating open circuits and representing charges and currents as emerging from the field rather than material sources is crucial.

To achieve that end two important elements were required. The first concerns the motion of idle wheels that represented electricity as governed by Ampere's law relating electric flux and magnetic intensity (curl $\mathbf{H} = 4\varpi\mathbf{J}$, where $\mathbf{H}$ is the magnetic field and $\mathbf{J}$ is the electric-current density). A consequence of that law was that it failed to provide a mechanism for the accumulation of charge because it applied only in the case of closed currents. Consequently a term $\partial\mathbf{D}/\partial t$ had to be added to the current so that it

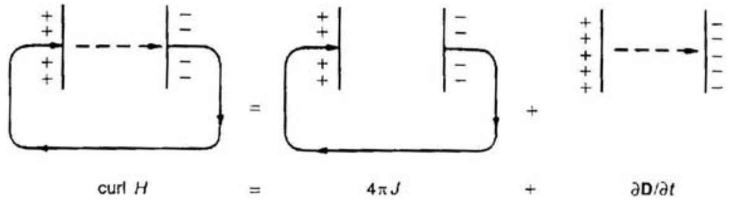

Figure 7.2 The displacement term $\partial \mathbf{D}/\partial t$ modified the original Ampere law where $\mathbf{D} = (1/c^2)E + 4\varpi P$ (the polarization vector).

was no longer circuital. As we saw earlier, and from the diagram in Figure 7.2, the dielectric between the coatings of a condenser fulfilled that need and was seen as the origin of the displacement current. The force of that current was seen as proportional to the rate of increase of the electric force in the dielectric and hence produced the same magnetic effects as a true current. Hence the charged current could be "regarded as" flowing in a closed circuit.[8] The modified term in Ampere's law took the value zero for the case of steady currents flowing in closed circuits and nonzero values for open circuits, thereby giving definite predictions for their magnetic effects.

But, as I noted earlier, there was a second feature concerning the account of displacement, namely, how the representation of electric current qua elastic *restoring* force (a crucial feature of the model) was used to represent open circuits. According to the mechanics of the model, the rotations push the small particles along giving rise to a current of magnitude $1/4\varpi$ **curl H** while the elastic distortions move the particles in the direction of the distortion, adding another contribution $\partial \mathbf{D}/\partial t$ to the current.[9] Maxwell had linked the equation describing displacement ($R = -4\varpi E^2 h$) with the ether's elasticity (modeled on Hooke's law) and also with an electrical equation representing the flow of charge produced by electromotive force. Hence R was interpreted as both an electromotive force in the direction of displacement and an elastic restoring force in the opposite direction. Similarly, the dielectric constant E is both an elastic coefficient and an electric constant, and h represents both charge per unit area and linear displacement. As a result, the equation served as a kind of bridge between the mechanical and electrical parts of the model. The electrical part included the introduction of the displacement current, the calculation of a numerical value for E (a coefficient that depended on the nature of the dielectric), and the derivation of the equations describing the quantity of displacement of an electric current per unit area. The mechanical part required that E represent an elastic force capable of altering the structure of the ether. The point that concerns us here, though, is exactly how the mechanical features of the model gave rise to electrical effects: that is, the relation between the mechanical distortion of the ether and the displacement current.

The answer lies in seeing how Maxwell's model represented the primacy of the field. As Seigel (1991, p. 99) points out, in Maxwell's equation of electric currents the **curl H** term (the magnetic field) appears on the right side with the electric current **J** on the left as the quantity to be calculated.

$$\mathbf{J} = 1/4\varpi\ (\mathbf{curl\ H} - 1/c^2\ \partial\mathbf{E}/\partial t)$$

To illustrate how this works mechanically, consider the following example: If we take a charging capacitor there is a growing electric field pointing from positive to negative in the space between the plates. This is the current owing to the solenoidal, closed loop **curl H** term. However, associated with this field there is a reverse polarization (Figure 7.3) due to the elastic deformation of the vortices acting on the particles. This gives rise to a reverse current between the plates that cancels the **curl H** term. This is because it is negative and points toward the positive plate as the capacitor is charging. The solenoidal term is incapable of producing accumulations of charge, so it is the reverse polarization that gives rise to charge on the capacitor plates and not vice versa. It is the constraint on the motion of the particles that reacts back as a constraint on the motion of the vortices that drives the elastic distortion of the vortices in the opposite direction. As a result, charge builds up through the progressive distortion of the medium. This elastic distortion is accompanied by a pattern of elastic restoring forces that correspond to the electric field *E*. Without this there would be no charge because it is responsible for relaxing the solenoidal property of the electric current.

