MODELING EPISTEMOLOGY:

EXAMPLES AND ANALYSIS IN COMPUTATIONAL PHILOSOPHY OF SCIENCE

Patrick Grim

Philosophy, Stony Brook
Center for Study of Complex Systems, Univ. of Michigan
Suite 700 Weiser, 500 Church St.
Ann Arbor, MI
patrick.grim@stonybrook.edu

ABSTRACT

What structure of scientific communication and cooperation, between what kinds of investigators, is best positioned to lead us to the truth? Against an outline of standard philosophical characteristics and a recent turn to social epistemology, this paper surveys highlights within two strands of computational philosophy of science that attempt to work toward an answer to this question. Both strands emerge from abstract rational choice theory and the analytic tradition in philosophy of science rather than postmodern sociology of science. The first strand of computational research models the effect of communicative networks within groups, with conclusions regarding the potential benefit of limited communication.   The second strand models the potential benefits of cognitive diversity within groups. Examples from each strand of research are used in analyzing what makes modeling of this sort both promising and distinctly philosophical, but are also used to emphasize possibilities for failure and inherent limitations as well.

Keywords: epistemology, computational modeling, philosophy of science, epistemic landscapes, agent-based

1   INTRODUCTION

Philosophy of science takes as its goal the understanding of the relation between scientific inquiry and the world, of scientific techniques, assumptions, methodologies, theory-formation, scientific testing, explanation, and confirmation. It is well-known, however, that contemporary philosophy has two distinct and sometimes warring traditions, often termed 'analytic' and 'continental.' The research surveyed here is entirely and unapologetically in the analytic tradition of philosophy of science, taking Carnap and Hempel as remote ancestors and perhaps naming Suppes and Skyrms as more immediate predecessors (Carnap 1928, Hempel 1966, Suppes 1993, Skyrms 1966, 1990). It has little contact with postmodern sociology of science, which might instead name Derrida, Lyotard and Latour as major figures, for example (Derrida 1976, Lyotard 1979, Latour and Woolger 1986, Latour 2007).

The work that follows is meant to highlight and analyze results from a recent 'computational turn' in analytic philosophy of science. Over the last ten or fifteen years, techniques of agent-based modeling and simulation have been applied to issues in philosophy of science in ways that only classical logic and Bayesian probability had been applied before (Grim 2006, 2009; Hegselmann & Krause 2006; Zollman 2007a, 2007b, 2010; Weisberg & Muldoon 2009; Muldoon 2013; Borg et al. 2018; Frey & Sessja 2018a, b). In a way that again distinguishes it from much postmodern philosophy of science, this work is distinctly realist rather than 'social constructivist' or anti-realist in character. The assumption is that there is indeed

a realm of scientific truth—the way the world is as a matter of fact—independent of us and of our efforts to find it. The central question being explored is one of optimal investigatory structure: what pattern of scientific communication and cooperation, between what kinds of investigators, is best positioned to lead us to that truth?

In what follows I use a string of examples to emphasize both the promise of this recent work and some limitations, and its fully traditional philosophical character in the sense outlined. Section 2 offers a recent history of social epistemology and traces a first strand in computational philosophy of science, one that models the effect of network structure among scientists in the exploration of epistemic landscapes. Section 3 pauses to offer some first philosophical reflections on both conclusions from that work and its general character. Section 4 introduces a second strand of research, emphasizing issues of diversity in scientific exploration. The lessons of failures and limitations in that second set of examples—suitably detected and documented within the research strand itself—are the subject of further philosophical reflection in section 5. Section 6 offers concluding reflections on the future of computational modeling in analytic philosophy of science.

2 COMPUTATIONAL PHILOSOPHY OF SCIENCE

2.1 The Advent of Social Epistemology

Epistemology is the study of knowledge. Traditional epistemology, however—the epistemology of Plato, Hume, Descartes, and Kant—has treated the acquisition and validation of knowledge on the individual level. The question for traditional epistemology was always how I as an *individual* can acquire knowledge of the objective world, when all I have to work with is my subjective experience.

