This article was downloaded by: [Australian National University]

On: 12 June 2014, At: 21:10

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK

Journal of Economic Methodology

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/rjec20>

Modeling herding behavior and its risks

Michael Weisberg ^a

^a Department of Philosophy , University of Pennsylvania , Philadelphia , USA

Published online: 25 Mar 2013.

To cite this article: Michael Weisberg (2013) Modeling herding behavior and its risks, Journal of Economic Methodology, 20:1, 6-18, DOI: [10.1080/1350178X.2013.774843](https://doi.org/10.1080/1350178X.2013.774843)

To link to this article: <http://dx.doi.org/10.1080/1350178X.2013.774843>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Modeling herding behavior and its risks

Michael Weisberg*

Department of Philosophy, University of Pennsylvania, Philadelphia, USA

Scientific communities can exhibit herding behavior when their members are reluctant to take research risks. This paper examines the epistemic consequences of herding behavior, as well as its surprising robustness. One way that herding behavior can be damped is by incentivizing research risks along the lines of recently announced National Science Foundation initiatives. Our models show that such incentives can have a small, but nevertheless significant impact on a scientific community's epistemic progress.

Keywords: cognitive labor; herding; agent-based model

Jel Codes: B41; C63

Philosophers of science have long been interested in the social structure of science and its impact on the scientific community's epistemic productivity. The philosophical literature refers to this topic as the *Division of Cognitive Labor*. Francis Bacon's analysis of how flourishing scientific communities should be structured is probably the earliest work on this topic (Bacon, 1620/2004). Bacon's project was thoroughly normative, with little attention paid to how science was actually practiced. With increased attention paid to these issues by historians and sociologists, and rethinking of traditional ideas about the role of community structure by philosophers, a more realistic picture of actual practice began to emerge.¹

Drawing on these insights, philosophers of science have developed new methods for studying the impact of community structure on scientific research. This paper is an attempt to apply some of these tools to the subject of this special issue: the risks of herding behavior in scientific communities. I will begin by outlining several new tools for studying the division of cognitive labor, including the one I will rely on in this paper. I will then discuss the issue of herding specifically, pointing to ways that this analysis could be extended to other issues.

1. Three approaches to cognitive labor

The recent philosophy of science literature contains three approaches to analyzing the division of cognitive labor: the *marginal contribution/reward* (MCR) approach, the *epistemic networks approach* and the *epistemic landscape approach*. This section discusses the basic ideas of all the three approaches and how they might be applied to the relevant issues.

1.1. *Marginal contribution/reward*

The earliest models of the division of cognitive labor were produced by Kitcher (1990, 1993) and Strevens (2003, 2006). In these models, which I call the MCR approach, the division of cognitive labor is treated as a community-level resource allocation problem. Imagine that the scientific community is trying to find the metabolic pathway of an important biological process and that the community could take several approaches to

*Email: weisberg@phil.upenn.edu

discover this pathway. To maximize its chances of discovering the pathway in the minimum amount of time, the community needs to find a way to divide its most important resource, scientists, across these different projects. The optimal allocation will be the one that maximizes the probability of finding the pathway, or of finding it in the shortest span of time or for the least cost.

While Kitcher and Strevens discuss how the community might calculate the optimum distribution of cognitive labor, the most interesting part of their analysis is conducted from the individual scientist's point of view. Strictly speaking, they model a representative agent, but it is easiest to think about the analysis from the point of view of a scientist newly entering the field. In their accounts, the scientist knows the current distribution of scientists to projects, as well as the success function for each project. Success functions represent the ability of the project to transform the cognitive resources of scientists into successful outcomes.

With the success function and the current distribution of cognitive labor, the model scientist can calculate its marginal contribution to the success of the different projects. In other words, the model scientist can figure out how much more probable a project's success would be if it were to join the project.

