Handshaking Your Way to the Top: Simulation at the Nanoscale

Eric Winsberg†‡

Should philosophers of science be paying attention to developments in "nanoscience"? Undoubtedly, it is too early to tell for sure. The goal of this paper is to take a preliminary look. In particular, I look at the use of computational models in the study of nano-sized solid-state materials. What I find is that there are features of these models that appear on their face to be at odds with some basic philosophical intuitions about the relationships between different theories and between theories and their models. My conclusion is that developments in nanoscience are not an unlikely place for novel insights in the philosophy of science to emerge.

1. Introduction. All of the pundits, prognosticators, and policy makers are in agreement: research into the science and technology of the nanoscale is going to be one of the hot scientific topics of the twenty-first century. According to the Web page of the National Nanotechnology Initiative, moreover, this should make nanotechnology and nanoscience "of great interest to philosophers." Admittedly, the kind of philosophers being imagined by the authors of the initiative Web page are most likely something like the nanotechnological analogues of bioethicists—not the kind of philosophers that typically convene at the meeting of the Philosophy of Science Association. But what about us?

Should we philosophers of science, those of us who are interested in methodological, epistemological, and metaphysical issues in the sciences, be paying any attention to developments in nanoscale research? Undoubtedly, it is too early to tell for sure. But arguably, the right prima facie intuition to have is that we should. After all, major developments in the history of the philosophy of science have always been driven by

†To contact the author, please write to: Eric Winsberg, Department of Philosophy, University of South Florida, FAO 226, 4202 Fowler Ave., Tampa FL 33620; e-mail: winsberg@cas.usf.edu.

‡I thank the many attendees at PSA 2004 and at a presentation at the University of Bielefeld for helpful comments.

Philosophy of Science, 73 (December 2006): 582–594. 0031-8248/2006/7305-0011$10.00
Copyright 2006 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/518328 Published online by Cambridge University Press

major developments in the sciences. It is true that, historically, most of those scientific developments have involved revolutionary changes at the level of fundamental theory (especially, of course, the revolutionary changes in physics at the beginning of the twentieth century). It's also true that nanoscience is unlikely to bring about innovations in fundamental theory. But surely there is no reason to think that new experimental methods, new research technologies, or innovative ways of solving a new set of problems within existing theory could not have a similar impact on philosophy. And it is not altogether unlikely that some of the major accomplishments in the physical sciences to come in the near future will have as much to do with modeling complex phenomena within existing theories as they do with developing novel fundamental theories.

So far, none of this is meant to be an argument, but simply an impressionistically motivated suggestion that nanoscience *might* be something of philosophical interest. The project of this paper is to look and see, and to try to give a more definitive answer. Because of my past work, the place that I am inclined to do that looking is in aspects of model building, especially methods of computer simulation, that are employed in nanoscience. What I find is that it does indeed look as if there are good prospects for philosophers of science to learn novel lessons, especially about the relations between different theories, and between theories and their models, by paying attention to developments in simulation at the nanoscale.

To begin, what exactly is "nanoscale science"? No precise definition is possible. But intuitively, it is the study of phenomena and the construction of devices at a novel scale of description: somewhere between the strictly atomic and the macroscopic levels. Theoretical methods in nanoscience, therefore, often have to draw on theoretical resources from more than one level of description.

Take, for example, the field of nanomechanics. "Nanomechanics" is the study of solid-state materials that are too large to be manageably described at the atomic level and too small to be studied using the laws of continuum mechanics. As it turns out, one of the methods of studying these nano-sized samples of solid-state materials is to simulate them (i.e., study them with the tools of computer simulation) using hybrid models constructed out of theories from a variety of levels (Nakano et al. 2001). As such, they create models that bear interestingly novel relationships to their theoretical ancestors. So a close look at simulation methods in the nanosciences offers novel insights into the kinds of relationships that exist between different theories (at different levels of description) and between theories and their models.

