

Rigorous Results, Cross-Model Justification, and the Transfer of Empirical Warrant: The Case of Many-Body Models in Physics

Axel Gelfert
Department of Philosophy
National University of Singapore
3 Arts Link, 117570 Singapore
email: axel@gelfert.net

June 4, 2007

Abstract

This paper argues that a successful philosophical analysis of models and simulations must accommodate an account of mathematically rigorous results. Such rigorous results may be thought of as genuinely model-specific contributions, which can neither be deduced from fundamental theory nor inferred from empirical data. Rigorous results provide new indirect ways of assessing the success of models and simulations and are crucial to understanding the connections between different models. This is most obvious in cases where rigorous results map different models on to one another. Not only does this put constraints on the extent to which performance in specific empirical contexts may be regarded as the main touchstone of success in scientific modelling, it also allows for the transfer of warrant across different models. Mathematically rigorous results can thus come to be seen as not only strengthening the cohesion between scientific strategies of modelling and simulation, but also as offering new ways of indirect confirmation.

1 Introduction

This paper argues that the philosophical analysis of models and simulations can be advanced by reconsidering the role of mathematically rigorous results in their evaluation and confirmation. It focuses on a class of models intended to describe the physical behaviour of systems that consist of a large number of interacting particles. Such *many-body models* are typically employed in order to account for a range of complex phenomena such as magnetism, superconductivity, and other phase transitions. Because of the dual role of many-body models as models of *physical systems* (with specific empirical phenomena as their explananda) and as *mathematical structures*, they form a sub-class of scientific models, from which one can arrive at general conclusions about the role of mathematical relationships in constructing and assessing models. Since many-body models lend themselves to computational evaluation via a range of techniques (e.g., Monte Carlo simulations, Green's function techniques, etc.), they are of special significance when it comes to analysing the relation between mathematical rigour and computer simulations.

The structure of the paper is as follows: First, an attempt is made to clarify the relation between models and simulations by drawing on recent work in the philosophy of models. While models have rightly earned their place in philosophical analyses of science, philosophical work on simulations is still in its relative infancy. It is argued (Section 2) that a proper appreciation of simulations in science requires a shift in focus from questions of representation to analyses of how science generates results. The question of how models can be used to generate specific results has sometimes been overlooked; in the case of mathematical models, it has often been regarded as unproblematic. This calls for an analysis of the role of mathematics in models and simulations more generally (Section 3). Section 4 argues that, beyond qualitative results and numerical outcomes, there exists the important class of mathematically rigorous results and relations, which often play a vital role as 'benchmarks' even when they lack an empirical interpretation. At a descriptive level, it may be noted that such 'benchmarking' often makes up the better part of research in computational physics (Section 5). However, as argued in Section 6, rigorous results and relations also allow for specific ways of indirect confirmation and transfer of empirical warrant. This point is elaborated on, both by engaging with the contemporary philosophical debate about indirect confirmation and epistemic coherence, and by drawing on an example from the recent history of physics, known as the Mott transition.

2 Simulating Models: From Representation to Results

While scientific models have enjoyed a long history as objects of philosophical inquiry, simulations have only recently begun to garner serious interest from philosophers of science. (See, for example, refs. (38), (20) and works quoted therein.) The relationship between models and simulations is by no means self-

evident; within science, too, there is considerable vacillation as to the order of dependency between models and simulations. Scientists frequently speak of ‘simulating a model’ (in the sense of running a computer simulation of the dynamic behaviour of a model given a certain input, e.g. consisting in parameter values and boundary conditions), but they also profess to engage in ‘modelling a simulation’: that is, using numerical techniques to construct a *computational* model whose only *raison d’être* is its ability to generate sets of simulated data. Whether a ‘model’ is being ‘simulated’, or a ‘simulation’ is being ‘modelled’, is – at least in actual scientific usage – often not so much a question of logical order as of temporal order: new computational techniques may be applied to well-established theoretical models, and vice versa. To complicate matters further, there exists some ambiguity in the term ‘simulation’, which may be understood as referring either to the process of applying a set of (usually numerical) techniques, or to the output generated by such a procedure. Given that actual scientific usage offers no clear verdict on the matter, a first task of any prospective philosophical analysis of simulation will have to be the clarification of the logical order and the conceptual relationship between models and simulations; this, or at least one step in this direction, is the task of the present section.

Recent philosophical interest in the use of models in science is in large part due to an approach, pioneered by Margaret Morrison and Mary Morgan, which views models as ‘mediating instruments’ (25, p. 10). As I shall argue later, the ‘models as mediators’ view, while by no means the only well-developed view of scientific models, has special affinities with the topic of simulation. Therefore, in the present context, it will serve to set the agenda for the discussion of scientific models. According to the mediator approach, models are to be regarded as more than mere unavoidable intermediary steps in applying our best scientific theories to specific situations. Rather, as ‘mediators’ between our theories and the world, models inform the interpretation of our theories just as much as they allow for the application of these theories to empirical phenomena. Models, it is claimed, ‘are *not* situated in the middle of an hierarchical structure between theory and the world’, but operate outside the hierarchical ‘theory-world axis’. (25, p. 17f.) Traditionally, unless their role was seen as merely heuristic, models were to be judged by how well they fit with the fundamental theory and the empirical data, or, more specifically, how well they explain the data *by the standards of the fundamental theory*: ideally, a model should display a tight fit both with theory and with empirical data. Indeed, on certain accounts of the formal relation between theory and data, any application of a theory to empirical phenomena – that is, any use of a theory that goes beyond the mere deduction of further theoretical statements – must necessarily happen via models embedded in the semantic structure of a theory. Without taking a stance on this issue, it seems obvious that such a view places a rather heavy theoretical load on the concept of ‘model’ – more, perhaps, than the notion of a scientific model (as employed in scientific strategies of modelling and simulation) can bear.

The ‘models as mediators’ approach, by contrast, insists that any scientific account of specific processes and phenomena necessarily depends on consider-

ations that are extraneous to fundamental theory. In this respect, Morrison and Morgan argue, the role of scientific models is similar to that of tools and scientific instruments; indeed, it is a crucial assumption of the mediator view that model building involves an element of creativity and skill – it is ‘not only a craft but also an art, and thus not susceptible to rules’ (25, p. 12). By focusing more on the process of model construction than on the logical relations into which models can enter with other abstract structures, such as theories, one might argue that the mediator view already displays a natural affinity towards questions of simulation. One of the mediator view’s fundamental tenets is the thesis that models ‘are made up from a *mixture* of elements, including those from outside the domain of investigation’ (25, p. 23). Of these varied elements, some will typically derive from theory, whereas others may originate from extra-theoretical considerations: ‘model construction involves a complex activity of integration’ (24, p. 44). More often than not, this integration is neither perfect nor complete. When certain elements of a model are incompatible, the integration cannot be perfect. This, for example, is the case in the Bohr model’s conflicting demands that the electrons in an atom should be conceived of as orbiting the nucleus on circular paths without losing energy, while at the same time viewing them as objects of classical electrodynamics. Integration may also remain incomplete for the simple reason that not all features of a system are eventually reflected in the model. As Daniela Bailer-Jones argues, ‘selection of aspects for the purposes of modelling is an accepted and well-practised creative strategy’ (1, p. 66). Which aspects are deemed relevant may depend on a range of criteria, including such factors as computational accessibility or explanatory interest, which themselves are determined less by theoretical first principles than by contingent facts of scientific practice. This conception of scientific models has sometimes been characterised as being fuelled by anti-theoretical sentiments, as in Nancy Cartwright’s emblematic pronouncement that ‘theories in physics do not generally represent what happens in the world – only models represent in this way’ (7, p. 180); this view is seconded by Morrison who writes that ‘the proof or legitimacy of the representation arises as a result of the model’s performance in experimental, engineering and other kinds of interventionist contexts’ – not by reference to theory. (23, p. 81) It is noteworthy, though, that neither Cartwright nor Morrison call into question the overall epistemic goal of all modelling and theorising: namely, to *represent*. Indeed, their claim is not that representation itself is unattainable, but rather that it can only be attained by means of models, not theories.

