New Approaches to the Division of Cognitive Labor

Michael Weisberg
University of Pennsylvania

forthcoming in *New Waves in Philosophy of Science*

Scientists are not lone agents, cut off from the outside world, responding only to information generated in their own laboratories. Rather, they make decisions about what to investigate by integrating what they discover for themselves with what they learn from others. They also take into account external factors such as grants, prizes, and prestige. These sources of feedback lead scientists to coordinate and divide their resources among differing approaches to the research domain. This coordination seems to enhance the success of scientific communities, but this coordination is neither planned nor explicit. Philip Kitcher has called this fact about scientific communities *the division of cognitive labor*.

The division of cognitive labor is one of the most striking features of modern scientific communities and has been argued to be a key component in their epistemic success (Gerson, 2008; Kitcher, 2002, 2003; Hull, 1988; Rosenberger, Grim, Anderson, Rosenfeld, & Eason, ms; Solomon, 2001; Strevens 2003, 2006; Weisberg and Muldoon, forthcoming; Zollman 2007, forthcoming). However, relatively little is known about the connection between these cognitive and social facts about science and the epistemic productivity of scientific communities. Research questions concerning the division of cognitive labor are both empirical and theoretical. Empirically, we might ask: What is the actual division of labor in particular research communities? How much overlap and repetition of research is found in the most successful communities? What is the normal extent and pattern of communication among scientists? How do scientists make choices about dividing their cognitive labor? What is the connection between this division and the division of expertise documented by psychologists (Keil, et. al., 2008)? While philosophers might contribute to these questions, they tend to fall into the domains of psychology and sociology of science.

But there are theoretical questions concerning cognitive labor that are more easily addressed with philosophical methods: What different types of divisions of cognitive labor are possible? How effective are these divisions for achieving scientific goals? Are there tradeoffs among these divisions? What kinds of individual motivations can lead to these divisions? How do restrictions of information and resources affect these choices and the division of labor which is an outcome of the choices? What kinds of incentives or structural features might the scientific community adopt to achieve better divisions of cognitive labor?

These questions are at the heart of new work in philosophy of science that investigates the social structure of science. This work examines how the division of cognitive labor contributes to science's epistemic power, and how it can create epistemic and social problems. A truly interdisciplinary part of philosophy of science, this work draws together philosophers, historians, sociologists, psychologists, computer scientists, and mathematicians. In this paper I will discuss three recent approaches to studying cognitive labor: Kitcher's and Strevens' *marginal/contribution reward*

(MCR) approach, Zollman's and Grim's *epistemic networks* approach, and the *epistemic landscape* approach I developed with Ryan Muldoon.

1. Marginal Contribution/Reward Approach

The best known approach for modeling cognitive labor has been championed by Philip Kitcher (1992, 1993) and Michael Strevens (2003, 2006). Kitcher and Strevens represent the division of cognitive labor 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 they can take several approaches to discovering 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 amount 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 point of view of the individual scientist. 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 account, 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, it can figure out how much more probable a project's success would be if it were to join the project.

On the basis of such models, Kitcher and Strevens make an important argument: Classical epistemic norms will lead scientists to misallocate their cognitive labor. If a classically rational, truth-seeking agent followed the procedure above, then she would join the project with the highest probability of success. But this isn't always what the scientific community as a whole wants to see happen. Maximizing the chance at success might involve distributing scientists across projects, not just to the projects with the best chance of success.

Returning to our example, imagine that one of the ways of investigating the biochemical pathway had a high probability of success when a reasonable number of scientists works 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 works on the project. In a case like this one, 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. However, this would not happen if everyone followed the classical epistemic norm: "take the approach most likely to lead to the truth."

To explain this discrepancy between classical epistemic norms and real scientific communities, Kitcher and Strevens draw on analyses from the history and the sociology of science. Scientists, they argue, are often motivated by more than just epistemic factors. Prestige, credit, money, jealousy, greed, and so forth, are all factors that may affect how scientists choose to divide their labor.

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 towards the project with the overall highest probability of success, but the second consideration pushes scientists towards 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.

