Paul Humphreys

Computational Science and its Effects

Zusammenfassung Der Computer und die Computerwissenschaften finden eine enorme Verbreitung in den Wissenschaften. Was bedeutet das aus wissenschaftsphilosophischer Perspektive? Paul Humphreys verwirft in seinem Text die Rede von einer Revolution und diagnostiziert stattdessen eine tiefgreifende Veränderung in der Bedeutung von ›Techniken‹ (*techniques*). Computermethoden bedeuten eine grundlegende philosophische Neuerung in folgendem Sinne: Sie schieben den Menschen aus dem Zentrum des Prozesses der Erkenntnisgewinnung. Damit setzen sie zwar eine historische Entwicklung fort, die mit dem Gebrauch von Instrumenten wie Uhren, Teleskopen oder Mikroskopen eingesetzt hat. Der Unterschied besteht aber darin, dass nicht mehr nur menschliche Wahrnehmungskapazitäten erweitert werden, sondern das Überlegen und Schließen selbst vom menschlichen Intellekt zur Maschine verlagert wird.
Bei diesem Beitrag handelt es sich um die gekürzte Fassung seines Vortrags *Methodological Novelties in Computer Simulations* im Rahmen der Abschlusstagung der Forschungsgruppe ›Wissenschaft im Anwendungskontext‹ im September 2007 (vgl. ZiF: *Mitteilungen* 1/2008, 24ff.).

This paper is the shortened version of the lecture Paul Humphreys gave on the final conference of the ZiF Research Group *Science in the Context of Application* in September 2007 (cf. ZiF: *Mitteilungen* 1/2008 pp. 24ff.).

Prof. Paul Humphreys
Corcoran Department of Philosophy, University of Virginia
pwh2a@virginia.edu

1. **Introduction**

The rise of computational science, which can be dated, somewhat arbitrarily, as beginning around 1945 – 1946[1], has had effects in at least three connected domains—the scientific, the philosophical, and the socio-technological context within which science is conducted.[2] Some of these effects are secondary, in the sense that disciplines such as complexity theory would have remained small theoretical curiosities without access to serious computational resources. Other effects, such as the possibility of completely automated sciences, are longer term and will take decades to alter the intellectual landscape. I shall provide here some examples of fine-grained philosophical effects as well as examples of more sweeping social and intellectual consequences that will suggest both the different ways of thinking that these methods require and a hint at how far-reaching they are.

First, we need a framework. In their paper 'Complex Systems, Modelling, and Simulation', SYLVAIN SCHWEBER and MATTHIAS WÄCHTER [2000] suggested that the introduction and widespread use of computational science constitutes what they call a 'HACKING Revolution' in science and that HACKING's use of 'styles of reasoning', a concept which originated with the historian of science A.C. CROMBIE, can give us useful insights into these methods. SCHWEBER and WÄCHTER have many useful things to say about simulations and related methods, but HACKING's framework does not sit well on computational science. Let me say why.

HACKING revolutions have four principal characteristics: First, 'they transform a wide range of scientific practices and they are multi-disciplinary'. In this, they are different from the more familiar KUHNian revolutions or the shifts in theoretical research programs suggested by IMRE LAKATOS, which tend to be limited to single scientific disciplines. Computational science satisfies this first condition, most notably because its methods are largely trans-disciplinary. Secondly, a HACKING revolution leads to new institutions designed to foster the new practices. The Santa Fe Institute is an example of this second feature.[3] The third characteristic is that 'the revolution is linked with substantial social change'. Changes in the social structure of science are hard to separate from more general societal changes introduced by computers, but it is true that the social structure of science has been affected by the easy electronic exchange of ideas, the dominant role of programmers in a research group, and remote access to supercomputers. The fourth characteristic is that 'there can be no complete, all-encompassing history of such revolutions'.

Although there is merit in the concept of a HACKING Revolution, I shall not use it here for

---

[1] I identify its origins with the use of electronic computers to perform Monte Carlo calculations at Los Alamos and JOHN MAUCHLEY's suggestion that ENIAC could be used for difference equation simulations, rather than for just routine arithmetical calculations. See METROPOLIS [1993], p. 127 for the second point. I do not vouch for the accuracy of METROPOLIS's recollections on this point although the exact historical turning point, if indeed 'exact' ever makes sense in historical claims, is unimportant. For those interested in technoscience, I note that the innovation had its origins at Los Alamos and other military research institutions rather than in industrial applications.