We can now state the relation between charge and the field: Charge is the center of elastic deformation that gives rise to a pattern of electromotive forces, which constitutes the field—the energy of deformation is the electric field energy. Put simply, it is the fields that give rise to charges and currents. The magnetic field gives rise to the solenoidal in both the wire and space between the plates and the changing electric field gives rise to

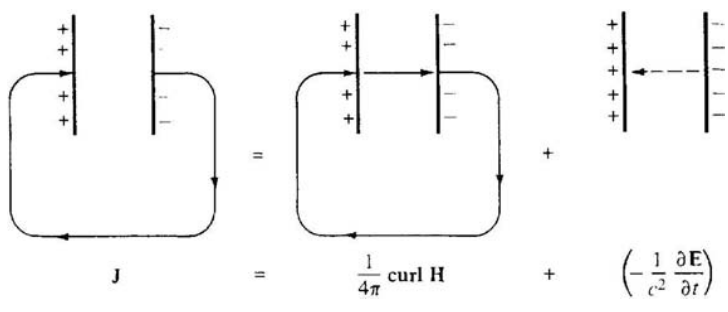

Figure 7.3 The reverse (polarization) current term has a negative sign, which cancels the **curl H** term between the plates.

a reverse current in the space between the plates. The sum of these two currents yields the conduction current **J**—an open circuit for the true current.

What this shows is how a completely fictional model provided an account of how electromagnetic forces could be produced in a mechanical ether. But, more significantly, it furnished the all important field theoretic picture that would become the basis for the modern account we have today. Once the displacement current secured the basis for a field theory, it was in some sense a rather small step to construct the *electromechanical* theory of light. I say electromechanical because this particular model didn't identify light with electromagnetic waves; rather, it was hypothesized that both originated from the same medium or field. Instead of approaching the problem as a mathematical one and obtaining solutions for the equations for the **E** or **H** fields in the form of transverse electromagnetic waves, the model occupied center stage with the mathematical account emerging from the model.

Very briefly, the rest of the 1861–1862 story goes something like this.[10] The velocity of propagation (V) of transverse torsion waves in the medium was given by the torsion modulus m, an elastic constant that controls the strength of the electric forces, divided by the mass density of the medium ρ_M which controls the strength of magnetic forces. Maxwell set the values of these parameters through a chain of linkages between the mechanical, electrical, and magnetic aspects of the model, which resulted in $V = c$ where c is the ratio of units—a measure of the relative strengths of electrostatic forces and electromagnetic forces. This ratio of electric to magnetic units depended on a quantity that had the dimensions of a velocity. There were five different methods for determining that velocity and using these experimental results Maxwell obtained a value for c that was very close to the velocity of light. Consequently the velocity of waves in the magnetoelectric medium was also roughly equivalent to the velocity of light.[11] But, and this is the important point, the rationale for setting the parameters to the values Maxwell chose, together with the equivalence between V and the ratio of units, followed directly from the mechanical and electromagnetic connections that emerged from the model. In some sense the model, like an experimental novel, had played out its own story; the succession of facts depended on the constraints imposed by the phenomena themselves via their place in the model.[12]

The problem of course was that no one, especially Maxwell, thought this model was anything but an elaborate fiction. But because the numerical relations between the optical and electromagnetic phenomena were too spectacular to ignore, he needed some way of showing that the existence of electromagnetic waves propagating with velocity c followed from electrical equations themselves, divorced from the mechanical underpinnings of the model. The answer came in 1865 via a purely dynamical theory and the introduction of electromagnetic variables into the equations of dynamics. However, with the abandonment of the mechanical model he also, and perhaps most importantly, needed a different justification for introducing displacement into his equations, for without it there was literally no field theory.

A number of assumptions made this possible. Based on facts about light he assumed that space or the ether must be capable of storing up energy, and from the forms of the equations of electricity and magnetism he showed the forms that energy must have if it is disseminated through space. If the medium is subject to dynamical principles then it will be governed by Lagrange's equations, and if we apply dynamics to electromagnetism we can substitute magnetic and electric energy for kinetic and potential energy. According to the Faraday–Mossotti theory of dielectrics, when a potential difference is applied across a dielectric it becomes polarized—molecules are positive at one end and negative at the other. This entails not only that electric energy is stored in the dielectric but that a transient electric current flows in it. If one assumed that space can store up energy it could also become polarized by a potential difference, and a changing potential difference would thereby produce a changing current (displacement) associated with a changing magnetic field. Once the displacement current was introduced Maxwell was able to deduce the properties of electromagnetic waves that accounted for light.

All of this was done in a paper entitled "A Dynamical Theory of the Electromagnetic Field." Of course, this is not to say that the new "dynamical theory" was without problems. Because there was no mechanical model it became difficult to conceive how exactly displacement operated. If electricity was being displaced, how did this occur? The problem was magnified because charge was interpreted as a discontinuity in displacement. And although displacement appears in the fundamental equations, divorced from its mechanical foundation, electromagnetism takes on the appearance of a phenomenological theory.[13] The continuation of this story and the efforts to unite mechanics and electromagnetism is both long and rather complicated.[14] Although it has interesting implications for debates about the relationship between fictional representations, mathematical representations, and concrete knowledge, those are not issues I can adequately deal with here. Instead I want to conclude this section by focusing on the problem I raised at the beginning, namely, what specific features can we isolate as playing a role in the transmission of information from fictional models?