However, there remains a gap between, on the one hand, the aspirations of the mediator view to solve the problem of scientific representation and, on the other hand, the way it assesses the success of scientific models. Merely asserting that models are instruments for intervening in the world, and that their representational success is to be assessed by their performance in ‘interventionist contexts’ leaves open how we derive knowledge from their application. As Bailer-Jones puts it, ‘[i]f one chooses to interpret “representation” in the way Morrison does, then there still remains a gap between good performance “in experimental, engineering and other kinds of interventionist contexts” and “giving

useful information”.’ (1, p. 67) The gap is not closed by simply insisting that models are ‘inherently intended for specific phenomena’ (33, p. 75), and that models are superior to theories because ‘they provide the kinds of details about specific mechanisms that allow us to intervene in the world’ (23, p. 83).

If one holds that models derive their justification exclusively from instrumental success in specific empirical phenomena, then what is needed is a measure for empirical success, which typically will hinge on comparison of the model’s predictions with empirical measurements, both at the quantitative level of numerical results and at the qualitative level of system behaviour. An uninterpreted model, however, does not in and of itself, without numerical evaluation, deliver quantitative or qualitative predictions about specific empirical phenomena. This is why one might say that, at least across much of the so-called ‘hard’ sciences, ‘the proper object of epistemic evaluation is a model in conjunction with a numerical method’ (14, p. 736). Scientifically important questions of accuracy and prediction are not exhausted by a philosophical analysis of whether or not a model stands in a representational relationship to certain aspects of reality. Primarily for this reason, what is needed in addition to an epistemology focused on representation, which has long been at the heart of the philosophical debate about models, is an ‘epistemology of results’, as it were. It is at this level that simulation gains significance: often, especially in the case of complex models, it is via the use of simulation techniques that specific numerical results and predictions are being derived from models. Somewhat similar to the way observation and measurement techniques furnish *empirical* data, simulation techniques generate specific instances of *simulated* data. Lest it be blind to this analogy, the philosophy of models, with its emphasis on representation, needs to be complemented by a philosophy of simulation, which takes due account of the non-trivial nature of generating results from models. As Eric Winsberg puts it, ‘we need an epistemology of simulation because simulation modeling is a set of scientific techniques that produces results.’ (37, p. 276) Quantitative agreement between simulated results and empirical data, though, may not be the most important criterion for assessing the adequacy of a model (or of simulation techniques, for that matter). As Sang Wook Yi points out, such quantitative agreement is usually ‘taken to be an extra virtue’, but is ‘not considered a necessary feature of model-based understanding’ – nor, one might add, of simulation-based understanding. (39, p. 84)

To the extent that one purpose of models is to enhance our scientific understanding of a phenomenon, numerical exactness and quantitative predictive success may thus become secondary considerations – though, no doubt, they will often be of independent interest, not least to those carrying out actual experiments on the kind of physical systems that are being simulated. As Robert Batterman has argued, in the case of models there may even be a direct tradeoff between quantitative exactness and explanatory value: ‘[T]he more one builds into one’s model – the more “exactly” it represents the phenomenon – the more likely it will be that solutions to the equations [of the mathematical model] will require methods of approximation’. On Batterman’s account, ‘what one would like is a good minimal model’ – one that captures the essential physics

rather than the ‘fine details’, even if this means that the model is a caricature of the ‘full’ system. Once such a model has been found, ‘[t]he adding of details with the goal of “improving” the minimal model is self-defeating – such improvement is illusory’. (3, p. 22) This calls for new ways of assessing and comparing the relative standing of models and simulations in general, beyond the criterion of quantitative accuracy. How, then, is one to evaluate models on the basis of their performance in simulations, given that agreement between simulated results and empirical data is at best one consideration amongst others? In Section 4, I shall discuss one important factor which, I want to suggest, no complete account of models and simulations can afford to neglect, namely the existence of benchmarks – more specifically, of mathematically rigorous results and relations (which, more often than not, may lack a straightforward empirical interpretation) – which allow for distinctive ways of assessing models and simulations across different contexts. Given the nature of such mathematically rigorous results and relations, let us first turn briefly to the role of mathematics in modelling and simulations more generally.

3 Mathematical and Computational Models

Mathematical models can take different forms and serve different purposes. They may be limiting cases of a more fundamental, analytically intractable theory, for example in the case of modelling planetary orbits as if planets were independent mass-points revolving around an infinitely massive sun. Sometimes, models connect different theoretical domains, as is the case in hydrodynamics, where Prandtl’s boundary layer model interpolates between the frictionless ‘classical’ domain and the Navier-Stokes domain of viscous flows. (See ref. (24).) In both cases, models allow for good quantitative predictions despite the intractability of the full theory. Even where a fundamental theory is lacking, mathematical models may be constructed, for example by fitting certain dynamical equations to empirically observed causal regularities (as in population cycles of predator-prey systems in ecology) or by analysing statistical correlations (as in models of stock-market behaviour).

This diversity of examples of scientific models notwithstanding, one can nonetheless single out a number of characteristic features of *mathematical models*. The mere fact that a model is expressed in mathematical form is clearly not enough. After all, a model is not merely a set of (uninterpreted) mathematical equations, theorems and definitions, as this would deprive it of its representational power: a set of equations cannot properly be said to ‘model’ anything, neither a specific phenomenon nor a class of phenomena, unless one gives some of the variables an interpretation that connects them with (some aspects of) empirically observable phenomena. After all, one of the key motivations for constructing a model, at least in cases where a ‘full’ theory is presumed to hold ‘in principle’, is the recognition that theories are about idealized objects (e.g., ‘mass points’) rather than real objects (e.g., planets).¹ In any case, it is impor-

¹Cf. Ronald Giere, who argues that there are good reasons to regard ‘Newton’s laws as

tant to acknowledge that mathematical models cannot merely be uninterpreted mathematical equations *if* they are to function as mediators of any sort; that is, if they are to *model* a case that, for whatever reason, cannot be calculated or described in terms of theoretical first principles.

The fact that mathematical models, like other kinds of models, require background assumptions for their interpretation, of course, does not rule out that in each case there may be a core set of mathematical relationships that model *users* regard as definitive of the mathematical model in question. In fact, where these mathematical features are not merely ‘inherited’ from an underlying fundamental theory, they may provide a mathematical model with precisely the autonomy and independence (from theory and data) that its role as mediator requires. Such autonomy differs from other, more radical kinds of independence from theory, which can sometimes be found in numerical modelling, at least when this involves fitting a set of – sometimes quite arbitrary – equations to empirical data. In cases with considerable uncertainty about the causal processes and dynamics laws governing a system (e.g., in analyses of the stock market), such mathematically informed ‘curve-fitting’ may be the theoretician’s last resort. It does, however, lie at the extreme end of possible ways of constructing mathematical models and, importantly, differs radically from simulation. Whereas ‘curve-fitting’ typically accommodates existing (past) data to an, often crude, mathematical model, simulation is essentially about the generation of new ‘data-like’ material – that is, of *simulated data* that were not antecedently available, neither via empirical observation nor via theoretical derivation.

Furthermore, though it may be true that, as Giere puts it, ‘[m]uch mathematical modeling proceeds in the absence of general principles to be used in constructing models’ (15, p. 52), it makes good sense to speak of a *mathematical model of* a phenomenon (or a class of phenomena) only if the mathematics employed (i.e., the *kind* of mathematical techniques and concepts) is in some way sensitive to the *kind* of phenomenon in question. For example, while it may be possible, if only retrospectively, to approximate the stochastic trajectory of a Brownian particle by a highly complex deterministic single-particle function, for example a Fourier series of perfectly periodic functions, this would hardly count as a good mathematical model: there is something about the phenomenon, namely its stochasticity, that would not be adequately reflected by a set of deterministic single-particle equations; such a set of equations would quite simply not be a mathematical model of *Brownian motion*.²

In addition to the requirement that the core mathematical techniques and concepts be sensitive to the kind of phenomenon that is being modelled, a further condition can be imposed on what should count as a *mathematical* model. Loosely speaking, the mathematics of the model should do some work in integrating the model’s various other elements; after all, it follows from the dis-

defining idealized abstract objects rather than as *describing* real objects’ (15, p. 52)

