

COMPUTER SIMULATIONS IN METAPHYSICS: POSSIBILITIES AND LIMITATIONS¹

BILLY WHEELER

<https://orcid.org/0000-0001-7017-2674>

*Sun Yat-Sen University
Department of Philosophy
Zhuhai, Guangzhou
China*

billy.wheeler@cantab.net

Article info

CDD: 110

Received: 05.07.2019; Revised: 13.09.2019; Accepted: 13.09.2019

DOI: <http://dx.doi.org/10.1590/0100-6045.2019.V42N3.BW>

Keywords

Computer modeling

Computer simulation

Methods in metaphysics

Humean Supervenience

Nomic Necessity

Abstract: Computer models and simulations have provided enormous benefits to researchers in the natural and social sciences, as well as many areas of philosophy. However, to date, there has been little attempt to use computer models in the development and evaluation of metaphysical theories. This is a shame, as there are good reasons for believing that metaphysics could benefit just as

¹ Earlier versions of this paper were presented at the Forum on Philosophical Methods at Sun Yat-Sen University, Zhuhai as well as the 8th Asia-Pacific Conference on Philosophy of Science at Fudan University, Shanghai. I thank the participants for their helpful comments. I am also grateful to two anonymous reviewers from *Manuscripto* who also provided invaluable feedback on an earlier version of this paper.

much from this practice as other disciplines. In this paper I assess the possibilities and limitations of using computer models in metaphysics. I outline the way in which different kinds of model could be useful for different areas of metaphysics, and I illustrate in more detail how agent-based models specifically could be used to model two well-known theories of laws: David Lewis's "Best System Account" and David Armstrong's "Nomic Necessitation" view. Some logically possible processes cannot be simulated on a standard computing device. I finish by assessing how much of a threat this is to the prospect of metaphysical modeling in general.

1 INTRODUCTION

Philosophers have been aware of the importance of models in the development and evaluation of scientific theories for some time. One important type of modeling in science involves computer simulation. Simulations are believed to have many advantages over physical and conceptual models. Firstly, by using the memory and processing power of a digital computer, they can perform calculations much faster than the human mind. As a result, more complex physical phenomena (such as those that have continuous rates of change or multiple variables) can be modeled. Secondly, the output of the data that is processed by a computer can be displayed using different visual media. Often this is in graphical form, but other more "realistic" 3D representations can be provided with a desktop monitor—and increasingly—by the use of a virtual reality headset. This provides heuristic benefits because it allows scientists to "see" the implications of their theory—revealing unforeseen consequences and possibly new emergent phenomena (Gould, Tobochnik & Christian 2006, 5).

Philosophers too have begun to realize what a powerful tool computer simulation can be in the development and

evaluation of their theories. Computer models and simulations have already been used in a number of different disciplines within philosophy, including: philosophy of science (Kevin Zollman 2007, 2010; Weisberg and Muldoon, 2009; De Langhe, 2014; Alexander, Himmelreich & Thompson, 2015), political philosophy (Über & Hartmann 2016; Beisbart & Hartmann 2011; Hahn, von Sydow & Merdes, 2019); social epistemology (Christoph Merdes, 2018; Hong & Page, 2004) and rational choice theory (Klein, Marx and Scheller, 2019).

Many of the computer modeling techniques used so far in philosophy have tended to be applied to social phenomena, such as group agency and rationality. But this leaves many corners of philosophy unexplored. This raises the question: can computer simulations be utilized elsewhere in philosophy in order to help develop and evaluate theories? There are many areas of philosophical research that have little to do with social processes. One obvious area is metaphysics which aims to describe the most basic constituents of reality. Theories of time, necessity, events, causation, lawhood, properties, and the mind-body problem—for example—might fruitfully be modeled on a computer. Can we expect the same benefits of simulations here that have been found elsewhere in philosophy?

This paper will attempt to answer this question by exploring the different methods that are available for modeling metaphysical objects and processes on a computer. I will discuss how rival metaphysical models ought to be evaluated as well as the limitations of using computers to model metaphysical phenomena.

The rest of the paper will have the following structure. I start in section 2 by looking more generally at the process of how computer models and simulations are created from scientific theories and what analogous methods can be followed to produce models from metaphysical theories.

Here I identify five different types of modeling practice that are relevant, two of which seem to be unique to metaphysical modeling. In section 3 I provide an illustration of the steps taken in creating a computer model of a metaphysical phenomenon. This is carried out using two well-known theories of lawhood: David Lewis's best systems account and David Armstrong's theory of nomic necessitation. Section 4 discusses two forms of computer model evaluation that are frequently found in the literature: validation and verification, and shows how these can also be applied to metaphysical models. Finally, in section 5, I discuss whether or not the non-computability of some logical processes poses a challenge to the use of models in metaphysics, given that most metaphysical theories are not limited in this way.

2 PROGRAMMING FOR METAPHYSICIANS

2.1. *The Simulation Process*

Computer models and simulations in science are as old as the computer itself and were arguably one of the main reasons for the rapid rise of computing technologies in the 1940s and 1950s (Humphreys 2004, 49). Despite this, there is no straightforward path from theory to simulation. The process is highly creative and depends on a good deal of background knowledge and expertise from the scientist in question. That said, it is possible to make some general remarks concerning the core features that simulations share. A useful conception has been provided by Eric Winsberg (2010, 10). His 5-stage account of computer simulation can be summarized using the following flowchart:

Theory → Model → Treatment → Solver → Results

Let me say a bit more about each of these stages. Most computer simulations or models begin with a scientific *theory*. I say most, because there are some exceptions, such as James Conway's *Game of Life* (Gardner, 1970) which are mainly exploratory in nature. Suppose the phenomenon to be simulated is the motion of a pendulum. Then the underlying theory will consist of Newton's laws of motion and gravity. Next comes the building of the *model*. Here Winsberg is not using the term "model" to refer to the computer model itself, but to some set of equations and initial conditions that satisfy the theory (2010, 10). In the case of the pendulum, the model will consist of a set of differential equations and an abstract description of the system such as "a point mass suspended from a massless string". The third stage, *treatment*, involves assigning initial values to the pendulum, such as its angle of displacement, the force due to gravity, any dampening forces, etc. The next stage is arguably the most difficult and requires the scientist to design a computer program that *solves* the equations for the initial conditions and subsequent states of the system. Here some level of accuracy is pre-determined and approximations in calculations are made in order for the results to be computable in a finite amount of time. Finally, the *results* of the calculations are recorded and displayed numerically or by a 3D virtual representation of the target physical system.