What is the sense of representation that is important here and how does it emerge from the fictional model? What is especially significant is that the development of the field equations did not proceed through the introduction of a new term called the displacement current as a mathematical modification to Ampere's law. Instead, what I tried to show is how the foundation for electromagnetism emerged from the molecular vortex model and was in fact determined by it. But the important issue here is not that Maxwell was capable of deriving a set of field equations from a false model, but rather what it was about the model that underscored the applicability of the equations. Put differently, what was the conceptual or theoretical content represented in the model that formed the basis for his electromagnetic worldview?

The answer centers, of course, on the displacement current. Maxwell knew that the ether or field did not consist of rotating vortices and idle wheels, but he also knew that in order to represent the Faraday picture of electromagnetism there had to be some account of how electricity could travel in free space and charge could be build up without material bodies. Consequently, he needed a way of representing this mechanically if he was to derive a set of equations applicable to this new field theoretic picture. The mechanical model became the focal point for an understanding of how charges and currents could be understood as field theoretic phenomena and the formulation of a mathematical account of those processes. Part of that account involved a modification of the relation between between Ampere's law, Coulomb's law, and the equation of continuity for open circuits, something that was indicated by the structure of the model. Once the basic mechanical features of the model were in place they constrained both the physical and mathematical descriptions of electromagnetic forces. In the same way that character development in a novel determines, to some extent, how the story will play out, the features of the model restrict the way that certain physical relations can be represented. By the time the *Treatise on Electricity and Magnetism* was completed in 1873, displacement had taken on the role of a primary current responsible for the transmission of electricity through space. Although the mature theory was an extension of the dynamical approach, neither would have been possible without the fictional model that provided a physical conception of how an electromagnetic field theory might be possible.

What this extended discussion hopefully shows is not only how certain information emerged from the fictional model but also the need for examining the model in detail in order to show exactly how this happens. Although I have only addressed one example here, the point I want to make is a general one concerning fictional models. To say that fictional models are important sources of knowledge in virtue of a particular kind of similarity that they bear to concrete cases or systems is to say virtually nothing about how they do that. Instead, what is required is a careful analysis of the model itself to uncover the kind of information it yields and the ways in which that information can be used to develop physical hypotheses. There are various ways that fictional models may be able to accomplish this, but each one will do so in a way that is specific to that particular model. The situation is radically different from the case of idealization where an analysis of the methods employed for both the idealizing process and its corrections will typically cover a variety of different cases. Consequently, idealization becomes a relatively easy category to define, but it still presents some challenges for uncovering how idealized models relate to real systems.

4. FROM THE IDEAL TO THE CONCRETE

So, how exactly does the fictional case differ from an idealized law/model where there are specific unrealistic conditions required for the law to hold, conditions that are false but nevertheless capable of approximating real physical systems? The Hardy–Weinberg law is a simple example of a fundamental law that makes assumptions not realized in any natural population. In what sense do we want to say that this law provides us with information about these populations? One of the things that renders laws explanatory, as highlighted by the D-N model, is the fact that they are general enough to apply to a diverse number of phenomena. In other words, they enable us to understand specific features of phenomena as similar in certain respects; for example, universal gravitation shows that both terrestrial and celestial bodies obey an inverse square force law. Nancy Cartwright (1983) claims that this generality is a reason for thinking fundamental laws like these are false; their generality results in their being unable to fully describe the situations they reportedly cover, or they deliberately omit aspects of the situation that are not relevant for the calculation at hand. In that sense they don't accurately describe concrete situations and are true only of objects in our models. The problem then is how could they possibly provide knowledge of concrete physical systems, or in this case, populations?

Part of Cartwright's reason for claiming that covering laws are false is to contrast them with phenomenological laws (or models) that supposedly do give us more accurate descriptions of the physical world. However, what the Hardy–Weinberg law shows is that embedded in what Cartwright would call a 'false' law is a great deal of accurate information about biological populations, information that was crucial in the synthesis of Mendelian heredity and Darwinian natural selection. To that extent it serves as an example of how mathematical idealization (and abstraction) can enhance our understanding far beyond simple predictive capabilities.[15] As we shall see later in this chapter, the Hardy–Weinberg law enables us to understand fundamental features of heredity and variation by establishing a mathematical relation between allele and genotype frequencies that embodies the very gene conserving structure that is the essential feature of Mendelism. What is important for my purposes here is to show why the unrealistic nature of its assumptions does not affect the significance of either the conclusions it provides or the information implicit in its formulation. Moreover, it is a nice example of the differences I want to highlight between idealization and the account of abstraction I described in the introduction.