²Classical Brownian motion can, of course, occur in globally deterministic systems; however, at the single-particle level, the influence of the environment on a given particle *P* can only be included in the particle’s equation of motion in the form of an effectively random force acting on *P*.

cussion in the previous section that, for a mathematical construct to count as a *model* of a phenomenon or process, it must extend beyond its formal mathematical representation as a set of uninterpreted equations. Extra-theoretical considerations as well as background assumptions that do not lend themselves to formalisation must all be in place for a model to be a tool of scientific inquiry. A bare mathematical structure alone does not lend itself to application to individual cases. The perhaps vague demand that the mathematical aspects of a model should contribute to the integration of all, or at least *a wide range*, of the model's elements, can be given a concrete interpretation by way of example. If, say, a mathematical model employs the calculus of partial differential equations, then it should also indicate which (classes of) initial and boundary conditions need to be distinguished. Through specifying dynamic equations and their initial and boundary conditions, mathematical models can efficiently subsume different domains under the same basic structure. Consider, for instance, Prandtl's boundary-layer model of fluid dynamics, which, by specifying boundary conditions along with the model equations, succeeds in integrating not only different spatial domains (the boundary layer surrounding an object, and the infinite flow into which it is immersed), but also different domains of dynamic behaviour (laminar versus turbulent flow), as well as various background assumptions (Bernoulli's 'no-slip' condition, Helmholtz's principles etc.). Prandtl's model, it has been argued, is successful precisely because of its capacity to integrate different elements: '[I]f unification is taken to mean a close relationship among the elements used – which one could call structural coherence – then “unification” would indeed be the right expression to characterize Prandtl's advance over the rational mathematicians and especially over his predecessor Helmholtz.' (17, p. 58) Other authors have referred to this capacity of mathematical models to successfully integrate different elements, or different aspects of the same phenomenon, as 'mathematical moulding':

Mathematical moulding is shaping the ingredients in such a mathematical form that integration is possible, and contains two dominant elements. The first element is moulding the ingredient of mathematical formalism in such a way that it allows the other elements to be integrated. The second element is calibration, the choice of the parameter values, again for the purpose of integrating all the ingredients.
(6, p. 90)

Calibration is essential to the function and functioning of models. However, as will be argued in the next section, it should not be understood in the narrow sense of fixing the parameter values of a given model; rather, calibration may also take place across different models, by inquiring into their quantitative and qualitative behaviour as well as into the non-empirical relationships that hold between them. Such calibration, as the example of mathematically rigorous results and relations in the following sections will show, is complementary to, yet nonetheless distinct from, say, the issue of validation in the case of computational models. Computational models are typically implemented in the form of

an algorithm, either on a computer or on a network of computers. Their main structural and computational features are determined by such factors as network topology, numerical methods and algorithms used, computing power etc. In this regard, they differ from mathematical models, which are typically represented in an analytically closed form by a set of equations and whose structural characteristics are determined by mathematical constraints, not by constraints of realising a technological implementation. While computational models are often a crucial step in the actual implementation of computer simulations, it makes sense not to conflate them with simulations either. As R.I.G. Hughes has urged, one ought to distinguish between ‘the use of computer techniques to perform calculations, on the one hand, and computer simulation, on the other’. (19, p. 128)

4 Rigorous Results as Benchmarks for Simulations

The present section aims to apply the general framework outlined above to a particular class of mathematical models intended to describe and explain the physical behaviour of systems that consist of a large number of interacting particles. Such models, usually characterised by a specific Hamiltonian (energy operator), are frequently employed in condensed matter physics in order to account for phenomena such as magnetism, superconductivity, and phase transitions. Many-body models are particularly suitable as an example in the present context, since they form a class of models that, on the one hand, picks out a wide, yet well-defined range of physical phenomena as their explananda and, on the other hand, can be characterised mathematically by a narrow range of representational techniques (e.g., the formalism of second quantization). Many-body systems are also among the systems most widely studied using computer simulation, and it is the use of certain mathematical features of many-body models as benchmarks for the simulation of many-body systems, which will serve as a tool by which to analyse the interplay between models, and simulations of mathematical models, more generally.

Recall the idea of ‘mathematical moulding’ mentioned in the previous section: namely, the capacity of mathematical models to integrate, in drawing on certain mathematical techniques, diverse elements – some deriving from fundamental theory, others of non-theoretical origin – and subsume them under one mathematical structure. This capacity is essential for the applicability of mathematical models to specific scientific problems. The mathematics of a model, however, does not merely serve the ‘sanitary’ purpose of integrating already existing elements into a coherent formal structure; it also contributes new elements. By virtue of their mathematical structure, mathematical models possess stable features and intrinsic characteristics which may be independent of the model’s function of representing a physical system. Importantly, they can stand in a formal relation to other mathematical models, even when these are mod-

els of different physical systems. Thus, a mathematical model may contribute new elements to the theoretical description of the physical system, or class of systems, under consideration – elements which are not themselves part of the fundamental theory, but which may, in turn, take on an interpretative or otherwise explanatorily valuable role.

One important class of examples of such newly contributed elements are *rigorous results and relations* in statistical physics and many-body physics. Over the years, these have attracted considerable attention and have even given rise to a special branch of theoretical physics that concerns itself with rigorous results. (For a summary of some groundbreaking earlier developments, see Baxter 1982 and Griffiths 1972; for a philosophical case study see Gelfert 2005.) The term ‘rigorous result’ calls for some clarification. What makes a result ‘rigorous’ is not its qualitative or numerical accuracy as measured against experiments. In fact, the kind of ‘result’ in question will often have no immediate connection with the empirical phenomenon (or class of phenomena) the model or theory is supposed to explain. (In this regard, the derivation of rigorous results is unlike, say, the simulation-aided generation of [data-like] results discussed towards the end of section 2.) Rather, it concerns an exact *mathematical* relationship between certain mathematical variables or structural components of the mathematical model, which may or may not reflect an empirical feature of the system that is being modelled. An extreme example would be an ‘exact solution’ to a model, in which a complete evaluation of every variable would be possible for all parameter values. Other examples of rigorous results include, but are not limited to, conditions on the asymptotic behaviour in special limiting cases (some of which may be ‘unphysical’ in the sense that they do not correspond to actual physical scenarios – such as the limit of ‘infinitely strong’ interaction among particles in a system), symmetry requirements for particular mathematical elements of a model, ‘impossibility theorems’ that rule out certain kinds of macroscopic or dynamic behaviour of a model, and so forth.

The ‘active’ contribution of the model – that is, its contributing *new elements* rather than merely *integrating* theoretical and experimental (as well as further, external) elements – is not only relevant to interpretative issues, but also has direct consequences for assessing the techniques used to evaluate the model in specific circumstances, either by computing observable quantities or by simulating possible scenarios using a range of techniques. This is particularly salient in the case of the rigorous results mentioned in the preceding paragraph. Rigorous results are exact results that are true *of a model* (or a class of models) rather than of a theory. They often take the form either of exact relations holding between two or more quantities, or of lower and upper bounds to certain observables. If, for example in a model of a magnetic phase transition, the order parameter in question is the magnetization, then rigorous results – within a given model – may obtain, dictating the maximum (or minimum) value of the magnetization or the magnetic susceptibility permitted by the model. Frequently, rigorous results and relations provide a partial mapping of a model’s mathematical structure onto relationships between observables. By checking the results of computer simulations against those relationships, one can

then hope to find out whether a given simulation technique respects the model's fundamental features.

