Durham University

Insights

Institute *of* Advanced Study

A Function for Fictions: Expanding the Scope of Science

Eric Winsberg

Volume 1 **2008** **Number 8**

ISSN 1756-2074

About Insights

Insights captures the ideas and work-in-progress of the Fellows of the Institute of Advanced Study at Durham University. Up to twenty distinguished and 'fast-track' Fellows reside at the IAS in any academic year. They are world-class scholars who come to Durham to participate in a variety of events around a core inter-disciplinary theme, which changes from year to year. Each theme inspires a new series of Insights, and these are listed in the inside back cover of each issue. These short papers take the form of thought experiments, summaries of research findings, theoretical statements, original reviews, and occasionally more fully worked treatises. Every fellow who visits the IAS is asked to write for this series. The Directors of the IAS – Ash Amin, Michael O'Neill, Susan J. Smith and Colin Bain – also invite submissions from others involved in the themes, events and activities of the IAS.

About the Institute of Advanced Study

The Institute of Advanced Study, launched in October 2006 to commemorate Durham University's 175th Anniversary, is a flagship project reaffirming the value of ideas and the public role of universities. The Institute aims to cultivate new thinking on ideas that might change the world, through unconstrained dialogue between the disciplines as well as interaction between scholars, intellectuals and public figures of world standing from a variety of backgrounds and countries. The Durham IAS is one of only a handful of comparable institutions in the world that incorporates the Sciences, Social Sciences, the Arts and the Humanities.

The focal point of the IAS is a programme of work associated with, but not exclusive to, an annual research theme. At the core of this work lies a prestigious Fellowship programme. This programme gathers together scholars, intellectuals and public figures of world standing or world-promise to address topics of major academic or public interest. Their mission is to anticipate the new and re-interpret the old, communicating across and working between disciplinary boundaries.

Every year, the Institute invites as many as twenty highly creative individuals to spend up to three months in Durham. They are located in Cosin's Hall, a magnificent and spacious 18th century mansion which, together with Durham Cathedral and Durham Castle, forms part of Palace Green, dominating the World Heritage Site of Durham Peninsula. During their stay, Fellows engage with departments and colleges, deliver public lectures and seminars, and, above all, join an international community of researchers to address the theme selected for that year. Further details of the IAS and its Fellowship programme can be found at www.durham.ac.uk/ias/fellows

Copyright

A FUNCTION FOR FICTIONS: EXPANDING THE SCOPE OF SCIENCE

To a first approximation, fictions are representations that do not concern themselves with truth. Science, to be sure, is full of representations. But the representations offered to us by science, or so we are inclined to think, are supposed to aim at truth, (or at least one of its cousins: approximate truth, empirical adequacy, reliability.) If the proper and immediate object of fictions is contrary to the aims of science, what role could there be for fictions in science? This paper will argue for at least one important role for fictions in science, especially in the computationally intensive sciences of complex physical systems – in computer simulation. Fictions, I will argue, are sometimes needed for extending the useful scope of theories and model-building frameworks beyond the limits of their traditional domains of application. One especially interesting way in which they do this is by helping to enable model builders to sew together incompatible theories and apply them in contexts in which neither theory by itself will do the job.

Fictions vs. Models

The history of discussions of fictions in the philosophy of science goes back at least to Hans Vaihinger's famous book, *The Philosophy of 'As If'* (Vaihinger, 1911). On Vaihinger's view, science is full of fictions and any representation that contradicts reality is a fiction; or, at least, a semi-fiction.[1] (A full fiction, according to Vaihinger, is something that contradicts itself.) And so, on this view, most ordinary models in science are a kind of fiction.

'Fictional,' however, is not the same thing as 'inexact' or 'not exactly truthful.' Not everything, I would argue, that diverges from reality, or from our best accounts of reality, is a fiction. Many of the books to be found in the non-fiction section of your local bookstore will contain claims that are inexact, or even false. But we do not, in response, ask to have them reshelved in the fiction section. An article in this morning's newspaper might make a claim, 'The green zone is a 10km^2 circular area in the center of Baghdad,' that is best seen as an idealization. And though I live at the end of a T-intersection 'Google maps' show the adjacent street continuing on to run through my home. Still, none of these is a fiction.