Although only a brief characterization, Winsberg's 5-stage account provides a good starting point for developing computer simulations of metaphysical phenomena.² We

² It might be said that much model construction in scientific practice starts with a real-world target system, and the model is built from this rather than from theory (see e.g. Morgan & Morrison, 1999). I agree this could be a fruitful alternative way of constructing computer models in metaphysics as well. However, in this paper I will follow closely Winsberg's 5-stage account that starts with theory, since the aim here is not to actually build

begin—as in the scientific case—with a theory, such as of time, properties, laws of nature, mental states, etc. Then an abstract model is created that satisfies the main components of the theory. It might be reasoned that this aspect of the process finds little counterpart in academic metaphysics which, for the most part, proceeds by conceptual analysis and definition by necessary and sufficient conditions. But this is not entirely true. It has been argued by Peter Godfrey-Smith (2006, 2012), Laurie Paul (2012) and Timothy Williamson (2018) that much metaphysical work carried out involves a form of model-building. Take as an example Lewis's (1983, 1986a and 1994) model of Humean Supervenience—the so-called "Humean Mosaic" of localized properties at space-time points spread over actual and possible worlds. Or Armstrong's (1983) model of lawhood as "N(F, G)" consisting of a relation of necessity N holding between two universals F and G. Other examples that might be given include Nancy Cartwright's (1999) nomological machines; John McTaggart's (1908) A and B time series; the so-called "Garden of Forking Paths" model of free will (McKenna & Coates, 2015); and David Chalmers's (1996) philosophical zombies. All these seem to count as examples of conceptual models which, could in principle, be modeled using a computer simulation.

The treatment and solver stages of the simulation process need to be implemented somewhat differently for metaphysical theories. Metaphysical theories, on the whole, do not contain algebraic formulae: instead they use sets of necessary and sufficient conditions or some other set of acceptability criteria. In terms of assigning initial values, the job of the simulator will be to create a computer program

computer models, but illustrate how they might be built given familiar metaphysical theories.

that satisfies the logical implications of the theory as best it can. The computer model will need to be functionally equivalent (with regard to the necessary and sufficient conditions, acceptability criteria, etc.) as much as possible with the conceptual model associated with the theory.

In section 4 I will illustrate in more detail what I mean by "functional equivalence", but for now let us return to Lewis's example of the mosaic. Here an agent-based model looks best suited to preserve many of the relations inherent in this idea. Let us set the initial state of the model to consist of a grid of n squares or "tiles". Each square represents a single space-time point. The cell, as in typical examples of cellular automata, can be in one of many different states. These states represent the "perfectly natural properties" that for Lewis can be instantiated at a space-time point (1983, 345). If we so wish, we can also include possible worlds as alternative arrays with different properties instantiated at counterpart space-time points. (This is necessary, for example, if we want to include Lewis's theory of counterfactuals, natural properties and causation.)

Finally, the results of the theory or simulation can be displayed in a way which is methodologically useful for the metaphysician. In what ways might a computer model or simulation be useful? Just like computer simulations in science, metaphysical simulations can be useful for creating instances of a theory that are more complex than the human mind can imagine or process. This might then reveal hidden structures or implications of the theory previously unforeseen. There are also evaluative benefits that can be had from computer simulations. One way this is likely to play out will be in terms of assessing the logical consistency of different parts of a single theory. Once one has a computer model it is possible to perform "metaphysical tests" on it by changing initial values or running additional programs. For example, one can add a program to the model of Lewis's

Humean mosaic to see if the laws which emerge really are axioms in a "best system" balancing simplicity and strength. (I will come back to this example in more detail in sections 3 and 4).

2. 2. *Types of Simulation*

We have seen that the stages of developing a computer model or simulation in metaphysics could follow a similar pattern to that found in science. Scientists use different types of simulation to model natural phenomena, depending on the kind of theory they are using and the aspect of the world they want to represent. In this section, I will outline five different types of simulation and explain the ways in which they can be used for building metaphysical models. What follows is not meant to be an exhaustive list, but captures the most likely ways computer simulations could be of benefit to metaphysicians.

(i) Equation-Based Models

The most common type of computer simulation found in scientific investigation is an *equation-based* model. This is typically constructed on the basis of a well-known, previously established formula for some physical system. The role of these is to model the dynamical nature of a system, and frequently involves multiple parameters with varying rates of change. Because they involve differential equations of one type or another, the treatment and solver stages use approximation techniques. One important method is "discretization", where instead of calculating values of physical quantities over the entire range of real numbers, the solution space is fixed to a finite range of values (Wolfson and Pert 1999, 27). The degree of approximation

and accuracy that is made can depend on a number of different factors, e.g. on how many resources are at the modeler's disposal, the greater the computing power one wants to spend, the less discretized the solutions, and the more accurate the final representation that is required.

It is difficult to tell whether this method of simulation will have much use in metaphysics.³ Most metaphysical theories, even when they are about physical phenomena, do not express their central ideas using numerical formulae. One way in which equation-based models could be of benefit to metaphysicians is in the validation of their theories (see section 4 for more on the process of validation). Most metaphysicians, even those who do not strictly identify as a "metaphysical naturalist", would adhere to the principle that a metaphysical theory that is consistent with our best current scientific theories is preferable to one that is not. This suggests a role for equation-based models in the evaluation of metaphysical theories. If an equation-based simulation of a physical phenomenon can be run alongside the program for some metaphysical theory, then it shows that the two are logically consistent with each other.

(ii) Agent-Based Models

As mentioned in the introduction, agent-based models (also known as "particle methods" and "atomic methods") form the majority of computer simulations currently used by

³ In some way many of the models created will have some form of an equation or formula in them. However, what I mean by an equation-based model here is one which takes a previously established set of physical formulas such as Newton's laws of motion and gravity. It is best to think of all models as being essentially rule-based or algorithm-based in the more broad sense of requiring formulas in their construction.

philosophers. They are typically the preferred choice of model for scientists studying social phenomena, but they have also been used to model atomic and molecular interactions (Than & Büttgenbach, 1994), galaxy formation (Farouki & Shapiro, 1980) and a range of meteorological phenomena (McGuffie & Henderson-Sellers, 2005). Their versatility makes them suitable for metaphysical modeling too. As suggested in the case of Humean Supervenience, they provide a way to model the properties of localized individuals and events. However, this example of an agent-based model fails to make the most of them, because their true value comes when they evolve over time according to a program. Here we see they have the potential to model a number of theories in metaphysics that involve time-dependent evolution: events, personal identity, lawhood, causation and free will—are all topics which could fruitfully be modeled using this variety of simulation.

