Penn Libraries
UNIVERSITY of PENNSYLVANIA

University of Pennsylvania
ScholarlyCommons

Departmental Papers (Philosophy)   Department of Philosophy

4-2009

Epistemic Landscapes and the Division of Cognitive Labor

Michael Weisberg
University of Pennsylvania, weisberg@phil.upenn.edu

Ryan Muldoon
University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/philosophy_papers

Part of the Philosophy Commons

Epistemic Landscapes and the Division of Cognitive Labor

Abstract

Because contemporary scientific research is conducted by groups of scientists, understanding scientific progress requires understanding this division of cognitive labor. We present a novel agent-based model of scientific research in which scientists divide their labor to explore an unknown epistemic landscape. Scientists aim to find the most epistemically significant research approaches. We consider three different search strategies that scientists can adopt for exploring the landscape. In the first, scientists work alone and do not let the discoveries of the community influence their actions. This is compared with two social research strategies: Followers are biased toward what others have already discovered, and we find that pure populations of these scientists do less well than scientists acting independently. However, pure populations of mavericks, who try to avoid research approaches that have already been taken, vastly outperform the other strategies. Finally, we show that, in mixed populations, mavericks stimulate followers to greater levels of epistemic production, making polymorphic populations of mavericks and followers ideal in many research domains.

Disciplines
Philosophy

Epistemic Landscapes and the Division of Cognitive Labor*

Michael Weisberg and Ryan Muldoon[†‡]

Because contemporary scientific research is conducted by groups of scientists, under-standing scientific progress requires understanding this division of cognitive labor. We present a novel agent-based model of scientific research in which scientists divide their labor to explore an unknown epistemic landscape. Scientists aim to find the most epistemically significant research approaches. We consider three different search strat-egies that scientists can adopt for exploring the landscape. In the first, scientists work alone and do not let the discoveries of the community influence their actions. This is compared with two social research strategies: Followers are biased toward what others have already discovered, and we find that pure populations of these scientists do less well than scientists acting independently. However, pure populations of mavericks, who try to avoid research approaches that have already been taken, vastly outperform the other strategies. Finally, we show that, in mixed populations, mavericks stimulate followers to greater levels of epistemic production, making polymorphic populations of mavericks and followers ideal in many research domains.

1. Introduction. The complexity of contemporary science far exceeds the physical and cognitive resources of any individual scientist. Because of this, scientific research is almost always tackled by communities of sci-entists of varying size and degree of coordination.[1] In other words, modern science requires the division of cognitive labor.

1. Such coordination can include minimal consultation about techniques and materials, diving up tasks within a research team, or massive community-based projects where no individual understands the whole project or the contribution of each member.

While these facts about the nature of contemporary science are well-known to philosophers, having been discussed by Kuhn and Lakatos, among others, surprisingly little has been written about the epistemology of divided cognitive labor and the strategies scientists do and should use in order to divide their labor sensibly.[2] What makes such work especially challenging is the need to simultaneously keep track of the actions of individual scientists and of the epistemic progress of a scientific community. We have to account for how divided knowledge among individuals can give rise to unified knowledge in the community. At the same time, optimal epistemic behavior of individual scientists can give rise to undesirable collective behavior, and the extent to which cognitive labor is divided may affect this.

The complexity and philosophical richness of these issues provides many avenues for investigation. Much of the extant literature, however, has focused on two closely related questions: What are the optimal distributions of cognitive labor? How can we make incentives for scientists to divide themselves in the ways most beneficial to the progress of science? These questions have primarily been addressed with the use of marginal contribution/reward (MCR) models.

When applied to the division of cognitive labor, the MCR framework envisions the following scenario: scientists have to choose a project to work on. In order to do so, they calculate their marginal contribution to the epistemic success of this project and also their potential reward, based on the reward scheme in place. The most well-known of these models, those proposed by Philip Kitcher (1990, 1993) and Michael Strevens (2003), show that optimal distributions of cognitive labor can be achieved even if all scientists acted in self-interested ways, at least under a certain set of assumptions.

While we have been somewhat critical about the MCR approach's assumptions elsewhere (Muldoon and Weisberg 2008), we believe that these studies reveal a very important insight: scientists' micromotives can

---

2. Much has been written on the more general topic of the social nature of scientific knowledge. 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 our aims is Solomon's (2001) work on social empiricism, which argues that normative scientific epistemology requires the assessment of internal as well as social factors. Our focus in this article is somewhat different; we focus on how individual scientists make choices about dividing their labor, not the diversity of perspectives required to put scientific knowledge on firm epistemic footing. 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).

look epistemically impure or shortsighted, yet these motives can actually help the community as a whole make rapid progress toward finding out the truth. Thus a core tenant of strategic models about the division of cognitive labor is that what is epistemically good for individuals may differ from what is epistemically good for the community.

This article embraces this insight but develops models of the division of cognitive labor in a considerably different manner from Kitcher and Strevens. Rather than employing the MCR framework, we develop an agent-based approach, where individual scientists adopt strategies to explore what we call the 'epistemic landscape'. As we will show throughout the article, modeling the division of cognitive labor in this way allows for greater representational flexibility for modeling epistemic situations that are common to modern scientific communities.

We will argue that to be maximally effective, scientists need to really divide their cognitive labor, coordinating in such a way to take account of what other scientists are doing. We also show, albeit in a preliminary way, that a mixed strategy where some scientists are very conservative and others quite risk taking leads to the maximum amount of epistemic progress in the scientific community.

2. Science as a Landscape. While almost all modern scientific research takes place in scientific communities, not every division of cognitive labor is the same. In one kind of scenario, scientists choose between different approaches, all of which aim at the same narrow goal. Such situations, like the race to find the structure of DNA, to synthesize Taxol, and to complete the human genome project, have a special kind of epistemic structure: From the point of view of the community as a whole, the thing that matters most is getting to the final answer as quickly as possible. Once we know the structure of DNA, finding it out a second time isn't very useful. Repetition of crucial elucidation experiments may be important, but only insofar as they help establish that the structure was elucidated correctly. These types of cases with definitive goals and end points are especially well-suited to be studied with the MCR approach and have been analyzed in detail by Kitcher and Strevens.