This is only a sample of the utility of Kitcher's and Streven's framework for studying cognitive labor, but it illustrates the kind of question that it can be most usefully applied to. When a fixed number of projects are available, and the incentives, success functions, and distribution of cognitive labor are known to all scientists, analytical results can be derived in this framework. Many of these results show that the scientific community is often right to hedge its bets by distributing cognitive labor. One way this can be encouraged is if the community organizes incentives so that self-interested individual scientists can lead the community as a whole to better fulfill its epistemic goals.

Productive as it has been, the MCR framework has been less successful in modeling situations of imperfect information and bounded rationality. For example, the models require that every agent know what every other agent is doing. The next approach to cognitive labor is specifically designed to investigate situations where this is not the case.

2. Epistemic Network Approach

A second approach to modeling cognitive labor focuses on how the social structure of science affects learning, confirmation, and the propagation of error. Famous laboratories such as the wartime Los Alamos nuclear weapons laboratory (Rhodes, 1987) and MIT RADAR laboratory (Galison, 1997) planned ways for scientists to continuously integrate their findings and share ideas with one another. Technological innovations such as the Internet and rapid forms of electronic publication make this possible on a much wider and geographically distributed scale. But is this high degree of collaboration an unambiguously good thing? Might such communication lead to the propagation and fixation of errors as well as knowledge?

The *epistemic network approach* investigates these questions by combining ideas from confirmation theory and network theory. 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 work by Grim (Rosenberger, Grim, Anderson, Rosenfeld, & Eason, ms) and Zollman (2007; forthcoming), lines of communication between scientists are represented by network graphs, such the ones seen in Figure 1. Each node of these graphs represents a scientist and each edge a communication channel. By altering the connectivity of the graph, from the minimally connected cycle to the maximally connected complete graph, Zollman and Grim have simulated different communication structures in science.

[Figure 1]

Zollman's recent work provides both an illustrative example of the potential of this work, and also a surprising and startling conclusion. He asks us to consider the following abstract scenario: scientists are trying to determine whether the world is in state φ_1 or φ_2. He likens this to trying to determine the causative agent for some disease. Are peptic ulcers, for example, caused by hyperacidity or bacteria? In standard Bayesian fashion, every scientist has a prior probability on the hypotheses that φ_1 is true and that φ_2 is true, where no scientist attaches probability 1 or 0 to either hypothesis and where these two hypotheses are mutually exclusive.

Now imagine that data start coming in by experiment. As each scientist begins experimenting, she receives information about the world. This information is drawn from a distribution centered around the objectively true answer, taking in to account the fact that experiments produce noisy data. Scientists update their beliefs according to the data that they generate, as well as the data generated by the peers to whom they are connected in a social network.

Zollman took this basic idea and applied it to network models suitable for investigation by simulation. He iterated the scenario described above until the community converge either on the correct answer or on another state where the population is no longer able to take in new information. The scenario was studied with randomized initial priors and repeated 10,000 times per condition.

Although there are many subtleties to Zollman's results, two especially interesting conclusions can be drawn from this work. First, scientists connected in a cycle converged to the truth more often than scientists connected in a wheel or in a fully connected graph. This suggests the unintuitive conclusion that careful limiting of information available to scientists may have certain advantages. Or to put the point more bluntly, less well-informed scientists might have an advantage over more well-informed ones, if the goal is to minimize error. The second result, however, suggests an advantage of the more highly connected communities. When these communities converge to the truth, they do so more rapidly. For the ten-scientist communities Zollman studied, those on complete networks converged about five times faster than those in the cycle network.

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 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.

Despite its advantages, the epistemic networks approach still requires a lot of initial fixity. The hypotheses under consideration as well as the network structure need to be fixed ahead of time and do not change. In addition, the models are homogeneous with respect to scientists. Every scientist is epistemically equivalent to all the others, the only difference being in connectivity. The final approach to cognitive labor allows us to have heterogeneous populations and allows the space of projects and hypotheses to change "on-line," a more accurate reflection of the way science proceeds.

3. Science on an Epistemic Landscape

A third approach to modeling the division of cognitive labor is the *epistemic landscape* 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 considered in this paper, the epistemic landscapes 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, and so forth. 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.

Let's start from a concrete example. Consider a psychologist interested in how children's understanding of pretense develops. Given that this is a new area, there are few guidelines about what specific projects she might choose to work on. She does know that another lab has demonstrated the basic developmental milestones of spontaneous pretend play, but this only constrains her next move very modestly. Assuming that she is early in her career, a reasonable first move is to start off by replicating the previous finding or by choosing a very similar project. Then, as she makes small changes from this initial approach, she can determine the extent to which the resulting findings are epistemically significant.