Perennial questions of individual epistemology remain, but the last few decades have seen the rise of a very different form of epistemology as well. Anticipated in early work by Philip Kitcher and Alvin I. Goldman, *social* epistemology is now evident both within dedicated journals and across philosophy quite generally (Kitcher 1993, Goldman 1999, Goldman & Whitcomb 2011, Goldman & Blanchard 2018). I acquire my knowledge of the world as a member of a social group: a group that includes those inquirers that constitute the scientific enterprise, for example. In order to understand the acquisition and validation of knowledge we have to go beyond the level of individual epistemology: we need to understand the social structure, dynamics, and process of scientific investigation.

It is within this social turn in epistemology in the analytic tradition that the tools of computational modelling—agent-based modeling in particular—become particularly appropriate. The idea is to not an epistemological examination *of* agent-based modeling (though that is indeed of continuing interest—see Frey and Šešelja (2018a, b)) but an attempt to put agent-based modeling to work in understanding the optimal structure and dynamics of groups of epistemic investigators.

2.2 Networks, Communication, and the Exploration of Epistemic Landscapes

We start with a set of examples from a first strand of computational philosophy of science, emphasizing the role of networks and epistemic landscapes.

One might think that access to more data by more investigators would inevitably optimize the truth-seeking goals of community of investigators. On that intuition, faster and more complete result—the contemporary science of the internet—would allow a faster, more accurate, and more exploration of nature. Surprisingly, however, this first strand of modeling offers robust conclusions regarding the potential benefits of *limited* communication.

2

In the spirit of rational choice theory, much of this work was inspired by analytical work in economics on infinite populations by Venkatesh Bala and Sanjeev Goyal (1998), computationally implemented for small populations in a finite context and with an eye to philosophical implications by Kevin Zollman (2007a, 2010). In Zollman's model, Bayesian agents choose between a current method φ_1 and what is set as a better method φ_2, starting with random beliefs and allowing agents to pursue the investigatory action with the highest subjective utility. Agents update their beliefs based on the results of their own testing results—drawn from a distribution for that action—together with results from the other agents to which they are communicatively connected. A community is taken to have successfully learned when all agents converge on the better φ_2.

Zollman's results are shown in Figure 2 for the three simple networks shown in Figure 1. The communication network which performs the best is not the fully connected network in which all investigators have access to all results from all others, but the maximally distributed network represented by the ring. As Zollman also shows, this is also that configuration which takes the longest time to achieve convergence.


Figure 1: A 10 person ring, wheel, and complete graph. After Zollman (2007a).


Figure 2: Learning results of computer simulations: ring, wheel, and complete networks of Bayesian agents. Adapted from Zollman (2007a).

The further concept of an *epistemic landscape* has emerged as of central importance in this strand of research. Analogous to a fitness landscape in biology, an epistemic landscape offers an abstract representation of ideal data that might in principle be obtained in testing a range of hypotheses (Grim 2006, 2009; Weisberg & Muldoon 2009; Hong & Page 2004, 2007). Figure 3 usees the example of data that might be obtained by testing alternative medical treatments. In such a graph the xz plane represents particular hypotheses regarding the most effective combination of radiation and chemotherapy. Graph height on the y axis represents some measure of success: the percentage of patients with 5-years survival

3

on that treatment, for example. If an investigator uses radiation therapy at rate x, with chemotherapy at rate z, the result will be the proportion of successful treatments represented on the y axis hovering over that point.


Fig. 3: A three-dimensional epistemic landscape. Points on the xz plane represent hypotheses regarding optimal combination of radiation and chemotherapy; graph height on the y axis represents some measure of success.

An epistemic landscape is intended as an abstract representation of the real world that is being explored. The key word, of course, is 'abstract': few would argue that such a model is fully realistic either in terms of the simplicity of limited dimensions or the precision in which one hypothesis has a distinctly higher value than a close neighbor. As in all modeling, the goal is to represent as simply as possible those aspects of a situation relevant to answering a specific question (Grim et al. 2007, 2013): in this case, the question of optimal scientific organization. Epistemic landscapes—even those this simple—have been assumed to offer a promising start. As outlined in section 4.2, however, one of the deeper conclusions that has emerged is how sensitive results can be to the specific topography of the epistemic landscape.