The Hardy–Weinberg law is often described as a consequence of Mendel's law of segregation, or a generalization of Mendel's laws as applied to populations. It relates allele or gene frequencies to genotype frequencies and states that in an infinite, random mating population in the absence of external factors such as mutation, selection, drift, and migration, one generation of random mating will produce a distribution of genotypes that is a function solely of allele frequencies. Moreover, this distribution does not change over subsequent generations, provided all conditions are held constant. In other words, if we have a pair of alleles Aa at a particular gene locus and the initial ratio of A to a is p to q, then for every succeeding generation the ratio will be p to q. Regardless of the distribution of genotypes in the initial generation the distribution for all succeeding generations will be

$$p^2A_1A_1 + 2pqA_1A_2 + q^2A_2A_2$$

where p^2 is just the probability of getting an A_1A_1 homozygote, which is the probability that the egg is A_1 times the probability that the sperm is A_1 (by the product rule for independent events). Both of these probabilities are p because in its simplest form the law assumes that the species is hermaphroditic. Because the heterozygote can be formed in two different ways the probability is $2pq$ (by the addition rule for mutually exclusive events). So, if you know the value for p then you know the frequencies of all three genotypes.

Because random mating does not change allele frequencies, all one needs to calculate the genotype frequencies after a round of random mating is the allele frequencies before random mating. In populations where each individual is either male or female with different allele frequencies it will take two generations to reach Hardy–Weinberg equilibrium. One can see then the relation between the stability of the frequencies and Mendel's law of segregation. With random cross-fertilization there is no disappearance of any class whatever in the offspring of the hybrids, and each class continues to be produced in the same proportion.[16]

But, and here is the important point, what is significant about the Hardy–Weinberg law is not so much the binomial form of the genotype frequency and the prediction of genotypes based on the stability of the population, but rather what the stability actually shows or presupposes. Despite the idealizing assumptions, the stability allows us to understand something about Mendelian populations that is significant for understanding heredity and variation. In other words, certain conditions must be present for the stability to be possible. Thus, the predictive success of the law is intimately connected with certain basic claims about genetic structure that are presupposed in its formulation. What the Hardy–Weinberg law says is that if no external forces act, then there is no intrinsic tendency for the variation caused by the three different genotypes that exist in a population to disappear. It also shows that because the distribution of genotype frequencies is independent of dominance, dominance alone cannot change genotype frequencies. In other words, there is no evidence that a dominant character will show a tendency to spread or a recessive one to die out. Instead, the genotype frequencies are maintained in constant proportions. The probabilistic genetic structure is conserved indefinitely; but should it be influenced by an outside force, such as mutation, the effect would be preserved in a new stable distribution in the succeeding generation.

This was crucial for understanding the problems with blending inheritance as advocated by the Darwinians, and to that extent the claim that the law is false in some sense misses the point if our concern is conveying information. Under blending inheritance, variation was thought to decrease rapidly with each successive generation, but Hardy–Weinberg shows that under a Mendelian scheme it is maintained. This pointed to yet another fundamental aspect of Mendelism, namely, the discontinuous nature of the gene, and why it was crucial for the preservation of variation required for selection. How was it possible for the genetic structure to be maintained over successive generations? The reason for the stability could be traced directly to the absence of fusion, which was indicative of a type of genetic structure that could conserve modification. This condition was explicitly presupposed in the way the law was formulated and how it functioned.[17] In that sense one can see the Hardy–Weinberg (H-W) law as the beginning of a completely new explanation of the role of mutation and selection and how they affect our understanding of evolution.

What I have focused on thus far has been the information about the nature of heredity that is embedded or presupposed in the structure of this law. But what about the so-called "falsity" of the assumptions under which it is alleged to hold? Although the "model population" specified by the H-W law bears little, if any, relation to actual human populations, we saw that the law was an important source of theoretical information. What happens when assumptions like random mating and infinite populations are replaced with more realistic assumptions true of actual populations, assumptions like assortative mating and small populations? Although there are many forms of nonrandom mating, what is crucial from the point of view of "unrealistic assumptions" is that in many cases it is possible to show that given a parental and daughter generation allele frequencies remain the same in each generation; hence genetic variation is maintained—the fundamental conclusion of the H-W law. However, because heterozygote frequency is less than that applying in random mating populations, the variation is in some sense cryptic; that is, you get the same allele frequencies but different genotype frequencies. But after one generation of random mating, the H-W genotype frequencies would be immediately restored. Similarly, in the case of infinite populations, once we relax this assumption we find that mean heterozygosity decreases very slowly with time as a result of the sampling drift implicit in the process. This slow loss can be understood as the stochastic analogue of the "variation-preserving" property of infinite populations described by H-W. Although a violation of these conditions destroys H-W equilibrium, we nevertheless learn some useful information about the population.