If we are looking for an example of a simulation model likely to stimulate those sorts of insights, we need look no further than so-called parallel

https://doi.org/10.1086/518328 Published online by Cambridge University Press

multiscale (or sometimes "concurrent coupling of length scales," CLS), methods of simulation. These methods were developed by a group of researchers interested in studying the mechanical properties (how they react to stress, strain, and temperature) of intermediate-sized solid-state materials. The particular case that I will detail below, developed by Farid Abraham and a group of his colleagues, is a pioneering example of this method.[1] What makes the modeling technique "multiscale" is that it couples together the effects described by three different levels of description: quantum mechanics, molecular dynamics, and continuum mechanics.

2. Multiscale Modeling. Modelers of nanoscale solids need to use these multiscale methods—the coupling together of different levels of description—because each individual theoretical framework is inadequate on its own at the scale in question. The traditional theoretical framework for studying the mechanical behavior of solids is continuum mechanics (CM). CM provides a good description of the mechanics of macroscopic solids close to equilibrium. But the theory breaks down under certain conditions. CM, particularly the flavor of CM that is most computationally tractable—linear elastic theory—is no good when the dynamics of the system are too far from equilibrium. This is because linear elastic theory assumes that materials are homogeneous even at the smallest scales, when in fact we know this is far from the truth. When modeling large samples of material, this approximation works, because the sample is large enough that one can effectively average over the inhomogeneities. Linear elastic theory is in effect a statistical theory. But as we get below the micron scale, the fine-grained structure begins to matter more. When the solid of interest becomes smaller than approximately one micron in diameter, this "averaging" fails to be adequate. Small local variations from mean structure, such as material decohesions—an actual tearing of the material—and thermal fluctuations, begin to play a significant role in the system. In sum, CM cannot be the sole theoretical foundation of "nanomechanics"—it is inadequate for studying solids smaller than 1 micrometer in size (Rudd and Broughton 2000).

The ideal theoretical framework for studying the dynamics of solids far from equilibrium is classical molecular dynamics (MD). This is the level at which thermal fluctuations and material decohesions are most naturally described. But computational issues constrain MD simulations to about 10^7–10^8 molecules. In linear dimensions, this corresponds to a constraint of only about 50 nanometers.

1. Good review literature on parallel multiscale simulation methods for nanomechanics can be found in Abraham et al. (1998), Broughton et al. (1999), Rudd and Broughton (2000), and Chelikowski and Ratner (2001).

https://doi.org/10.1086/518328 Published online by Cambridge University Press

So MD methods are too computationally expensive, and CM methods are insufficiently accurate, for studying solids that are on the order of 1 micron in diameter. The good news is that the parts of the solid in which the far-from-equilibrium dynamics take place are usually confined to regions small enough for MD methods. So the idea behind multiscale methods is that a division of labor might be possible—use MD to model the regions where the real action is, and use CM for the surrounding regions, where things remain close enough to equilibrium for CM to be effective.

There is a further complication. When cracks propagate through a solid, this process involves the breaking of chemical bonds. But the breaking of bonds involves the fundamental electronic structure of atomic interaction. So methods from MD (which use a classical model of the energetic interaction between atoms) are unreliable right near the tip of a propagating crack. Building a good model of bond breaking in crack propagation requires a quantum mechanical (QM) approach. Of course, QM modeling methods are orders of magnitude more computationally expensive than MD. In practice, these modeling methods cannot model more than 250 atoms at a time.

The upshot is that it takes three separate theoretical frameworks to model the mechanics of crack propagation in solid structures on the order of 1 micron in size. Multiscale models couple together the three theories by dividing the material to be simulated into three roughly concentric spatial regions. At the center is a very small region of atoms surrounding a crack tip, modeled by the methods of computational QM. In this region, bonds are broken and distorted as the crack tip propagates through the solid. Surrounding this small region is a larger region of atoms modeled by classical MD. In that region, material dislocations evolve and move, and thermal fluctuations play an important role in the dymanics. The far-from-equilibrium dynamics of the MD region is driven by the energetics of the breaking bonds in the inner region. In the outer region, elastic energy is dissipated smoothly and close to equilibrium, on length scales that are well modeled by the linear-elastic, continuum dynamical domain. In turn, it is the stresses and strains applied on the longest scales that drive the propagation of the cracks on the shortest scales.