Another type of scenario in which scientists divide their cognitive labor involves research on the same topic broadly construed, but with small differences in the activities and goals of particular scientists. For example, within the research program of synthetic biology (Benner 2003), a group of chemists successfully synthesized novel DNA nucleotides that function analogously to naturally occurring DNA bases. This initial synthesis by one group of scientists (Liu et al. 2003) led another to incorporate these bases into a strand of DNA, creating what they called 'xDNA' (Gao, Liu, and Kool 2005). Still another group studied the electronic properties and

stability of these molecules using computational methods (McConnell and Wetmore 2007). These individual episodes of research were independent, but they built off of one another. Progress by one research group affected the research trajectory of others, but unlike in Watson and Crick's elucidation of DNA's structure, a significant discovery made by one did not signal the end of the specific research topic.

This kind of division scenario seems to us to be more common, the sort that makes up the bulk of scientific progress. Even highly significant findings of the sort reported each week in *Science* and in *Nature* often result from research organized in this way. Our models are designed to capture and analyze this division by mapping information about the microstructure of scientific research to spatial components of an epistemic landscape. Division of labor is represented as the distribution of agents throughout the landscape and scientific change as the exploration of the landscape. We now turn to these details.

2.1. Components of the Epistemic Landscape. A single epistemic landscape corresponds to the research 'topic' that engages a group of scientists. Research topics can be individuated at broad and narrow scopes, but we will focus on relatively narrow scopes. The scope of our models approximately corresponds to the topic that a specialized research conference or advanced level monograph might be devoted to. For example, in psychology, the study of young children's abilities to engage in and reason about pretend play is a topic of the scope we have in mind. Similarly, the study of opiod receptors in chemical biology, critical phenomena in statistical physics, or plant chemical communication in biology are all topics of the appropriate scope for our models.

The second conceptual component of epistemic landscapes are 'approaches'. These are narrow specifications of how an individual scientist or research group investigates the topic. An approach includes:

1. the research questions being investigated,
2. the instruments and techniques used to gather data,
3. the methods used to analyze the data, and
4. the background theories used to interpret the data.

For example, among the researchers studying children's ability to engage in pretend play (a single topic), classes of approaches might involve investigating the developmental time course, the differences between individual and group play, how children play with peers versus adults, and so on. Within these classes of approaches would be the specifics of the research method, including where the population was drawn from, direct manipulation versus observational approaches, the props used to initiate pretend play, and the like. Finally, the same techniques aimed at the same

questions may yield different results in light of the background theories used to interpret the data. We mean for approaches to be individuated quite finely with respect to each of these four aspects.

The final main component of an epistemic landscape is the 'epistemic significance' of the results yielded by following a particular approach. Here we follow Kitcher (1993) in claiming that finding out true things about the world is extremely easy—all the null results collected in every laboratory tell us true things about the world, but many of these null results are not very scientifically interesting. What scientists really care about are significant true things.[3]

An important and foundational debate in philosophy of science concerns the source of scientific significance. A classical perspective holds that some facts have intrinsic scientific significance. A radical alternative holds that all judgments of scientific significance are merely the result of dominant ideologies and other political and social forces that influence scientists and scientific consumers as much as anyone else. Moderate positions acknowledge both the social origin of much of what we take to be important in scientific knowledge, but also that some questions and answers have significance internal to the goals and structures of science. Our model makes no commitment about the source of significance judgments. It only requires that the community of scientists working on the same topic would make the same or nearly the same judgments about significance.[4]

We now have the basic components necessary to construct the landscape: topics, approaches, and the significance of research conducted with these approaches. The boundaries of the landscape are delimited by the topic; the coordinates of the landscape correspond to approaches; and the topography of the landscape corresponds to significance. The epistemic landscape represents the hypothetical research approaches that could be taken with respect to a given research topic and the epistemic significance that research employing a particular approach can yield.

Strictly speaking, epistemic landscapes only need two dimensions, one for approaches and one for significance. But a realistic epistemic landscape

3. Technically, the significance in our model should be thought of as significance of the truth that is uncovered by employing a given approach.

4. In most real cases, this is an idealization, which we made here in order to keep the models simple and to bring them in line with Kitcher's and Strevens's. Real judgments about epistemic significance are likely to be colored by background beliefs and non-epistemic goals. In future, more complex landscape models, other kinds of judgments of significance such as 'technological significance', or monetary value might be represented as additional dimensions on the landscape. The scientists' overall judgments about the significance of the truths yielded by an approach will be a function of all of these dimensions and might vary between scientists.

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

would be of high dimensionality because approaches can vary along many dimensions, as discussed above. HIV drug researchers, for example, might share the research topic of treating HIV and the dimensions of their landscape might correspond to different drugs, to different treatment cycles, to different assessment protocols, and so forth. Each point in the landscape would represent a particular choice about these aspects of the approach along with the epistemic significance that employing the approach yields.

Although realistic landscapes are of high dimensionality, we will employ a three-dimensional[5] landscape in the models discussed in this article. The x and y coordinates of points on this landscape will correspond to aspects of the approach and the z coordinate will correspond to the epistemic significance of the truths yielded by adopting that approach. Further, to make our model manageable, we will discretize the topography, describing 'patches' centered around integer coordinates for x and y and having a particular significance value. Figure 1 shows such a landscape.

5. Our primary motivation for adopting the three-dimensional landscape was conceptual clarity and computational simplicity. More complex landscapes can be generated easily and they often yield additional local maxima. However, without making more specific real-world commitments about what the topography of a particular landscape represents, we believe that the prudent course is to keep the landscapes simple. Future investigations could profitably explore landscapes of higher dimensionality and greater ruggedness.

2.2. Scientist Agents. So far, we have described the scientist-independent parts of our model—the structure of the epistemic landscape and the information encoded in it. One of the major advantages of our models when compared to MCR models is that we can more realistically represent the actual epistemic situation of scientists who have limited knowledge about the landscape. Scientists do not see the whole landscape at the beginning of the simulation; they learn about the landscape by exploring it or observing others.

As these models are agent-based models, individual scientists or research groups are explicitly represented as individuals. Each scientist will have a series of agent variables including its position in the epistemic landscape, memory about the patches already explored, and a variable which describes the algorithm it uses to explore the epistemic landscape. In more complex models, the scientists can also have variables corresponding to individual utility functions, sets of skills, level of talent, prestige, resources, and so on.

How do scientists move through the epistemic landscape? This is one of the major areas of flexibility in this family of models. Different exploration rules can be explored, as can mixed strategies where subpopulations employ different rules.

To fix ideas, we begin by describing the simplest movement rule we employ called **HE**, for hill climbing with experimentation. Scientist agents following this rule, whom we call 'controls', only keep track of the significance of their current location and the significance of their previous location on the epistemic landscape. These scientists make all of their decisions on their own, evaluating how to find a better approach given where they have been. We consider these agents to be control agents because they make all of their decisions independently, as if they were the only ones engaged in research. They also resemble the epistemicaly pure agents in classical discussions of scientific rationality.