The partial independence of rigorous results from fundamental theory, and the fact that they are model-specific, makes them interesting 'benchmarks' for the numerical and analytical techniques of calculating observable quantities from the model. R.I.G. Hughes notes this, albeit only in passing, in his case study of one of the first computer simulations of the Ising model: 'In this way the verisimilitude of the simulation could be checked by comparing the performance of the machine against the exactly known behaviour of the Ising model.' (19, p. 123) The significance of 'benchmarks' for the purposes of simulations can hardly be overestimated. As Winsberg emphasises, simulations are often performed to investigate systems for which data are sparse; hence, 'comparison with real data can never be the autonomous criterion by which simulation results can be judged' (37, p. 287). If empirical data are not available, other reliable means of calibration must be found as a substitute. This is where rigorous results play an important role, and indeed may be crucial to the assessment of a simulation's success, given that '[t]he first criterion that a simulation must meet is to be able to reproduce known analytical results' (37, p. 288). Rigorous results thus can be seen to play an essential role in the verification of a simulation, where *verification* 'is taken to mean the testing of the model in relation to existing analytical solutions [...] as a benchmark' (as opposed to a simulation's *validation* against empirical data). (17, p. 59)

5 The Cohesion of Simulative Practice: A Descriptive Perspective

Strategies of modelling and simulating physical systems raise a number of justificational questions. Are the methods that are being used reliable? Does the outcome successfully describe reality? Do models and simulations enhance our understanding of the phenomena that are being studied? As discussed above, such questions have typically been discussed in terms of whether or not the model in question is a faithful representation of the physical system. However, it is by no means obvious how, in practice, the representational relationship between a model and reality could be assessed globally. At best, one can hope to probe this relationship locally and test the model's performance in specific circumstances. The mediator view of models argues that it is a model's performance in specific 'interventionist contexts', in connection with 'specific phenomena', which is the main source of justification and determines the model's validity. On this account, the specific outcomes in different instances of employing a model determine its instrumental value which, in turn, is considered a measure of the model's justification. However, in the present section I want to suggest that, from a descriptive perspective, such a purely 'outcome-based' perspective does not exhaust the range of actual sources of justification in our modelling and simulation practices.

It should be emphasized that the outcome-based approach is not limited to comparison with empirical data. As the more and more widespread use of simulation techniques suggests, the data-like results of computer simulations can take on a similar role for the purposes of assessing a model's performance. At a descriptive level, it is by no means clear that the main activity of researchers consists in assessing the model's performance in experimental or other empirical contexts. At least an equal amount of work goes into comparing and calibrating different methods of numerical evaluation against each other. That is, the calibration often takes place not between models and empirical data, but amongst different methods of numerical evaluation, irrespective of *their* empirical accuracy. This can be made particularly salient in the case of the many-body systems referred to earlier. When it comes to the use of many-body models in solid-state and condensed matter physics, it is not uncommon to come across whole papers on, say, the problem of 'magnetism in the Hubbard model', which do not make a single reference to empirical data. (As an example, see Tusch, Szczech, Logan 1996.) Even in cases where quasi-exact numerical results are obtainable for physical observables (for example by Quantum Monte Carlo simulations), these will often be compared not to empirical data but instead to other results arrived at by other numerical methods. As the authors of an early (1976) review of simulation techniques in condensed-matter physics put it, computer simulation 'has served to take the place of experimental measurements in many a theoretical paper discussing properties where the experiments had not been performed, were not possible to perform, had been performed on systems not corresponding to the theoretical Hamiltonian, or simply had not been performed to the accuracy and definition of the computations' (32, p. 169). Rather than adjust the parameters of the model to see whether the behaviour of a specific physical system can be modelled with empirical accuracy, the parameters will be held fixed to allow for better comparison of the different evaluative techniques with one another, often singling out one set of numerical outcomes (e.g., those calculated by Monte Carlo simulations) as authoritative. However, such comparison between different numerical methods, rather than with actual empirical data, does not yet constitute an extra source of justification for the model or simulation in question – it has merely replaced empirical with simulated outcomes.

The existence of rigorous results and relations as benchmarks, as discussed in the previous section, adds to the use of computer simulation a host of new considerations that go beyond mere comparison with empirical data obtained from actual measurements. For example, it is standard practice to run simulations for special limiting cases and compare the numerical results thus obtained with mathematically known exact results – even when these involve 'unphysical' assumptions, such as letting a (finite) physical variable mathematically diverge to infinity. While the capacity to numerically reproduce such exact results in limiting cases is no direct indicator of a model's empirical validity, it points to the robustness and stability of the numerical method that is being employed. Other bench-marking techniques involve reproducing exact *solutions* of a model (which are typically obtainable only for restricted model systems consisting of a small number of interacting particles, or systems that are confined to low

dimensions, such as a one-dimensional ‘chain’ of atoms as opposed to a fully extended three-dimensional lump of material), or they may require demonstrating that a simulation is invariant with respect to simple transformations (e.g., of coordinate systems). Scientists typically refer to benchmarks of this kind as a way of ‘testing’ a numerical technique, and as way of ‘improving our understanding’ of the corresponding model.³ If a numerical technique meets a large number of benchmarks associated with rigorous results about a model, then one may say that, in a loose sense of the word, both the model and the simulation technique display a significant degree of ‘cohesion’, irrespective of the empirical significance of the model.⁴

Given that cohesion between rigorous results about a mathematical model and the numerical methods applied to it, does not guarantee that a model is an empirically adequate representation of a real system, one might be tempted to doubt the significance of this weak kind of ‘cohesion’. One might insist that, while it may be true that a good deal of preliminary benchmarking and cross-checking has to happen before a model’s predictions can be compared with empirical data, nonetheless empirical success is the ultimate goal. The view has some legitimacy, though it hinges on a rather narrow interpretation of scientific modelling: In many cases it is precisely the activity of cross-checking and ‘benchmarking’ that drives research and makes up the better part of it. At the very least, it must be acknowledged that some of the most heavily researched models typically are not being assessed by their performance in specific empirical contexts. In part, this is because many models never were intended for *specific* phenomena in the first place, but for a qualitative understanding of a range of physical systems.⁵ In many areas of research, as R.I.G. Hughes points out with respect to the physics of critical phenomena, ‘a good model acts as an exemplar of a universality class, rather than as a faithful representation of any one of its members’ (19, p. 115). The notion of ‘universality class’ in particular, which makes it possible to give a comprehensive classification of models independent of the dynamical details of the physical systems in question, is directly associated with rigorous results and relations (concerning the critical exponents in the vicinity of a phase transition). (See ref. (9) for a detailed survey.) The mathematical framework of renormalization group theory, too, has given rise to exact mappings of certain kinds of mathematical models on to one another, whereby, as Cyril Domb puts it, certain models ‘ceased to be an *odd man out*, and could be fitted into a general framework’ (9, p. 244-5.). It is this process of fitting *prima facie* disparate models into a general framework, typically with the aid of rigorous results and relations among models, whereby cohesion of simulative practice is achieved and model-based understanding in science is deepened.

³See refs. (34) and (11), the latter of which makes use of all of the afore-mentioned benchmarking strategies.

⁴Here, as in the title of this section, the term ‘cohesion’ is chosen, in line with standard usage in English, to indicate a weaker kind of ‘sticking-together’ than would be implied by (epistemic) ‘coherence’. For a somewhat different distinction between ‘cohesion’ and ‘coherence’, see Olsson (2005), Chapter 5.

⁵On this point, see also Yi (2002).

However, as I shall argue in the next section, such apparent cohesion at the descriptive level of scientific practice is underpinned by a more specific kind of (epistemic) cross-model justification.

6 Coherence and Cross-Model Justification

In the previous section, I argued that the way in which scientists employ, and numerically evaluate models, is characterised by a high degree of cohesion in the sense that the activity of cross-checking and benchmarking allows for the comparison between known exact results and the outcomes of different numerical methods across a range of models. To paraphrase Domb (1996), individual models may thus be tied into a general methodological framework, whereby they become part of a more general fabric of mathematical and numerical techniques, rather than being considered in isolation. This sense of ‘cohesion’, however, while no doubt suggestive, may be considered too weak to sanction any *specific* appeal to individual rigorous results as a source of justification. In the present section, I shall make a significantly stronger point regarding the epistemic role of mathematically exact relations between models. In particular, I shall argue that epistemic warrant may flow from one model to another in virtue of their being linked by an appropriate, mathematically rigorous relation. Not just any kind of rigorous relation will do, but by way of example I intend to show that, given certain general conditions, there is a straightforward way in which the empirical success of one model warrants belief in the adequacy of another, quite different model, provided both are linked by an appropriate mathematically rigorous relation. First, however, I shall contrast the kind of cross-model justification proposed here with more established versions of indirect justification.