Is there a form of scientific communication which optimizes its truth-seeking goals in exploration of a landscape? In a series of agent-based models starting with Grim (2006) and more fully developed in Grim et al. (2014), agents are dropped as communicatively linked explorers on an epistemic landscape. In such a design simulation can be used to explore the effect of network structure, the topography of the epistemic landscape, and the interaction of the two.

The simplest form of the results (Grim 2006) echo the pattern seen in different forms in Bala and Goyal (1998) and in Zollman (2007a, 2010), here played out on epistemic landscapes. Agents start with random hypotheses as points on the x-axis of a two-dimensional landscape. They compare their results (the height of the y axis at that point) with those of the other agents to which they are networked. If a networked neighbor has a higher result, the agent moves toward an approximation of that point (in the interval of a 'shaking hand') with an inertia factor (generally 50%, or a move halfway). The process is repeated by all agents, progressively exploring the landscape in attempting to move toward more successful results (for full details see Grim (2006)).

On 'smooth' landscapes of the form of the first two graphs in Figure 4, agents in any of the networks shown in Figure 5 succeed in finding the highest point on the landscape. Results become much more interesting for epistemic landscapes that contain a 'needle in a haystack' as in the third graph in Figure 4.

4

Fig 4: Two-dimensional epistemic landscapes. Adapted from Grim (2009).

ring radius 1   small world   wheel   hub   random   complete

Fig 5: Sample networks.

Using an inertia of 50% and a 'shaking hand' of .08, 50 agents networked as a ring with radius 1 converge on the global maximum in the 'needle in the haystack' landscape in 66% of runs. For a ring with radius 2, in which agents are connected to two others on each side, results drop immediately to 50% of runs in which agents find the global maximum. With a 9% probability of rewiring, the success rate of small worlds drops to 55%. Wheels and hubs have a 42% and 37% success rate, respectively. Random networks with a 10% probability of connection between any two nodes score at 47%. The worst performing communication network on a 'needle in a haystack' landscape is the 'internet science' of a complete network in which everyone instantly sees everyone else's result.

An extension of these results appear in Grim & Singer et al. (2013). There a small sample of landscapes is replaced with a quantified 'fiendishness index,' roughly representing the extent to which a landscape embodies a 'needle in a haystack.' Higher fiendishness quantifies a lower probability that hill-climbing from a randomly chosen point 'finds' the landscape's global maximum (for details see Grim & Singer et al. (2013)). Landscapes, though still two-dimensional, are 'looped' so as to avoid edge-effects also noted in Hegselmann and Krause (2006).

Here again results emphasize the epistemic advantages of ring-like or distributed network over fully connected networks in the exploration of intuitively difficult epistemic landscapes. Distributed single rings achieve the highest percentage of cases in which the highest point on the landscape is found, followed by all other network configurations. Total or completely connected networks show the worst results over all. Times to convergence are shown to be roughly though not precisely the inverse of these relationships.

3   LANDSCAPE MODELS: A FIRST PHILOSOPHICAL ANALYSIS

This first series of examples is offered in an attempt to show the structure, some of the promise, and some of the limitations of computational modeling in analytic philosophy of science.

The goal is an understanding of truth-finding maximization at that most abstract and general level, with the models functioning as philosophical arguments—something like thought experiments—rather than showing any attempt at prediction from real data. What these models suggest is that it is distributed

5

networks of communication between investigators, rather than full and immediate communication between all, that will—or at least *can*—maximize accuracy of findings: a more real grasp of the real world.