It is the interactions between the effects on these different scales that lead students of these phenomena to describe them as "*inherently* multiscale" (Broughton et al. 1999, 2391). What they mean by this is that there is significant feedback between the three regions. All of these effects, each one of which is best understood at its own unique scale of description, are strongly coupled together. Since all of these effects interact simultaneously, it means that all three of the different modeling regions have to be coupled together and modeled simultaneously. The fact that three dif-

https://doi.org/10.1086/518328 Published online by Cambridge University Press

ferent theories at three different levels of description need to be employed makes the models "multiscale." The fact that these different regions interact simultaneously, that they are strongly coupled together, means that the models have to be "parallel multiscale."

An instructive way to think about the meaning of the phrase "parallel multiscale" is to compare two different ways of going about integrating different scales of description into one simulation. The first more traditional method is what Abraham's group, in keeping with their computational background, label "serial multiscale." The idea of serial multiscale is to choose a region, simulate it at the lower level of description, summarize the results into a set of parameters digestible by the higher level description, and then pass those results up to a simulation of the higher level.

But serial multiscale methods will not be effective when the different scales are strongly coupled together: "There is a large class of problems for which the physics is inherently multiscale; that is the different scales interact strongly to produce the observed behavior. It is necessary to know what is happening simultaneously in each region since one is strongly coupled to another" (Broughton et al. 1999, 2391). What seems to be required for simulating an inherently multiscale problem is an approach that simulates each region simultaneously, at its appropriate level of description, and then allows each modeling domain to continuously pass relevant information back and forth between regions—in effect, a model that seamlessly combines all three theoretical approaches. Sticking to language borrowed from computer science, Abraham's group refers to this method as "parallel multiscale" modeling. They also refer to it as "concurrent coupling of length scales." What allows the integration of the three theories to be seamless is that they overlap at the boundary between the pairs of regions. These boundary regions are where the different regions "shake hands" with each other. The regions are called the "handshaking regions" and they are governed by "handshaking algorithms." We will see how this works in more detail in the next section.

The use of these handshaking algorithms is one of the things that make these parallel multiscale models interesting. Parallel multiscale modeling, in particular, appears to be a new way to think about the relationship between different levels of description in physics and chemistry. Typically, after all, we tend to think about relationships between levels of description in mereological terms: a higher level of description relates to a lower level of description more or less in the way that the entities discussed in the higher level are made up out of the entities found in the lower level. That kind of relationship, one grounded in mereology, accords well with the relationship that different levels of models bear to each other in what the Abraham group label serial multiscale modeling. But parallel multiscale

https://doi.org/10.1086/518328 Published online by Cambridge University Press

models appear to be a different way of structuring the relationship between different levels of description in physics and chemistry.

I would like to offer a little bit more detail about how these models are put together, and, in particular, to say a bit more about how the handshaking algorithms work—in effect, to illustrate how one seamless model can integrate more than one level of description. To do this, though, I have to first of all say a bit more about how each separate modeling level works. I turn to that in the next section.

3. Three Theoretical Approaches.

3.1. Continuum Mechanics (Linear Elastic Theory). The basic theoretical background for the model of the largest scale regions is linear elastic theory, which relates, in linear fashion, stress—a measure of the quantity of force on a point in the solid—with strain—a measure of the degree to which the solid is deformed from equilibrium at a point. Linear elastic theory, combined with a set of experimentally determined parameters for the specific material under study, enables you to calculate the potential energy stored in a solid as a function of its local deformations. Since linear elastic theory is continuous, in order for it to be used in a computational model it has to be discretized. This is done using a "finite element" method. This technique involves a "mesh" made up of points that effectively tile the entire modeling region with tetrahedra. Each mesh point is associated with a certain amount of displacement—the strain field. At each time step, the total energy of the system is calculated by "integrating" over each tetrahedron. The gradient of this energy function is used to calculate the acceleration of each grid point, which is in turn used to calculate its position for the next time step. And so on.