Although this sounds straightforward in the abstract, there are actually many degrees of freedom making the choice of the next project non-trivial. Even if the scientist starts by investigating the developmental trajectory of pretend play and wants to "take the next step," what direction should she take it in? This is where computational models can help us start to get a handle on the process of project choice.

3.1 Epistemic Landscape

Epistemic landscape models begin by postulating a set of *approaches*, narrow specifications of how a research topic is investigated. 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.

In the psychology example described above, different approaches might involve investigating these phenomena in children of different ages, looking at the differences between individual and group play, considering how children play with peers vs. adults, etc. 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. Approaches are narrowly individuated in all of these respects and represented spatially in our models.

Scientists 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's confine ourselves to what Kitcher (2001) 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's further suppose that all of our scientists are sufficiently talented that adopting an approach will always yield whatever significant truths could be discovered by employing that approach.

With these assumptions, we can construct an *epistemic landscape*. The dimensions of this landscape correspond to aspects of approaches, along with an additional dimension of epistemic significance. Each point is thus an approach with a corresponding degree of epistemic significance.

Epistemic landscapes are the scientist-independent part of this representational framework. Here, only objective information about the world is represented. It takes the form: If such and such an investigation (narrowly construed) is performed by a competent scientist, then such and such results will obtain. All of this information is meant to be independent of any particular scientist.

In order to make the landscape tractable to simulation, we need to take three further, pragmatic steps. First, we discretize the landscape so that "patches" instead of points correspond to approaches. Second, we reduce the dimensionality of the landscape from the hundreds of dimensions in a realistic landscape to three: two for approach and one for significance. This reduction can be thought of as either "bundling" aspects of approaches together, or more plausibly as a subspace of the larger approach space. Finally, 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 2, which is the landscape used for the tests described in this chapter.

Figure 2

3.2 Scientist Agents

The second half of the framework brings us back to the scientist' 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 have taken and how successful they were. What a scientist doesn't know is the entire topography of the epistemic landscape. She can't 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 she has been before, so that she can determine later if she 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, learning 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.

The models Muldoon and I (forthcoming) have developed focus on approach choice, much like MCR models. Once our model scientists can determine the significance of the truths discovered with their current approach and their neighbors' approaches, they must decide what approach to try next. A major advantage of epistemic landscape models is that these decisions can be made differently by different agents; each can have its own strategy. Our preliminary analyses have investigated three types of agents with different in strategies: *controls*, *mavericks*, and *followers*.

3.3 Controls

We begin by considering model scientists with extremely limited knowledge: They only know about what approaches they have taken in the past and the significance of the truths generated by taking these approaches. These scientists are unaware of or uninterested in the rest of the community. Their only interest is in surveying the epistemic landscape in order to find areas of high significance. In addition, the scientists will engage in some degree of experimentation: behavior that deviates from their current research trajectory. This will keep the agents from getting marooned in zero significance areas. Because such scientists do not take other's discoveries into account, they will serve as our baseline case for investigating the division of cognitive labor.

In our epistemic landscape models, scientists choose their approaches by employing strategies. The controls deploy the following, very simple strategy:

CONTROL

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 its significance equal to the previous patch?

If Yes: With probability *e*, 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.

3. With probability *e*, change to a random heading and move forward.

As discussed above, this strategy restricts the agents' knowledge to what they can detect themselves about the significance of a patch. They will not even notice if another agent is currently on the same patch. The rule ensures that their movement will always be in the direction of increasing significance if they are on a gradient. If they get stuck in a low significance area, the experimentation probability will eventually cause them to move in a new direction and find a more significant area. Scientists following this search pattern are guaranteed to find local significance maxima in finite time. However, the experimentation procedure is not strong enough to knock scientists off of local maxima, unless the peak is constrained to a single patch. This means that the controls are guaranteed to find at least one approach of peak significance, but not necessarily both.