[2] There are other domains it has affected, but I shall restrict my discussion to these three.

[3] Although the Institute has recently announced that because complexity science is now well established, it must move in new directions.

two reasons. The first is that HACKING revolutions share their second and third components with KUHNian revolutions (because of the tight link between the intellectual and sociological aspects of KUHN's position, these components are satisfied almost by default in a KUHNian revolution). And the fourth condition is almost trivially true of any such historical episode. This leaves only the multi-disciplinary aspect, which is important but lacks fine structure. Secondly, let me make a distinction between *replacement revolutions* and *emplacement revolutions*. Replacement revolutions are the familiar kind in which an established way of doing science is overthrown and a different set of methods takes over. Emplacement revolutions occur when a new way of doing science is introduced which largely leaves in place existing methods. The introduction of laboratory experimentation was an emplacement revolution in the sense that it did not lead to the demise of theory or of observation. Similarly, the rise of computational science constitutes an emplacement revolution. This is not to say that theory and experiment are not affected by computational approaches, because certain theoretical methods have now been taken over by computational methods, and many experiments are now computer assisted, but theory and experiment have not been abandoned and considered scientifically unacceptable in the way that the replacement revolutions of COPERNICUS over PTOLEMY, NEWTON over DESCARTES, or DARWIN over gradualism resulted in the untenability of the previous approaches.

What of styles of reasoning? Here are six cases, originally identified by CROMBIE, that are cited by HACKING, as examples of the genre:

(a) The simple method of postulation exemplified by the Greek mathematical sciences.
(b) The deployment of experiment both to control postulation and to explore by observation and measurement.
(c) Hypothetical construction of analogical models.
(d) Ordering of variety by comparison and taxonomy.
(e) Statistical analysis of regularities of populations, and the calculus of probabilities.
(f) The historical derivation of genetic development.
(HACKING [1992], p.4)

HACKING then says: "Every style of reasoning introduces a great many novelties including new types of objects; evidence; sentences, new ways of being a candidate for truth or falsehood; laws, or at any rate modalities; possibilities. One will also notice, on occasion, new types of classification, and new types of explanations....Hence we are in a position to propose a necessary condition for being a style of reasoning: each style should introduce novelties of most or all of the listed types and should do so in an open-textured, ongoing, and creative way" (ibid, pp. 11 – 12, slightly reformatted). One could squeeze computational science into this framework because three of the five criteria are satisfied—novelties of evidence, sentences, and possibilities—but as I shall argue, laws are the wrong vehicle for understanding what is distinctive about computational science, and the novel objects are better understood as novel representations. Moreover, 'style of reasoning' has an anthropocentric flavor that is best avoided in this context. So 'style of reasoning' is not a good fit. Instead, I shall use the term 'technique' in what follows.

2. The Main Issue

Let me put the principal philosophical novelty of these methods in the starkest possible way: Computational science introduces new issues into the philosophy of science because it uses methods that push humans away from the centre of the epistemological enterprise. In doing this, it is continuing a historical development that began with the use of clocks and compasses, as well as the optical telescope and microscope, but it is distinctively different in that it divorces reasoning, rather than perceptual, tasks from human cognitive capacities. There were historical ancestors of computational science, such as astrolabes and orreries, but their operation was essentially dependent upon human calculations.

Until recently, science has always been an activity that humans carry out and analyze. It is also humans that possess and use the knowledge produced by science. In this, the philosophy of science has followed traditional epistemology which, with a few exceptions such as the investigation of divine omniscience, has been the study of human knowledge. LOCKE's *Essay Concerning Human Understanding*, BERKELEY's *A Treatise Concerning the Principles of Human Knowledge*, HUME's *A Treatise of Human Knowledge*, REID's *Essays on the Intellectual Powers of Man* are but a few examples; the CARTESian and KANTian traditions in their different ways are also anthropocentric.[4] In the twentieth century, the logical component of logical empiricism broke free from the psychologism of earlier centuries, but the empiricist component prevented a complete separation.[5] Two of the great alternatives to logical empiricism, QUINE's and KUHN's epistemologies, are rooted in communities of human scientists and language users. Even constructive empiricism and its successor, the empirical stance, are firmly anchored in human sensory abilities. (VAN FRAASSEN [1980], [2004]). There are exceptions to this anthropocentric view, such as POPPER [1972] and FORD, GLYMOUR, and HAYES [2006], but the former's World 3 is too abstract for our concerns and the latter's artificial intelligence orientation does not address the central issues of computational science.[6]