In the seventeenth century, scientific results were exchanged slowly, from person to person, in the form of individual correspondence. In today's science results are instantly available to everyone. In headline form, the clear conclusion from these models is that the communication mechanisms of seventeenth century science may be more reliable in exploring the real world than the internet communications of today. Zollman draws the corollary conclusion that loosely connected communities made up of less informed scientists might be more reliable in seeking the truth than communities of more informed scientists that are better connected (Zollman 2007a, p. 13).

Simple philosophical models of this sort are intended to offer not specific predictions but a general *understanding*: not only *that* certain exploratory structures can maximize navigation on the landscapes, but *why*. In all the models surveyed, more connected networks produce inferior results because agents move too quickly to salient but sub-optimal results: to local rather than global maxima. In Zollman's Bayesian models, the fact that updating is primarily local serves as a buffer against misleading results. Distributed networks are more robust to the occasional string of bad results because their effect is restricted to small regions rather than immediately communicated to everyone in the network. In the landscape models surveyed, connected networks result in all investigators moving toward the same point, currently announced to everyone as highest, skipping over large areas in the process—precisely where the 'needle in the haystack' might be hidden. In more distributed networks, local action results in a far more even and effective exploration of widespread areas of the landscape, with a predictably higher probability of finding any hidden needle. In terms of the conceptualization offered by John Holland, distributed epistemic networks optimize exploration rather than exploitation (Holland 1975).

In order to get a real grasp on the real world—in order to most effectively find the truth—how *should* we structure the funding and communication structure of our scientific communities? It is clear both from these results in their current form, and in further work along these general lines, that the answer may well be 'landscape'-relative: it may well depend on what kind of question is at issue what form scientific communication ought to take. It may also depend on what desiderata are at issue. The models surveyed emphasize accuracy of results, abstractly modeled. All those surveyed concede that there is a clear trade-off between accuracy or results and speed of community consensus (Zollman 2007a, 2007b, Grim et al. 2014). For many purposes, ethical as well as practical, it may often be far better to work with a result that is only roughly accurate but available today than to wait 10 years for a result that is many times more accurate but arrives far too late.

A range of further examples is offered in the next section, emphasizing several ways in which philosophical computational models of this sort can fail, even on their own terms.

4 MODELING DIVERSITY AND EXPERTISE

4.1 Division of Scientific Labor in Exploring Epistemic Landscapes

A second tradition of work in computational philosophy of science also uses epistemic landscapes, but attempts to model the effect not of network structure but of the division of labor and diversity within scientific groups. An influential but unfortunately flawed precursor in this tradition is the work of Weisberg and Muldoon (2009).

Two views of Weisberg and Muldoon's landscape appear in Figure 6. In their treatment, points on the base plane of the landscape represent 'approaches'—abstract representations of the background theories,

6

methods, instruments and techniques used to investigate a particular research question. Heights at those points are taken to represent scientific significance (following Kitcher 1993).


Fig. 6. Two visions of Weisberg and Muldoon's landscape of scientific significance (height) at different approaches to a research topic.

The agents that traverse this landscape are not networked, as in the earlier studies noted, except to the extent that they are influenced by agents near them on the landscape. What is significant about the Weisberg & Muldoon model, however, is that their agents are not homogeneous. Two types of agents play a primary role.

'Followers' take previous investigation of the territory by others into account in order to follow successful trends. If any previously investigated points in their Moore neighborhood have a higher significance than the point they stand on, they move to that point (randomly breaking any tie). Only if no neighboring investigated points have higher significance and uninvestigated point remain, followers move to one of those.

'Mavericks' avoid previously investigated points much as followers prioritize them. Mavericks choose *un*explored points in their neighborhoods, testing significance. If higher than their current spot, they move to that point.

Weisberg and Muldoon measure both the percentages of runs in which groups of agents find the highest peak and the speed at which peaks are found. They report that the epistemic success of a population of followers is increased when mavericks are included, and that the explanation for that effect lies in the fact that mavericks can provide pathways for followers: "[m]avericks help many of the followers to get unstuck, and to explore more fruitful areas of the epistemic landscape" (for details see Weisberg & Muldoon (2009), pp. 247 ff). Against that background they argue for broad claims regarding the value for an epistemic community of combining different research strategies. The optimal division of labor that their model suggests is "a healthy number of followers with a small number of mavericks."