The reason this is important is because it is not in the interest of the scientific community, or the beneficiaries of scientific knowledge, for all scientists to join the project with the highest probability of success. By way of example, Kitcher imagines a scenario where there are two different approaches to finding the structure of an important molecule. One approach has a probability of success when a reasonable number of scientists work on the project. At the same time, there is another approach that has a relatively low probability of success, but this probability of success could be realized if a small number of scientists work on the project. In this case, assuming diminishing marginal returns for increasing numbers of scientists, the community would be better off if a small number of scientists worked on the second approach.

But if the scientists know their marginal contribution and if this somehow factors into a reward scheme, then incentives for a more optimal allocation can be developed. Consider the simple case, where scientists are motivated by credit alone. Assuming that the lion's share or all of the credit goes to the scientist who makes a discovery first, scientists will want to take into account both the probability of success of the project and the probability that they will be the first one to complete the project. The first consideration pushes scientists toward the project with the overall highest probability of success, but the second consideration pushes scientists toward projects that have fewer scientists working on them. This fact, Strevens (2003) has argued, explains why the scientific community has adopted the priority rule, the rule that whoever discovers something first gets all the credit.

1.2. *Epistemic networks*

The second approach to modeling cognitive labor focuses on how the social structure of science affects learning, confirmation, and the propagation of error. Scientists attach probabilities to hypotheses using information that they discover for themselves and information that they learn from other members of their community. In new works by Zollman (2007, 2010) and Grim, Rosenberger, Anderson, Rosenfeld, and Eason (forthcoming), lines of communication between scientists are represented by network graphs. Each node of these graphs represents a scientist and each edge represents a communication channel. By altering the connectivity of the graph, from the minimally

connected cycle to the maximally connected complete graph, Zollman and Grim et al. have simulated different communication structures in science.

The epistemic networks approach promises to be extremely helpful in studying divisions of cognitive labor that result in limited information to individual scientists. Zollman and Grim et al. have primarily used this method to address issues of confirmation, but one could also imagine studies of project choice in this framework as well. One could start with the MCR framework, but then include network structure to explicitly limit the information available to scientists in making a choice about how to divide their cognitive labor.

1.3. *Epistemic landscapes*

The third approach to modeling the division of cognitive labor is the epistemic landscapes approach, which I have developed along with Ryan Muldoon. This approach incorporates aspects of MCR and epistemic networks, but envisions scientific research as being more like foraging than either microeconomic optimization or Bayesian reasoning. Like epistemic network models, this approach severely limits the information available to scientists. The models focus, however, on MCR-like questions about how scientists choose to divide their cognitive labor. The approach starts with agents of very limited knowledge and rationality and tests how interesting dynamics can arise from the bottom up.

Like the other approaches to cognitive labor mentioned in this section, the epistemic landscape approach begins from the individual scientist's point of view. We start by considering the epistemic situation of individual scientists – what information they have, what their motivations are, etc. We then try to work out what rules scientists might follow when making decisions about which research projects to pursue and implement these rules in the models.

Epistemic landscape models are *agent-based*, which means that a community of scientists is represented as individuals. These individuals can be heterogeneous with respect to their knowledge and strategies for learning about the world. The remainder of this paper is about the epistemic landscapes approach.

2. **Constructing an epistemic landscape**

Epistemic landscape models begin by postulating a set of *research approaches*, narrow specifications of how a research topic is investigated. Research approaches specify the research questions being investigated, the instruments and techniques used to gather data, the methods used to analyze the data, and the background theories used to interpret the data.

Researchers make strategic choices when they choose or modify their approaches. In particular, they want to choose approaches that generate results of equal or greater significance to their current approach. So we need a way of representing the connection between approaches and the significance of the knowledge generated by using that approach.

For present purposes, let us confine ourselves to what Kitcher (1993) has called the *epistemic significance* of scientific knowledge. This is the purely scientific value of a result that the community agrees on. Financial and other pragmatic values are excluded for present purposes. Let us further suppose that all scientists are equally talented and that adopting an approach will reveal the approach's significant truths.