(iii) *Monte Carlo Models*

Monte Carlo methods of modeling are used by scientists to mimic random processes or to quickly calculate an approximation to an unknown quantity. Not everybody agrees that these models should be considered a separate form of simulation in themselves (see Grüne-Yanoff & Weirich 2010, 30). Whilst they can be used to simulate *indeterministic* processes—such as quantum decay (Gordon & Gordon, 2007)—for the most part they are used as a method of approximation for *deterministic* processes, where one or more values of a physical quantity are either unknown or too difficult to compute (Woolfson & Pert 1999, 24-27).

Because Monte Carlo methods can be used to model indeterministic processes, this suggests an important role for them in simulating a range of metaphysical objects and processes. For example, theories of consciousness and

mental states which make an appeal to indeterministic quantum phenomena, such as those given by Penrose (1989) or Chalmers (1995), might make use of Monte Carlo methods to simulate the emergence of mental states. Theories of free will which account for incompatibilist freedom using indeterminism as a fundamental process in nature (Laura Ekstrom, 2000 and 2003; Alfred Mele, 2006) would likewise find Monte Carlo a useful method for modeling human agency.

As in the case of physical simulations, metaphysicians could also use Monte Carlo methods even if their underlying theories are purely deterministic. Whenever we want to quickly approximate some value, a Monte Carlo method is a good option. For example, suppose we want to know the distance between two possible worlds P_1 and P_2 . Assuming some finite field of propositions which are either true or false, a random sample of their truth values will give us an approximation of the distance between P_1 and P_2 .

(iv) Dualist Models

The three types of modeling looked at so far all have instances in the natural and social sciences. It is, after all, where they originated from. But this doesn't mean that philosophers ought to be limited to only these kinds of model. In fact, given that scientists and metaphysicians study different phenomena, it stands to reason that metaphysicians will use types of model that are unique to philosophical investigation.

One such example is a *dualist model*. These are useful for studying the properties of objects or events that take place in "different worlds"—that is, places which have different physical properties and laws. The most recognizable form of dualism in philosophy is substance dualism in the mind-body problem. A dualist model can work in a number of different

ways. One kind of dualist model would be *program-led* which takes two different simulations—with different objects and laws—and compares the ways they change and evolve over time. Naturally, from a substance dualist perspective, one important area of investigation is the possibility of causal interaction between these two worlds.

A different kind of dualist model would be *agent-led*. Here one world is designed using a computer program that allocates a role for *input* that ultimately derives from the choices of an agent. This agent could be an artificial agent, although for simple design it could also be an actual human who provides input using a suitable device. Computer games and training simulators provide examples of this latter kind of dualist model, which have already gained some interest by philosophers (see e.g. Cogburn & Silcox, 2009). Here the “world” produced by the program is such that it is impossible to tell at a given time prior to agent interaction what the state of the system will be after agent interaction. These kinds of model are useful for advocates of incompatibilist free will (e.g. Meghan Griffith 2005, 2007, and 2010; Jonathan Jacobs and Timothy O'Connor, 2013) who believe that an agent acts freely provided they cause events in a world but are not themselves caused by prior events in that world.

(v) *Divine Intervention Models*

Divine intervention models are similar to agent-led dualist models in that events inside a model or virtual world can be determined by agents in another world. However, whereas agent intervention models involve only small inputs of data (such as, for example, playing a computer game)—a divine intervention model specifically allows for an agent to change the very algorithms by which the model is run (i.e. its program).

In this respect, divine intervention models closely resemble models for theories of supernatural agency. Theories of miracles, for example, could be modeled using this kind of computer simulation. If we assume something like Hume's definition of a miracle as a "violation of a law of nature", then a simulation can be designed which sees an agent external to the program intervening in the very writing of the program itself, changing its basic operation. This could involve simply adding a new feature to the output of the model which was not a result of its prior state, or it might involve a rewriting of the algorithms themselves that govern the evolution of the simulation.

There are other topics in the philosophy of religion that also benefit from divine intervention models. Debates surrounding creation and conservation, for example, could look at the different ways a model is generated by a program and the input needed by the "designer" of the program over time. Interestingly, by using computer simulations already provided by the sciences, philosophers can test various theses about the relationship between religion and science, including God's role in the design of physical laws, evolution, and the possibility of life after death. There is also scope for simulating the so-called antinomies of God: could a program be designed, for example, that always maximizes the utility of the agents within it? Could the designer know everything about the program whilst simultaneously allowing for free will of the agents modeled? These, and other questions relating to the properties of God and his role in creation, could be explored using this type of model.

3. AN ILLUSTRATION: HUMEAN SUPERVENIENCE VS. NOMIC NECESSITATION

In this section I will give an example of how a metaphysical phenomenon can be modeled using one of the types of computer simulation. The metaphysical object in question is *lawhood* and I will explore how two different theories—Humean Supervenience and Nomic Necessitation—form the basis for a computer model of laws. I am not a computer scientist, and the accounts I outline here will neither specify a programming language nor a formal algorithm for generating the model. Instead the point is to give an intuitive, informal, characterization of the steps involved in adapting a metaphysical theory to a computer model. Having two examples at hand will be useful in the next section when illustrating how metaphysical models should be evaluated.

3.1. *Simulating Laws 1: Humean Supervenience*


I have already hinted at how the general framework of Humean Supervenience, as given chiefly by Lewis, might be simulated using an agent-based model. Here I want to give more concrete detail into how this can be used to go beyond local properties and bring about other metaphysical objects, including laws.

According to Lewis:

Humean Supervenience ... is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another. ... We have geometry: a system of external relations of spatiotemporal distances between points. ... And at those points we have local qualities:

perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that. (1986a, ix).

Agent-based models provide a natural way of representing his idea. Because for Lewis everything supervenes on properties at space-time points, a model of Humean Supervenience need not evolve with time. Indeed, space-time is a particular, and so reality is one giant "mosaic" of space-time points. This brings to mind the layout of a chessboard, with each square representing a single space-time point. In our agent-based model, we can decide how many squares (space-time points) we want. For simplicity's sake, let us consider an agent-based model with just one-dimension:



In this model each of the squares represents a different space-time point and is the possible bearer of a quality: what Lewis identifies with as the "perfectly natural properties". These properties can be monadic or an n-part relation, in which case they are instantiated by more than one of the space-time points. In our model we can change the value of the squares in our grid to represent the qualities at space-time points. Again, for simplicity's sake, let us assume that there are only two different monadic properties: "on" which is represented by a dark square, and "off", which is represented by an empty square.

A world in a Lewisian sense is then provided by a distinct pattern of "on" and "off" values for squares in the grid.

These worlds can be more or less regular depending on the arrangement of the properties that are instantiated. For example, if A and B are two different worlds, then A clearly exhibits less order than B:



The order or pattern within a world is very important for Lewis because it is ultimately this which laws, counterfactuals, and causation depend upon. Since I am concerned to model laws, I will focus on these only.