On standard accounts, indirect justification is a matter of a piece of evidence E confirming one hypothesis H^* by way of confirming another hypothesis H , which is more directly related to E than is H^* (for example because E is logically entailed by H , but not by H^*). If H and H^* are related in such a way that any epistemic support for H may be taken as boosting confidence in H^* , then it is plausible, following the aforementioned terminology, to speak of a relation of *indirect confirmation* between E and H^* .⁶ A historical example is discussed by Laudan and Leplin (1991), who describe how empirical evidence of magnetic striping in magmatic material on ocean floors was regarded as providing indirect confirmation of the hypothesis that continents had undergone significant climate changes throughout geological history. The two propositions, at first sight, seem to have little to do with each other, but they can be linked by Alfred Wegener’s continental drift thesis and the theory of plate tectonics: If continental shelves do indeed move relative to one another along suboceanic faults where new geological material is constantly being formed, then one would

⁶It must be noted, however, that speaking of indirect confirmation as a *relation* may be misleading in some contexts, since the term ‘indirect confirmation’ does not refer to a logical relation between propositions expressed by evidential statements and hypotheses, but instead merely expresses a methodological connection.

expect to find both magnetic striping on ocean floors (caused by the – independently known – multiple reversals the Earth’s magnetic field has undergone in the past) *and* variation in a continent’s climate as it moves across different latitudes. The tectonic drift theory thus both accounts for the formation of new geological material and predicts that continents move across different climate zones. The evidence of magnetic striping, E , directly confirms the hypothesis (H) that geological material is being formed in an ongoing process. By identifying a possible mechanism underlying continental drift, it also lends support to tectonic drift theory as a whole and, by entailment, to the thesis (H^*) that continents have undergone major climate changes in the past. The example at hand instantiates what is perhaps the most elementary formal account of indirect confirmation: if two hypotheses, H and H^* , are entailed by a comprehensive theory – in the present case, the theory of plate tectonics – and if a given empirical finding E confirms H , then E also confirms H^* . While this picture appears intuitive, not least thanks to the choice of historical example, Okasha (1997) has convincingly argued that it runs into trouble as a formal account of indirect confirmation, since it gives rise to a version of Hempel’s paradox: namely, that any contingent statement confirms any contingent statement. This suggests that more than merely attending to the logical relations between evidential statements, hypotheses, and theories is needed in order to make sense of indirect confirmation.⁷ For example, one might try to find a way of restricting indirect confirmation in such a way that a given hypothesis may receive indirect support only from a subset of other hypotheses and empirical findings, and not from just any evidential statement. This might help to avoid Hempel’s paradox, but it requires an account of when indirect confirmation is possible in principle.

The possibility of indirect confirmation, as a relationship between a piece of evidence and different hypotheses (or conjunctions of hypotheses), has recently been elaborated on, and has been given a formal underpinning, by work on the (purported) truth-conduciveness of coherence. (See refs. (22) and (8).) As Franz Dietrich and Luca Moretti (2005) point out, while it is doubtful that coherence is truth-conducive *in general* (in the sense that, given certain non-arbitrary background assumptions, a set of hypotheses S that is more coherent is always more probable), it is nonetheless possible to specify conditions under which coherence is *confirmation-conducive*. These conditions can be made mathematically precise on any coherence measure that satisfies certain non-trivial confirmation transmission properties. More specifically, it can be shown rigorously that, given an appropriate coherence measure (typically defined relative to a probability function, as a mapping of sets of statements onto real numbers indicating the degree of coherence among the set’s members), there exist finite (non-trivial) coherence thresholds, such that, for any set S that meets the threshold, the following principles hold:⁸ (i) evidence for any one member of S

⁷Indeed, this observation may provide additional motivation for the project, to be pursued later in this section, of shifting attention from relations between scientific theories and hypotheses to relations between scientific models.

⁸The condition of ‘non-triviality’ in relation to the coherence threshold c demands that, for

also confirms each member of S , and (ii) evidence for any one member of S also confirms the logical conjunction of all members of S .⁹ More recently, Moretti (2007) has exploited the formal symmetry between evidential statements and hypotheses to construct two analogous ‘evidence-gathering properties’, to stand alongside the confirmation transmission properties associated with (i) and (ii). The upshot of these findings is the vindication of two standard practices of scientific confirmation: the confirmation of a theory (regarded as a conjunction of statements) by confirmation of its parts, and the method of indirect confirmation of one hypothesis H^* by evidence for another, related hypothesis H . Provided that the two general confirmation-transmission properties are satisfied, if H and H^* are sufficiently strongly coherent (in the sense specified by the coherence threshold for the set S comprising H and H^* as its members), then evidence for H flows to H^* , as well as to the conjunction $H \wedge H^*$. Note that, unlike the elementary formal account of indirect confirmation derived from Laudan and Leplin (1991), the present account does not suffer from Hempel’s paradox. By making the possibility of indirect confirmation conditional on the existence of strong coherence between two hypotheses, the account places limits on which hypotheses can lend support to each other, thus avoiding the otherwise paradoxical conclusion that every contingent statement would appear to confirm every other.

The significance of these findings is considerable, not least since they provide a clear sense in which coherence is epistemically advantageous. There does, however, remain the question of what coherence amounts to: what do we mean when we say that one belief (hypothesis, statement, proposition) coheres with another, or indeed coheres better with one belief (or hypothesis, etc.) than with another. Intuitively, two beliefs are coherent if they ‘hang together’ sufficiently well.¹⁰ Different proposals are available as to how best to cash out this intuitive understanding of coherence. Thus, Olsson (2002: 249) defines coherence of a set of beliefs as proportional to the degree to which the contents of those beliefs agree and proposes a coherence measure based on this idea. By contrast, Fitelson (2003) regards coherence as an extension of the notion of (deductive) logical coherence of sets of statements; hence, he introduces a coherence measure which essentially averages over all reciprocal relations of support (confirmation) that hold between any statement of a given set and any subset of other statements of that set. Different analyses of our pre-theoretical understanding of coherence may be pursued, and a number of other proposals have by now been formalised. Interestingly, the confirmation-transmission properties discussed above have been shown to hold for several of these formal coherence measures (includ-

a given coherence measure C , the value of c be less than the maximal coherence level across all permissible sets S of statements (that is, of finite, nonempty sets S containing statements H with $p(H) > 0$), where the maximal coherence level will be the one assigned by C to any set of pairwise equivalent statements or formulae. (See Dietrich and Moretti 2001, p. 413.)

⁹The former principle may also be weakened to allow for the pairwise consideration of scientific hypotheses, i.e. by considering binary subsets S^* containing hypotheses H and H^* instead of the full set S ; see (Moretti 2007).

¹⁰This locution is due to (Tomoji 1999) and has recently been elaborated on by (Moretti 2007).

ing those proposed by Olsson and Fitelson), and one might even speculate that the satisfaction of the conditions associated with the possibility of confirmation transmission is itself a strong desideratum for any viable coherence measure. Unlike (some forms of) traditional coherentism, there is no suggestion in any of these proposals that coherence alone can somehow guarantee the likely truth of a set of propositions – that it can generate justification ‘from scratch’, as it were¹¹ – but instead what is shown is how empirical confirmation may be transmitted across different (sets of) statements. This observation is of direct relevance to the study of models and simulations, since it serves as a reminder that more than mere ‘cohesion’ of scientific practice alone (in the loose sense discussed in the previous section) is required, or at least desirable, in order to warrant trust in the results of modelling and simulation techniques. It also makes plain why it is useful and important to study empirical cases of transfer of warrant across models, especially in such instances where, as I claim, mathematically rigorous formal relations between models suffice to enable such cross-model justification. By way of example, it may be possible to test the suggestion that one possible way in which mathematical models can be seen to be coherent is by analysing those mathematically rigorous relations that hold between them and link them together. Indeed, given that certain mathematical relations, such as mappings from one domain to another, simultaneously connect a large number of elements to each other, it appears plausible that mathematically rigorous results may provide an extra dimension along which models can ‘hang together well’.