With these assumptions in place, we can construct model epistemic landscapes. The dimensions of such landscapes correspond to aspects of approaches, along with an

additional dimension of epistemic significance. Each point corresponds to an approach and a degree of epistemic significance. An example epistemic landscape is shown in Figure 1.

The second half of the framework brings us back to the scientist's point of view. In the real world, a scientist knows what approaches she has taken in the past and the success of these approaches. She also knows what approaches some of her colleagues had taken and how successful they were. What a scientist does not know is the entire topography of the epistemic landscape. She cannot be certain how significant untried approaches will prove to be. All she can do is to make inferences about these approaches and then try them out. These facts are all reflected in the epistemic landscape framework.

In epistemic landscape models, a scientist's current approach is represented spatially: Scientists' locations on the landscape correspond to their current approaches. The model scientist maintains some memory about where it has been before, so that it can determine later if it is going in a plausible direction. Further, by adopting a particular approach, scientist agents are able to determine the significance associated with their current approach. This corresponds to a real-world scientist conducting successful research and determining how significant her results actually are.

Since epistemic landscapes are supposed to be models of the social structure of science, agents need to have information about the activities and progress of other scientists. Real scientists devote time to reading the literature, attending conferences, and communicating with colleagues in order to learn what has been tried and what has been successful. In the models described below, this will be represented in a simplified manner, which roughly corresponds to reading the scientific literature or going to conferences. Agents will be able to see a limited range of other approaches, determine whether these approaches have been tried and, if so, discover the degree of significance associated with these approaches.

3. Spontaneous herding on epistemic landscapes

This section will introduce a very simple epistemic landscape model, which is described in detail in Weisberg and Muldoon (2009). Neither the landscape in this model nor the types of scientist agents are highly realistic. Nevertheless, as we will see, we can learn a lot by starting very simple and then making the model more complex. Indeed, starting with such a simple model will reveal the main issue to be considered in this paper: herding behavior.

Figure 1. An example epistemic landscape of the form used for the models in this paper.

3.1. Two Gaussian, three-dimensional landscape

Our first model landscape will be constructed in three dimensions. Two of the dimensions will correspond to two aspects of research approaches. The third dimension will correspond to epistemic significance. As many research domains contain more than one region of epistemic significance, we will use two Gaussian functions centered at different locations to calculate the significance for each research approach. In order to make the landscape tractable to simulation, we need to take two further pragmatic steps. First, we discretize the landscape so that *patches* instead of points correspond to approaches. Second, we wrap the landscape on a torus so that it has no edges. When we do this, we get a landscape much like the one depicted in Figure 1, which is the landscape used for the tests described in this article.

Although individual scientists will be seeking research approaches of high epistemic significance, it is useful to have a global measure of the community's progress. Let us call the ratio of explored approaches to total research approaches of non-zero significance the community's *epistemic progress*. Efficient, successful scientific communities will maximize epistemic progress in minimal time.

3.2. Three strategies

The next elements of our simple model are the scientist agents. Just as the representation of research approaches is highly simplified, so too will be our scientist agents. We will assume the following:

1. Agents are only interested in seeking epistemic significance.
2. All agents are equally talented at research, so they always learn the correct level of significance of their current approach.
3. Agents' memory lasts only one model cycle, so they remember what research approaches they adopted in the previous model cycle and the significance of that approach.
4. Agents can see the significance of any previously explored approach in their *Moore neighborhood*, the eight approaches adjacent in the epistemic landscape.
5. Agents have no way of estimating the significance of previously unexplored approaches.
6. Agents are initially distributed randomly in the low significance regions of the epistemic landscape.

With these general assumptions, it will be possible to formulate three exploration strategies.

3.2.1. Controls

The first strategy is called *control*. Agents following this strategy do not take into account what others are doing at all. So this represents the extreme case where scientists do not really divide their cognitive labor; each is trying to find the most epistemically significant approach for itself.