Appealing to the abstraction/idealization distinction I introduced at the beginning can further clarify our understanding of how deviations from the conditions or assumptions specified by the law affect its applicability. Essentially we can divide the assumptions associated with the Hardy–Weinberg law into two groups. The first involves assumptions that don't allow for relaxation without violating H-W equilibrium, such as infinite population size and random mating. The second includes the absence of selection, migration, and mutation. These assumptions affect allele frequencies but not random mating. For example, selection may be taking place in a population that is nevertheless breeding randomly. Violations of these latter conditions will not rule out H-W proportions; instead, the allele frequencies will change in accordance with the changing conditions. In other words, these conditions function as idealizations that may or may not hold but whose effect on the system can be straightforwardly calculated. Put differently, we can think of them as external factors that isolate basic features of a Mendelian system that allow us to understand how variation could be conserved.

Contrast that situation with the requirements of infinite populations and random mating. Infinite populations are crucial in that one must be able to rule out genetic drift, which is a change in gene frequencies that results from chance deviation from expected genotypic frequencies. That is, we must be able to determine that detected changes are not due to sampling errors. Although random mating seems like the kind of restriction that is typically violated, we can see how its violations affect gene frequencies: In the case of assortative mating there will be an increase in homozygosity for those genes involved in the trait that is preferential such as height or eye color. Traits such as blood type are typically randomly mated. Similarly, in the case of inbreeding there will be an increase in homozygosity for all genes. Because both of these assumptions are necessary for H-W equilibrium, they cannot, in general, be corrected for and in that sense are necessary features for the applicability of the law. In other words, they ought to be considered abstractions rather than idealizations because they describe situations that cannot approximate real-world situations through the addition of correction factors.[18]

My account of abstraction is somewhat different from that typically described in the literature. Although abstraction and idealization are sometimes conflated, Cartwright (1989) has distinguished them in the following way: Idealization is a process where one starts with a concrete object and then mentally rearranges some of its (inconvenient) features or properties. This enables us to write down a law describing its behavior in certain circumstances. In some cases it is possible to just omit factors that are irrelevant to the problem, but for the factors that are relevant they are sometimes given values that are not, strictly speaking, accurate but allow for ease of calculation. The idealizations presupposed by the ideal gas law are an example (infinitesimal size of molecules and absence of intermolecular forces). In these cases we sometimes know the degree to which the idealization is a departure from the real situation and if necessary its effect can be estimated.

Abstraction presents a different scenario; here what Cartwright calls the "relevant features" have been genuinely subtracted. The example she uses to illustrate the point is the comparison between a law and a model. Many of the effects studied in modern physics make a very small contribution to the total behavior of a system and hence the laws governing these systems do not really approximate, in any real sense, what happens in concrete cases. These laws are instances of abstractions. The model, on the other hand, takes the relevant factors and assigns them convenient values in order to facilitate calculation. Although the latter may be unrealistic in the sense that it gives an idealized representation of particular properties, it still makes contact with the world insofar as it includes properties that are relevant to the system's behavior. Abstract laws do not literally describe the behavior of real systems because (1) they subtract features in order to focus on a single set of properties or laws as if they were considered in isolation and (2) no amount of theory will ever allow us to complete the process of concretization. Laws that govern the laser abstract from its material manifestations to provide a general description that is "common to all, though not literally true of any" (1989, p. 211). By contrast, the assumption of infinitesimal size for molecules can be corrected to make the ideal case more realistic.

Although I think Cartwright is essentially right in her claim that we can never "concretize" abstractions, the question that interests me is why that is the case. In her discussion of abstraction Cartwright mentions Duhem, who is also concerned with the notion of abstraction in physics. Because physics needs to be precise, it can never fully capture the complexity that is characteristic of nature. Hence, the abstract mathematical representations used by physics do not describe reality but are better thought of as imaginary constructions—neither true nor false. Laws relate these symbols to each other and consequently are themselves symbolic. Although these symbolic laws will never touch reality, so to speak, they do involve approximations that are constantly undergoing modification due to increasing experimental knowledge (Duhem, 1977, p. 174). In that sense, Duhem's account of abstraction seems also to incorporate elements of what Cartwright would call idealization. It also seems clear from his account that the gap between symbolic laws and reality is due, essentially, to an epistemic problem that besets us in the quest for scientific knowledge. Laws are constantly being revised and rejected; consequently, we can never claim that they are true or false. In addition, because of the precise nature of physics we must represent reality in a simple and incomplete way in order to facilitate calculation. This describes both idealization and abstraction, depending on how one chooses to simplify.