To assess the research potential of a community of controls, Muldoon and I asked 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 scientists increases?
2. 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 2, but varying the initial positions and numbers of scientists. With respect to the first question, we found that ten controls found both peaks 95% of the time, but that the time it took to find these peaks varied from about 550 to about 43,000 cycles of the model. This variance seems to be explained solely as a function of the initial distribution of scientists to approaches, which were randomly assigned, although agents were restricted to low significance areas. Increasing the number of agents ensured that both peaks were found (as happens 100% of the time in our simulations with 20 agents) and that the mean time to finding these peaks was rapidly diminished with increasing numbers of controls.

While it is obviously important for the scientific community to find the epistemic landscape's peaks, much important research also happens using non-maximally significant approaches. To investigate this, let *epistemic progress* be the fraction of patches with significance greater than zero that have been visited by the community of scientists.

Employing the same epistemic landscape as before, we discovered a linear relationship between the number of controls and the average epistemic progress of the community. For any fixed number of model cycles, increasing the number of controls gives a linear increase in the average epistemic progress of the community. Fixing the number of controls and increasing the number of model cycles also results in a linear increase of epistemic progress. In more recent work, my students and I have shown that this result is robust across many changes to the epistemic landscape (Werner, Naecker, Muldoon, and Weisberg, 2008). Specifically, changing the number, location, and size of the hills and peaks do not change these results in any appreciable, qualitative way.

Despite their modest success, a community of controls 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. An obvious reason for this is that controls cannot learn from one another or take into account what other scientists are doing when they plan their next moves. Communities of controls do not divide cognitive labor; each member of the community is acting as if it were the only member.

3.4 Getting Social: Followers and Mavericks

In the last section, I discussed the behavior of a classically rational agent: one that simply tries to find approaches yielding truths of the highest significance. This involves no division of cognitive labor at all. Each agent acts as if it was the only one. Understanding the behavior of these agents first provides a baseline against which we can now compare the behavior of scientists who do divide their cognitive labor.

Followers divide their cognitive labor by reasoning that the best way to find more significant truths about the world is to find the approach which has yielded the highest significance so far, and move in that direction. This is simulated in several steps. At the beginning of each cycle of the model, followers examine the patches in their Moore neighborhood, the 8 patches immediately adjacent to the one on which they are currently located. The agents 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

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.

A parallel set of simulations to the ones described above reveals that followers are surprisingly bad at finding peaks and making epistemic progress. With 200 followers a single peak (approach of maximum significance) was found 60% of the time, with both approaches being found only 12% of the time. In addition, as we can see in Figure 3, 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. This contrasts poorly with communities of controls. It took fewer than 300 controls to reach the same significance achieved by 400 followers.

FIGURE 3

These results suggest that the high degree of coordination and learning exhibited by followers is not an optimal strategy, at least when it is shared by the entire community. This result might lead one to think that learning from others inevitably leads to bad epistemic outcomes — indeed, this is one lesson we might learn from Zollman's models. But is the problem coordination with other agents, or the way that followers coordinate? The third strategy will let us investigate this question.

Like followers, *mavericks* pay attention to what others are doing, but they use this information differently. Instead of moving towards approaches yielding high significance, mavericks move away from explored territory.

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

MAVERICK

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.

Mavericks' performance is vastly superior to that of controls and followers. The maximally significant peaks are always found, even with very small populations. As with the controls and followers, we examined populations of 10 to 400 mavericks in increments of 10. 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. Mavericks are thus extremely efficient at finding peaks and, due to the methods they use to find the peaks, they also make excellent and rapid epistemic progress. Figure 3 makes the contrast between mavericks and followers clear and shows that, 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.

3.5 Polymorphic Populations

It seems highly unlikely that a real scientific community would be composed exclusively of mavericks or followers. Strategies for scientific inquiry vary between scientists and within careers. Some scientists are more individualistic and maverick-like; others follow trends and are follower-like. So what happens in the epistemic landscape model when populations are mixed?

An initial result for mixed populations is that the substitution of a *single* maverick for a follower results in a statistically significant increase of epistemic progress made by a community composed of followers. This increase of epistemic significance, however, is small. To investigate this effect further, Muldoon and I studied populations of scientist-agents with fixed size, but varied the ratio of mavericks to followers. We found that even small numbers of mavericks mixed in with followers causes the population to make considerably more epistemic progress.