3.2. Molecular Dynamics. In the medium-scale regions, the basic theoretical background is a classical theory of interatomic forces. The model begins with a lattice of atoms. The forces between the atoms come from a classical potential energy function for silicon proposed by Stillinger and Weber (Stillinger and Weber 1985). The Stillinger-Weber potential is much like the Leonard-Jones potential in that its primary component comes from the energetic interaction of nearest neighbor pairs. But the Stillinger-Weber potential also adds a component to the energy function from every triplet of atoms, proportional to the degree to which the angle formed by each triplet deviates from its equilibrium value. Just as in the finite element case, forces are derived from the gradient of the energy function, which are in turn used to update the position of each atom at each time step.

https://doi.org/10.1086/518328 Published online by Cambridge University Press

3.3. Quantum Mechanics. The very smallest regions of the solid are modeled as a set of atoms whose energetic interaction is governed, not by classical forces, but by a quantum Hamiltonian. The quantum mechanical model they use is based on a semi-empirical method from computation quantum chemistry known as the "Tight Binding" method. It begins with the Born-Oppenheimer approximation. This approximation separates electron motion and nuclear motion and treats the nuclei as basically fixed particles as far the electronic part of the problem is concerned. The next approximation treats each electron as basically separate from the others and confined to its own orbital. The semi-empirical part of the method is to use empirical values for the matrix elements in the Hamiltonian of these orbitals. For example, the model system that Abraham's group has focused on is solid-state silicon. So, the values used for the matrix elements come from a standard reference table for silicon—derived from experiment. Once again, once a Hamiltonian can be written down for the whole system, the motions of the nuclei can be calculated from step to step.

4. Handshaking between Theories. Clearly, these three different modeling methods embody mutually inconsistent frameworks. They each offer fundamentally different descriptions of matter and they each offer fundamentally different mathematical functions describing the energetic interactions among the entities they describe.

The key to building a single coherent model out of these three regions is to find the right handshaking algorithm to pass the information about what is going on in one region that will affect a neighboring region into that neighbor. One of the difficulties that beset earlier attempts to exchange information between different regions in multiscale models was that they failed, badly, to conserve energy. The key to Abraham's success in avoiding this problem is that his group constructs their handshaking algorithms in such as way as to define a single expression for energy for the whole system. "The overarching theme is that a single Hamiltonian is defined for the entire system" (Broughton et al. 1999, 2393). The expression is a function of the positions of the various "entities" in their respective domains, whether they be mesh elements, classical atoms, or the atomic nuclei in the quantum mechanical region.

The best way to think of Abraham's handshaking algorithms then, is as expression that defines the energetic interactions between, for example, the matter in the continuum dynamical region with the matter in the molecular dynamical regions. But this is a strange idea indeed—to define the energetic interactions between regions—since the salient property possessed by the matter in one region is a (strain) field value, while the other is the position of a constituent particle, and in the third it is an electron

https://doi.org/10.1086/518328 Published online by Cambridge University Press

cloud configuration. To understand how this is possible, we have to simply look at the details in each case.

4.1. Handshaking between CM and MD. To understand the CM/MD handshaking algorithm, first envision a plane separating the two regions. Next, recall that in the finite element method of simulating linear elastic theory, the material to be simulated is covered in a mesh that divides it up into tetrahedral regions. One of the original strengths of the finite element method is that the finite element mesh can be varied in size to suit the simulation's needs, allowing the simulationists to vary how fine or coarse the computational grid is in different locations. When the finite element (FE) method is being used in a multiscale model, this feature of the FE mesh becomes especially useful. The first step in defining the handshake region is to ensure that as you approach the plane separating the two domains from the finite element side, the mesh elements of the FE domain are made to coincide with the atoms of the MD domain. (Farther away from the plane, the mesh will typically get much coarser.)

The next step is to calculate the energy of the "handshake region." This is the region between the last mesh point on one side and the first atom on the other. The technique that Abraham's group use is essentially to calculate this energy twice, once from the perspective of FE and once from the perspective of MD, and then average the two. Doing the first of these involves pretending that the first row of atoms are actually mesh elements, doing the second involves the opposite—pretending that the last row of mesh elements are atoms.