The agents start out distributed randomly through zero significance areas of the landscape and facing a particular direction, which we call their 'heading'. Controls employ the following movement rule each cycle of the model:

HE Rule:
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 2% probability, 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.

The **HE** rule is a very basic hill-climbing algorithm with the addition of an experimentation rate. Scientists moving around the epistemic landscape rely only on what they can detect themselves about the significance of a patch; indeed, they do not even notice if another agent is currently on the same patch. They move in the direction of increasing significance and if they get stuck in a low significance area, they will ultimately move in an experimental new direction and find a more significant area. The experimentation rate is not strong enough, however, to knock scientists off of a local maximum, unless the peak is constrained to a single patch. Agents following a hill-climbing algorithm of this type are guaranteed to find at least a local maximum in finite time (Russell and Norvig 1995, 111).

What does it mean for a scientist to visit a patch in these models? In the most abstract sense, it means that the scientist has explored that portion of the epistemic landscape. There are some good reasons to leave the interpretation at this level of abstraction, because there really is no additional structure in the model to guide a more concrete interpretation. However, because we want the models reported here to ultimately form the base of more realistic models, we believe further interpretation of scientists visiting patches is needed.

To give further interpretation, we need to be very clear about what is not included in the model. There is no notion of a research cost in the model. In each model cycle, every scientist is permitted to move. Similarly, there is no notion of the differential time it might take to fully investigate any patch of the epistemic landscape. Whether the patch has been previously visited or not and whether investigating the patch could yield significant truths or not does not constrain movement through the epistemic landscape. Finally, there is no notion of changing significance on the basis of what has happened in previous cycles of the model. In real science, when a highly significant part of the epistemic landscape has been well-explored, there is little to be gained by scientists further exploring that exact region. But in the models reported here, no aspect of this phenomenon—such as a finite number of possible publications per patch—will be accounted for.

Given this background, we interpret visiting a patch as follows: when a scientist visits a patch, this means that she tries to determine whether there is a significant truth to be determined at the patch. In other words, she tries to determine whether a particular approach will yield a significant truth. This might be accomplished by reading the literature, doing an experiment, or communicating with other scientists. Our models do not

distinguish between these possibilities and since research is costless, the yield of any of these approaches is equivalent. Further, our models assume that scientists are extremely talented at laboratory and library research. A visit to a patch will always yield a truth of the objective significance value associated with that particular approach.[6] Higher values of significance correspond to the discoveries of more epistemically significant truths. In real science, this probably corresponds to higher prestige publications or grants. But in our model, significance is simply a value[7] that can be collected.

The models we will analyze in subsequent sections are thus very idealized. We believe that this is the appropriate way to begin agent-based investigations of the division of cognitive labor and will allow us to get a handle on the basic dynamics of cognitive labor represented on epistemic landscapes. Despite these idealizations, we can make considerable progress in studying strategies for the division of cognitive labor and how these strategies interact when the community is polymorphic in strategy. Nevertheless, it would be premature to draw quantitative conclusions from the models before future work can investigate the robustness of the result upon relaxing these idealizing assumptions.

One of the most important differences between the epistemic landscape approach and the approach employed by Kitcher and Strevens is that in the MCR framework, information about the potential success of a project is embedded in the calculation performed by scientists. In our model, scientists do not have a global view of the landscape. They can only see the parts that they have explored, as well as any information they get from the exploration of others. This fundamental difference will become clearer as we now describe the behavior of the scientists in simulations which employ this framework.

3. Hill Climbing with Experimentation. By way of initial analysis of our models, we will describe a first series of simulations involving both simple

6. In a more complex model, one could more fully distinguish between landscape exploration and landscape exploitation. The question of which approaches yield significant truths would be settled by exploration, but actually generating those significant truths would involve some kind of investment into research using that approach which we call 'exploitation'. The current models only keep track of exploration, because keeping track of exploitation would require introducing a cost for doing research and a less trivial connection between finding a highly significant patch and actually yielding a significant, publishable truth.

7. From the point of view of the models' agents and from our own point of view, all that matters is the ordering of the significance values. Since the models described in this article are not intended to be correlated to any actual scientific topic, the absolute values are not particularly meaningful.

epistemic landscapes and scientists following the **HE** rule.[8] The epistemic landscapes used in this study are built on a 101 patch by 101 patch toroidal grid. Significance is determined by two Gaussian functions generated with similar parameter sets.[9] The baseline significance for a given grid patch is 0. At the boundary of a set of significant patches this jumps from 0 to 50, signaling entry to an area of epistemic significance.[10] From there, the significance grows according to the Gaussian function. Figure 1 is a three-dimensional representation of such a landscape, but in our subsequent discussion, it will be more straightforward to examine two-dimensional contour plots, where height is represented with lighter shades of gray. Figure 2 is such a representation.

This epistemic landscape is not meant to model any particular target scientific domain; however, we believe that it has several features which are common to many kinds of domains we wish to study. First, most domains have multiple approaches that will yield significant truths. Second, the approaches likely to yield significant results cluster together, and are not scattered randomly through the epistemic landscape. Third, there is likely to be more than one cluster of promising approaches in a given topic domain. It is thus necessary to represent multiple promising approaches, but we believe that two peaks are sufficient in these basic models.

The epistemic landscape is populated with control agents, scientists who follow the **HE** rule. As we have already discussed, when a population follows this rule, all of the agents will eventually find their way to one of the peaks. So while it is important for model validation purposes to ensure that all agents eventually find a peak on the epistemic landscape during a simulation, the most interesting things we can learn are about the short- and medium-term behavior of the model. Specifically, we will ask the following questions:

1. How fast does a community of controls find the two peaks of the epistemic landscape? How does this scale up as the number of sci-

---

8. All of the simulations described in this article were carried out using models constructed with Netlogo 4.0 beta (Wilensky 1999). The behavior of these models was subsequently verified in Netlogo 4.0. Code for the models, as well as example parameter sets, can be found from the authors.

9. Two-dimensional Gaussian functions have the form $f(x, y) = A \exp(-a(x - x_0)^2 + b(x - x_0)(y - y_0) + c(y - y_0)^2)$, where A is the amplitude, (x_0, y_0) the center, and the parameters a, b, and c control the spread of the function in three dimensions. For the studies described in this article, we used the parameter set $A = .75$, $a = .02$, $b = .01$, and $c = .02$ for the Gaussian centered at $(25, 25)$ and $A = .7$, $a = .01$, $b = .01$, and $c = .01$ for the Gaussian centered at $(-5, -5)$

10. Any boundary change yields the same behavior. We use the large jump for ease of visualization.

Figure 2. Two-dimensional representation of epistemic landscape. In the simulation, this grid is wrapped around a torus so that there are no edges. Lighter shades of gray correspond to greater significance.

entists increases?