Control agents follow a simple *hill climbing with experimentation* procedure. This procedure can be described as follows:

1. Move forward one patch.
2. Ask: Is the patch I am investigating more significant than my previous patch?
If Yes: Move forward one patch.

If No: Ask: Is it equally significant as the previous patch?

If Yes: With probability 0.02, move forward one patch with a random heading. Otherwise, do not move.

If No: Move back to the previous patch. Set a new random heading. Begin again at Step 1.

This strategy is the baseline to compare our next two strategies, which add in a division of cognitive labor.

3.2.2. *Follower*

One extreme way for a scientific community to divide its labor is for each agent to imitate the best work it knows about, and then move on from there. A simple way to implement this rule is the *follower* strategy. Followers behave as follows:

Ask: Have any of the approaches in my Moore neighborhood been investigated?

If yes: Ask: Is the significance of any of the investigated approaches greater than the significance of my current approach?

If yes: Move toward the approach of greater significance. If there is a tie, pick randomly between them.

If no: If there is an unvisited approach in the Moore neighborhood, move to it, otherwise, stop.

If no: Choose a new approach in the Moore neighborhood at random.

Followers will thus follow the ‘trails’ of successful agents until they are surrounded by approaches of lower significance than the one they are currently adopting.

3.2.3. *Maverick*

The third initial strategy is the maverick strategy. Unlike followers who try to figure out what others are doing so that they can imitate the best, mavericks try to figure out what others are doing so they can *avoid* everyone else. These agents are mavericks in the sense that they will always choose to go it alone.

The maverick strategy can be described as follows:

Ask: Is my current approach yielding equal or greater significance than my previous approach?

If yes: Ask: Are any of the patches in my Moore neighborhood unvisited?

If yes: Move toward the unvisited patch. If there are multiple unvisited patches, pick randomly between them.

If no: If any of the patches in my neighborhood have a higher significance value, go toward one of them, otherwise stop.

If no: Go back to 1 patch and set a new random heading.

3.3. *Results*

The landscape and agents described above are identical to the ones Muldoon and I analyzed in our 2009 paper. Here is a brief summary of the results:

To assess the research potential of a community of controls, Muldoon and I focused on epistemic progress. Specifically, we asked: How much epistemic progress does the community of scientists make? How does this scale up as the number of scientists increases?

These questions were addressed by repeatedly employing the landscape shown in Figure 1, but by varying the initial positions and numbers of scientists. Figure 2 shows the results of these simulations. In it, we can see a number of interesting trends. First, epistemic progress of a community of controls increases linearly with increasing numbers of agents. It also increases linearly with simulation time, but this is not shown in this graph.

While we expected communities of followers to do better than communities of controls, because they imitate the best research approach, this is not at all that happens. Followers always do worse than controls.

Even more surprisingly, mavericks, who try to avoid all other agents, do the best by far are the most successful agents. As a comparison, with 100 mavericks, the community achieves 0.55 epistemic progress after 200 model cycles. With 400 mavericks, they achieve epistemic progress of 0.90 in the same time. This stands in stark contrast to followers, where the average epistemic progress of a community of 400 followers is only 0.17 on our diagnostic landscape, which means that only 17% of the significant approaches were discovered.

Before delving deeper, we can see the first lesson: While it may be useful for scientists to take into account what other have discovered, this information can be detrimental to progress when taken into account in the wrong way. Zollman (2007) draws a similar conclusion on the basis of his epistemic network models.

3.4. Herding

Why do the followers make so little epistemic progress? One way to answer this question is to look at the behavior of a population of followers in a typical simulation. Figure 3 depicts the equilibrium state of one run of the model with 200 followers. In this figure, we can see the paths followed by the individual agents. It shows that the vast majority of followers get stuck in low significance regions. Moreover, followers get stuck in low significance regions because their natural tendency is to cluster together, or, more colorfully, they exhibit *herding* behavior.