Suppose for example that there is an atom on the MD side of the border. It looks over the border and sees a mesh point. For the purpose of the handshaking algorithm, we treat that mesh point as an atom, calculate the energetic interaction according to the Stillinger-Weber potential, and divide it by two (remember, we are going to be averaging together the two energetics.) We do this for every atom/mesh point pair that spans the border. Since the Stillinger-Weber potential also involves triples, we do the same thing for every triple that spans the border (again dividing by two.) This is one half of the "handshaking Hamiltonian." The other half comes from the continuum dynamics' energetics. Whenever a mesh point on the CM side of the border looks over and sees an atom, it pretends that atom is a mesh point. Thus, from that imaginary point of view, there are complete tetrahedra that span the border (some of whose vertices are mesh points that are "really" atoms.) Treating the position of that atom as a mesh point position, the algorithm can calculate the strain in that tetrahedron and integrate over the energy stored in the tetrahedron. Again, since we are averaging together two Hamiltonians, we divide that energy by two.

https://doi.org/10.1086/518328 Published online by Cambridge University Press

We now have a seamless expression for the energy stored in the entire region made up of both the continuous solid and the classical atoms. The gradient of this energy function dictates how both the atoms and the mesh points will move from step to step. In this way, the happenings in the CM region are automatically communicated to the molecular dynamics region, and vice versa.

4.2. Handshaking between MD and QM. The general approach for the handshaking algorithm between the quantum region and the molecular dynamics region is similar: the idea is to create a single Hamiltonian that seamlessly spans the union of the two regions. But in this case, there is an added complication. The difficulty is that the tight binding algorithm does not calculate the energy locally. That is, it doesn't apportion a value for the energy for each interatomic bond; it calculates energy on a global basis. Thus, there is no straightforward way for the handshaking algorithm between the quantum and MD region to calculate an isolated quantum mechanical value for the energetic interaction between an outermost quantum atom and a neighboring innermost MD atom. But it needs to do this in order to average it with the MD value for that energy.

The solution that Abraham and his group have developed to this problem is to employ a trick that allows the algorithm to localize that QM value for the energy. The trick is to employ the convention that at the edge of the QM region, each "dangling bond" is "tied off" with an artificial univalent atom. To do this, each atom location that lies at the edge of the QM region is assigned an atom with a hybrid set of electronic properties. In the case of silicon, what is needed is something like a silicon atom with one valence electron. These atoms, called "silogens," have some of the properties of silicon and some of the properties of hydrogen. They produce a bonding energy with other silicon atoms that is equal to the usual Si-Si bond energy, but they are univalent like a hydrogen atom. This is made possible by the fact that the method is semi-empirical, and so fictitious values for matrix elements can simply be assigned at will. This makes it such that the silogen atoms don't energetically interact with their silogen neighbors, which means that the algorithm can localize their quantum mechanical energetic contributions. Finally, once the problem of localization is solved, the algorithm can assign an energy between atoms that span the threshold between regions that is the average of the Stillinger-Weber potential and the energy from the Hamiltonian in the tight-binding approximation. Again, this creates a seamless expression for energy.

5. Three Questions. In the sequel, I will suggest that there are features of these multiscale models—with their integration of different levels of description, their "handshaking algorithms," and their silogens—that ap-

https://doi.org/10.1086/518328 Published online by Cambridge University Press

pear on their face to be at odds with some basic philosophical intuitions about the relationships between different theories and between theories and their models. But before I begin to draw any philosophical conclusions, I think it is important to note that this area of research—nanomechanics in general and these multiscale methods in particular—is in its relative infancy. And while Abraham and his group have had some success with their models, researchers in these areas are still facing important challenges. It is probably too early to say whether or not this particular method of simulation will turn out, in the great scheme of things, to be the right way to go about predicting and representing the behavior of "intermediate-sized" samples of solid-state materials. Hence, it is probably also too early to be drawing conclusions, methodological or otherwise, from these sorts of examples.