Lewis's account of laws is given in his celebrated "Best System Account" (1973, 1983, 1986a, 1994). Although versions of this view had been suggested previously by Mill (1895) and Ramsey (1927), it is Lewis's version that is the most developed. According to Lewis, laws are axioms in a deductive system that capture the true propositions at a world. So for example, if the first square in the A-world is "on", a statement expressing this fact is a truth at that world. Of course, one way to systematize all these truths is just to list them one-by-one: "square-1 is on", "square-2 is off", "square-3 is on", etc. This would provide a very informative summary of the truths, but it wouldn't be very simple. According to Lewis, the laws are the axioms in the "best deductive system"—one which is *best balanced* between informativeness and simplicity:

Take all deductive systems whose theorems are true. Some are simpler, better systematized than others. Some are stronger, more informative than others. These virtues compete: An uninformative system can be very simple, an unsystematised compendium of miscellaneous information can be very informative. The best system is one that strikes as good a balance as truth will allow between simplicity and strength. How good a balance that is will depend on how kind nature is. A regularity is a law iff it is a theorem of the best system. (1994, 478)

This raises our first big treatment and solver issue in creating a metaphysical model for Humean Supervenience. Lewis never specifies just how the best system emerges from the local matters of fact. He certainly believes it is mind-independent (1973, 73). Given that the world is effectively one static arrangement of properties, the best system is one among many sets of propositions that are made true by that arrangement. So we now know that in Lewis's view there must exist more than just the arrangement of local properties: there exist also the propositions *made true* by those properties. These propositions are abstract objects, and Lewis tends towards realism about them (1986b, 54). They do not, therefore, merely exist in the mind and are not reducible to the space-time points and their properties.⁴

⁴ A common way of thinking about propositions in the context of Humean Supervenience is as the class of possible worlds where the propositions are true. This approach still requires us to be realist about sets, which exist as abstract objects. It is no help, therefore, in the current context, and would still need a program to identify them and print them as output. For simplicity's sake, I will therefore continue just to talk of "propositions" leaving open what they amount to in a Humean metaphysics.

Our model of Humean Supervenience needs to contain then both concrete and abstract objects. Of the concrete objects there is just one, namely the Humean mosaic of space-time points and their properties. There are many abstract objects and these supervene on the concrete Humean mosaic. However, in order to get them into our model we need to design a program—a "Supervenience Program"—if you will, that generates them from the mosaic. In the case of the true propositions, such a program would not be difficult to write. For every square in the mosaic, we merely need to design a program which completes the following task:

1. Start with square $n=1$.
2. If n is "on" write $\langle n \text{ is on} \rangle$, otherwise write $\langle n \text{ is off} \rangle$.
3. Move on to square $n=2$.
4. Stop when there are no more squares.

Completing this task will give us a list of true propositions for the world. To get to the laws we now need to find a way of capturing the same information (or less if the gains in simplicity are worth it) using fewer resources. How can we interpret the properties "strength" and "simplicity" of a system in a language which is friendly to computer programmers? The strength criterion is perhaps easier to understand. Strength, or amount of information, can be measured in terms of *how many* of the true propositions for the world are contained in the program that produces the best system. If such a program reproduces all the

propositions, then it is maximally strong.⁵

Simplicity is a little more difficult to interpret because of the different meanings and ways this term is used in ordinary language. I have argued elsewhere that Lewis's original version of the best system account can be improved if the idea of simplicity is interpreted in terms of "algorithmic complexity" (Wheeler 2016, 2017 and 2018). This is a well-known measure of complexity in the computer and information sciences and has been put on a rigorous mathematical footing (Li & Vitanyi, 2008). In its simplest terms, the complexity of a string of symbols, or other information structure, is equal to its length in some predetermined computing language. Because strings of symbols which contain regularity can be algorithmically compressed, the same amount of information can be represented in a simpler way.

Algorithmic complexity, therefore, provides a suitable measure to compare the length of strings and programs and to determine which one is "simpler" (by being shorter in length) than another.

As it is possible to measure both strength and simplicity using a computer program, so it is now possible—in theory—to design a program whose job is to search for the "best system"—a set of propositions that provides the best balance between these two virtues. Whether such a program could actually be constructed and whether or not it would return the right result is something that needs to be tested in practice.

In the next section I will highlight some of the ways in

⁵ It is possible such a program might produce propositions that are not contained in the theoretically best system possible, i.e. it produces too many. In this case we would have to conclude that the program is overly complex and therefore fails the simplicity requirement.

which this model is likely to fail in adequately representing Humean Supervenience and what conclusions can be drawn about the plausibility of Humean Supervenience, in general, as a theory of lawhood.

3.2. *Simulating Laws 2: Nomic Necessitation*


A well-known alternative to Lewis's Humean Supervenience account of lawhood comes from David Armstrong (1983). According to Armstrong, laws do not supervene on regularities among intrinsic, non-modal properties at space-time points. Instead, laws exist in the world as fundamental features, and play an important role in the evolution of the physical world, determining which events occur. This is explained by postulating a new kind of relational property, what he calls "contingent" or "nomic necessity":

Suppose it to be a law that Fs are Gs. F-ness and G-ness are taken to be universals. A certain relation, a relation of non-logical or contingent necessitation, holds between F-ness and G-ness. This state of affairs may be symbolized as "N(F, G)". Although N(F, G) does not obtain of logical necessity, if it does obtain then it entails the corresponding Humean or cosmic uniformity: $(x) (Fx \supset Gx)$. That each F is a G, however, does not entail that F-ness has N to G-ness. (1983, 85)

The idea is that some—but possibly not all—first-order universals are connected by a second-order universal "N". The characteristic feature of this property is that when two or more universals are connected by N, the occurrence of

one brings about the occurrence of the other.

We might try to model the core aspect of this theory using agent-based models once again. Because the instantiation of the universal N between two properties is precisely what brings one about on the occurrence of the other, it suggests using a dynamical model. Each state of the system at one time t_1 determines the state of the system at the next instant t_2 . This time, let us design our universe as involving two dimensions as follows:




Here each square represents a particular and as before can come in different values depending on the universals that it instantiates. We also need to recognize that some of these universals themselves have properties, namely nomic necessitation relations N. Let us suppose that in our model there are only three different kinds of first-order property: *whiteness*, *greyness*, and *blackness*, each being represented by a square shaded in that color. We also recognize one second-order property N(*greyness*, *blackness*) that relates the universals *greyness* and *blackness* with the relation of nomic necessity.