For comparison, Figure 4 shows the paths followed by 50 mavericks. All mavericks find the peaks, and it is clear from their paths that they take very different paths up to the peak,

Figure 2. Comparison of the epistemic progress of controls, followers and mavericks. Controls and mavericks measured after 200 cycles, followers after 1000.

Figure 3. Equilibrium state of 200 followers on an epistemic landscape (represented in two dimensions). Depicted paths are the paths each follower took to its equilibrium state.

which means they get a higher score for epistemic progress. The mavericks also clearly display anti-herding behavior. By virtue of their strategy, they avoid one another altogether.

Real scientists behave neither like pure followers nor like mavericks, yet scientific behavior is follower-like more often than not. Building on the research of other scientists is

Figure 4. Equilibrium state of 50 mavericks on an epistemic landscape. Depicted paths are the paths each maverick took to its equilibrium state.

essential for further progress and, perhaps more importantly, the only way to secure funding. True mavericks would need to continuously start from scratch in their research, never building off of anyone else's work nor sharing other people's results. And they would very rarely be awarded grants, as funding agencies require that a successful track record be demonstrated.

While follower-like behavior is common and encouraged, the herding behavior it encourages is worrisome. It seems that there is a substantial risk that followers will not find the most significant research approaches, and certainly will miss many other promising ones. So, the next question we must ask is whether this behavior depends on the particular idealizing assumptions we made, or is it an inherent tendency of anything like the follower strategy. Such questions are best analyzed by robustness analysis (Levins, 1966; Weisberg, 2006; Weisberg & Reisman, 2008; Wimsatt, 1981, 2007). In this case, we will engage in structural robustness analysis, to try to see if we can break the herding behavior by modifying the model.

4. Is herding a robust property of followers?

To investigate the robustness of herding behavior among followers, we can make modifications to the follower strategy that will allow followers to explore beyond their immediate neighborhoods. My first modification along these lines is to give the followers more information about other approaches.

4.1. Increased community size

In the simple version of the model, followers only have information about the visited approaches in their Moore neighborhoods, their immediately past approach, and their current approach. We can make a small modification to the model and change the definition of the neighborhood to something more flexible. For our new definition, I will rely on what is called an $r(n)$ neighborhood, all of the approaches in radius n from the current approach. Call the value of n the *community size*. The Moore neighborhood corresponds approximately to a community size of 1.4.

If we fix the number of followers to 150, but on successive runs allow n to range from 1 to 20, we can observe the effect of increasing community sizes on epistemic progress and herding behavior. In Figure 5, I plot the epistemic progress of follower communities with increased neighborhood sizes.

As we can see from this graph, there is an overall trend of linearly increasing epistemic progress with increasing community size. While populations of 150 followers with a community size of the Moore neighborhood achieve mean epistemic progress of about 0.10, community sizes of 20 achieve mean epistemic progress of 0.4, a four-fold increase.

There are some other interesting features of this graph. Although the overall trend is a linear increase, there is a significant reduction in epistemic progress with an initial increase in community size. This is most likely due to the same underlying herding tendency, which is magnified by the increased information set available with expanded communities. What happens is that when the community size is increased moderately, the typical follower's information set is still about low significance regions of the landscape. This has the effect of further trapping most followers in low significance regions. Only as the community size gets larger does the typical follower gain information about higher significance regions. This can be seen in Figure 6, which shows the behavior of followers with increasing community size.

Figure 5. Mean epistemic progress of populations of 150 followers with differing community sizes.

The final aspect of this graph to consider is the mean epistemic progress of populations at the higher end of the graph. Since the grid used in this simulation is 101×101 , a community size of 20 represents each scientist being able to see a little more than 12% of the total approaches from their current approach. Even with this degree of knowledge of the epistemic landscape, followers only manage to find about 40% of the significant approaches. This is about the same degree of epistemic progress exhibited by 50 mavericks. So, while followers' epistemic progress can be improved with increased community size, they never achieve a high degree of epistemic progress.