To better understand why this happened, we adopted a second diagnostic measure, the *total progress* of the community. Total progress corresponds to the fraction of approaches investigated, whether these approaches are significant or not. This measure allows us to see how much total activity is being performed by the scientific community.

figure 4

Figure 4 summarizes the results, and we can see that 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. This massive increase in total progress is primarily a result of the increased stimulation of followers by mavericks. As mavericks pass through a region where followers are located, they effectively "unlock" approaches for investigation that the followers would have otherwise avoided. This means that in addition to their own direct effects on epistemic and total progress, mavericks make a very important indirect contribution to epistemic and total progress by stimulating followers.

Despite the fact that mavericks stimulate the followers to make considerable epistemic and total progress, pure populations of mavericks still perform better than equivalently sized mixed populations. Scientist for scientist, mavericks make more epistemic progress and more total progress than populations of followers. Yet as I mentioned above, scientific communities are typically polymorphic in strategy. Why might this be?

One possibility is that the maverick strategy is more costly both to individuals and to the community. In real communities, this could be due to the strategy's extreme anti-conservatism. Mavericks cannot learn from their neighbors, or even borrow their neighbor's techniques, equipment, background research and the like. 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 when resources are finite, optimum research communities should be composed primarily of followers, but with some mavericks included. In a series of new simulations, I have investigated this possibility by adding two additional aspect to the model: research resources and research cost.

I start off with a fixed pool of available resource units. Research costs are represented as the consumption of units of this resource per agent per cycle of the model. To investigate situations where mavericks are more costly than followers, the followers consume one unit per cycle, while mavericks consume ten units per cycle. If polymorphism in real populations of scientists is explained by research costs, then imposing these restrictions ought to reveal a tradeoff: In a fixed population of scientists with modest resources, there will be some optimum balance of mavericks and followers, mostly consisting of followers.

Investigation of the new, resource-limited models is ongoing, but two preliminary results are worth noting. The first result is that mavericks are always better at making epistemic progress than followers. My students and I expected that imposing modest resource limits and making each maverick more expensive than a follower would result in a tradeoff where a mixed population of mavericks and followers made maximal epistemic progress; we were incorrect. Mavericks are so much better than the followers that in most scenarios, pure populations of mavericks are better than mixed populations, despite the fact that they have many fewer cycles to investigate the landscape.

Figure 5

Only when we imposed severe restrictions — ones that caused a pure population of mavericks to only run for about 25 cycles — did we start to see the tradeoff (Figure 5b). Even here, the optimum balance of mavericks and followers is shifted much further in the direction of mavericks than we expected. While these models are still relatively simple, these results suggest that higher research costs cannot be the whole explanation for polymorphism in research communities.

Epistemic landscape models' flexibility for representing cognitive labor should be evident, even from this brief presentation. To date, however, they still lack some of the positive features of the other approaches. Unlike the MCR approach, epistemic landscape models do not provide an easy way of determining the *community-level* optimum distribution of labor. Different distributions need to be tried by simulation, upon specification of the measure of interest. Similarly, landscape models have not yet represented community structure other than Moore neighborhoods of approaches. It is easy to generalize the Moore neighborhood concept to bigger areas, but this is still a neighborhood of approaches, not of agents.

4. Connecting the Three Approaches

The models described in this chapter are all extremely simple. Each one of them can be extended in myriad ways, and some of their key results might be tested by looking at sociological and psychological data. Rather than describe these possible extensions and empirical tests, I want to close with a few comments about how the central ideas of the three modeling approaches might inform one another in future work.

One way of uniting the approaches is to take conclusions generated by MCR models, and investigate the conditions under which they could arise using network and landscape models. For example, say we wanted to discover the conditions under which the scientific community would adopt the Priority Rule. An agent-based model, similar to the ones used in the epistemic landscape approach, could be used to relax some of Strevens' key assumptions, seeing under what conditions the Priority Rule produces the optimal distribution of labor. Another model might incorporate an evolutionary dynamic (with mentors "giving birth" to students) and show under what conditions the Priority Rule would spontaneously arise. The approaches might also make contact when additional sources of value or significance (practical, monetary, etc.) are included in the landscape. Here one might investigate the interaction of reward schemes and strategies with different epistemic landscapes.