We can design a program which evolves the simulation as follows:

1. For a square n , if n contains a first-order universal not connected by N , then repeat its value.
2. For a square n , if n contains a first-order universal connected by N , change its value to the universal it is connected to.
3. Repeat stages 1-2 for every instant of time.

It is a common complaint against Armstrong's account that he never specifies the conditions under which the obtaining of a relation N between two universals forces the first to bring about the occurrence of the second (Lewis 1983, 366; van Fraassen 1989, 86). Here, let us assume that there is some program which does this work and that it brings it about in a time-wise manner at every instant.

Between two instances on the clock of the computer (t_1 and t_2), we then have the following two states:



The simulation will look very similar to other kinds of cellular automata models of the universe, such as those investigated by Conway (Gardner, 1970) and Stephen Wolfram (2002). However, there are some differences.

Firstly, the system does not evolve according to some pre-determined laws which are programmed into the computer. The only program we need is one which searches for relations of N and then fulfils the required instantiation of the consequent universal in the next time instant. In this sense, each N is like a “mini-program” or subroutine, which is brought about whenever the main program (as given in the steps above) finds it. Secondly, unlike ordinary cellular automata, the state-evolution of the system depends on more than just spatial and temporal relations. Indeed, on this model of Armstrong’s universe, the value of any given square could depend on the value of any other square. It need not be limited to locality conditions of space and time. The advantage of this is that it allows Armstrong’s theory to simulate a wider variety of physical phenomena, including quantum entanglement.

4 THE VALIDATION AND VERIFICATION OF METAPHYSICAL SIMULATIONS

The evaluation of computer models and simulations in science is frequently said to proceed via two different strands: validation and verification. These two aspects are summarized by Winsberg⁶:

The epistemology of simulation can be cleanly

⁶ It must be noted that Winsberg himself thinks the distinction between verification and validation in model evaluation is an over simplification of actual practice. See his 2010 (19-25).

divided into two components: so-called verification and validation. *Verification*, on this conception, is the process of determining whether or not the output of the simulation approximates the true solutions to the differential equations of the original model. *Validation*, on the other hand, is the process of determining whether or not the chosen model is a good enough representation of the real-world system for the purpose of the simulation. (2010, 19-20)

The terms validation and verification are being used by simulationists in a way that is almost the reverse of how they are ordinarily understood by philosophers. The validation of a model does not concern its logical consistency, but rather the degrees to which the model "saves the phenomena" in the required respects, i.e. how much the model conforms to the data gained through observation and experiment.

Conversely, the verification of a model or simulation is not a matter of its empirical content, but rather, how well the model lives up to the expectations of its underlying scientific theory. Put into practice, to validate a computer simulation of a pendulum requires comparing the values for various parameters produced by running the simulation against *benchmark data* gained from measurements of real pendulums. To verify the simulation requires showing that the behavior of the pendulum does not deviate too much from what the underlying theory (i.e. Newton's laws of motion and gravity) predicts.

Of course, computer simulations and models are *going to deviate* from their real-world counterparts and the predictions made by theories. For the most part, this is because of the discretization process employed to make them computationally tractable. The question is whether or not the

behavior deviates in a way that is explained by the limits of the computing process itself or whether it is some problem in the design of the program or the underlying theory which inspired it.

Can these two strands be used in the evaluation of models created from metaphysical theories? I will now argue that they can, and explain some of the ways in which metaphysical verification and validation is likely to differ from computer modeling in the natural and social sciences.

4.1. *Validation*

The validation of a computer simulation usually means checking it against data obtained from a real-world example of the phenomenon modeled. It is essentially a form of empirical evaluation. However, when it comes to metaphysical theories, conditions of acceptability require more than just empirical adequacy. Theories also need to fit our "pre-analytic intuitions" or "judgments" about the phenomenon in question. The role of *intuition* versus *observation* in the methodology of metaphysics is a contested issue, and it is not one that I will take a stand on here. In this case we can simply say that if a metaphysician is inclined more towards *naturalism*, then they will place greater emphasis on the empirical validation of models; on the other hand, if they are inclined more towards *a priori* or "armchair" methods, then they will place greater emphasis on conformity to our pre-analytic intuitions.⁷

Let us start with the empirical side of validation. Unlike scientists who have empirical data by which to check the

⁷ It also worth mentioning that advocates of extreme versions of methodological *a priorism* in metaphysics are unlikely to approve of using computer simulations and models in the first place.

output of a model or simulation, metaphysicians are in no such position. Of course, metaphysicians do have their own records of experience with everyday objects and processes, and traditionally these have been an important source of evidence in favor of, or against, a theory. But when we talk about the validation of a metaphysical theory in an empirical sense, we are likely to mean conformity to our current scientific understanding of the world. As mentioned, a metaphysical theory that is consistent with our best scientific theories is more likely to be preferred than one that is not.

This suggests a novel use of computers in the evaluation of metaphysical theories that has been little explored. A good, scientifically literate metaphysician can surely check the empirical foundations of their theory using their own mind. But there are two setbacks to this approach. Firstly, very few philosophers are this well versed in the fundamental theories of science, which are increasingly complex and mathematical in nature. Secondly, even a scientifically literate philosopher can make mistakes, and due to the complexity of the philosophical and scientific theory in question, fail to fully comprehend its consequences. Computer modeling helps alleviate these concerns and provides a means to compare the consistency of metaphysical theories with our best scientific understanding of the world.

This can be achieved by trying to model existing scientific results using, as its basic program of operation, some metaphysical theory. Going back to the examples based on Lewis and Armstrong, these can have their consistency with the fundamental laws of physics evaluated in the following way. In the case of Humean Supervenience, we first need to construct a grid which satisfies the four basic dimensions of space and time. Then the values, or states of the space-time points, need to instantiate fundamental properties identified by our physicists, such as: spin, charge, mass, momentum, strangeness, etc. The program in the model for Humean

Supervenience (as given in section 3) will then output two important sets of data: (i) the true propositions for this Humean world, and (ii) the best system, and hence the laws, for this world. The consistency of this metaphysical model with science depends on how well its output matches what our best scientific theories tell us. If the laws which emerge from the Humean mosaic are wildly at odds with the laws of fundamental physics, then clearly this model has failed the validation test.

In practice, the validation of a metaphysical model with a scientific theory will be more complicated than this. The mismatch between the output of a model and real-world physical phenomena might be the result of limitations in the computer itself, rather with than the underlying metaphysical theory. Even worse, there may be logical limits to what can be computed in a finite amount of time using finite resources, which affects the possibility of our metaphysical model ever being able to reproduce physical phenomena. I will come back to this issue in section 5.