I take as a starting point, therefore, the assumption that we ought to count as non-fictional many representations in science that fail to represent exactly; even representations that in fact contradict what our best science tells us to be the case about the world. Many of these kinds of representations are best captured by our ordinary use of the word 'model.' The frictionless plane, the simple pendulum, and the point particle all serve as good representations of real systems for a wide variety of purposes. All of them, at the same time, fail to represent exactly the systems they purport to represent, or, for that matter, any known part of the world. They all incorporate false assumptions and idealizations.

But, contra-Vaihinger, I urge that, because of their function (in ordinary contexts), we continue to call these sorts of representations 'models', and resist calling them fictions. It would seem to me to be simply wrong to say that ordinary models in science are not concerned with truth or any of its cousins. In sum, we do not want to get carried away. We do not want all (or almost all) of the representations in science, on maps, and in journalism, etc. to count as fictions. To do so risks not only giving a misleading overall picture of science ('all of science

is fiction!') but also of weakening to the point of emptiness a useful dichotomy between different kinds of representations – the fictional and the non-fictional. If only representations that are exact are non-fictions, then even the most ardent scientific realist will have to admit that there are precious few non-fictional representations in the world.

Fictions, then, are rarer in science than are models. That is because most models in science aim at truth or one of its cousins. But fictions do not. It might seem then, that on this more narrow conception of what it is for a representation to be a fiction, there will probably turn out to be no fictions in science. Part of my goal, then, is to show that there are. To do this, I will have to define my more limited conception of what it is to be a fiction. This will involve being clearer about what it means to 'aim at truth or one of its cousins.'

So, how should we proceed in demarcating the boundary between fictions and non-fictions? The salient difference between a fictional and a non-fictional representation, it seems to me, rests with the proper *function* of the representation. Indeed, I argue that I should count any representation – even one that misrepresents certain features of the world (as most models do) – as a non-fiction if we offer it *for* the sort of purpose for which we ordinarily offer non-fictional representations.

I offer, in other words, a pragmatic, rather than a correspondence, conception of fictionality. What, then, is the ordinary function of non-fictional representations? I would suggest that, under normal circumstances, when we offer a non-fictional representation, we offer it for the purpose of being a 'good enough' guide to the way some part of the world is, for a particular purpose. Here, 'good enough' implies that the model is accountable to the world, (in a way that fictions are not) in the context of that purpose. On my account, to hold out a representation as a non-fiction is ipso facto to offer it for a particular purpose and to *promise* that for that purpose, the model 'will not let you down,' when it comes to offering guidance about the world for that purpose. In short, non-fictional representations promise to be *reliable* in a certain sort of way.

But not just in any way. Consider an obvious fiction: the fable of the grasshopper and the ant. This, you may recall, is the story of the grasshopper who sings and dances all summer while the ant toils at collecting and storing food for the coming winter. When the winter comes, the ant is well prepared, and the grasshopper, who is about to starve, begs for his charity. This is what we might call a didactic fiction. The primary function of the fable, one can assume, is to offer us important lessons about the way the world is and how we should conduct ourselves in it. For the purpose of teaching children about the importance of hard work, planning and the dangers of living for today, it is a reasonably reliable guide to certain features of the world (or so one might think).

So why is it a fiction? What is the difference, just for example, between didactic fictions and non-fictions if both of them can serve the purpose of being reliable guides? To answer this, I think we need to consider the representational targets of representations. The fable of the grasshopper and the ant depicts a particular state of affairs. It depicts a land where grasshoppers sing and dance, insects talk, grasshoppers seek charity from ants, etc. If you read the fable incorrectly, as a non-fiction, you will think the fable is describing some part of the world – its representational target. And if you read it as a non-fiction, then you will think it is meant to be a reliable guide to the way this part of the world – this little bit of countryside where the grasshopper and the ant live – is. In short, the fable is a useful guide to the way the world is in some general sense, but it is not a guide to the way its prima facie representational target is. And that is what makes it, despite its didactic function, a fiction.

Non-fictions, in other words, are not just reliable guides to the way the world is in any old way. They describe *and point to a certain part of the world* and say 'if you want to know *about that part of the world I am pointing to*, for a certain sort of purpose, I promise to help you in that respect and not let you down.' The importance of this point about the prima facie representation targets of representations will become clear later.