At this point I need to draw a distinction. Call the current situation within which humans deal with science that is carried out at least in part by machines the *hybrid scenario*, and the more extreme situation of a completely automated science replacing the science conducted by humans the *automated scenario*. This distinction is important because in the hybrid scenario, one cannot completely abstract from human cognitive abilities when dealing with representational and computational issues. In the automated scenario one can, and it is for me the more interesting philosophical situation, but in the near term we shall be in the hybrid scenario and so I shall restrict myself here to that case. It is because we are in the hybrid scenario that computational science constitutes an emplacement revolution. If the automated scenario comes about, we shall then have a replacement revolution.

For an increasing number of fields in science, an exclusively anthropocentric epistemology is no longer appropriate because there now exist superior, non-human, epistemic

---

[4]A KANTian approach can be generalized to non-human conceptual categories, although the extent to which humans could understand those alien categories is then a version of one philosophical challenge faced by computational science.

[5]CARNAP's *Aufbau* (CARNAP [1928]) allows that a physical basis could be used as the starting point of the reconstruction procedure, but adopts personal experiences as the autopsychological basis. The overwhelming majority of the literature in the logical empiricist tradition took the human senses as the ultimate authority.

[6]One can usefully borrow POPPER's thought experiment in which all of the world's libraries are destroyed and ask how much of contemporary science would be affected if neutron bombs shut down all of the world's computers. Much of 'big science', especially in physics and astrophysics, would be impossible to carry out.

authorities. So we are now faced with a problem, which we can call the *anthropocentric predicament*, of how we, as humans, can understand and evaluate computationally based scientific methods that transcend our own abilities and operate in ways that we cannot fully understand. Once again, this predicament is not entirely new because many scientific instruments use representational intermediaries that must be tailored to human cognitive capacities. With the hybrid situation, the representational devices, which include simulations and computationally assisted instruments such as automated genome sequencing, are constructed to balance the needs of the computational tools and the human consumers. We can call the general problem of inventing effective intermediaries the *interface problem* and it is a little remarked upon aspect of scientific realism when we access the humanly unobservable realm using instruments. Just as scientific instruments present philosophy with one form of the metaphysical problem of scientific realism and its accompanying epistemological problems, so computational science leads to philosophical problems that are both epistemological, a feature that has been emphasized by ERIC WINSBERG and JOHANNES LENHARD[7], and metaphysical.

3. What is Metaphysically Different About Computational Science?

The essence of computational science is providing computationally tractable representations; objects that I have elsewhere called computational templates.[8] It is an important feature of templates that they are trans-disciplinary. The philosophical literature on scientific laws, with its emphasis on counterfactuals, nomological necessity, logical form, and so on, often does not stress the fact that the fundamental laws of a science are uniquely characteristic of that science. Although NEWTON's laws applied to any material object in the eighteenth century, they did not characterize biological objects qua biological objects in the way that they did characterize what it was to be a physical object. Nowadays, the HARDY-WEINBERG law is a characteristic feature of population biology, and it makes no sense in chemistry or physics.[9]

I mentioned above that laws are the wrong vehicle for understanding computational science. The reason for this is connected with the point just noted that scientific laws are intimately tied to a particular science and its subject matter, whereas the emphasis of computational science is on trans-disciplinary representations. There are some candidates for laws of this trans-disciplinary type in complexity theory, such as ZIPF's Law, a power law that reasonably accurately describes the distribution of city sizes, network connection densities, the size of forest fires, and a number of other phenomena that are the result of scale-invariant features. Just as theory and experiment involve techniques that are to a greater or lesser extent subject matter independent, so too does computational science. This cross-disciplinary orientation has at least two consequences that are worth mentioning. First, it runs counter to the widely held view that models are local representations. It is, of course, true that many models are far less general than theories, but the existence of widely used computational templates suggests

---

[7]See e.g. WINSBERG [2001], [2003], LENHARD [2007].

[8] See HUMPHREYS [2002], [2004] Chapter 3, [2008].