Unfortunately, the model has turned out to be flawed, most disappointingly in terms of simple implementation errors in which >= was used in place of > (Alexander et al. 2015). Weisberg and Muldoon's software agents do not in fact operate in accord with their outlined strategies. As implemented, their followers tend to get trapped into oscillating between two equivalent spaces (often of value 0). When followers are properly implemented, it turns out that mavericks help the success of a community solely in terms of discovery by the mavericks themselves, not by getting followers 'unstuck' who shouldn't have been stuck in the first place. As support for the claim that division of labor and strategic diversity are important epistemic drivers, then, the philosophical argument of the Weisberg and Muldoon model proves inadequate.

7

4.2 Diversity and Differing Epistemic Landscapes

Alexander et al. never deny the general *conclusion* that Weisberg and Muldoon draw: that cognitive diversity or division of cognitive labor can favor social epistemic outcomes. What they deny is that the Weisberg and Muldoon model adequately *supports* that conclusion. There is a different modeling approach, built on a different model of diversity, that does support that conclusion (Hong & Page 2004, Page 2007). But it also supports a point that Alexander et al. emphasize: that the advantages of cognitive diversity can very much depend on the epistemic landscape being explored (Grim et al. 2019).

Lu Hong and Scott Page work with a landscape of 2000 points, wrapped around as a loop. Each point is assigned a random value between 1 and 100. Their epistemic individuals explore that landscape using heuristics composed of three ordered numbers between, say, 1 and 12. An example helps. Consider an individual with heuristic <2, 4, 7> at point 112 on the landscape. He first uses his heuristic 2 to see if a point two to the right—at 114—has a higher value than his current position. If so, he moves to that point. If not, he stays put. From that point, whichever it is, he uses his heuristic 4 in order to see if a point 4 steps to the right has a higher peak, and so forth. An agent circles through his heuristic numbers repeatedly until he reaches a point from which none within reach of his heuristic offers a higher value (Fig. 7).


Fig. 7 An example of exploration of a landscape by an individual using heuristics as in Hong and Page (2004).

Hong and Page score individuals on a given landscape in terms of the average height they reach across all 2000 points. But their real target is the value of diversity in groups. With that in mind, they compare the performance of (a) groups composed of the 9 individuals with highest-scoring heuristics on a given landscape with (b) groups composed of 9 individuals with random heuristics on that landscape. In each case groups function together in what has been termed a 'relay.' For each point on the 2000-poiint landscape, the first individual of the group finds his highest reachable value. The next individual of the group starts from there, and so forth, circling through the individuals until a point is reached at which none can achieve a higher value (Fig. 8). The score for the group as a whole is the average of values achieved in such a way across all of the 2000 points.


Fig. 8 An example of exploration of a landscape by the first two individuals of a group, using their heuristics in relay as in Hong and Page (2004).

8

What Hong and Page demonstrate in simulation is that groups with random heuristics routinely outperform groups composed entirely of the 'best' individual performers. They christen their findings the 'Diversity Trumps Ability' result. In a replication of their study, the average maximum on the 2000-point terrain for the group of the 9 best individuals comes in at 92.53, with a median of 92.67. The average for a group of 9 random individuals comes in at 94.82, with a median of 94.83. Across 1000 runs in that replication, a higher score was achieved by groups of random agents in 97.6% of all cases (Grim et al. 2019). Hong and Page also offer a mathematical theorem as a partial explanation of such a result (Hong and Page 2004). That component of their work has been attacked as trivial or irrelevant (Thompson 2014), though the attack itself has come under criticism as well (Kuehn 2017, Singer 2019).

The Hong-Page model solidly demonstrates a general claim attempted in the flawed Weisberg-Muldoon model: cognitive diversity can indeed be a social epistemic advantage. But the Hong-Page model also proves limited in ways that have not always been acknowledged in application (Anderson 2006, Landemore 2013, Gunn 2014, Weymark 2014). In particular, it proves sensitive to the precise character of the epistemic landscape employed.