What Duhem's view captures is a philosophical problem that focuses on the gap between reality and our representation of it. As we saw earlier, the account of abstraction I am concerned with has its basis in this gap as well, but is motivated by particular kinds of abstract representations that are required for dealing with certain kinds of systems. Put differently, both Duhem's and Cartwright's accounts of abstraction describe physics more generally and the problems that beset a realist interpretation of theories, laws, and models. Although those issues have some bearing on my point, my overall concern stems from the way that we are *constrained* to represent certain kinds of systems in a mathematically abstract way if we are to understand how they behave. Because this goes beyond problems of calculation to issues about explanation, 'abstraction as subtraction' is not a useful category for my purposes. In my account one of the things that makes abstraction especially interesting is that it is the mathematical representation that provides the foundation for understanding causal features of the system. The physics and mathematics are inextricably intertwined, making our very characterization of these systems mathematical abstractions.

5. BEYOND FICTIONS: THE NECESSITY OF ABSTRACTIONS

In the preceding discussion we saw how infinite populations were an important constraint in the operation of the Hardy–Weinberg law. They are necessary for eliminating the chance or random influences on gene frequencies from one generation to the next, something that is common in small populations. Deviations from infinite population size can, of course, be handled because the kinds of populations to which one applies these models are never infinite. But in moving to smaller populations one must recognize that as population size decreases the effects of drift will become more predominant, thereby making it difficult to determine whether particular features of the population are the result of drift or selection. In other words, we cannot determine whether the population is undergoing evolutionary change. In that sense infinite population size (along with random mating) is a necessary condition for H-W equilibrium to be maintained and for determining how deviations from it (selection, mutation, etc.) are to be understood.

Other cases where abstract representations are crucial for understanding how the system in question behaves include phase transitions. There are a variety of physical systems that fall into this category: superconductivity, superfluidity, magnetism, crystallization, and several others. In each of these cases there is a spontaneous symmetry breaking associated with a phase transition that explains the occurrence of the superconducting or magnetic state of matter. Very briefly, the situation is as follows: In thermodynamics (TD), phase transitions are accounted for in terms of discontinuities in the thermodynamic potentials. However, once we move to statistical mechanics (SM) the equations of motion that govern these systems are analytic and hence do not exhibit singularities. As a result, there is no basis for explaining phase transitions in SM. In order to recover the TD explanation we need to introduced singularities into the equations and thus far the only way to do this is by assuming the number of particles in the system is infinite. That is, we need to invoke the "thermodynamic limit," $N \rightarrow \infty$, which assumes that the system is infinite in order to understand the behavior of a real, finite system. Note that the problem here isn't that the limit provides an easier route to the calculational features associated with understanding phase transitions; rather, the assumption that the system is infinite is *necessary* for the phase transitions to occur. In other words, we have a description of a physically unrealizable situation that is *required* to explain a physically realizable one (the occurrence of phase transitions).

Given this situation, how should we understand the relation between the abstract description and the concrete phenomena it supposedly explains? Although one might want to claim that within the *mathematical* framework of SM we can causally account for (explain) the occurrence of phase transitions by assuming the system is infinite, it is nevertheless tempting to conclude that this explanation does not help us to *physically* understand how the process takes place because the systems that SM deals with are all finite. Similar doubts have been expressed by Callender (2001) and Earman (2003), who argues against taking the thermodynamic limit as a legitimate form of idealization: "a sound principle of interpretation would seem to be that no effect can be counted as a genuine physical effect if it disappears when the idealizations are removed" (p. 21). Both claim that, we shouldn't assume phase transitions have been explained (or understood) if their occurrence relies solely on the presence of an idealization. Initially this seems an intuitive and plausible objection, but if we reflect for a moment on the way that mathematical abstraction is employed it becomes clear that this line of reasoning quickly rules out explanations of the sort we deem acceptable in other contexts. Here the distinction I mentioned at the beginning between idealization and abstraction becomes especially important. Specifically, we need to distinguish between the kind of abstraction that is, in some sense, dictated by our models and the more straightforward kinds of mathematical idealizations that are used simply to facilitate calculation. In the former case the abstraction becomes a fundamental part of how the system is modeled or represented and consequently proves crucial to our understanding of how the system/phenomena behave.

For example, consider the intertheoretical relations that exist in fluid mechanics between Navier–Stokes equations and the Euler equations or between theories like wave optics and ray optics, and classical and quantum mechanics. Because of the mathematical nature of physical theories the relations between them will typically be expressed in terms of the relations between different equations/solutions. In each case we are interested in certain kinds of limiting behavior expressed by a dimensionless parameter δ. In fluid dynamics δ is equal to 1/Re (Reynolds number), and in quantum mechanics it is Planck's constant divided by a typical classical action ($\hbar/S$). But, in fluid mechanics (as in the other cases listed earlier) the limit $\delta \rightarrow 0$ is singular and it is this singularity that is responsible for turbulent flows.