[9] To prevent misunderstanding, I note that although the term 'law' is used for such things as the weak and strong laws of large numbers in probability theory, because these are purely mathematical results this is a courtesy use of the term 'law'. They lack at least the nomological necessity possessed by scientific laws.

that the disunity of science thesis that often accompanies the 'models are local' thesis is simply wrong about the areas of contemporary science that lend themselves to the successful use of such templates. Secondly, it runs orthogonally to the traditional reductionist approach to understanding. Reduction suggests to us that we can better understand higher level systems by showing how they can be reduced to, how they can be explained in terms of, lower level systems. Computational templates suggest that we can gain understanding of systems without pursuing reduction by displaying the common structural features possessed by systems across different subject domains. In saying this, I am not claiming that these trans-disciplinary representations did not exist prior to the introduction of computational science. What the latter development did was to allow the vastly increased use of these techniques in ways that made their application feasible.

I can illustrate the issue involved using as an example agent based simulations. Agent based simulations are in certain ways very different from what one might call equation-based simulations. It is a common, although not universal, feature of agent based models that emergent macro-level features appear as a result of running the simulation, that these features would not appear without running the simulation, that new macro-level descriptions must be introduced to capture these features, and that the details of the process between the model and its output are inaccessible to human scientists. No traditional modeling methods address the first, second, and fourth features of these simulations. Let me elaborate a little on how the third point plays out in this context. The situation has been nicely captured by STEPHEN WEINBERG: "After all, even if you knew everything about water molecules and you had a computer good enough to follow how every molecule in a glass of water moved in space, all you would have would be a mountain of computer tape. How in that mountain of computer tape would you ever recognize the properties that interest you about the water, properties like vorticity, turbulence, entropy, and temperature?" (WEINBERG [1987], p. 434). Many of the 'higher level' conceptual representations needed to capture the emergence of higher level patterns do already exist in other theoretical representations; they are the starting point for what ERNEST NAGEL called inhomogeneous reductions. With other agent based models the situation is different because the simulation itself will, in some cases, construct a novel macro-level feature. It is this constructivist aspect of simulations, one that runs in the opposite direction to the traditional reductionist tendency of theories, that is a characteristic feature of agent based models in particular, although it also can be a focus of equation based models. Constructivism was memorably described in ANDERSON [1972] and is a key element of the arguments presented in LAUGHLIN and PINES [2000].[10] These emergent patterns in computer simulations form the basis for what MARK BEDAU has characterized as "weak emergence" (BEDAU [1997]) and traditional human modeling techniques will not generate them from the agent base. They can only be arrived at by simulation.

This emphasis on higher level patterns is not restricted to computational science or to emergence. It is a feature of multiply realizable systems and of physical systems in which universality is exhibited. As another example, NIKLAS LUHMANN, the German sociologist, has

---

[10]The use of generative mechanisms as an element of constructivism is noted in KÜPPERS and LENHARD [2006].

persuasively argued for the irrelevance of individual humans in various functional systems.[11] For example, within consumer economies, it is irrelevant who purchases the pack of cigarettes—they can be male, female, Chilean or Chinese, middle-aged or old, white collar or blue collar—all that matters is that the relevant economic communications take place. Indeed, LUHMANN's work is a striking example of a research program within which the importance of humans as individuals is severely diminished and the emphasis placed on the autonomy of higher level features. LUHMANN was an early advocate of autopoiesis, a process that leads to self-organizing systems. One of the core features of self-organizing systems is that there is no central organizing force controlling the system. The American Stryker forces that are currently operating in Iraq and Afghanistan are a contemporary example of the movement towards engineering systems of this kind. Every member of a squad is issued a radio or other communications device, with the result that information is no longer concentrated in and processed through a central command system, and the lowest ranking infantryman will often be better informed of the dynamically evolving state than will the commanding officer. Because the command hierarchy is still in place, the tension between the two is understandably the subject of much debate.

Computational science can also produce significant shifts in specific sciences. For example, general equilibrium theory, which dominated neo-classical economics for decades, is now being challenged by rival approaches such as agent based micro-economics and evolutionary game theory. These developments are sensible because humans tend to have a good insight into the nature of social and economic relations between individuals and much less of a firm grip on the kind of grand hyper-idealized theory that was once dominant.