The other aspect of validation that is unique to philosophical modeling involves consistency with our pre-analytic intuitions. Presumably, the underpinning theory from which the model is constructed will already have passed some testing in this area. This is likely to have been largely conceptual in nature: does the theory describe time, properties, free will, etc., in a way which satisfies my already existing ideas about these processes? But computer modeling can provide a new way to validate theories against our intuitions. Here the medium through which the theory is tested is not conceptual but visual. The question becomes: can the theory be used to create a world which, when experienced, matches my expectations about how the objects ought to behave?

The experiences can be generated through a simple desktop monitor, but more realistic and interesting

experiences can be generated using virtual reality. In both cases the individual becomes immersed in the model or simulation and can use this to judge its consistency with their intuitions. In the case of the model for Armstrong's universe, would our experience of a virtual world governed by this program produce the same ideas about lawhood, accidental regularities, and determinism that are ordinarily produced by our experiences of the physical world? One of the most striking examples of this in practice comes from early work on Conway's *Game of Life*. This "model" of the universe, which as we have already remarked is different from Armstrong's in a number of ways, nonetheless produces objects which behave in ways strikingly similar to biological organisms. Simply put: these objects behave in ways that we expect they *ought* to behave. If the same type of results can be generated using a metaphysical model, then this provides a good test of validation against our pre-analytic intuitions.

4.2. *Verification*

The verification of a computer model involves comparing the results (output) of the simulation to expected values predicted by its underpinning theory. In the case of scientific models, this means producing numerical values via the simulation for a range of magnitudes and comparing their accuracy against examples calculated analytically. A good test of verification results in values that do not deviate too wildly—a bad test would result in values that do. But things are not as simple as this. As I have remarked above, no practicing scientist would expect the results of their computer model to match those predicted by their theory completely. In the case of the pendulum, the theory predicts that it will change acceleration at every infinitesimal displacement from its starting point, but because the

computer can only make a finite number of calculations, its results will not reflect this.

There is, therefore, a great deal of approximation involved in the verification process and how much approximation is permitted very much depends on the scientists in question and the intended application of the model.

What can be said about the verification of computational models for metaphysical theories? Again, we ought to say that a good test of a model's verification is if it reproduces the results of the theory well—and a bad test is if it does not. The results of a metaphysical theory are not numerical and so we cannot compare values for magnitudes in the same way as a scientific theory. The “results” of a metaphysical theory are its logical implications that are intended to hold true for some phenomenon in question. In Armstrong's theory of lawhood, it is intended to be a logical consequence of $N(F, G)$ that the regularity $(x)(Fx \supset Gx)$ will hold. Verifying a computer model of Armstrong's theory would then presumably require checking whether or not a program that mimics $N(F, G)$ is also one that produces the regularity $(x)(Fx \supset Gx)$. If another program needs to be added to produce $(x)(Fx \supset Gx)$ on top of $N(F, G)$, then this would count as an instance of an *ad hoc* addition, and ought to count against the computer model.

Whereas scientific models need to be numerically consistent with their theories, I shall say that metaphysical models need to be *functionally equivalent* to their theories. In other words, they need to be designed in such a way that the program produces the same class of intended consequences for the model as there are in the theory. In order to verify a metaphysical model, the model needs to be created and actually run on a computer. All that can be done here is identify some of the intended consequences of a theory that should be checked once the simulation is up and running.

Let us return to the computer models of laws outlined in section 3. There are a number of implications of Lewis's Best System Account that ought to be checked when verifying the model. Below is a list of the most important:

- No modal relations
- Emerging true propositions
- A measure of simplicity
- A measure of strength
- The “best balance” between simplicity and strength
- Emerging laws of nature

Lewis's complete Humean Supervenience program is vast, and if one wanted to use this model to simulate all the metaphysical phenomena it aims to account for, then you could add to this list: counterfactuals, causal relations and necessary truths. A number of items on this list will be worth checking in detail because there are *prima facie* reasons for thinking the model will not produce them in the right respects. One, which we have looked at already, is the emergence of true propositions. There is evidence for thinking Lewis is a realist about these and so they exist in the Humean world—although as abstract, not concrete, objects. Nonetheless, we saw that in our model they do not emerge naturally from the Humean mosaic. Because the Humean mosaic is not the result of a program, but rather a “given” in the system, there is nothing the model can do to produce the propositions. (Things are not improved either by having more mosaics, one for each possible world.) To get these we have to add an additional program which checks all the squares in the grid one-by-one, registers their value, and prints the output <P>, where “<P>” stands for the proposition in question. Is this program an *ad hoc* addition to the model? There seems to be no corresponding sign of it in Lewis's theory.

Another worry, and one that is perhaps more important in the debate about laws, is whether or not a program for the best system can be designed that produces the right laws. In Lewis's theory, the laws are the axioms in all the systems best balanced between simplicity and strength. If we measure strength by the number of true propositions and simplicity in terms of algorithmic complexity, this will leave many options available. Sacrificing strength for simplicity is a lot like lossy data compression (Braddon-Mitchell, 2002) where large gains in simplicity can be had by "idealizing" the data—making it more regular and orderly than it actually is. But how much idealization one performs depends on how much strength one is willing to sacrifice. Lewis's theory calls for the "best balance" but it's not obvious at all what kind of compression-to-idealization ratio this would require.

If both these points (and possibly others beside) are seen to be problems for the model, what can a metaphysician conclude? The first step would be to follow ordinary scientific practice and try creating an improved computer model. After all, the models sketched in section 3 are only a first attempt. It's highly likely that better models will be designed and it is a natural part of computer simulation in science that models should go through many rounds of evaluation and redesign. Perhaps, for example, a program can be designed which produces a ratio between compression and idealization that is easy to implement and produces the same kind of results as Lewis's best balanced axioms.

On the other hand, if after repeated attempts no acceptable model can be constructed, what then should be concluded? It is reasonable that at this stage the metaphysician should look back at the theory itself and consider whether or not it needs to be amended or abandoned completely. However, when it comes to evaluating the truth of a metaphysical theory based on the

success or failure of a computer model, it raises a fundamental question: should metaphysical theories be beholden to computational implementation? As we will now see, there are some reasons for thinking that this whole endeavour is unsuitable for the evaluation of metaphysical theories.

5 FUNDAMENTAL LIMITS TO COMPUTER MODELING

The design of a computer model or simulation is limited in the sense that we cannot design programs for which we do not have the requisite hardware to manage. Hardware is limited in most instances by speed and storage. A program which requires more storage capacity than can be provided, or requires a very large amount of time to execute, will not be very useful. Such limitations are accounted for in the verification stage of evaluation. In scientific examples, careful model design can alleviate some of these problems. But there are limitations to what can be effectively computed that go beyond the physical properties of our chosen computing device. There are “fundamental limits” to computing which seem to place logical constraints on the very kind of programs—and therefore models—we can design.