2. If epistemic progress can be approximated as the percentage of significance yielding approaches discovered, how much epistemic progress does the community of scientists make? How does this scale up as the number of scientists increases?

To answer the first question, we ran a simulation experiment where 10 controls were placed randomly in zero significance areas of the epistemic landscape and their initial headings were also randomized. They were allowed to move around the landscape according to the **HE** rule. The simulation was cycled[11] continuously until each of the two peaks had been found by at least one scientist or else a time limit had elapsed. We set the time limit to 50,000 cycles, which pilot simulations suggested were long enough to ensure that the scientists landed on a peak. Based on this pilot data, we interpret a simulation that runs to the elapsed time to mean that all of the controls had piled on to a single peak, leaving the second one unvisited (which is a possible equilibrium state of our model). Moreover, despite being noncommittal on the amount of research time that one cycle

11. One cycle corresponds to each scientist agent following its rule set one time.

of the simulation corresponds to, we think that 50,000 cycles is far longer than the lifetime of most research topics. For the sake of generating an intuition about the time scales, imagine that a single model cycle corresponded to an average day of research, then 50,000 cycles would almost be 137 years.

The simulation was repeated 100 times using 10 controls and the standardized epistemic landscape described above. After each simulation, the random number generator was reseeded. We found that 95 times out of 100, the 10 controls successfully found both peaks of maximum significance. That is, they found the two most significant patches out of the 10,201 patches on the landscape. Among these 95 successful simulated communities, the time to finding the two significant peaks varied considerably from a maximum of 43,004 cycles to a minimum of 553 cycles. The mean for these runs was 6,075 with standard deviation 8,518 and the median was 2,553. More importantly, the length of runs is distributed in a heavy-tailed distribution, with 60% of the runs being completed in 4,000 cycles and 80% being completed in 10,000 cycles.

These results suggest that a small population of controls can run into trouble in a number of ways. First, it is simply the luck of the draw whether a small population will find the patches of maximum epistemic value in short order, after a grueling, long period of time, or ever. While the community will eventually find at least one peak, it may converge to a suboptimal situation, finding only one peak, and be stuck there forever. And if the community is especially unlucky, converging on this single peak may take a very long time. There will be huge variance in these facts and it can only be explained by random factors.

The next step is to see how these results change with increasing numbers of scientists working on the same research topic and hence located in the same epistemic landscape. We ran a second set of 100 simulations, increasing the number of controls to 20. This has two dramatic effects: it nearly ensures that the scientific community finds both peaks[12] and it halves the median time for the community to find both peaks. We continued analyzing the time to convergence by rerunning the simulation adding 10 scientists at a time. The result of these simulations is shown in Figure 3.

From these simulations, we learn that the probability of finding the approaches of maximal significance in a timely manner is strongly dependent on the number of independently working scientists in the community. Further, we learn that there are diminishing marginal returns for adding scientists. With this particular landscape, the differences between

12. While there is no guarantee that they will reach both peaks, it was true for all of our simulations and seems to be a high probability outcome.

Median Time to Find Approaches of Maximum Signifcance

Figure 3. Median time for the scientific community to find the two maximally signif-icant approaches on a two-Gaussian landscape. Each simulation was run 100 times.

groups of 10 scientists—after the model is already populated with 30 or more—are relatively small.

So far we have only looked at the community's ability to find the peaks of the epistemic landscape. While it is obviously important for the scientific community to find these peaks, much important research also happens on the slope of the peaks with significant, but nonmaximally significant approaches. The next step is to consider how controls fare in exploring these nonmaximal, but nevertheless significant portions of the epistemic landscape.

To determine this, we will define 'epistemic progress' to be the percentage of patches with significance greater than zero that have been visited by the community of scientists. Employing the same epistemic landscape as before, we will examine a series of small communities of scientists and determine how much epistemic progress these scientists make over set periods of time. The results of these simulations can be found in Figure 4.

As we can see in Figure 4, there is a linear relationship between the number of controls and the average epistemic progress of the community. For any of the fixed lengths of time that we measured, we can see that increasing the number of controls gives a linear increase in the average
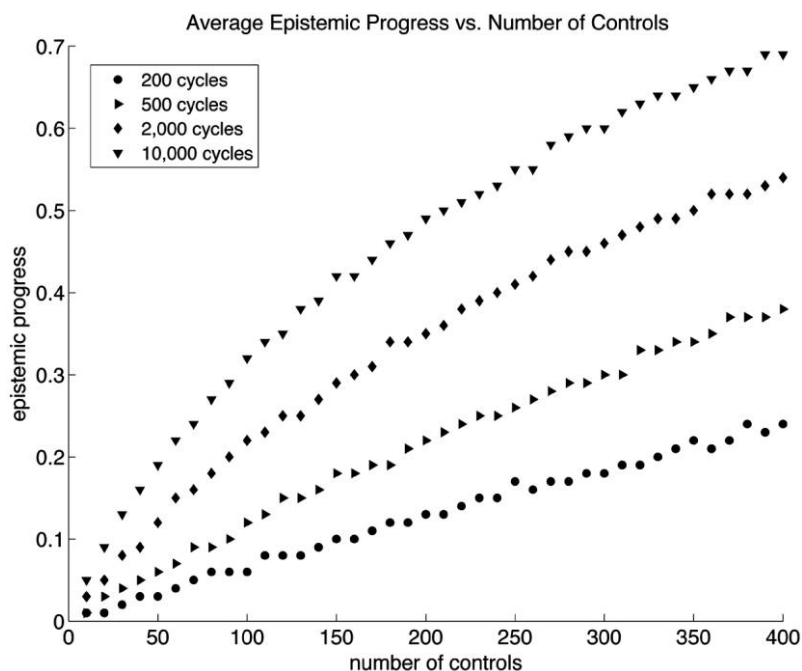
Figure 4. The average epistemic progress of scientific communities following the HE rule.

epistemic progress of the community. As we might expect, for any given number of controls, the longer we wait, the greater the epistemic progress.