Before considering a detailed example, it is important to point out some of the distinctive features of models and to consider what, specifically, transmission of confirmation would amount to in this case. In the historical scenario discussed by Laudan and Leplin, where geological evidence of magnetic striping, E , indirectly confirms (H^*), the climate change hypothesis, by confirming (H) the existence of ongoing formation of geological material (and thereby identifying a possible mechanism underlying the tectonic drift theory), there is a clear sense in which both H and H^* are part of, or deducible from, a more general theory. It is against the backdrop of such a theory, in conjunction with independently established auxiliaries, that one can, for example, draw on the background knowledge that there have been multiple reversals of the Earth’s magnetic field throughout geological history.¹² The existence of an overarching theory, which not only comprises the hypotheses to be confirmed but also places constraints on which empirical findings constitute confirming evidence, is a feature of most examples of indirect confirmation that are discussed in the literature. Typically, then, indirect confirmation of a hypothesis H^* is the result of identifying evidence in support of one of the more central tenets H of a scientific theory, and recognising that H and H^* are sufficiently coherent for the evidence to also confirm H^* .¹³ In the case of models, however, this account

¹¹On this point, see also (Shogenji 2001, p. 98).

¹²It is noteworthy in this context that Laudan and Leplin’s concern with indirect confirmation is only secondary, insofar as their main concern is with individuating, and choosing between, empirically equivalent theories. For criticism see Okasha (1997).

¹³That confirmation of more central tenets of a theory precedes the testing of more pe-

needs to be modified, for several reasons. First, models often include simplifying assumptions and idealizations which may or may not be limiting cases of an underlying fundamental theory. Demanding that two models be subsumable under the same general theory may be too restrictive and may exclude legitimate instances of cross-model support. Second, the notion of confirmation is not as straightforward for models as it is for individual hypotheses. On the standard picture in terms of conditional probabilities, a given piece of evidence E confirms a hypothesis H iff $p(H|E) > p(H)$. For models, the relationship between empirical evidence and confirmation will often be more tentative, insofar as a model is typically only intended to explain certain empirical aspects of a phenomenon; hence, which empirical data are considered relevant may be subject to revision. However, this by no means invalidates the idea of empirical confirmation *broadly construed*; it merely cautions against reducing a model's empirical success to, say, the quantitative accuracy of its predictions. This applies in particular to areas where models are employed in order to make sense of the qualitative features of a system, as in the study of complex phenomena such as phase transitions. In such cases, a model may be considered well-confirmed if, for example, it helps to identify a (possible) microscopic mechanism that would explain a certain (observed) macroscopic behaviour, and if further empirical evidence supports the existence of the postulated microdynamics. Whether a model is considered successful or not is not so much a matter of quantitative precision, or of reproducing particular empirical data (and certainly not just any empirical data!), but of whether it stands in an adequate representational relation to a physical system or class of physical systems. As Batterman points out, a model may well be considered successful when it gives a 'caricature' of a physical system rather than an accurate portrayal of all its details. Despite these differences between the appraisal of models and confirmation of (ordinary) hypotheses, it is nonetheless possible to press the analogy with the standard picture of confirmation of a hypothesis H by evidence E , as long as one recognises that what should properly take the place of H is a hypothesis *about the model* – namely, the hypothesis that it is an adequate representation of a physical system – rather than just any statement that has been derived from the model, whether predictive or otherwise.¹⁴ An interesting corollary of this account concerns the possibility of a transfer of empirical warrant from one model to another, namely when a well-confirmed model M is related to another model M^* in such a way that the very fact of M 's being an empirically successful representation of a physical system at the same time guarantees the (likely) representational success of M^* as a model of *another*, not necessarily similar physical system. In this case, it appears perfectly legitimate to speak of a case of indirect confirmation of one model, M^* , by (the empirical success of) another, M , given that

ripheral hypotheses, if true, is merely a historically contingent fact; nothing in the present argument hinges on this temporal sequence.

¹⁴This is not so far-fetched a construal as it may seem: after all, in Bayesian epistemology too, it is not the objective probability of a hypothesis given certain evidence that is at stake, but 'roughly', as Ellery Eells and Branden Fitelson put it, 'the degree of belief the agent would have in H were the agent to learn that E '. (10, p. 663)

the relation between the two furnishes extra reasons for the hypothesis that M^* is indeed a representation of a physical system, when previously this may have been a matter of mere stipulation. Such a case may well turn out to conform to the criteria of confirmation-transmission reported earlier, and indeed, as we shall see, appears *prima facie* vindicated by the example to be discussed in the remainder of this section.

In support of this general picture of indirect confirmation between models, I shall develop in detail a case study from the recent history of physics. In particular, I intend to show how empirical warrant may flow from one model to another, quite different model in virtue of their being linked by an appropriate, mathematically rigorous relation. This will substantiate the point, made above, that rigorous results and relations may give rise to a specific kind of (epistemic) cross-model justification that goes beyond mere cohesion of scientific practice. The case to be discussed concerns the so-called Mott metal-insulator transition. (For a review see ref. (13).) As it will turn out, attempts to explain this phenomenon (more about which shortly) with the aid of many-body models, rely heavily on the empirical success of other, quite dissimilar models as a source of justification. In particular, the many-body model now most widely believed to be capable of explaining the Mott transition – the Hubbard model – was historically recognised to be able to fill this role, only because of its standing in an adequate, mathematically rigorous relation to another widely discussed many-body model, the Heisenberg model. Before discussing in detail the character of the models involved in the example, and how the mathematically rigorous relation between them can function as a channel of empirical warrant, a brief description of the phenomenon and its history is in order. Its first description dates back to the 1930s, when experimental research on the properties of metals and semiconductors flourished, not least thanks to the successful band theory of electronic semiconductors, proposed by A.H. Wilson in 1931. In 1937, two Dutch industrial scientists, J.H. de Boer and E.J.W. Verwey, called attention to an anomaly of the conduction properties of certain transition metal oxides, such as nickel oxide. As De Boer and Verwey pointed out, its incomplete $3d$ band of electrons means that nickel oxide should be an excellent conductor instead of the extreme high-resistance insulator it actually is. According to Wilson’s band theory of conduction, even the smallest amount of energy should be enough to excite electrons into conducting states, whereas experiments showed nickel oxide to display a high electric resistivity. Nevill Mott (1937), replying to De Boer and Verwey’s conference paper, suggested that the unexpected insulating behaviour of nickel oxide and similar transition metal oxides might be caused by the repulsive interactions among the electrons themselves, thus rendering them unable to carry electricity. The suggestion was speculative at the time, but one of the predictions that followed for such ‘Mott insulators’, as they became known, was a sharp transition from insulating to conducting behaviour as the interaction strength between the electrons is being manipulated (e.g., by ‘doping’ nickel oxide with lithium impurities). Experimental evidence for the occurrence of such metal-insulator transitions was found soon thereafter, yet for a long time no adequate theoretical model of how such transitions come about

could be identified.

In parallel to the study of the electric properties of metals and semiconductors, significant research efforts were also expended on the study of magnetic properties, both of insulators and of ferromagnetic materials. Two models are especially significant in this context: the Heisenberg model and the Hubbard model. Werner Heisenberg, in 1928, proposed a model on the basis of the emerging quantum mechanics, which aimed to describe magnetic behaviour in systems with fixed spins. Some 35 years later, John Hubbard employed the by then far more sophisticated formalism of second quantisation, to propose a model of magnetism for the radically different case of systems with itinerant (electron) spins. The physical interpretations of the two models differ starkly: In the former case, spins – the ‘elementary magnets’, so to speak – are spatially located at atoms in a crystal lattice, whereas in the latter case they are associated with freely moving, delocalized electrons. The Heisenberg model, in particular, was highly successful in explaining a number of empirical laws and phenomena for insulators with localized magnetic moments, such as Bloch’s $T^{3/2}$ law, the Curie-Weiß law, the occurrence of spin waves, and critical phenomena in the vicinity of phase transitions. The model’s empirical success was helped partly by its amenability to techniques of analytical approximation and numerical computation; indeed, the various results and predictions of the Heisenberg model are generally regarded as ‘so convincing that one does not expect very much additional information from [its] unknown exact solution’ (27, p. 126). The Hubbard model’s empirical success in modelling the ferromagnetic behaviour of metals is somewhat more ambiguous, mainly due to its more complex mathematical structure (which results from a mathematical term representing the collective electron-electron interaction), which makes the derivation of quantitative predictions and relationships (such as the various empirical laws mentioned in the case of the Heisenberg model) more difficult. Hubbard himself was able to show that a ferromagnetic phase transition was not ruled out by the model and that certain (qualitative) conditions would be favourable for ferromagnetism. It was not until the advent of quasi-exact simulation techniques, however, that the Hubbard model was vindicated as indeed capable of representing ferromagnetic behaviour.