Fictional representations, on the other hand, are not thought to be good enough guides in this way. They are offered with no *promises of a broad domain of reliability*. Unlike most models in science, fictions do not come stamped with promissory notes that say something like 'in these respects and to this degree of accuracy (those required for a particular purpose), some domain of the world – the domain that I purport to point to – is like me, or will behave like me.'

So, crucial to understanding the difference between fictional and non-fictional representations is to understand the different functions that they are intended to serve. But intended by whom? A brief note about intentionality is in order: It probably sounds, from the above, as though in my view whether or not a representation counts as a fiction depends on the intention of the author of the representation. On such a view, if the author *intends* for the representation to carry with it such a promissory note, then the representation is non-fictional. This is close to my view. But I prefer to distinguish fictions from non-fictions without reference to the intention of the author. If we find, someday, the secret diaries of James Watson and Francis Crick, and these reveal to us that they intended the double helix model as a planned staircase for a country estate, this has no bearing on the proper function of the model. On my view, what a representation is for depends not on the intention of the author, but on the community's norms of correct use.

Consider the famous tapestry, hanging in the Cloisters museum in New York, the 'Unicorn in Captivity'. This is a good example of a representation. Quite possibly, it was the intention of the author that the tapestry be taken as a non-fictional representation belonging to natural history. More importantly, his contemporaries probably believed in unicorns. Had there been museums of natural history at the time, this might have been where the tapestry would have hung. But today, the tapestry clearly belongs in a museum of art. In deciding whether the tapestry is a fiction or a non-fiction does not depend on the intention of the author, it depends on what, based on the community's norms, we can correctly take its function to be – on what kind of museum the community is likely to display it in. Once upon a time, that role might have been in providing a guide to the animal kingdom. But it no longer is.

Fictions in Science

It is clear that science is full of representations that are inexact, and of representations that contradict what we know or believe to be the case about the world. It is clear, in other words, that science is full of models. But are there fictions in science? If a representation is not even meant to serve as a reliable guide to the way that the world is, if it is a fiction, would it not necessarily fall outside of the enterprise of science?

Despite the rather liberal constraints I want to impose on being a non-fiction, I believe there are a variety of roles that fictions can play in science. I want to outline one such role here. To do so, I turn to two examples of computer simulations in the physical sciences where, I argue, fictions are employed. The first example comes from simulation methods in a field called 'nano-mechanics,' and the second comes from computational fluid dynamics. In both

of these examples, model builders add fictions to their models in order to extend the useful scope of the model-building frameworks they employ beyond the limits of their traditional domains of application.

The first example comes from a set of methods in nano-mechanics called 'parallel multiscale' simulations, (or sometimes 'concurrent coupling of length scales').[2] 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 bits of solid state silicon. What makes these modeling techniques 'multiscale' is that they seamlessly couple together the effects described by three different levels of description: quantum mechanics, molecular dynamics, and continuum mechanics. By 'sewing' together these three model-building frameworks, model builders are able to go beyond the useful scope of any one of the three alone.

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. It is an idealization. When modeling large samples of material, this idealization 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, 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 'nano-mechanics' – it is inadequate for studying solids smaller than one 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 fifty nanometers.

So MD methods are too computationally expensive, and CM methods are insufficiently accurate, for studying solids that are on the order of one micron in diameter. On the other hand, 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 – using MD to model the regions where the real action is and 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, it 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 two hundred and fifty 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 one 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. Here, 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. This is where material dislocations evolve and move, and thermal fluctuations play an important role in the dynamics. 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.

Researchers describe these sorts of phenomena as 'inherently multiscale.' This means that what is required for simulating these phenomena is an approach that simulates each region simultaneously, at its appropriate level of description, and then allows each modeling domain to pass relevant information continuously back and forth between regions – in effect, a model that seamlessly combines all three theoretical approaches. 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.'