Similarly, in the ray limit where geometrical optics accurately describes the workings of telescopes and cameras the wavelength $\lambda \to 0$. Because ψ is nonanalytic at $\lambda = 0$, the wave function oscillates infinitely rapidly and takes all values between +1 and–1 infinitely often in any finite range of x or t. A good deal of asymptotic behavior that is crucial for describing physical phenomena relies on exactly these kinds of mathematical abstractions. What we classify as "emergent" phenomena in physics, such as the crystalline state, superfluidity, and ferrogmanetism, to name a few, are the results of singularities and their understanding depends on just the kinds of mathematical abstractions already described.

How then should we think about these kinds of mathematical abstractions and their relation to physical phenomena? My point is that the abstract (mathematical) representations supplied by our models are what forms our understanding of these systems. In the case of phase transitions there are formal accounts or definitions that appeal to zeros in the partition function, changes in symmetry and orderliness, and the existence of fixed points.[19] In each of these cases a sharp phase transition is possible—the transition temperature is well defined and the appearance of new orderliness is abrupt. Similarly, these formal features function as indicators of the kind of phenomena we identify with phase transitions—a sharp change in specific volume or density, a change in symmetry or the scaling of properties as measured by critical exponents and correlation functions. In other words, the mathematics provides not only a representation and precise meaning for phase transitions, but it also enables us to associate that representation with dynamical behavior such as symmetry breaking and the appearance of order. The abstract model illustrates the essential features of the phenomenon in question. In that sense the mathematics and the physics are crucially intertwined.

This also has implications for experimental practice. In cases where $N < \infty$ a phase transition is recognized by a finite change in a property like density or magnetization for an infinite change in another property like temperature or the magnetic field (as in the case of permanent magnetization). Yet, in a theory that has only finite volume or finite N we can't be sure that we are identifying a phase transition because the formal continuity of the pressure-volume curve is guaranteed by the analyticity in the activity for finite N. In these cases any discontinuity or unsmoothness is rounded off or smeared. When we do an experiment we are looking for jumps or discontinuities in the data that cannot be smoothed over. Although these jumps and curves are in the phenomena themselves, it is the job of our models (like the Ising model) to tell us the exact form they will take, for example, logarithmic singularity.[20] This is why the appeal to infinitely large systems is crucial; only there will the appropriate kinks and jumps emerge! Similarly in population genetics we need the assumption of infinite populations to determine whether changes in gene frequencies are the result of selection.

If one accepts my characterization of abstraction, then subscribing to Callender's and Earman's account would be tantamount to ignoring large portions of both the mathematical and physical foundations of our theories. On their idealization view we need an answer to why cases like the thermodynamic limit is considered illegitimate when sending distances and time intervals to zero in the use of differential equations is not. Here my distinction between idealization and abstraction offers some help. We typically think of an idealization as resembling, in certain respects, the concrete phenomenon we are trying to understand. We usually know how to correct or compensate for what we have left out of, or idealized in, the description. That is, we know how to add back things like frictional forces that may not have been needed for the problem at hand. Or, we know how to change the laws that govern a physical situation when we introduce more realistic assumptions, as in the move from the ideal gas law to the van der Waals law. These cases are best thought of as idealizations that represent a physical situation in a specific way for a specific purpose.

Constrast this with the kinds of mathematical abstractions described earlier. In the case of the thermodynamic limit, we don't introduce abstractions simply as a way of ignoring what is irrelevant to the problem or as a method for calculational expediency. Instead, the mathematical representation functions as a necessary condition for explaining and hence understanding the phenomena in question.[21] Thinking of abstraction in this way sheds a light on the problem Callender and Earman mention because the problem is very different from the more straightforward cases of idealization. The importance of mathematical abstraction in these contexts requires us to think differently about what constitutes explanation and understanding, but that challenge seems unavoidable. If physical phenomena are and in some cases *must* be described in terms of mathematical abstractions, then it seems reasonable to expect that their explanations be given in similar terms.

We can see how this situation differs significantly from the way fictional models are used to convey information. Although fictional models may constrain the physical possibilities once the model structure is in place, there is typically a choice about how to construct the model and how to represent the system/phenomenon. There is less freedom of movement with idealizations in that specific kinds of approximation techniques inform and determine the way the system is represented. That said, we often have a choice about which parameters we can leave out or idealize, given the problem at hand. In cases of mathematical abstraction, however, we are completely constrained as to how the system is represented, and the abstraction is a necessary feature of our theoretical account.