These limits were discovered in an attempt to show whether or not mathematics is complete—that is, whether or not every theorem stateable in the language of mathematics could be proven. As is well known, Alan Turing (1937) answered this question in the negative and his method for answering it was to formalize the idea of an “effective computing procedure”. Now known as the Turing Machine (TM), it provides a paradigm mechanical computer that is comprised of tape, a read-write head, state register, and table of rules. Essentially, for each symbol it reads on the tape

(input), it looks up the rule it should follow whilst in its current state, and prints a new symbol (output). It then moves on to a new state and a new symbol. Once all the symbols have been read (i.e. there is only blank tape left) it halts.

Showing whether or not some theorem in mathematics is provable is therefore equivalent to showing that a TM could follow its rules and halt. Turing was able to show that for a class of numbers (the “computable reals”), a TM could be designed such that it was impossible to tell whether or not it had halted after some time or whether it was continuing to calculate the next digit in the sequence of the number.⁸ In other words, there are some problems or theorems in mathematics which are not decidable using a standard computing device.

This problem, which has some structural analogies to Gödel’s incompleteness theorems and Russell’s paradox, shows that there are absolute limits to what can be expected from any computing device. And the limit comes about, not because of the make-up of the computing device itself, but because of the *content* of the problem that it is trying to solve. Most of these problems arise because the computer is asked to do something which involves reference to itself and/or to irrational numbers. Given that the problem arises because of the content of what the computer is asked to compute, it is entirely feasible that some metaphysical models could find themselves up against such limits.

Suppose we wanted to develop the model of Humean Supervenience from section 3 to include counterfactuals. Now, according to Lewis, a counterfactual such as “ $P \square \rightarrow Q$ ” is true if, and only if, in the nearest possible world where

⁸ For the details of Turing’s argument, see Cockshott, McKenzie & Michaelson (2012, 67-73).

P is true, so is Q (1973, 8). To generate the class of true counterfactuals we need to add an array of grids whose space-time points take alternative values to those designated for the actual world. In every conceivable way a space-time point could change its value, a new possible world needs to be created. This creates a level of complexity in the model which allows for the possibility of self-reference. If we also allow for states of space-time points which take irrational values, then Turing's undecidability result looms large.

What could a metaphysician say in response? There are proposals for models of computation that aim to go beyond the "Turing Barrier"—in other words, computers that can complete an infinite number of tasks. One such idea has been developed by Jack Copeland (2002). Copeland notes that in Turing's original undecidability result, the problem emerges because it is assumed that it will take a TM an infinite amount of time to complete an infinite number of tasks. But by appealing to an idea first suggested by Russell, this need not necessarily be the case. Provided the length of time it takes to complete each task reduces at every instant, the sum total length of an infinite task could be a finite amount of time, and therefore, computable:

Imposing the same temporal patterning upon a Turing Machine produces what I have termed an *accelerating* Turing Machine. These are Turing Machines that perform the second primitive operation called for by the program in half the time taken to perform the first, the third in half the time taken to perform the second, and so on. Let the time take to perform the first primitive operation called for by the program be one "moment". Since:

$$1/2 + 1/4 + 1/8 + \dots + 1/2^n + 1/2^{n+1} + \dots$$

is less than 1, an accelerating Turing Machine—or a human computer—can perform infinitely many primitive operations before two moments of operating time have elapsed. (2002, 283)

If a simulation of Lewis's counterfactuals does not halt on a standard TM, a TM which accelerated its performance in the way Copeland suggests might be able to decide whether or not an arbitrary counterfactual “ $P \square \rightarrow Q$ ” was true, even if it made reference to itself or involved irrational magnitudes.

Whether or not proposals for hypercomputation such as Copeland's accelerating TM are feasible is highly contested. For example, even if a computer could be built which reduced its operating time by half for every task, the machine would eventually need to send signals faster than the speed of light, and would produce so much heat in the process it would turn itself into plasma in a very little amount of time (Cockshott, McKenzie & Michaelson 2012, 188). It might not halt because the TM is destroyed in the process!

But these are *physical* limitations and it remains open to the metaphysician to argue that provided hypercomputation is possible in some possible world (with a different physics to our own), the metaphysical feasibility of the model is sound. The fact we cannot actually build it with our own equipment becomes another physical limitation, and as we have seen, this can be accounted for in the verification stage of evaluation.

6 CONCLUSION

There are real benefits metaphysicians can get from using computer models in their work. I hope to have shown that given the different variety of model and simulation types, there are many possible avenues for metaphysical research that could benefit from using models. There are genuine logical concerns facing the implementation of metaphysical systems on a computer too. Whether or not these are fatal depends on the feasibility of forms of hypercomputation that avoid Turing's undecidability results. However, given that these problems are likely to inflict only a small subset of all conceivable models, this ought not to deter metaphysicians from adopting computer models and simulations in the short term.

REFERENCES

- ALEXANDER, J., HIMMELREICH, J. & THOMPSON, C. Epistemic Landscapes, Optimal Search, and the Division of Cognitive Labor. *Philosophy of Science*, 82(3), 424-453, 2015.
- ARMSTRONG, D. *What is a Law of Nature?* Cambridge: Cambridge University Press, 1983.
- BEISBART, C. & HARTMANN, S. Computersimulationen in der Angewandten Politischen Philosophie – Ein Beispiel. *Kolloquium*, 21(1), 1153-1162, 2016.
- BRADDON-MITCHELL, D. Lossy Laws. *Nous*, 35(2), 260-277, 2001.
- CARTWRIGHT. *The Dappled World*. Cambridge: Cambridge University Press, 1999.

- CHALMERS, D. Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200-219, 1995.
- _____. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press, 1996.
- COCKSHOTT, P., MACKENZIE, L. & MICHAELSON, G. *Computation and its Limits*. Oxford: Oxford University Press, 2012.
- COGBURN, J. & SILCOX, M. *Philosophy through Computer Games*. London: Routledge, 2009.
- COPELAND, J. (2002). Accelerating Turing Machines. *Mind and Machines*, 12(2), 281-301, 2002.
- DE LANGHE, R. A Unified Model of the Division of Cognitive Labor. *Philosophy of Science*, 81(3), 444-459, 2014.
- EKSTROM, L. *Free Will: A Philosophical Study*. Boulder, CO: Westview, 2000.
- _____. *Free Will, Chance, and Mystery*. *Philosophical Studies*, 113(2), 153-80, 2003.
- FAROUKI, R. & SHAPIRO, S. Computer Simulations of Environmental Influences on Galaxy Evolution in Dense Clusters. *The Astrophysical Journal*, 24(1), 928-945, 1980.
- GARDNER, M. *Mathematical Games – The Fantastic Combinations of John Conway's New Solitaire Game "Life"*. *Scientific American*, 223, 120-123, 1970.
- GODFREY-SMITH, P. *Theories and Models in Metaphysics*. *The Harvard Review of Philosophy*, 14(1), 4-19, 2006.
- _____. *Metaphysics and the Philosophical Imagination*. *Philosophical Studies*, 160(1), 97-113, 2012.