4.2. Experimentation

Another way to investigate the robustness of herding behavior is to introduce the possibility of followers engaging in risky research, which I will call *experimentation*. These agents calculate the standard next move of a follower, but push a little further into unexplored territory. In my implementation, the agents are still fundamentally followers because most of the time they engage in standard follower behavior. Moreover, even when they experiment, these agents will experiment along the lines of previously successful research.

Figure 6. Populations of 150 followers with increasing community sizes (r).

This perturbation to the original model is motivated by a program recently introduced by the National Science Foundation (NSF). The Creative Research Awards for Transformative Interdisciplinary Ventures (CREATIV) program sets aside \$24 million of the NSF's \$5.5 billion budget to fund potentially transformative research. If a program officer identifies a high risk, but potentially high payoff proposal, he or she can directly recommend the proposal for funding, bypassing the normal review process, which is known to error on the side of funding conservative, safe research (Mervis, 2011).

To implement experimentation in our low-dimensional epistemic landscape, we can introduce a new type of agent with the following properties: engage in standard follower behavior most of the time, but with some probability p , 'jump' two approaches in the direction that currently looks most promising. By varying the community size and the experimentation probability p , we can simulate followers with increasing experimentation rate.

As a baseline, let us consider the epistemic progress of 200 followers with community size of 5. This community's average epistemic progress is 0.10, meaning just more than 10% of the significant approaches were investigated after 500 model cycles. Holding the community size fixed, but increasing the experimentation probability to $p = 0.02$ resulted in a doubling of epistemic progress (0.20). This is a substantial increase, but still yields only about 20% exploration of the significant approaches.

To probe at the extremes, we can double the community size and further increase the experimentation probability to 0.1, which will result in quite a lot of experimentation in a community of 200 members (see illustrations of these studies in Figure 7). These changes result in a further improvement to the average epistemic progress, which is increased to 0.24. While significant, this is not much higher than what we see with the smaller experimentation rate of 0.02.

A more comprehensive picture of the role of experimentation can be seen in Figure 8. This graph makes clear that the majority of the improvement to epistemic progress happens with the introduction of a small amount of experimentation. There is only marginal further improvement after the initial introduction.

At the high end of experimentation, where followers engage in risky research 10% of the time, we only get about 24% of the significant approaches being explored. Recall that 50 mavericks with a small community size could get to epistemic progress above 0.4. This suggests that the herding property of followers and follower-like strategies is extremely robust and extremely hard to break. Increasing community size, hence increased

Figure 7. Three populations of followers with community size 5 after 500 model cycles. The first engages in no experimentation, the second experiments with probability 0.02 and the third experiments with probability 0.1.

Figure 8. Average epistemic progress for a population of 200 followers who engage in experimentation with community size 5.

communication, can help a little. Increasing experimentation through CREATIV-like programs can help communities make progress, but only to a small degree. This helps illustrate the robustness of herding behavior, and the systemic and epistemic risks associated with follower-like behaviors.

5. Future directions

The models described in this paper are extremely simple, and I have just scratched the surface on how they can be used to study the phenomenon of herding in economics, as well as many other aspects of the social structure of economics and other sciences. There are many ways in which these models could be developed, including taking into account actual sociological and psychological data that might be relevant to the phenomenon of herding. One might also work toward the construction of higher dimensional realistic landscapes based on real citation data. But I want to focus on a few purely theoretical extensions that might be of interest.

One interesting direction would be to combine some of the ideas from Zollman's epistemic network models into epistemic landscape models in order to create more realistic versions of the follower and maverick strategies. The work of some scientists will be particularly salient to others because of fame, skill, friendship, or other factors. It is reasonable to assume that followers only follow some scientists, and not necessarily just the ones employing nearby approaches. Similarly, mavericks may be comfortable following some very important leaders in their field, but otherwise they want to avoid what most other scientists are doing.