4. What is Epistemically New About Computational Science?

The rise of computational science has allowed an enormous increase in scientific applications. But this expansion has also been accompanied by a shift in emphasis from what is possible in principle to what is possible in practice, with the countervailing result that the domain of science has also shrunk. Let me explain.

4.1 *In Practice, Not In Principle.*

One feature of computational science is that it forces us to make a distinction between what is applicable in practice and what is applicable only in principle. Here the shift is first, from the complete abstraction from practical constraints that is characteristic of much of traditional philosophy of science, and second from the kind of bounded scientific rationality that is characteristic of the work of SIMON and WIMSATT, within which the emphasis tends to be on accommodating the limitations of human agents. Ignoring implementation constraints can lead to inadvisable remarks. It is a philosophical fantasy to suggest, as MANFRED STÖCKLER does that "In principle, there is nothing in a simulation that could not be worked out without computers" ([2000], p. 368).[12]

In saying this I am not in any way suggesting that in principle results are not relevant in

---

[11] LUHMANN's culminating work is LUHMANN [1997], which is not yet available in an English translation. I am grateful to TIHA VON GHYCZY for conversations about various aspects of LUHMANN's thought.

[12] The first versions of THOMAS SCHELLING's agent based models of segregation, and the first versions of Conway's Game of Life were done 'by hand', but almost all contemporary simulations require abilities that go so far beyond what is possible by the unaided human intellect.

some areas. They clearly are; there are also other issues to which the philosophy of science needs to devote attention. One of the primary reasons for the rapid spread of simulations through the theoretically oriented sciences is that simulations allow theories and models to be applied in practice to a far greater variety of situations. Without access to simulation, applications are sometimes not possible; in other cases the theory can be applied only to a few stylized cases.

Within philosophy, there is a certain amount of resistance to including practical considerations, a resistance with which I can sympathize and I am by no means suggesting that the investigation of what can (or cannot) be done in principle is always inappropriate for the philosophy of science. One source of resistance to using in practice constraints is already present in the tension between descriptive history of science and normative philosophy of science, and in the tension between naturalistic approaches (which tend to mean different things to different people) and more traditional philosophy of science. But the appeal to in principle arguments involves a certain kind of idealization, and some idealizations are appropriate whereas others are not. A long-standing epistemological issue involves the limits of knowledge. Are there things that we cannot know, and if so, can we identify them? There surely cannot be any question that this is a genuine philosophical problem. Of course, it is not new—KANT famously gave us answers to the question. The question of what we can know, or more accurately, what we can understand, has been transformed by the rise of computational science and it is partly a question of what idealizations can legitimately be used for epistemic agents. We already have experience in what idealizations are appropriate and inappropriate for various research programmes. The move away from hyper-rational economic agents in micro-economics to less idealized agents mentioned earlier is one well-known example. For certain philosophical purposes, such as demonstrating that some kinds of knowledge is impossible even in principle, in principle arguments are fine. But just as humans cannot in principle see atoms, neither can humans in principle be given the attributes of unbounded memory and computational speed. This is the reason underlying epistemic opacity, one of the key epistemological features of the new methods.

4.2 *Epistemic Opacity*

One of the key features of computational science is the essential epistemic opacity of the computational process that leads from the abstract model underlying the simulation to its output. Here a process is epistemically opaque relative to a cognitive agent X at time t just in case X does not know at t all of the epistemically relevant elements of the process. A process is essentially epistemically opaque to X if and only if it is impossible, given the nature of X, for X to know all of the epistemically relevant elements of the process.[13] For a mathematical proof, one agent may consider a particular step in the proof to be an epistemically relevant part of the justification of the theorem, whereas to another, the step is sufficiently trivial to be eliminable. In the case of scientific instruments, it is a long-standing issue in the philosophy of science whether the user needs to know details of the processes between input and output in order to know that what the instruments display accurately represents a real entity.

---

[13]In my 2004, I used only the straightforward 'epistemically opaque' terminology. I now think that distinguishing between the weaker and stronger senses is useful. It is obviously possible to construct definitions of 'partially epistemically opaque' and 'fully epistemically opaque' which the reader can do himself or herself if so inclined. What constitutes an epistemically relevant element will depend upon the kind of process involved.