6. CONCLUSIONS

One of my goals in this chapter was to introduce distinctions among the processes involved in constructing a fictional model, an idealization, and mathematical abstraction. Each of these categories transmits knowledge despite the presence of highly unrealistic assumptions about concrete systems. Indeed, examples of the sort I want to classify as abstractions require these types of assumptions if we are to understand how certain kinds of phenomena behave. Although there is a temptation to categorize any type of unrealistic representation as a "fiction," I have argued that this would be a mistake, primarily because this way of categorizing the use of unrealistic representations tells us very little about the role those representations play in producing knowledge. That said, fictional models can function in a significant way in various stages of theory development. However, in order to uncover the way these models produce information we need to pay particular attention to the specific structure given by the model itself. There is no general method that captures how fictional models function or transmit knowledge in scientific contexts; each will do so in a different way, depending on the nature of the problem. By separating these models from other cases of abstraction and idealization we can recognize what is distinctive about each and in doing so understand how and why unrealistic representations can nevertheless provide concrete information about the physical world.[22]

NOTES

1. For an extended discussion see Morrison (2002).
2. For more discussion of nuclear models see Morrison (1998) and Portides (2000).
3. In other words, is the presence of idealized assumptions different from assumptions that are false in the sense of being deliberately false, like assuming that the ether is made up of rotating vortices of fluid or elastic solid particles? No amount of correction changes the form of these latter assumptions.
4. Although we typically say that models provide us with representations, I think we can also extend that idea to laws. By specifying the conditions under which the law holds (even if they are just straightforward ceteris paribus conditions) we have specified a scenario or a context that defines the boundaries for the operation of the law. This specification can be understood as a representation of the context under which the law can be assumed to hold.
5. Cartwright (1989), in particular, characterizes abstraction as the "taking away" of properties that are part of the system under investigation. For example, when modeling a superconductor one abstracts the type of material the superconductor is made of.
6. As we shall see later, it isn't really the presence of idealizing conditions that is responsible for transmitting information; rather, it is the what is presupposed in the binomial formula for predicting genotype frequencies.
7. All references to Maxwell's papers are contained in the 1965 edition of collected papers.
8. This is what Seigel (1991) refers to as the "standard account." He also remarks (p. 92) that according to the standard account the motivation for introducing the displacement current was to extend Ampere's law to open circuits in a manner consistent with Coulomb's law and the continuity equation. Although

I agree with Seigel that this was certainly not Maxwell's motivation, I think one can incorporate aspects of the standard account into a story that coheres with the historical evidence. My intention here is not to weigh in on this debate but simply to show how various aspects of the field equations emerged from the model.

9. It is perhaps interesting to note here that in the modern version the electric current **J** appears on the right because it is regarded as the "source" of the magnetic field **H**. This is because charges and currents are typically seen as the sources or causes of fields. This, however, was introduced by Lorentz, who combined the Maxwellian field theory approach with the continental charge-interaction tradition, resulting in a kind of dualistic theory where charges and currents as well as electric and magnetic fields are all fundamental, with the former being the source of the latter.
10. See Seigel (1991, pp. 130–135).
11. In addition to the agreement between the velocity of propagation of electomagnetic waves and light there were also two other connections between electromagnetic and optical phenomena that emerged from the model; one was the Faraday effect and the other involved refractive indices of dielectric media.
12. In the end, however, there was no real explanation of the way that elastic forces produced electric lines of force. Instead, he was able to calculate the resultant electrical force without any precise specification of how it arose (i.e., there was no calculation of a stress tensor of the medium, as in the magnetic case, from which he could then derive the forces).
13. For an extended discussion of the differences in the two accounts of displacement see pages 146–147 and 150–151 in Seigel (1991).
14. See Morrison (2000) for a detailed account of the development of the electromagnetic theory and the unification of electromagnetism and optics.
15. Why this is an instance of both abstraction and idealization will be discussed later.
16. The law of segregation refers to the fact that the characters that differentiate hybrid forms can be analysed in terms of independent pairs; that is, each analagen acts separately—they do not fuse. We can also understand this as stating that any hybrid for a given character produces an offspring distributed according to definite proportions. If the pure parental forms are A and a and the hybrid Aa, then the offspring of the hybrid will be distributed according to the ration 1A:2Aa:1a. Pearson (1904) was probably the first to show the relation between the law of segregation and the stability of a population in the absence of selection.
17. The notion of presupposed that I have in mind here is the same as the one connected to the ideal gas law—in order for the law to hold one must presuppose that the molecules of the gas are infinitesimal in size and have no forces acting between them.
18. In addition to the presence of unrealistic assumptions about the kinds of populations in which the law holds, the stability embedded in the structural form of the law is also crucial in explaining how/why it works. The variation preserving feature associated with the stability requires that the gene have a discrete, atomistic character. The importance of the "unrealistic" constraints is to highlight the impact of different factors and the degree to which they affect the evolutionary processs. Because the law represents a stable stationary state for a sexually reproducing population, it enables us to judge the effects of selection, preferential mating, etc. on the homogeneous field of allele and genotype frequencies. The *fundamental assumption* required for this picture to work is that of a structure that preserves modifications.

19. The most important information in the renormalization group flow is given by its fixed points, which give the possible macroscopic states of the system at a large scale.
20. Similarly, the uniformity of convergence defines what we mean phenomenologically by a phase of matter.
21. For an excellent discussion of issues surrounding the notion of reduction and the thermodynamic limit see Batterman (2005).
22. Support of research by the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.