Hong and Page's landscape is one in which each of 2000 points is given a random value between 1 and 100: a purely random landscape. One consequence of that fact is that the group of 9 best heuristics on different random Hong-Page landscapes have essentially no correlation: a high-performing individual on one landscape need have no carry-over to another. Grim et al. (2019) expands the Hong-Page model to incorporate other landscapes as well, in ways which challenge the general conclusions regarding diversity that have been drawn from the model but which also suggest the potential for further interesting applications.

An easy way to 'smooth' the Hong-Page landscapes is to assign random values not to every point on the 2000-point loop but every second point, for example, with intermediate points taking an average between those on each side. Where a random landscape has a 'smoothness' factor of 0, this variation will have a randomness factor of 1. A still 'smoother' landscape of degree 2 would be one in which slopes are drawn between random values assigned to every third point. Each degree of smoothness increases the average value correlation between a point and its neighbors. Grim et al. consider landscapes of varying 'smoothness' along roughly these lines, though with a randomization that avoids the lock-step intervals suggested (for details see Grim et al. 2019).

Using Hong and Page's parameters in other respects, it turns out that the 'Diversity Trumps Ability' result holds only for landscapes with a smoothness factor less than 4. Beyond that point, it is 'ability'—the performance of groups of the 9 best-performing individuals—that trumps 'diversity'—the performance of groups of random heuristics.

The Hong-Page result is therefore very sensitive to the 'smoothness' of the epistemic landscape modeled. As hinted in section 2.2, this is an indication from within the modeling tradition itself of the danger of restricted and over-simple abstractions regarding epistemic landscapes. As Grim et al. (2019) demonstrate, moreover, sensitivity is not to that modeling factor alone: social epistemic success depends on the pool of numbers from which heuristics are drawn as well, with 'diversity' showing strength on smoother landscapes if the pool of heuristics is expanded as well. Results also depend on whether social interaction is modeled using of Hong-Page's 'relay' or an alternative dynamics in which individuals collectively (rather than sequentially) announce their results, with all moving to the highest point announced by any. Different landscape smoothnesses, different heuristic pool sizes, and different interactive dynamics will favor the epistemic advantages of different compositions of groups, with different proportions of random and best-performing individuals (for details see Grim et al. 2019).

5   FAILURES AND LIMITATIONS

Rodney Brooks is quoted as saying 'the problem with simulations is that they are doomed to succeed.' That form of unfalsifiability would indeed be a scientific failing.

Beyond their inherent interest, what the examples of the preceding section emphasize is several ways in which computational philosophy *can* fail, like simulations generally, even on their own terms. It is much to the credit of this research tradition that those failures have been detected and documented within the tradition itself.

One ground for failure is failure of 'verification.' A model may not in fact do what it is 'supposed' to do: it may not successfully model the dynamics for which it conceived. That is the sorrow of the Weisberg and Muldoon model. As Alexander et al. demonstrate, its implementation did not match its conception. We have emphasized that philosophical models operate as something like thought experiments, within a context of something like logical argument. If the computational model does not operate as intended or represented, the thought experiment misfires. The conclusion offered for the argument, as Alexander et al. emphasize, is not in fact supported by the model.

Another ground for potential failure, even for fully verified models, is lack of robustness. This is clearly relative to the conclusions drawn from the model. The Hong-Page model has been widely appealed to as general support for general cognitive diversity in groups in general. It has been presented to NASA, has been cited in support of expected institutional effects of diversity requirements at UCLA, and has been appealed to in a brief before the Supreme Court in support of promoting diversity in the armed forces (Fisher v. Univ. of Texas 2016). As Grim et al, demonstrate, the model is simply not robust enough across its several parameters to support sweepingly general claims regarding diversity and ability or expertise. Is that a problem internal to the model, or an external matter of its interpretation or application? There is much to be said for the latter alternative. The model is and remains an interesting one—interesting often in the ways in which it *does* show sensitivity to different parameters. In that regard it is a model more sensitive to the complexity of epistemic groups and environments than the 'one size fits all' manner in which it has often been presented.