The basic idea behind the handshaking algorithms is rather simple in principle, but becomes rather complicated in practice. Each of the three basic regions has its own computational scheme based on a function that gives the energy of interaction between the simplest 'elements' of the scheme. In the CM region, the elements are the vertices of the grid over which the continuous equations have been discretized. The energy function comes from the elastic forces. In the MD region, the elements are molecules, and the energy function comes from a classical force function. In the QM region, the elements are atoms, and the energy function is a quantum Hamiltonian. What the handshaking algorithms do is to match up the elements from the neighboring regions and, when calculating the energetic interaction of, say, a vertice and an atom, take the average of the value that each of the two different schemes provides.

Suppose, for example, that there is an atom on the MD side of the border. It looks over the border and sees a grid point. For the purpose of the handshaking algorithm, we treat that mesh point as an atom, calculate the energetic interaction according to the classical 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. 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.

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 algorithm used to calculate the energy in the QM region does not calculate the energy locally. That is, it does not apportion a value for the energy for each inter-atomic 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 researchers 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 values for matrix elements can simply be assigned at will and put in 'by hand.' This makes it such that the silogen atoms do not 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 spans 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.

Silogen atoms, I want to argue, are fictions. To see this, we need to look at their function. But we need to be careful. If we view the overall model that drives the simulation as a whole, it is clearly non-fictional. The representational target of the Abraham model is micron-sized pieces of silicon and this model is meant to be a reliable guide to the way that such pieces of silicon behave. To endorse such a model in the relevant respect is *to promise that* for the purpose of designing nano-electromechanical systems (NEMS), *the model will not let you down*. It will be a good enough guide to the way these pieces of silicon behave, accurate to the degree, and in that respect, necessary for NEMS design. Though it contradicts reality in more ways than we could probably even count, Abraham's model is a non-fiction.

But within the simulation, we can identify components of the model that, prima facie, play their own local representational role. Each point in the QM region appears to represent an individual atom. Each point in the MD region, and each tetrahedron in the CM region, has an identifiable representational target. Some of these points, however – the ones that live in the QM/MD handshaking region – are special. These points, which represent their representation targets as silogen atoms, do not function to be reliable guides to the way those atoms behave. Silogens are used *for* tying off the energy function at the edge of the QM region. They are used *for* getting the whole simulation to work properly, not for depicting the behavior of the particular atom being modeled. We are deliberately getting things wrong locally so that we get things right globally. The silogen atoms are fictional entities that 'smooth over' the

inconsistencies between the different model-building frameworks, and extend their scope to domains where they would individually otherwise fail. In a very loose sense, we can think of them as being similar to the fable of the grasshopper and the ant. In the overall scheme of things, their grand purpose is to inform us about the world. But if we read off from them their prima facie representational targets, we are pointed to a particular domain of the world about which they make no promises to be reliable guides.

The second example comes from a method in computational fluid dynamics (CFD) used to model the supersonic flow of fluids with strong shocks.[3] Shock waves are notoriously difficult to simulate numerically. The reason is that the computational models involved use spatial 'grids,' and an abrupt shock will fall between the grid lines. This, in turn, makes the energy of the simulation blow up. John von Neumann came up with the idea of trying to blur the shock over a few grid cells. He did this by making up an 'artificial viscosity' for the fluid. The trick is to treat the fluid as extremely viscous right where the shocks are. This makes the shocks thicker than a few grid cells, and the simulation will no longer blow up.

Many simulations in CFD employ methods of discretization. This begins with a set of mathematical equations that depict the time-evolution of the system being studied in terms of rules of evolution for the variables of the model. In the case of CFD, the equations in question are a version of either the Euler or Navier-Stokes equations, depending on the factors to be included, and the coordinate system to be employed. To study a particular flow problem, it would be desirable to find solutions to the equations of those models. With most of the systems we encounter in fluid dynamics, however, the models that are suggested directly by theory consist of second-order, non-linear differential equations. It is often impossible, even in principle, to find closed-form solutions to these equations.

What a simulationist must do, therefore, is to find a replacement for the model hatched out of the theory – one that consists of a set of equations that can be iterated numerically, using step by step methods. When a theoretically motivated model is thus 'discretized,' and turned into a simulation model, the original differential equations are transformed into difference equations, and crafted into a computable algorithm that can be run on a digital computer. The finite 'difference' between the values of spatial and temporal variables is sometimes called the 'grid size' of the simulation, since one can imagine a spatio-temporal grid on which the simulation lives.