A scientific community that adopts the **HE** rule as its way of exploring the epistemic landscape is neither very effective nor very efficient. Large populations of controls can achieve high degrees of epistemic progress, but it takes a considerable amount of time for this to happen. One reason for this, which we believe is revealed in the simulation data presented above, is that the scientists in such communities cannot learn from one another. Scientists do not take into account what other scientists are doing when they plan their next moves. In some sense, the community that follows **HE** isn't really dividing cognitive labor among its members. Each member of the community is acting as if it were the only member of the community. The progress of the community is achieved simply by the wisdom of the crowd. In the next section, we describe two new rules for epistemic landscape exploration, which do take into account what other scientists are working on and are more accurately described as rules that divide cognitive labor.

4. Followers and Mavericks. Controls apparently suffer because they cannot learn from one another. In the next two strategies we will describe, scientists are very strongly influenced by what their neighbors are doing and attempt to learn from what their neighbors have previously discovered. In one of these strategies, scientists are strongly biased in favor of doing what others have done; in the second, they try to avoid what others have done.

There are two simple ways that agents can learn from what other agents have done. The first is for them to explicitly learn from other nearby agents and the second is for agents to leave markers in the epistemic landscape signifying that a particular approach has been explored. In the strategies discussed in this section, we opt for the latter approach because distance on our landscape is not really physical distance. When our agents are near one another on the landscape, this means that they are working on similar things, not that they are in physical or even communicative proximity. Real scientists working on similar projects may communicate in the short term through talking but leave 'markers' in the form of publications for posterity. The marks left in our epistemic landscapes correspond to publications, in an abstract way. This allows agents to communicate to one another about what regions have been successfully explored.

4.1. Followers. In the next set of simulations, the agents will employ the strategy we call **Follow**, and we will refer to these agents as *followers*. These agents attempt to take information about previously successful approaches and use it to find approaches of even greater significance. Specifically, at the beginning of each cycle of the model, followers examine the patches in their Moore neighborhood, the eight patches immediately adjacent to the one on which they are currently located. Followers will then move to the previously explored approach of maximum significance in their Moore neighborhood, if such an approach is available. More specifically, followers execute the following decision procedure:

> **Follow Rule:**
>   Ask: Have any of the approaches in my Moore neighborhood been
>   investigated?
>   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 towards 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.

As with the community of controls, we first asked how quickly a community of followers can converge on the approaches of maximum epistemic significance and then evaluate the epistemic progress of follower communities over time. In order to make comparisons with the previous set of simulations, we will continue to employ the same two-Gaussian epistemic landscape.

To examine the time to convergence on maximally significant approaches, we ran simulations where followers were placed randomly in zero significance areas of the epistemic landscape. The simulation was cycled until at least one scientist had found each of the two peaks or else a time limit of 1,000 cycles had elapsed.[13]

We ran 100 simulations with 10 followers, then in subsequent batches of 100 simulations, we increased this number by 10 up to 200 followers. After each simulation, the random number generator was reseeded. With only 10 followers, not a single population managed to find both approaches of maximum significance and only 3% managed to find at least one approach of maximum significance. At the high end of this simulation with 200 followers, a single approach of maximum significance was found 60% of the time, with both approaches being found only 12% of the time. However, when the populations of followers did find both peaks, this happened very rapidly with an average time to converge on the two peaks (among the populations that did converge) of 56 cycles, which suggests that the randomly placed agents were near the boundary of significance at the beginning of the simulation. The data for the entire batch of simulations is shown as a histogram in Figure 5.

Turning now to the epistemic progress of communities of followers, we followed a similar procedure to our analysis of the control group. We ran simulations of populations of 10–400 followers for 1,000 cycles. In every case, the population quickly converged to its final value for epistemic progress and remained stationary throughout many of the 1,000 cycles. A typical time-course for this population is recorded in Figure 6.

As we can further see in Figure 7, adding additional followers does result in the community of followers making great strides in their epistemic progress. The average epistemic progress of a community of 400 followers

---

13. We use a much shorter maximum time limit in this study because the **Follow** movement rule does not ensure that scientists will find one of the peaks and because pilot simulations revealed that the model reaches its steady state after a relatively small number of cycles.

Frequency of Followers Finding Peaks of Maximum Significance
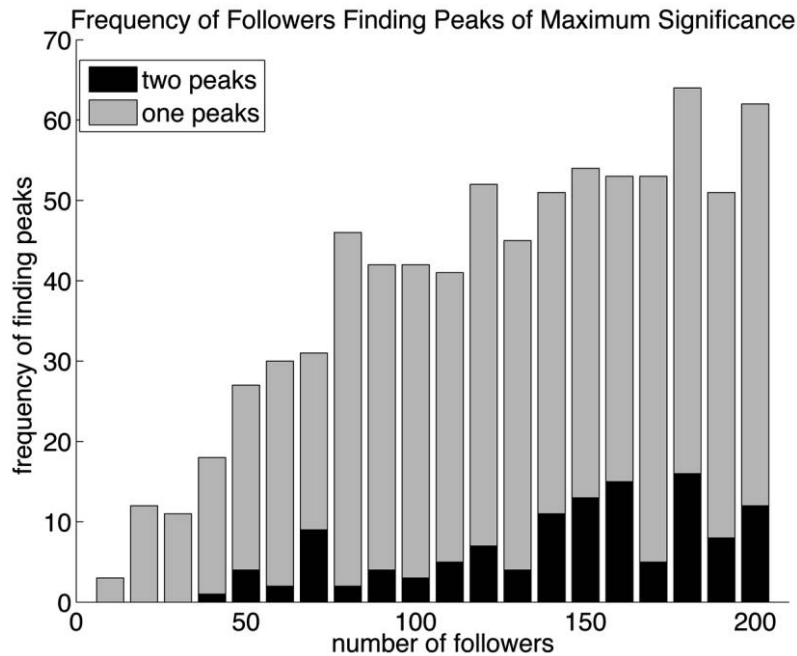

Figure 5. The frequency of convergence on approaches of maximum significance.

is 0.17, whereas the average epistemic progress for 10 followers was 0.0065. This contrasts poorly with communities of controls. In just 500 cycles, 400 controls progressed to epistemic significance level 0.24 and after 10,000 cycles, they reach 0.69. It took fewer than 300 control scientists 500 cycles to reach the maximum epistemic significance achieved by 400 followers in 1,000 cycles.