4   COMPUTATIONAL PHILOSOPHY OF SCIENCE: FUTURE PROSPECTS

Philosophical investigation is standardly marked by four characteristics. It is abstract, rather than concrete. It operates in terms of logical argument, rather than empirical data. Its goal is understanding at the most general level, rather than specific prediction or retrodiction. It is often normative, rather than descriptive. Analytic philosophers in particular can be expected to emphasize these characteristics.

There is room—indeed there is *need*—for forms of modeling that continue the tradition of philosophical modeling outlined in previous sections but that compromise at least one of these characteristics. The core change for which there is room—indeed for which there is a need—is a change in the second requirement. There is no reason why the questions we ask and the answers we seek are not questions and answers which we should approach using *both* empirical data and logical argument. What many of us would like to see is

Models in philosophy of science expanded to incorporate empirical data would admittedly be *less* philosophical. If they ended up being more informative—in whatever way they proved informative—that would seem a small cost. We have for too long been hampered by the borders of our disciplines. New techniques and methodologies, including computational techniques, need not come with any proprietary investment in established disciplinary boundaries.

REFERENCES

Alexander, J. M., J. Himmelreich and C. Thompson 2015. "Epistemic Landscapes, Optimal Search, and the Division of Cognitive Labor". *Philosophy of Science* vol. 82, pp. 424-453.

Anderson, E. 2006. "The Epistemology of Democracy." *Episteme* vol. 3, pp. 8-22.

Bala, V. and S. Goyal 1998. "Learning from neighbours". *Review of Economic Studies* vol. 65, pp. 565-651.

Carnap, R. 1928. *Der Logische Aufbau der Welt*. Leipzig: Felix Meiner Verlag. Trans as *The Logical Structure of the World*, R. A. George, Univ. of California Press 1967.

Derrida, J. 1976. *Of Grammatology*, tr. G. C. Spivak. Baltimore: Johns Hopkins University Press.

Frey, D. & D. Šešelja 2018a. "What is the Epistemic Function of Highly Idealized Agent-Based Models of Scientific Inquiry?" *Philosophy of the Social Sciences*. https://doi.org/10.1177/0048393118767085

Frey, D. & D. Šešelja 2018b. "Robustness and Idealizations in Agent-Based Models of Scientific Interaction." *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axy039

Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Goldman, A. I. and T. Blanchard 2015. "Social Epistemology". *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/epistemology-social/ Accessed August 22, 2018.

Goldman, A. I. and D. Whitcomb 2011. Eds,, *Social Epistemology: Essential Readings*. New York: Oxford University Press.

Gunn, P. 2014. "Democracy and Epistocracy." *Critical Review* vol. 26, pp. 59-79.

Grim, P. 2006. "Tangled Webs: The Philosophical Importance of Networks." The Marshall Weinberg lecture, University of Michigan.

Grim, P. 2009. "Threshold Phenomena in Epistemic Networks" *Proceedings, AAAI Fall Symposium on Complex Adaptive Systems and the Threshold Effect*, FS-09-03, AAAI Press.

Grim, P., R. Rosenberger, B. Anderson, A. Rosenfeld, and R. E. Eason 2013. "How Simulations Fail." *Synthese* vol. 190: 2367-2390.

Grim, P., R. Eason, E. Selinger and R. Rosenberger 2007. "What Kind of Science is Simulation?" *Journal of Experimental and Theoretical Artificial Intelligence* vol. 19, 19-28.

Grim, P. and Singer, D. J., S. Fisher, A. Bramson, W. Berger, C. Reade, C. Flocken and A. Sales 2014. "Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence". *Episteme* vol. 10, pp. 441-464.

Grim, P., D. J. Singer, A. Bramson, B. Holman, S. McGeehan and W. J. Berger 2019. "Diversity, Ability, and Expertise in Epistemic Communities," *Philosophy of Science* vol. 86, pp. 98-123.