On the other hand, it might *not* be too early to start thinking about what kinds of basic philosophical intuitions about science are likely to come under pressure—or to be informed in novel ways—if and when these scientific domains mature. So we might, at this stage, try to pinpoint some basic philosophical questions, questions whose answers are likely to be influenced by this kind of work. In other words, what I want to do here is simply to offer some ideas about what kinds of questions philosophers are likely to be able to shed light on, prospectively, if they keep an eye on what is going on in nanoscale modeling and simulation—especially with regard to multiscale methods—and to provide a sneak preview of what we might discover as the field progresses. Here are three such questions.

Question 1: What relationships are possible between levels of description? One issue that has received perennial attention from philosophers of science has been that of the relationship between different levels of description. Traditionally, the focus of this inquiry has been debate about whether or not, and to what extent or in what respect, laws or theories at higher levels of description are reducible to those at a lower level.

Underlying all of this debate, I believe, has been a common intuition: the basis for understanding interlevel interaction—to the extent that it is possible—is just applied mereology. In other words, to the extent that the literature in philosophy of science about levels of description has focused on whether and how one level is reducible to another, it has implicitly assumed that the only interesting possible relationships are logical ones—that is, intertheoretic relationships that flow logically from the mereological relationships between the entities posited in the two levels.[2]

2. An important exception is the recent work of Robert Batterman (2002).

https://doi.org/10.1086/518328 Published online by Cambridge University Press

But if methods that are anything like those described above become accepted as successful in nanoscale modeling, that intuition is likely to come under pressure. The reason is that so-called parallel multiscale modeling methods are forced to develop relationships between the different levels that are perhaps suggested, but certainly not logically detemined, by their mereology. Rather, developing the appropriate relationships "requires physical insight" (Broughton et al. 1999).

What this suggests is that there can be a substantial physics of interlevel interaction; a physics which is guided, but by no means determined by either the theories at each level or the mereology of their respective entities. Indeed, whether or not the relationships employed by Abraham and his group will turn out to be the correct ones is an empirical/physical question and not a logical/mereological one.

Question 2: How important is the consistency of a set of laws? This is an issue that has begun to receive attention only recently, particularly in the work of Mathias Frisch (2004). Using classical electrodynamics (CED) as an example, Frisch has challenged a common philosophical intuition about scientific theories: that the internal consistency of its laws is a necessary condition that all successful theories have to satisfy. I want to make a similar point here. In this case, the example of multiscale modeling seems to put pressure on a closely related, if somewhat weaker, intuition: that an inconsistent set of laws can have no models.

In a formal setting, this claim is obviously true; indeed it is true by definition. But rarely in scientific practice do we actually deal with models that have a clear formal relationship to the laws that inspire them. Most likely, the intuition that inconsistent laws cannot produce a coherent model in everyday scientific practice rests as much on pragmatic considerations as it does on the analogy to formal systems: how, in practice, could mutually conflicting sets of laws guide the construction of a coherent and successful model?

We can start by looking at what we learn from Frisch. In CED the strategy is usually to keep the inconsistent subsets of the theory properly segregated for a given model. "The Maxwell-Lorentz equations can be used to treat two types of problem. We can appeal to the Maxwell equations to determine the fields associated with a given charge and current distribution; or we can use the Lorentz force law to calculate the motion of a charged particle in a given external electromagnetic field" (Frisch 2004, 529). In other words, in most models of CED, each respective model draws from only one of the two mutually inconsistent "sides" of the theory. This technique works for most applications, but there are exceptions where the method fails. Models of synchrotron radiation, for example, necessarily involve both mutually inconsistent parts of the theory.

https://doi.org/10.1086/518328 Published online by Cambridge University Press

There are problems, in other words, that require us to calculate the field from the charges as well as to calculate the motion of the charges from the fields. But the solution method, even in the synchrotron case as Frisch describes it, is still a form of segregation. The segregation is temporal. You break the problem up into time steps: in one time step the Lorentz equation are used, in the next, the Maxwell equations, and so on.

A form of segregation is employed in multiscale modeling as well, but it is forced to break down at the boundaries. Each of the three theoretical approaches is confined to its own spatial region of the system. But the fact that there are significant simultaneous and back-and-forth interactions between the physics in each of these regions means that the strategy of segregation cannot be entirely effective. *Parallel* multiscale methods require the modeler to apply, in the handshaking region, two different sets of laws. The laws in Abraham's model, moreover, are each pair-wise inconsistent. They offer conflicting descriptions of matter and conflicting accounts of the energetic interactions between the constituents of that matter. But the construction of the model in the handshaking regions is guided by both members of the pair. When you include the handshaking regions, parallel multiscale models are—all at once—*models of an inconsistent set of laws*.