One of the earliest uses of finite difference simulations arose in connection with the Manhattan Project during World War II. John von Neumann and his group of researchers used finite difference computations to study the propagation and interaction of shock waves in a fluid, a subject crucial to the success of the atomic bomb.

We generally think of shock waves as abrupt discontinuities in one of the variables describing the fluid, but it was quickly recognized that treating them in this way would cause problems for any numerical solution. The reason is that a shock wave is not a true physical discontinuity, but a very narrow transition zone whose thickness is on the order of a few molecular mean-free paths. Even with today's high speed and high memory computers, calculating fluid flow with a differencing scheme that is fine enough to resolve this narrow transition zone is wildly impractical. On the other hand, it is well known that a simulation of supersonic fluid flow that does not deal with this problem will develop unphysical and unstable oscillations in the flow around the shocks. These oscillations occur because of the inability of the basic computational method to deal with the discontinuities associated with a shock wave – the higher the shock speed, the greater the amplitude of these oscillations becomes. At very high

speeds, such a simulation quickly becomes useless. To make it more useful and accurate, what simulationists need to do is somehow to dampen out these oscillations.

The generally accepted way to do this, originally devised by von Neumann and Richtmyer, while working at Los Alamos, is to introduce a new term, an 'unphysically large value of viscosity,' into the simulation, which is called 'artificial viscosity.' The inclusion of this term in the simulation is designed to widen the shock front, and blur the discontinuity over a thickness of two or three grid zones. This enables the computational model to calculate certain crucial effects that would otherwise be lost inside one grid cell; in particular, the dissipation of kinetic energy into heat (von Neumann and Richtmyer, 1950).

A standard equation for the value of artificial viscosity at a particular point in a simulation is given below:

$$Q = \begin{cases} l^2(\Delta x)^2 \, \rho(\frac{\delta v}{\delta x})^2 & (\frac{\delta v}{\delta x}) < 0 \\ 0 & \text{otherwise} \end{cases}$$

where *l* is a dimensionless constant, *delta x* is the grid size, *rho* is the density, and *v* is the local value of the velocity. This makes the value of viscosity proportional to the local vorticity of the fluid, and guarantees that it will be extremely high only in the local presence of a shock.

The high viscosity attributed to the fluid near the shocks is unrealistic. It contradicts our best accounts of what the world is actually like at the level of description in which fluids operate. In some respects, it is like the frictionlessness of an inclined plane, in that, just as it is literally false of any actual plane in the world that it is frictionless, it is literally false of any fluid that it has a viscosity that varies in proportion to its local vorticity. In that respect, von Neumann-Richtmyer models are not fictions. They are just models. They are representations that contain substantial inaccuracies, but that for certain purposes (including the purpose, unfortunately, of designing thermonuclear weapons), are 'good enough' guides to the way their representational targets (the entire fluid system in question) behave.

But models of fluids with shocks are unusual creatures in certain respects. Ever since Ludwig Prandtl introduced the idea of an aerodynamic boundary layer, where the flow of a fluid is divided into two flow fields, one inside the layer, where viscosity matters most, and one in the rest of the fluid, where viscosity can be ignored, fluid dynamicists have become accustomed to the idea that models of fluids are compartmentalized. And it is traditional to think of shocks in the same way. The model of a fluid with shocks is composed of a sub-model of the shock, and a sub-model of the rest of the fluid.

If we think of models of fluids in this way, then the model of the shock inside the von Neumann-Richtmyer (vNR) model is a fiction. No one thinks: 'treat the fluid inside a shock as having an enormous value for viscosity and you will not go wrong in understanding the behavior of the shock.' In fact, if you look at the behavior of the fluid in the regions where you have applied the high viscosity, you are getting things all wrong – deliberately so: you have virtually eliminated the existence of abrupt shocks. This is not a formula for gaining understanding of shocks. It is a formula for allowing the overall simulation to survive the computational disasters that shocks usually bring about. The hope is that it is okay. The hope is that by getting things dramatically wrong near where the shocks will be, you will prevent energy blowups in your algorithm, and you will get things reasonably right overall. If we think of the vNR method as something that we use to build models of fluids, it is a non-fiction, and artificial viscosity is a non-fictional property of the fluid. However if we think of the *models*

of shocks that we build with the vNR method, these are fictions, and the high viscosity is a fictional characteristic of the fluid in the shock.