That two many-body models may differ in their empirical success is of course to be expected, especially when their target phenomena are different. In the present example, the Heisenberg model had acquired considerable empirical confirmation as a model of the magnetic properties of insulators, long before the Hubbard model was shown to be capable of representing ferromagnetic behaviour in metals. However, as discussed earlier, comparison with empirical data alone does not exhaust the ways in which mathematical models may be studied: insight may also be gained from studying the formal relations that hold between models, qua their status as mathematical objects. In the case of the Heisenberg and the Hubbard model, the mathematically rigorous relations between them are especially intimate. For example, it can be shown that, under certain conditions (at half-filling – i.e., when half of the quantum states in the conduction band are occupied – and in the strong-coupling limit, when the parameter rep-

representing the relative strength U of the electron-electron interaction diverges) the Hubbard model can be mapped fully on to one version of the Heisenberg model (namely, the spin-1/2 antiferromagnetic Heisenberg model). Under the specified conditions, the two models, despite their different physical interpretations, are *de facto* isomorphic and display the same dynamic behaviour. In addition to such rigorous relations between models, there also exist rigorous results for certain limiting cases of each individual model. For example, in the case of vanishing electron-electron interaction ($U = 0$), the Hubbard model simplifies considerably and is mathematically equivalent to the much older, exactly solvable Fermi gas model (proposed by Enrico Fermi in 1935), which describes the statistics and conduction properties of an ideal metal in terms of a perfect electron gas.

As it turns out, when considered in conjunction, these mathematically rigorous findings combine to allow for the kind of transfer of empirical warrant that, I suggested earlier, could take place across models, even when these describe very different physical systems. In order to complete the example, what needs to be shown is how the Hubbard-Heisenberg mapping and the identity between the Hubbard model at $U = 0$ and the Fermi electron gas model combine to channel empirical warrant from the Heisenberg model to the Hubbard model. Interestingly, this transfer of warrant extends beyond the Hubbard model's original purpose of representing ferromagnetism, and instead applies to the Hubbard model *as a model for the Mott metal-insulator transition*: what will be vindicated is the hypothesis – which previously had only been stipulated – that the Hubbard model stands in a representational relation to a class of real physical systems (including compounds involving ferromagnets such as nickel), in this case comprising those substances that are *Mott insulators*. In order to see how the indirect confirmation of a representational relationship comes about, one needs to consider the rigorous results and relations discussed earlier. Recall that, at $U = 0$ (zero electron-electron interaction), the Hubbard model has been rigorously shown to describe a perfect conductor (ideal Fermi electron gas), whereas in the limit of $U \rightarrow \infty$ (diverging interaction strength) it maps exactly onto the antiferromagnetic Heisenberg model. The latter, in virtue of its empirical success, had long become ‘the “standard model” for the description of magnetic insulators’ (13, p. 75). One is thus faced with a model which, for vanishing parameter value $U = 0$, describes a conductor and, towards the other end of the spectrum of possible parameter values of U , describes an insulator. The former is a direct rigorous result about the Hubbard model considered in isolation, whereas the latter follows from the rigorous relation that maps the Hubbard model onto the Heisenberg model, coupled with the body of empirical evidence supporting the adequacy of the Heisenberg model as a model of antiferromagnetic insulators. Given these properties of the Hubbard models in the two limits of weak ($U = 0$) and strong ($U \rightarrow \infty$) electron-electron interaction, there are good grounds to believe that, at some finite intermediate level U_c a transition in the model's behaviour occurs from conductor to insulator – precisely what happens in the case of the Mott metal-insulator transition. Hence, as one leading physicist summarises the consensus of his peers, ‘we are confident that

the model is indeed capable of describing a Mott transition at a critical interaction strength U_c ' (13, p. 75). Given that the Hubbard model is formulated in terms of operators representing electronic behaviour only, this may be taken as a confirmation of Mott's original hypothesis that the repulsive interaction among otherwise free electrons may be responsible for the unexpected insulating behaviour. Note, however, that, in this case, whatever confirmation is conferred to Mott's original *hypothesis* is dependent on the evidence for the Hubbard model's success as a *representation* of a class of physical systems, where such evidence of representational success is itself partly constituted by mathematically rigorous results that allow for the flow of warrant from the empirically well-confirmed Heisenberg model to the initially more stipulative Hubbard model. In other words, in the example just discussed, empirical warrant first flows from one model to another, in virtue of their standing in an appropriate mathematically rigorous relation, from which it may then 'trickle down' to the level of individual hypotheses. Rigorous results and relations may thus come to be seen as playing an essential role in conferring cross-model justification.

Before concluding, I wish to briefly address two concerns that might stand in the way of recognising the capacity of rigorous results and relations to confer justification from one model to another. First, one might worry whether purely formal results can indeed confer justification. How can the mere fact that one model is related to another through a set of mathematically rigorous relations in itself be a source of justification? This worry is misplaced, however, in that it conflates two very different aspects of justification: its generation and transmission. The present account of cross-model justification only concerns the latter; that is, it proposes a novel way of how empirical warrant may flow from one model to another, and need not be committed to the more radical claim that formal results can somehow create justification 'from scratch'. Indeed, in this regard the situation is the same as in recent research into the epistemic role of coherence, reported earlier in this section, according to which one can save many of the intuitions concerning the epistemic utility of coherence without having to resort to the dubious claim that coherence is in itself always truth-conducive. A second worry concerns the status of those mathematically rigorous results and relations that hold only in 'unphysical' limiting cases. Consider the example discussed above, of the Hubbard-Heisenberg mapping in the limiting case of diverging interaction strength $U \rightarrow \infty$. Such a case cannot, strictly speaking, purport to describe any actual physical system, where the interaction is necessarily finite. However, the fact that certain mathematical limiting cases are empirically vacuous need not undermine their relevance, given that the validity of mathematically rigorous results and relations does not hinge on their empirical confirmation.¹⁵ Furthermore, given that the construction of models is itself heavily dependent on techniques of idealization, limiting procedures and 'asymptotic reasoning' (see Batterman 2001), there is no principled reason why the same techniques should not be used in comparing and evaluating such

¹⁵This is not to say that one cannot often give an entirely reasonable empirical interpretation of such limiting cases, for example by specifying which orders of magnitudes are to be considered relevant.

models. Indeed, if one takes one's lead from actual scientific practices of modelling and confirmation, and from the historical example of the Mott transition in particular, there seems little reason to doubt that cross-model justification is possible and that rigorous results and relations need not *themselves* be imbued with direct empirical significance in order to serve their auxiliary role of conferring empirical warrant from one model to another.

7 Conclusion

In this paper, I have argued that mathematically rigorous results and relations play a crucial role in many areas of scientific modelling. Not only do they serve as a means of coordinating and calibrating the numerical techniques and approximations employed in actual scientific practice; they also give rise to a specific kind of cross-model justification, in that they allow for the transfer of empirical warrant between models, by virtue of their standing in an appropriate, mathematically rigorous relation. Importantly, rigorous results are genuinely new contributions of a model: they are neither entailed by theoretical 'first principles', nor can they be inferred from empirical data. In this sense, rigorous results are internal to a model, or class of models, and cannot be assimilated to either theory or data. As such, they illustrate the capacity of models to take on roles beyond both fundamental theory and performance in specific empirical and interventionist contexts. It is such rigorous results, I claim, which guide much of research by providing the focus for modelling strategies and attempted refinements of evaluative techniques, whether by numerical means, analytical evaluation, or computer simulation. Characteristically, rigorous results provide non-empirical constraints, which may serve as general 'benchmarks' for model construction and computer simulation; it is such constraints and benchmarks which guide the process of model refinement. Since rigorous results generally are validated by a model's formal features as a mathematical object, they can also relate different models in quite unexpected ways. This allows for the transfer of empirical warrant from one model to another, even in cases where both represent different classes of physical systems, as in the mapping of the Hubbard model at half-filling onto the Heisenberg model. It is only in virtue of this mathematically rigorous mapping, coupled with the empirical success of the Heisenberg model, that the Hubbard model has been recognised as successfully representing a class of real physical systems, known as Mott insulators. By vindicating a model's representational success, based on the empirical success of another, quite different model, rigorous results and relations allow for a specific kind of indirect confirmation. This phenomenon of (epistemic) cross-model justification, I have argued, is not merely another form of 'moulding' a mathematical model to concrete empirical situations, nor is it exhausted by the merely descriptive cohesion of actual scientific practices of modelling and simulation. Rather, rigorous results and relations are novel contributions at the level of mathematical models and, as such, deserve closer philosophical analysis and attention.