Within the hybrid scenario, no human can examine and justify every element of the computational processes that produce the output of a computer simulation or other artifacts of computational science. This feature is novel because, prior to the 1940s, theoretical science had not been able to automate the process from theory to applications in a way that made the details of parts of that process completely inaccessible to humans. Many, perhaps all, of the features that are special to simulations are a result of this inability of human cognitive abilities to know in detail what the computational process consists in. The computations involved in most simulations are so fast and so complex that no human or group of humans can in practice reproduce or understand the processes. Although there are parallels with the switch from an individualist epistemology, within which a single scientist or mathematician can verify a procedure or a proof, to social epistemology, within which the work has to be divided between groups of scientists or mathematicians, so that no one person understands all of the process, the sources of epistemic opacity in computational science are very different.

One of the major unresolved issues in many areas of computational science is whether the invention of new mathematical techniques might eventually replace some of these computational methods. I have frequently heard the suggestion that if we introduced a new class of functions that were solutions to the existing, currently intractable model, this would not change the way the model relates to the world. In fact it would, because with the availability of analytic solutions, the epistemic opacity of the relation between the model and the application would disappear. Moreover, even if this were to happen, the fact that the computational methods are, during our era, an unavoidable part of scientific method makes them of philosophical interest, just as the use of the Ptolemaic apparatus for computing planetary orbits is still of philosophical interest.

There are aspects of computational science that are simply not addressed by either of the two traditional philosophical accounts of theories. The traditional syntactic account of theories distinguished between some types of theories; those that were recursively axiomatizable, those whose axioms sets are only recursively enumerable, and a few other types. Computer scientists have since added to this classification, in moving from the simple issue of (Turing) computability to measures of theoretical computational complexity, such as P, NP, P-SPACE, and many others. This refinement can be incorporated within the syntactic account of theories. Other issues about the power of different computational architectures that are also relevant to computational science cannot be so incorporated. It is possible that if operational quantum or biological computers are built, a number of scientifically intractable problems will become tractable, opening up new areas of research. This is not an issue that is in any way addressed by traditional modeling techniques and although philosophical discussions of quantum computing have not been motivated much by issues in the area of simulations, the area is novel and is relevant to computational science. (See e.g. MERMIN [2007]).

4.3 *The Link Between Science and Technology*

The final issue to be addressed is the way in which progress in various sciences is now tied to technological advances in ways that go beyond the dependencies produced by a reliance on instrumentation. Computer simulations are crucially dependent upon computational load issues, and science must often wait until the next generation of machines is developed for these load demands to be accommodated. Technological issues arise in other ways as well: there are problems of extending models when substantial chunks of existing code are written in b adaption to later research; the former may require obsolete hardware to run. Philosophers of science are free to abstract from these issues, but then in some areas of science their accounts will simply misrepresent how progress is made.

Even with idealizations, these computational features are relevant. Here is one particular example: Determining energy levels is a core interest for molecular chemists. Physical chemistry employs quantum mechanics as its basic theoretical apparatus, but ab initio calculations of the energy levels are impossible to carry out for any but the smallest molecules. The simple valence bond and molecular orbital models do not provide accurate predictions even for hydrogen molecules, so they have to be supplemented with dozens of extra terms to account for various features. They therefore employ multiple approximations and are heavily computational. So the approximations chosen in the Hartree-Fock self-consistent field approach, a standard method of calculating ground state energies in ab initio quantum chemistry, are inextricably linked with the degree to which those calculations can actually be carried out in practice. On the other side there is now a growing sense that a different problem has arisen; that new techniques need to be developed to effectively exploit the massive computational power that is now available in many areas.[14]

5. Conclusion

Although some scepticism has been expressed about the novelty of computer simulations and related techniques (e.g. STÖCKLER [2000], FRIGG and REISS [2008]; for a response see HUMPHREYS [2008]), there is more than enough evidence to support claims that they constitute an important addition to the techniques of science, on a par with theoretical representations and experiment. I hope that what I have said above gives some insight into the more general intellectual consequences of these new ways of doing science.

[14] 'Rationale for a Computational Science Center', unpublished report, University of Virginia, March 2007.

References

ANDERSON, PHILIP W. [1972]: More is Different, *Science* 177: 393 – 396.

BEDAU, MARK [1997]: Weak Emergence, *Philosophical Perspectives* 11, pp. 375 – 399.

CARNAP, RUDOLPH [1928]: *Der logische Aufbau der Welt*. Berlin. (English translation published as *The Logical Structure of the World*, ROLF GEORGE (translator). Berkeley: University of California Press, 1967.)