So, the function of artificial viscosity is not to be a reliable guide (even in certain limited respects and to some limited degree of accuracy) to the way fluids in shocks behave. What silogen atoms and artificial viscosity have in common (though one is an entity, the other a property) is just that: they both point at some part of the world, but neither is for getting the behavior of that part of the world right. They add to the overall reliability of the simulation by sacrificing the reliability of the simulation in getting things right about the part that they represent.

Despite the apparent conflict of the function of fictions with the aims of science, fictions, it would seem, play at least one important role in science. They help us to sew together inconsistent model-building frameworks, and to extend those frameworks beyond their traditional limits.

Notes

[1] My understanding of Vaihinger comes almost entirely from Fine (1993).

[2] My discussion of parallel multiscale simulations draws from several sources. Included are Abraham et al. (1998), Broughton et al. (1999), Chelikowski and Ratner (2001), Nakano et al. (2001), and Rudd and Broughton (2000). For more details, see Winsberg (2006b).

[3] My discussion of parallel multiscale simulations draws from the following sources: Campbell (2000), von Neumann and Richtmyer (1950), and Caramana et al. (1988). For more details, see Winsberg (2006a).

Reference List

Abraham, F. et al. (1998) Spanning the length scales in dynamic simulation. *Computers in Physics* 12:6: 538-46.

Broughton, J. et al. (1999) Concurrent coupling of length scales: methodology and application. *Physical Review* B 60:4: 2391-403.

Campbell, J. (2000) Artificial Viscosity for Multi-dimensional Hydrodynamics Codes. http://cnls.lanl.gov/Highlights/2000-09/article.htm

Caramana, E. J., Shashkov M. J. and Whalen P. P. (1988) Formulations of artificial viscosity for multi-dimensional shock wave computations. *Journal of Computational Physics* 144: 70–97.

Chelikowski, J. and Ratner, M. (2001) Nanoscience, nanotechnology, and modeling. *Computing in Science and Engineering* 3:4: 40-1.

Fine, A. (1993) Fictionalism. *Midwest Studies in Philosophy* 18: 1-18.

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

Rudd, R. E. and Broughton, J. Q. (2000) Concurrent coupling of length scales in solid state systems. *Physica Status Solidi*. B 217: 251–91.

Vaihinger, H. (1911) Die Philosophie des Als Ob. Berlin: Verlag von Reuther & Reichard. Published in English as *The Philosophy of 'As If'*. (Translated by Ogden, C.K.) London: Kegan Paul, 1923.

von Neumann, J. and Richtmyer, R. D. (1950) A method for the numerical calculation of hydrodynamical shocks. *Journal of Applied Physics* 21: 232–47.

Winsberg, E. (2006a) Models of success vs. the success of models: reliability without truth. *Synthese* 152: 1-19.

- - - . (2006b) Handshaking your way to the top: simulation at the nanoscale. *Philosophy of Science* 73: 582–94.

Backlist of Papers Published in Insights

2008 Volume 1

No.	Author	Title	Series
1	Boris Wiseman	Lévi-Strauss, Caduveo Body Painting and the Readymade: Thinking Borderlines	General
2	John Hedley Brooke	Can Scientific Discovery be a Religious Experience?	Darwin's Legacy
3	Bryan R. Cullen	Rapid and Ongoing Darwinian Selection of the Human Genome	Darwin's Legacy
4	Penelope Deutscher	Women, Animality, Immunity – and the Slave of the Slave	Darwin's Legacy
5	Martin Harwit	The Growth of Astrophysical Understanding	Modelling
6	Donald MacKenzie	Making Things the Same: Gases, Emission Rights and the Politics of Carbon Markets	Modelling
7	Lorraine Code	Thinking Ecologically about Biology	Darwin's Legacy

Insights

Insights is edited by Susan J. Smith, IAS Director and Professor of Geography.
Correspondence should be directed to Audrey Bowron (a.e.bowron@durham.ac.uk).