8 Acknowledgments

I would like to thank Barry Loewer and Imre Kondor for constructive discussions during my stay at Collegium Budapest (Institute for Advanced Study), where this paper was first conceived. I am also grateful to audiences at the CNRS/LSE conference on ‘Models and Simulations’ in Paris and at Tokyo Institute of Technology. Two anonymous referees for *Synthese* have offered exceptionally detailed and helpful comments, for which I am very grateful.

References

- [1] Daniela M. Bailer-Jones: “When scientific models represent”, *International Studies in the Philosophy of Science*, 17, 2003, 59-74.
- [2] Robert W. Batterman: *The Devil in the Details. Asymptotic Reasoning in Explanation, Reduction and Emergence*, Oxford: Oxford UP 2001.
- [3] Robert W. Batterman: “Asymptotics and the Role of Minimal Models”, *British Journal for the Philosophy of Science*, 53, 2002, 21-38.
- [4] Rodney J. Baxter: *Exactly Solved Models in Statistical Mechanics*, New York: Academic Press 1982.
- [5] J.H. de Boer and E.J.W. Verwey: “Semi-Conductors with Partially and with Completely Filled 3d-Lattice Bands”, *Proceedings of the Physical Society (Extra Part)*, 49, 1937, 59-71.
- [6] Marcel Boumans: “Built-in justification”, in Margaret Morrison and Mary S. Morgan (eds.): *Models as Mediators. Perspectives on Natural and Social Science*, Cambridge: Cambridge UP 1999, 68-96.
- [7] Nancy Cartwright: *The Dappled World. A Study of the Boundaries of Science*, Cambridge: Cambridge UP 1999.
- [8] Franz Dietrich and Luca Moretti: “On Coherent Sets and the Transmission of Confirmation”, *Philosophy of Science*, 72, 2005, 403-424.
- [9] Cyril Domb: *The Critical Point. A Historical Introduction to the Modern Theory of Critical Phenomena*, London: Taylor and Francis 1996.
- [10] Ellery Eells and Branden Fitelson: “Measuring Confirmation and Evidence”, *The Journal of Philosophy*, 97, 2000, 663-672.
- [11] Satoshi Ejima, Florian Gebhard, and Satoshi Nishimoto: “Tomonaga-Luttinger parameters for doped Mott insulators”, *Europhysics Letters* 70, 2005, 492-498.
- [12] Enrico Fermi: “Sulla quantizzazione del gas perfetto monoatomico”, *Rendiconti della Accademia Nazionale dei Lincei*, 6, 1935, 145-149.
- [13] Florian Gebhard: *The Mott Metal-Insulator Transition: Models and Methods. (Springer Tracts in Modern Physics, Vol. 137)*, Berlin: Springer 1997.

- [14] Axel Gelfert: “Mathematical Rigor in Physics: Putting Exact Results in Their Place”, *Philosophy of Science*, 72, 2005, 723-738.
- [15] Ronald N. Giere: “Using Models to Represent Reality”, in Lorenzo Magnani, Nancy J. Nersessian and Paul Thagard (eds.): *Model-Based Reasoning in Scientific Discovery*, New York: Plenum Publishers 1999, 41-57.
- [16] Robert B. Griffiths: “Rigorous results and theorems”, in Cyril Domb and Melville S. Green (eds.): *Phase Transitions and Critical Phenomena*, New York: Academic Press, 1972, 8-109.
- [17] Michael Heidelberger: “Applying Models in Fluid Dynamics”, *International Studies in the Philosophy of Science*, 20, 2006, 49-67.
- [18] John Hubbard: “Electron correlations in narrow energy bands”, *Proceedings of the Royal Society of London (Series A: Mathematical and Physical Sciences)*, 276, 1963, 238-257.
- [19] R. I. G. Hughes: “The Ising model, computer simulation, and universal physics”, in Margaret Morrison and Mary S. Morgan (eds.): *Models as Mediators. Perspectives on Natural and Social Science*, Cambridge: Cambridge UP 1999, 97-145.
- [20] Paul Humphreys: *Extending Ourselves*, Oxford: Oxford UP 2004.
- [21] Larry Laudin and Jarret Leplin: “Empirical Equivalence and Underdetermination”, *The Journal of Philosophy*, 88, 1991, 449-472.
- [22] Luca Moretti: “Ways in which coherence is confirmation conducive”, *Synthese* (forthcoming, as of 2007 available at Springer Online-First).
- [23] Margaret C. Morrison: “Modelling Nature: Between Physics and the Physical World”, *Philosophia Naturalis*, 35, 1998, 65-85.
- [24] Margaret Morrison: “Models as autonomous agents”, in Margaret Morrison and Mary S. Morgan (eds.): *Models as Mediators. Perspectives on Natural and Social Science*, Cambridge: Cambridge UP 1999, 38-65.
- [25] Margaret Morrison and Mary S. Morgan: “Models as mediating instruments”, in Margaret Morrison and Mary S. Morgan (eds.): *Models as Mediators. Perspectives on Natural and Social Science*, Cambridge: Cambridge UP 1999, 10-37.
- [26] Nevill Mott: “Discussion of the Paper by De Boer and Verwey”, *Proceedings of the Physical Society (Extra Part)*, 49, 1937, 72-73.
- [27] Wolfgang Nolting: “Ferromagnetism and Electronic Correlations”, in Ferdinando Mancini (ed.): *Lectures on the Physics of Highly Correlated Electron Systems (AIP Conference Proceedings Vol. 527)*, Melville: American Institute of Physics Publishing 2000, 118-225.
- [28] Samir Okasha: “Laudan and Leplin on Empirical Equivalence”, *British Journal for the Philosophy of Science*, 48, 1997, 251-256.

- [29] Erik J. Olsson: *Against Coherence: Truth, Probability, and Justification*, Oxford: Oxford UP 2005.
- [30] Tomoji Shogenji: “The Role of Coherence in Epistemic Justification”, *Australasian Journal of Philosophy*, 79, 2001, 90-106.
- [31] Mark Steiner: “The Application of Mathematics to Natural Science”, *The Journal of Philosophy*, 86 No. 9, 1989, 449-480.
- [32] M. Steiner, J. Villain, and C.G. Windsor: “Theoretical and experimental studies on one-dimensional magnetic systems”, *Advances in Physics*, 25, 1976, 87-209.
- [33] Mauricio Suárez: “Theories, Models, and Representations”, in Lorenzo Magnani, Nancy J. Nersessian and Paul Thagard (eds.): *Model-Based Reasoning in Scientific Discovery*, New York: Plenum Publishers 1999, 75-83.
- [34] A. Toschi, P. Barone, M. Capone, and C. Castellani: “Pairing and superconductivity from weak to strong coupling in the attractive Hubbard model”, *New Journal of Physics* 7 (7), 2005, 1-17.
- [35] Michael A. Tusch, Yolande H. Szczech, and David E. Logan: “Magnetism in the Hubbard model: An effective spin Hamiltonian approach”, *Physical Review B*, 53 No. 9, 1996, 5505-5517.
- [36] A.H. Wilson: “The theory of electronic semiconductors”, *Proceedings of the Royal Society of London (Series A: Mathematical and Physical Sciences)*, 133, 1931, 458-491.
- [37] Eric Winsberg: “Sanctioning Models: The Epistemology of Simulation”, *Science in Context*, 12, 1999, 275-292.
- [38] Eric Winsberg: “Simulated Experiments: Methodology for a Virtual World”, *Philosophy of Science*, 70, 2003, 105-125.
- [39] Sang Wook Yi: “The Nature of Model-Based Understanding in Condensed Matter Physics”, *Mind and Society*, 3, 2002, 81-91.