FORD, KENNETH, CLARK GLYMOUR, and PATRICK HAYES [2006]: *Thinking About Android Epistemology*. Menlo Park, CA: AAAI Press.

FRIGG, ROMAN and JULIAN REISS [2008]: A Critical Look at the Philosophy of Simulation, *Synthese*. [forthcoming]

HACKING, IAN [1992]: 'Style' for Historians and Philosophers, *Studies in History and Philosophy of Science* 23, pp. 1 – 20.

HUMPHREYS, PAUL [2002]: Computational Models, *Philosophy of Science* 69, pp. S1 – S11.

HUMPHREYS, PAUL [2004]: *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.

HUMPHREYS, PAUL [2008]: The Philosophical Novelty of Computer Simulation Methods, *Synthese.* [forthcoming]

KÜPPERS, GÜNTER and JOHANNES LENHARD [2006]: From Hierarchical to Network-Like Integration: A Revolution of Modeling Style in Computer-Simulation, pp. 89 – 106 in: *Simulation: Pragmatic Constructions of Reality—Sociology of the Sciences, Volume 25*, JOHANNES LENHARD, GÜNTER KÜPPERS, and TERRY SHINN (eds). Berlin: Springer.

LAUGHLIN, ROBERT B. and DAVID PINES [2000]: 'The Theory of Everything', *Proceedings of the National Academy of Sciences* 97, pp. 28 – 31.

LENHARD, JOHANNES [2007]: Computer Simulations: The Cooperation Between Experimenting and Modeling, *Philosophy of Science* 74, pp. 176 – 194.

LUHMANN, NIKLAS [1997]: *Die Gesellschaft der Gesellschaft*. Frankfurt/Main: Suhrkamp.

MERMIN, N. DAVID [2007]: *Quantum Computer Science*. Cambridge: Cambridge University Press.

METROPOLIS, NICHOLAS [1993]: The Age of Computing: A Personal Memoir, pp. 119 – 130 in: *A New Era in Computation*, NICHOLAS METROPOLIS and GIAN-CARLO ROTA (eds). Cambridge, Mass: The MIT Press.

POPPER, KARL [1972]: Epistemology Without a Knowing Subject, pp. 106 – 152 in: KARL POPPER, *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.

REDHEAD, MICHAEL [1980]: Models in Physics, *British Journal for the Philosophy of Science* 31, pp. 145 – 163.

SCHWEBER, SAM and MATTHIAS WÄCHTER [2000]: Complex Systems, Modeling and Simulation, *Studies in History and Philosophy of Modern Physics* 31, pp. 583 – 609.

SHAPIRO, STEWART (ed) [2005]: *The Oxford Handbook of Philosophy of Mathematics and Logic.* New York: Oxford University Press.

STÖCKLER, MANFRED [2000]: On Modeling and Simulations as Instruments for the Study of Complex Systems, pp.355 – 373 in: MARTIN CARRIER, GERALD MASSEY, and LAURA RUETSCHE (eds), *Science at Century's End: Philosophical Questions on the Progress and Limits of Science*. Pittsburgh: University of Pittsburgh Press.

SUPPES, PATRICK [1962]: Models of Data, pp. 252 – 261 in ERNEST NAGEL et al (eds), *Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress.* Stanford: Stanford University Press.

VAN FRAASSEN, BAS [1980]: *The Scientific Image*. Oxford: The Clarendon Press.

VAN FRAASSEN, BAS [2004]: *The Empirical Stance*. New Have: Yale University Press.

WEINBERG, STEPHEN [1987]: Newtonianism, Reductionism, and the Art of Congressional Testimony, *Nature* 330, 433 – 437.

WINSBERG, ERIC [2001]: Simulations, Models, and Theories: Complex Physical Systems and Their Representations, *Philosophy of Science* 68, pp. 442 – 454.

WINSBERG, ERIC [2003]: Simulated Experiments: Methodology for a Virtual World, *Philosophy of Science* 70, pp. 105 – 125.

WIMSATT, WILLIAM C. [1974]: Complexity and Organization, pp. 67 – 86 in: *PSA 1972: Proceedings of the 1972 Biennial Meeting of the Philosophy of Science Association*, KENNETH SCHAFFNER and ROBERT COHEN (eds). Dordrecht: D. Reidel Publishing Company.