Hegselmann, R. & U. Krause 2006. "Truth and Cognitive Division of Labour: First Steps toward a Computer Aided Social Epistemology". *Journal of Artificial Societies and Social Simulation* vol. 9. http://jasss.soc.surrey.ac.uki/9/3/10.html

Hempel, C. 1966. *Philosophy of Natural Science*. Englewood Cliffs, NJ: Prentice-Hall.

Holland, J. 1975. Adaptation in Natural and Artificial Systems. Cambridge, MA: MIT Press, Bradford Books.

Hong. L. & S. Page 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-ability Problem Solvrers." *Proceedings of the National Academy of Sciences* vol.101, pp. 16385-16389.

Kitcher, P. 1993. *The Advancement of Science*. Oxford: Oxford University Press.

Kuehn, D. 2017. "Diversity, Ability, and Democracy: A Note on Thompson's Challenge to Hong and Page". *Critical Review*, vol. 29, pp. 72–87.

Landemore, Hélène (2013). *Democratic Reason: Politics, Collective Intelligence, and the Rules of the Many*. Princeton, NJ: Princeton University Press.

Latour, B. 2007. Reassembling the Social: An Introduction to Agent-Network-Theory. Oxford: Oxford Univ. Press.

Latour, B. and S. Woolger 1986. Laboratory Life: The Construction of Scientific Facts. Princeton, NJ: Princeton Univ. Press.

Lyotard, J-F. 1979. *La condition postmoderne: rapport sur le savoir*. Paris: Minuit. Tran. as *The Postmodern Condition: A Report on Knowledge*, G. Bennington. Minneapolis, MN: Univ. of Minnesota Press, 1984.

Muldoon, R. 2013. "Diversity and the Division of Cognitive Labor". *Philosophy Compass*, vol. 8, pp.117–125.

Page, S. 2007. *The Difference*. Princeton University Press.

Singer, D. J. 2019. "Diversity, Not Randomness, Trumps Ability". *Philosophy of Science* vol. 86, pp. 178-191.

Skyrms, B. 1966. *Choice and Chance: An Introduction to Inductive Logic*. Dickenson 1966, 4th ed. Wadsworth 1999.

Skyrms, B. 1990. *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard Univ. Press.

Suppes, P. 1993. *Models and Methods in the Philosophy of Science*. Springer.

Thompson, A. 2014. "Does Diversity Trump Ability?: An Example of the Misuse of Mathematics in the Social Sciences." *Notices of the American Mathematical Society* vol. 61, pp. 1024-1030.

Weisberg, M. and R. Muldoon 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* vol. 76, pp. 225-252.

Weymark, John A. (2014) "Cognitive Diversity, Binary Decisions, and Epistemic Democracy," Vanderbilt University Center for the Study of Democratic Institutions, Working paper 5-2014.

Zollman, K. 2007a. "The Epistemic Benefit of Transient Diversity". Manuscript, University of California Irvine https://pdfs.semanticscholar.org/58f9/afc4df6d4c8d099698f9b4c9eaa65dfc1082.pdf

Zollman, K. 2007b. "The Communication Structure of Epistemic Communities", *Philosophy of Science* vol. 74, pp. 574-587.

Zollman, K. 2010. "Social Structure and the Effects of Conformity". *Synthese* vol. 172, pp. 317-340.

AUTHOR BIOGRAPHY

PATRICK GRIM is Distinguished Teaching Professor Emeritus with the Group for Logic & Formal Semantics, Dept. of Philosophy, Stony Brook University, and Philosopher in Residence with the Center for Study of Complex Systems at the University of Michigan. Grim publishes widely in modeling across a range of disciplines, is editor of the *American Philosophical Quarterly* and over 35 volumes of the *Philosopher's Annual*, and is co-author of *The Philosophical Computer: Exploratory Essays in Philosophical Computer Modeling*. His email address is patrick.grim@stonybrook.edu.