As can be seen from Figure 6, populations of followers tend to reach their equilibrium epistemic progress very rapidly. Once this is reached, the population ceases to move about the landscape. To further analyze this behavior, we traced the path of individual followers during the course of individual model runs, an example of which is shown in Figure 8. These plots show three behaviors of interest: first, clusters of followers who start out close to hills in the epistemic landscape follow each other up the hill. These are the only followers that ever make it on to the hills at all. Second, finding one's way onto the hill does not guarantee making it to the top, which strongly contrasts with the behavior of controls. If a control finds the edge of a hill, she will ultimately make it to the peak. However, if followers bump into each other on the way up, they can get stuck following

Figure 6. Typical time-course for 300 followers on a two-Gaussian epistemic landscape.

each other around on a suboptimal region of the hill. Finally, the vast majority of followers who start far away from the hills on the landscape never get close to the landscape because, if alone, they end up following their own trail. Or if around others, they end up circling around the trails each other make.

In the before and after pictures in Figure 8, we show how a population of 300 followers starts off and how it reaches its equilibrium. In the second picture, we have let each follower trace out its path. All three behaviors of interest are exhibited in these figures.

These three behaviors suggest that the high degree of coordination and learning from others exhibited by followers is simply not paying off. Populations of followers do not even make as much epistemic progress as the same sized population of controls. But is the problem coordination with other agents, or the way that followers coordinate? In the next section, we examine this question by analyzing a third strategy which we call the 'maverick' strategy.

4.2. Mavericks. Like followers, mavericks take into account which approaches have been previously explored and which ones were successful.

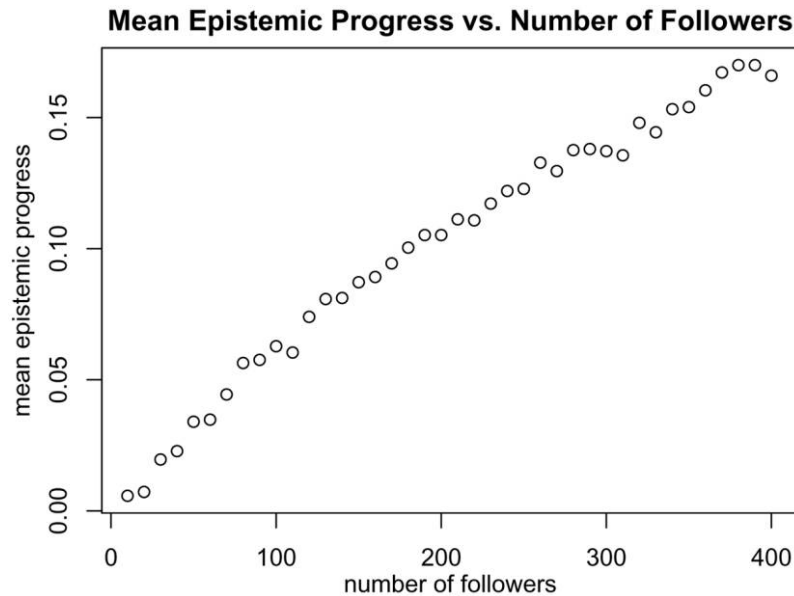**Mean Epistemic Progress vs. Number of Followers**

Figure 7. Epistemic progress of communities of followers after 1,000 cycles of the model.

However, unlike followers, mavericks avoid previously examined approaches, while followers emulate them.

At the beginning of each cycle of the model, mavericks examine the patches in their Moore neighborhood and execute the following decision procedure:

>**Maverick Rule:**
> 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 towards 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 towards one of them, otherwise stop.
> If no: Go back 1 patch and set a new random heading.

We first examined the mavericks' efficiency at finding the two approaches of maximum significance. Unlike in the case of the followers, 10 mavericks nearly always found both peaks (99% of the time), and 20

Figure 8. Exploration of the epistemic landscape by a community of 300 followers before (*A*) and after (*B*) movement begins. The tails behind agents are a plot of the paths they followed during the course of the simulation.

mavericks always found both peaks in our simulations. In addition, the mavericks are far more efficient at finding the peaks than controls. With 10 mavericks, the mean time to find both peaks was only 80 model cycles. With 100 mavericks, the mean time to find both peaks was 37 cycles and this is only slightly improved by adding 100 more mavericks to make a total of 200. With 200 mavericks, the average time to find both peaks is 33 cycles.

The mavericks are similarly impressive when we examine their epistemic progress: Large amounts of progress is made with very few agents in a very short amount of time. As with the controls and followers, we examined populations of 10–400 mavericks in increments of 10. We sampled the community's epistemic progress after 200, 500, and 2,000 cycles of the model.

As expected, the worst performance was with 10 agents and the shortest amount of time. The mean value for epistemic progress in this case was merely 0.10. In other words, 10% of the significant approaches had been found. After another 300 cycles, this hardly improves (0.12) suggesting that the source of this low value is actually the mavericks' efficiency at hill climbing. Populations of 10 mavericks find the peak approaches before they can explore a sufficient number of alternative approaches.

Increasing the number of mavericks drastically increases the epistemic progress of the community. With 100 mavericks, the community achieves 0.55 epistemic progress after 200 cycles. With 400 mavericks, they achieve epistemic progress of 0.90 after 200 cycles, meaning that nearly every significant approach has been explored.

Figure 9. Comparison of the epistemic progress of controls, followers, and mavericks. Controls and mavericks measured after 200 cycles; followers after 1,000.

As with the small number of mavericks, there is little change in epistemic progress after 200 cycles. With 100 mavericks, for example, the community gets to 0.63 from 0.55. With 200, it moves from 0.75 to 0.80. In all cases, populations of mavericks make the progress that they are going to make quickly and they find the maximally significant approaches quickly. We compare the epistemic progress of mavericks to followers and controls in Figure 9.

Having examined the performance of pure populations of controls, followers, and mavericks on the same epistemic landscape, we can now draw some preliminary conclusions. Mavericks are extremely efficient at finding peaks, and, due to the methods they use to find the peaks, they also make excellent epistemic progress. In contrast, followers are very poorly equipped to find the peaks of the landscape to make epistemic progress. The contrast between these two groups suggests that while it may be important to take into account information about what other scientists are doing, if one takes it into account in the wrong way, it can be disastrous. Like mavericks, populations of controls are pretty good at finding peaks; given enough time, they will always find at least one of the peaks. However, they take far longer to find the peaks and make far less epistemic progress per number of scientists then do mavericks.

So far we have only looked at pure populations, where all the scientists

follow the same strategy. In the next section, we report preliminary analysis of mixed populations. Of particular interest will be the effect that mavericks can have on populations of followers.

5. Mixed Populations of Mavericks and Followers. Our initial study of mixed maverick and follower populations[14] asks a very simple question: Does the addition of a single maverick to a large population of followers make a difference? Specifically, does this addition increase the epistemic progress of the community and does it alter the behavior of the followers in any other significant way?