Another possible development is to add multiple incentives to the models. Rather than only being able to explore the landscape in search of epistemic significance, real investigators are also motivated by prestige, practical value, and monetary value. Additional dimensions corresponding to these other sources of value might be added to an epistemic landscape, and then strategies might involve a utility function balancing these kinds of considerations.

These are only a few possibilities. The framework is intended to be extremely flexible and my collaborators and I hope that many more applications and developments will be found.

Note

1. A wide range of philosophers including Giere (1988), Hull (1988), Solomon (1992), Kitcher (1993) and Thagard (1993) have emphasized that science involves the coordinated cognitive effort of many scientists. Closer to the aims of this paper is Solomon's (2001) work on social empiricism, which argues that normative scientific epistemology requires the assessment of internal as well as social factors. Sociologists of science have also discussed this issue, but their primary focus has tended to be the incentive structure of science (e.g. Merton, 1957) or the ways in which scientists navigate the complex relationships created by the division of research labor (e.g. Gerson 2008).

References

- Bacon, F. (1620/2004). *The Oxford Francis Bacon: The Instauration magna part II: Novum organum and associated texts*, Vol. 11. New York: Oxford University Press.
- Gerson, E. M. (2008). Reach, bracket, and the limits of rationalized coordination: Some challenges for CSCW. In M. S. Ackerman, C. Halverson, T. Erickson, & W. A. Kellogg (Eds.), *Resources, co-evolution, and artifacts: Theory in CSCW* (pp. 193–220), London: Springer-Verlag.
- Giere, R. N. (1988). *Explaining science: A cognitive approach*. Chicago, IL: University of Chicago Press.
- Hull, D. L. (1988). *Science as a process: An evolutionary account of the social and conceptual development of science*. Chicago, IL: University of Chicago Press.
- Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy*, 87(1), 5–22.
- Kitcher, P. (1993). *The advancement of science*. Oxford: Oxford University Press.
- Levins, R. (1966). The strategy of model building in population biology. In E. Sober (Ed.), *Conceptual issues in evolutionary biology* (1st ed., pp. 18–27). Cambridge, MA: MIT Press.
- Merton, R. K. (1957). Priorities in scientific discovery. *American Sociological Review*, 22, 635–659.
- Mervis, J. (2011). NSF creates fast track for out-of-the-box proposals. *Science*, 334(6058), 883.
- Grim, P., Rosenberger, R., Anderson, B., Rosenfeld, A., & Eason, R.E. How Simulations Fail. *Synthese*, forthcoming.
- Solomon, M. (1992). Scientific rationality and human reasoning. *Philosophy of Science*, 59(3), 439–455.
- Solomon, M. (2001). *Social empiricism*. Cambridge, MA: MIT Press.
- Strevens, M. (2003). The role of the priority rule in science. *Journal of Philosophy*, 100, 55–79.
- Strevens, M. (2006). The role of the Matthew effect in science. *Studies in History and Philosophy of Science Part A*, 37(2), 159–170.
- Thagard, P. (1993). Societies of minds: Science as distributed computing. *Studies in the History and Philosophy of Science*, 24(1), 49–67.
- Weisberg, M. (2006). Robustness analysis. *Philosophy of Science*, 73, 730–742.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2), 225–252.
- Weisberg, M., & Reisman, K. (2008). The robust Volterra principle. *Philosophy of Science*, 75, 106–131.
- Wimsatt, W. C. (1981). Robustness, reliability, and overdetermination. In M. Brewer & B. Collins (Eds.), *Scientific inquiry and the social sciences* (pp. 124–163). San Francisco, CA: Jossey-Bass.
- Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings*. Cambridge, MA: Harvard University Press.
- Zollman, K. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.
- Zollman, K. (2010). Social structure and the effects of conformity. *Synthese*, 172(3), 317–340.