- GORDON, S. & GORDON, F. Random Simulations of Radioactive Decay. *Primus*, 3(3), 323-330, 1993.
- GOULD, H., TOBOCHNIK, J. & CHRISTIAN, W. *An Introduction to Computer Simulation Methods*. Reading MA: Addison-Wesley, 2006.
- GRIFFITH, M. Does Free Will Remain a Mystery? A Response to van Inwagen. *Philosophical Studies*, 124(3), 261–269, 2005.
- _____. Freedom and Trying: Understanding Agent-Causal Exertions. *Acta Analytica*, 22(1), 16–28, 2007.
- _____. Why Agent-Caused Actions are Not Lucky. *American Philosophical Quarterly*, 47(1), 43–56, 2010.
- GRUNE-YANOFF, T. & WEIRICH, P. The Philosophy and Epistemology of Simulation: A Review. *Simulation Gaming*, 41(1), 20-50, 2012.
- HAHN, U., VON SYDOW, M. & MERDES, C. How Communication can Make Voters Choose Less Well. *Topics in Cognitive Science*, 11(1), 194–206, 2019.
- HONG, L. & PAGE, S. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46), 16385-16389, 2004.
- HUMPHREYS, P. *Extending Ourselves: Computational Science, Empiricism and Scientific Method*. Oxford: Oxford University Press, 2004.
- JACOBS, J. & O'CONNOR, T. (2013). Agent Causation in a Neo-Aristotelian Metaphysics. In S. Gibb, E. J. Lowe, & R. D. Ingthorsson (Eds.), *Mental Causation and*

- Ontology (pp. 173-192). Oxford: Oxford University Press.
- KLEIN, D., MARX, J. & SCHELLER, S. (2019). Rational Choice and Asymmetric Learning in Iterated Social Interactions – Some Lessons from Agent-Based Modeling. In K. Marker, A. Schmitt, & J. Sirsch (Eds.), *Demokratie und Entscheidung* (pp. 277-294). Wiesbaden: Springer.
- LEWIS, D. *Counterfactuals*. Oxford: Blackwell, 1973.
- *New Work for a Theory of Universals*. *Australasian Journal of Philosophy*, 61(4), 343-377, 1973.
- *Philosophical Papers Volume II*. Oxford: Oxford University Press, 1986a.
- *On the Plurality of Worlds*. London: Blackwell, 1986b.
- *Humean Supervenience Debugged*. *Mind*, 103(412), 473-490, 1994.
- LI, M. & VITANYI, P. *An Introduction to Kolmogorov Complexity and Its Applications*. New York: Springer, 2008.
- MCGUFFIE, K. & HENDERSON-SELLERS, A. *A Climate Modelling Primer*. Chichester: Wiley & Sons, 2005.
- MCKENNA, M. & COATES, J. (2015). *Compatibilism*. (E. Zalta, Ed.) Retrieved from *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition): <https://plato.stanford.edu/archives/win2018/entries/compatibilism>
- MCTAGGART, J. *The Unreality of Time*. *Mind*, 17(68), 457-474, 1908.

- MELE, A. *Free Will and Luck*. Oxford: Oxford University Press, 2006.
- MERDES, C. *Strategy and the pursuit of truth*. *Synthese*, Forthcoming, 1-22, 2018. Retrieved from <https://doi.org/10.1007/s11229-018-01985-x>
- MILL, J. S. *A System of Logic*. London: Routledge, 1985.
- MORGAN, MARY S., and MARGARET MORRISON, eds. *Models as Mediators: Perspectives on Natural and Social Science*. N edition. Cambridge ; New York: Cambridge University Press, 1999.
- PAUL, L. *Metaphysics as modeling: the handmaiden's tale*. *Philosophical Studies*, 160(1), 1-29, 2012.
- PENROSE, R. *The Emperor's New Mind*. Oxford: Oxford University Press, 1989.
- RAMSEY, F. *Facts and Propositions*. *Aristotelian Society Supplementary Volume*, 7, 153–170, 1927.
- THAN, O., BUTTGENBACH & S. *Simulation of Anisotropic Chemical Etching of Crystalline Silicon using a Cellular Automata Model*. *Sensors and Actuators A: Physical*, 45(1), 85-89, 1994.
- TURING, A. *On Computable Numbers, with an Application to the Entscheidungsproblem*. *Proceedings of the London Mathematical Society*, 45(2), 230-265, 1937.
- UBLER, H. & HARTMANN, S. *Simulating Trends in Artificial Influence Networks*. *Journal of Artificial Societies and Social Simulation*, 19(1), 2016. Retrieved from <http://jasss.soc.surrey.ac.uk/19/1/2.html>
- VAN FRAASSEN, B. *Laws and Symmetry*. Oxford: Clarendon Press, 1989.

- WEISBERG, M. & MULDOON, R. Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science*, 76(2), 225-252, 2009.
- WHEELER, B. Simplicity, Language-Dependency and the Best System Account of Laws. *Theoria: An International Journal for Theory, History and Foundations of Science*, 31(2), 189-206, 2016.
- _____. Humeanism and Exceptions in the Fundamental Laws of Physics. *Principia: An International Journal of Epistemology*, 21(3), 317–337, 2017.
- _____. *Idealization and the Laws of Nature*. Geneva: Springer, 2018.
- WILLIAMSON, T. Model-Building in Philosophy. In R. Blackford, & D. Broderick (Eds.), *Philosophy's Future: The Problem of Philosophical Progress* (pp. 159-172). Oxford: Wiley-Blackwell.
- WINSBERG, E. *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press, 2010.
- WOLFRAM, S. *A New Kind of Science*. Champaign: Wolfram Media, 2002.
- WOOLFSON, M. & PERT, G. *An Introduction to Computer Simulation*. Oxford: Oxford University Press, 1999.
- ZOLLMAN, K. The Communication Structure of Epistemic Communities. *Philosophy of Science*, 74(5), 574-587, 2007.
- _____. The Epistemic Benefit of Transient Diversity. *Erkenntnis*, 72(1), 17-35, 2010.