In order to address these questions, we employed the same epistemic landscape as in the earlier studies with populations of 400 followers. We allowed the model to run for 1,000 cycles and measured the epistemic progress and the total number of approaches that were explored by the community. This was compared to a second set of populations, this time with 400 followers and the addition of a single maverick. The same measurements were taken.

For both measures, there was a significant difference between these samples. Adding a single maverick increased the epistemic progress of a population of followers by an average of 0.02 ($t = -5.74$, $p < 0.001$, two-tailed) and it increased the total number of approaches investigated by 232 ($t = -10.8$, $p < 0.001$, two-tailed). In contrast, the difference in average approaches investigated and average epistemic progress between a population of 400 followers and a population of 401 followers are not significant ($t = -1.47$, $p = 0.144$, two-tailed).

Since adding even a single maverick to a population of followers makes a significant difference, we conducted a series of mixed population studies to demonstrate the effect of systematically adding mavericks to populations of followers. Using the same epistemic landscape, we systematically studied populations of 10–400 followers, mixing in 10 mavericks at a time up to a maximum of 50 mavericks. After 500 cycles of the model, the epistemic progress was recorded. As expected, the added mavericks had several significant effects.

As we saw in Section 4, pure populations of followers make very little epistemic progress. Populations of 100 followers made epistemic progress of 0.07, and 400 followers only made an average progress of 0.17. In contrast, when just 10 mavericks are added, 100 followers improve to an average of 0.15 (a 214% increase), and 400 followers improve to 0.28 (a 165% increase). This suggests that even small additions of mavericks to

14. Although in this section we analyze mixed populations of agents employing pure strategies, this is equivalent to homogeneous populations employing mixed strategies (Maynard Smith 1982).

Figure 10. Epistemic progress of mixed communities of followers and mavericks after 500 model cycles.

populations of followers massively boosts the productivity of that population. We can see this result clearly in Figure 10, which plots the epistemic progress for different mixed populations of mavericks and followers.

We also note that the epistemic progress of these mixed populations is not due to the mavericks alone. If the progress were solely due to the mavericks, then the absolute increase in epistemic progress would be the same across mixed populations, with the one caveat that there would be a slight tendency to decrease progress as the number of followers increases. Instead we find the opposite trend: there is an increase of 0.08 in a population of 100 followers, whereas with 400 followers the increase is 0.11. This suggests there is an indirect stimulation of follower activity that accounts for the additional epistemic progress. Thus, the increase in productivity of the mixed population is due both to the direct actions of the mavericks, namely, their own efficiency at finding high significance approaches, and the effect that mavericks have on followers. Mavericks help many of the followers to get unstuck, and to explore more fruitful areas of the epistemic landscape.

Our final study of mixed populations of mavericks and followers examined what happens when the total number of scientists is held fixed,
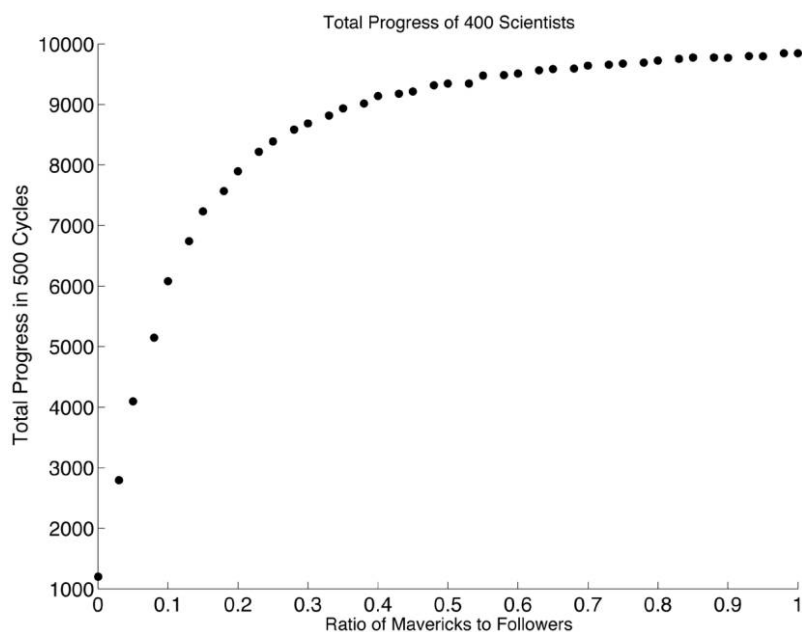
Figure 11. Average number of approaches investigated by mixed populations of followers and mavericks after 500 model cycles. The landscape consists of 10,201 approaches total. The population size is held fixed at 400 scientists, but the ratio of mavericks to followers is varied from 0 mavericks to 400 mavericks.

but the ratio of mavericks to followers was adjusted from 100% followers to 100% mavericks. This time, instead of just looking at what we have been calling 'epistemic progress', we consider the *total progress* of the community. We define 'total progress' as the total number of approaches investigated, whether significant or not. This measure allows us to see how much total activity is being performed by the scientific community, which we need to keep track of to fully understand the effect of strategy distributions in a population.

Figure 11 summarizes the results of this final analysis and the results are rather striking. The initial addition of mavericks (ratios of .02–.10) causes rapid tripling then quadrupling of the number of approaches investigated. Further small increases in the number of mavericks (0.10–0.40) take the population to around the 90% mark for the number of approaches explored in 500 cycles. Remembering that the mavericks quickly converge to the approaches of maximal significance themselves, this rapid increase in the number of projects explored is primarily a result of the increased stimulation of followers by mavericks briefly passing through their region.

Thus, in this set of simulations as in the others, we see the very significant indirect affect that mavericks have on research progress via their ability to stimulate the followers.

6. Dividing Cognitive Labor in Normal Science. The simulations described in this article only scratch the surface of what might be explored using epistemic landscape models. Landscapes can be made more rugged, they can contain more information, exploration strategies can take into account more information, an economy of money and credit can be included, and so forth. Much work remains to be done in realizing these possibilities, all of which we believe can be built within our existing framework.

Even with our current models and current landscape, we have observed a number of very interesting general trends about the division of cognitive labor. The first is a connection between cognitive labor and Thomas Kuhn's (1962) extensive discussions of 'normal science'.