The methods developed by these researches for overcoming these inconsistencies (the handshaking algorithms) may or may not turn out to be too crude to provide a reliable modeling approach. But by paying close attention to developments in the field of nanoscale modeling, a field in which the models are almost certainly going to be required to involve hybrids of classical, quantum, and continuum mechanics, philosophers are likely to learn a great deal about how inconsistencies are managed. In the process, we will be forced to develop richer accounts of the relationships between theories and their models—richer accounts, in any case, than the one suggested by the analogy to formal systems.

Question 3: How do models differ from ideal descriptions? (What role can falsehoods play in model building?) It has been widely recognized that many successful scientific models do not represent exactly. A simple example: the model of a simple harmonic oscillator can quite successfully predict the behavior of many real physical systems, but it provides at best only an approximately accurate representation of those systems. Nevertheless, many philosophers hold to the intuition that successful models differ from ideal descriptions primarily in that they include idealizations and approximations. Ronald Laymon has made this intuition more precise with the idea of "piecewise improvability" (Laymon 1985). The idea is that while many empirically successful models deviate from ideal descrip-

https://doi.org/10.1086/518328 Published online by Cambridge University Press

tion, a small improvement in the model (that is, a move that brings it closer to an ideal description) should always result in a small improvement in its empirical accuracy.

But what about the inclusion of "silogens" in multiscale models of silicon? Here, piecewise improvability seems to fail. If we make the model "more realistic" by putting in more accurate values for the matrix elements at the periphery of the QM region, then the resulting calculation of the energetic interactions in the handshake region will become less accurate, not more accurate, and the overall simulation will fail to represent accurately at all. The lesson of this and other examples is that models can sometimes successfully make use not only of approximations and idealizations but also outright "falsifications." False assumptions, it appears, can be systematically successful.[3] Nanoscale models, particularly simulation models, are likely to put pressure on the philosophical intuition that success and reliability always come from truth.

REFERENCES

Abraham, F., J. Broughton, N. Bernstein, and E. Kaxiras (1998), "Spanning the Length Scales in Dynamic Simulation," *Computers in Physics* 12 (6): 538–546.

Batterman, R. (2002), *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. New York: Oxford University Press.

Broughton, J., F. Abraham, N. Bernstein, and E. Kaxiras (1999), "Concurrent Coupling of Length Scales: Methodology and Application," *Physical Review B* 60 (4): 2391–2403.

Chelikowski, J., and M. Ratner (2001), "Nanoscience, Nanotechnology, and Modeling," *Computing in Science and Engineering* 3 (4): 40–41.

Frisch, M. (2004), "Inconsistency in Classical Electrodynamics," *Philosophy of Science* 71 (4): 525–549.

Laymon, R. (1985), "Idealization and the Testing of Theories by Experimentation," in P. Achinstein and O. Hannaway (eds.), *Observation, Experiment and Hypothesis in Modern Physical Science*. Cambridge: MIT Press, 147–173.

Nakano, A., et al. (2001), "Multiscale Simulation of Nano-systems," *Computing in Science and Engineering* 3 (4): 56–66.

Rudd, R. E., and J. Q. Broughton (2000), "Concurrent Coupling of Length Scales in Solid State Systems," *Physica Status Solidi B* 217: 251–291.

Stillinger, F. H., and T. A. Weber (1985), "Computer Simulation of Local Order in Condensed Phases of Silicon," *Physical Review B* 31: 5262–5271.

Winsberg, E. (2006), "Models of Success vs. the Success of Models: Reliability without Truth," *Synthese* 152: 1–19.

3. For other examples of "falsifications," as well as for a discussion of their implications for scientific realism, fundamentalism, and the status of "reliability" as a viable semantic notion, see Winsberg (2006).

https://doi.org/10.1086/518328 Published online by Cambridge University Press