Kuhn himself described normal science as 'puzzle solving', a way of articulating the details of a paradigm. One thing the behavior of agents in our model makes clear is that this is too simple a characterization for describing the division of cognitive labor in nonrevolutionary circumstances. All of our agents are doing normal science, yet some quickly converge to the maximally significant patches (and get papers in *Nature*), others find their way to significant areas, no doubt producing high quality, slightly derivative research, and others seem completely hopeless, marooned forever, employing approaches which can generate few results of significance. It is only these latter scientists who truly seem to be puzzle solving, at least in the most pejorative sense. The rest of the scientists are discovering significant truths, doing significant research. So one general lesson we might take away from this analysis is that one should make finer divisions among normal science activities if one is interested in cognitive labor.

Closely related to this is the differential suitability of the strategies for different kinds of normal science. Unsurprisingly, followers seem very well suited for puzzle solving—the simple articulation of details of a paradigm. Mavericks can partially fulfill this role, but their search patterns through the epistemic landscape are not particularly well suited for the kind of long term analyses required, for example, to add one more decimal place to a known constant. However, having a small population of mavericks in the midst of a larger population of followers helps the followers to puzzle solve. As we discussed in Section 5, even a few mavericks can cause the followers to explore a greater portion of the insignificant portion of their epistemic landscape, the regions we have associated with puzzle solving.

While the followers are good at puzzle solving, the mavericks are es-

pecially efficient at finding the peaks of maximum significance. As we showed, individual mavericks find the peaks extraordinarily quickly and indeed the whole population converges rapidly on those peaks. This means that if one wants to search the landscape rapidly for the most significant truths, one should employ a population of mavericks, at least as opposed to followers or controls. Even small populations of mavericks will be sufficient.

The maverick strategy of seeking out unknown epistemic territory has an important relationship to the class of problems that Kitcher and Strevens are most interested in. For those winner-take-all problems in which there is no particular value to discovering something a second time, mavericks have an important advantage. They converge on peaks very quickly because they do not duplicate the approaches of others. This strategy can thus be interpreted as a behavioral representation of the MCR approach favored by Kitcher and Strevens, insofar as both strategies seek out the greatest potential gains. Just as diminishing marginal returns discourage agents from joining projects that are already well populated, the maverick strategy avoids approaches that have already been tried. However, as our model can address a wider range of divisions of cognitive labor, we can also see how mavericks perform in more common scenarios.

We have also seen that in mixed populations, mavericks can provide pathways for followers to find the base of the peaks on the epistemic landscape. Once the followers find these bases, they are reasonably efficient at finding the tops. And mavericks can also stimulate followers to engage in pure puzzle solving, ensuring that the landscape is fully explored to find hidden significant approaches. Therefore, mixed populations of mavericks and followers are valuable divisions of cognitive labor.

The models presented in this article are, of course, highly idealized, even with respect to what we might accomplish in this framework. This makes drawing larger conclusions from them difficult, because we would want to know how robust the results we have discovered are to further perturbations and complexifications of the model. That said, we can draw some tentative conclusions about divisions of cognitive labor if we make one further assumption: Different strategies have differential costs. In particular, it is more costly to be a maverick than a follower.

That a maverick research strategy is more costly than a follower research strategy seems plausible because of the strategy's anticonservativism. Followers not only learn from their neighbors, but presumably they can borrow techniques, equipment, background research, and the like. They do not need to do everything for themselves. Mavericks, on the other hand, are studiously avoiding what has been done before and hence have to take a much larger research burden on themselves. Unless one had a very large research budget consisting of lots of money, supplies,

and helpers, it would be professionally, institutionally, and personally very costly to be a maverick.

If it is more costly to be a maverick, then optimum research communities are going to be composed of a healthy number of followers with a small number of mavericks. At this point, without considerably more detail added to our models, it is hard to say exactly what the optimum balance should be. The followers do the bulk of the puzzle solving, exploring every last corner of the epistemic landscape to make sure that there are no hidden patches of high significance. They also simply articulate the paradigm, which has an important role in science, even if it is not what garners one the most praise or glory. The mavericks have two roles. Though small in number, they are essential for stimulating the followers to expand their research horizons. They also do the majority of finding the most significant peaks, at least at first. But finding the most significant peaks is only part of the scientific enterprise; articulating the paradigm matters too. Thus a polymorphic population of research strategies thus seems to be the optimal way to divide cognitive labor.

REFERENCES

Benner, S. A. (2003), "Synthetic Biology: Act Natural", *Nature* 421: 118.
Gao, J., H. Liu, and E. T. Kool (2005), "Assembly of the Complete Eight-Base Artificial Genetic Helix, xDNA, and Its Interaction with the Natural Genetic System", *Angewandte Chemie International Edition* 44: 3118–3122.
Gerson, E. M. (2008), "Reach, Bracket, and the Limits of Rationalized Coordination: Some Challenges for CSCW", in M. S. Ackerman et al. (eds.), *Resources, Co-evolution, and Artifacts: Theory in CSCW.* Dordrecht: Springer-Verlag, 193–220.
Giere, R. N. (1988), *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
Hull, D. L. (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago: University of Chicago Press.
Kitcher, P. (1990), "The Division of Cognitive Labor", *Journal of Philosophy* 87: 5–22.
——— (1993), *The Advancement of Science*. Oxford: Oxford University Press.
Kuhn, T. S. (1962), *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
Liu, H., et al. (2003), "A Four-Base Paired Genetic Helix with Expanded Size", *Science* 302: 868–871.
Maynard Smith, J. (1982), *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
McConnell, T. L., and S. D. Wetmore (2007), "How Do Size-Expanded DNA Nucleobases Enhance Duplex Stability? Computational Analysis of the Hydrogen-Bonding and Stacking Ability of xDNA Bases", *Journal of Physical Chemistry B* 111 (11): 2999–3009.
Merton, R. K. (1957), "Priorities in Scientific Discovery", *American Sociological Review* 22: 635–659.
Muldoon, R., and M. Weisberg (2008), "Robustness and Idealization in Models of Cognitive Labor", manuscript.
Russell, S. J., and P. Norvig (1995), *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.
Solomon, M. (1992), "Scientific Rationality and Human Reasoning", *Philosophy of Science* 59: 439–455.

——— (2001), *Social Empiricism*. Cambridge, MA: MIT Press.
Strevens, M. (2003), "The Role of the Priority Rule in Science", *Journal of Philosophy* 100: 55–79.
Thagard, P. (1993), "Societies of Minds: Science as Distributed Computing", *Studies in History and Philosophy of Science* 24: 49–67.
Wilensky, U. (1999), "Netlogo", Center for Connected Learning and Computer-Based Modeling. Evanston, IL: Northwestern University, http://ccl.northwestern.edu/netlogo/.