A second potential connection between the three approaches would be to incorporate epistemic networks into MCR and landscape models. Their incorporation into MCR models would be one way to introduce bounded rationality into models that otherwise assume that representative agents have access to all relevant information. One might also imagine using epistemic networks in 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.

Finally, ideas from the landscapes approach might be fruitfully adopted in epistemic network models. While the connectivity in network models can be varied, the actual epistemic strategy of each scientist is homogenous and based on simple Bayesian reasoning. As the landscape approach shows, interesting results arise when strategies are varied and populations are polymorphic. Similar tests could be deployed using epistemic networks, where the agents at each node reason differently.

References

Hull, David (1988) *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*, (Chicago: University of Chicago Press).

Keil, F.C., Stein, C., Webb, L., Billings, V.D., & Rozenblit, L. (2008) 'Discerning the Division of Cognitive Labor: An Emerging Understanding of How Knowledge is Clustered in Other Minds,' *Cognitive Science, 32(2),*259-300.

Kitcher, Philip (1992) 'The Division of Cognitive Labor,' *Philosophy of Science*, 87, 5-22.

___ (1993) *The Advancement of Science*, (New York: Oxford University Press).

___ (2001) *Science, Truth, and Democracy*, (New York: Oxford University Press).

Galison, Peter Louis (1997) *Image and Logic: A Material Culture of Microphysics*, (Chicago: University of Chicago 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*, (Springer-Verlag), 193–220.

Rosenberger, Robert, Patrick Grim, Brian Anderson, Adam Rosenfeld, and Robb Eason, (ms) "Science in Simulation: Structure and Evaluation"

Rhodes, Richard (1986) *The Making of the Atomic Bomb*, (New York: Simon & Schuster).

Strevens, Michael (2003) 'The Role of the Priority Rule in Science,' *Journal of Philosophy*, 100:2, 55-79.

___ (2006) 'The Role of the Matthew Effect in Science,' *Studies in the History and Philosophy of Science*, 37, 159-170.

Solomon, Miriam (2001) *Social Empiricism*, (Cambridge, Mass: MIT Press).

Weisberg, Michael and Ryan Muldoon (forthcoming) 'Epistemic Landscapes and the Division of Cognitive Labor,' *Philosophy of Science*.

Werner, Mark, Benjamin Naecker, Ryan Muldoon, and Michael Weisberg (2008) 'Robustness Tests on Epistemic Landscapes,' [archive address].

Zollman, Kevin J.S. (forthcoming) 'Social structure and the effects of conformity,' *Synthese*.

Zollman, Kevin J.S. (2007), 'The communication structure of epistemic communities. *Philosophy of Science*74(5): 574-587.

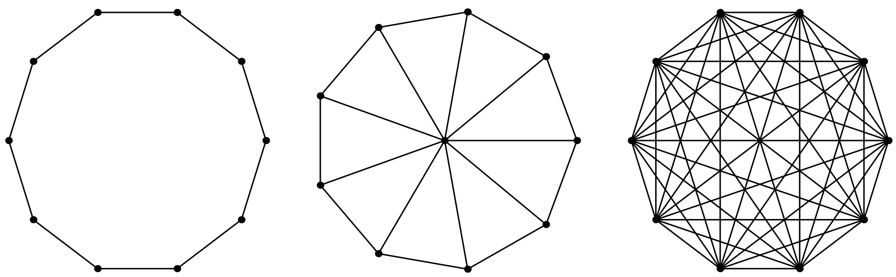

Figure 1: Zollman's three epistemic networks with 10 nodes. From left to right: the cycle, the wheel, and the complete network.

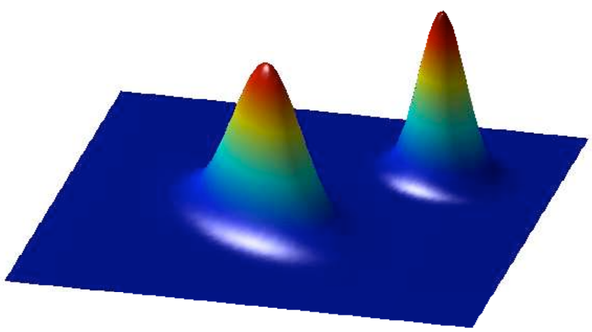

Figure 2: The example epistemic landscape used in the simulations described in this paper.

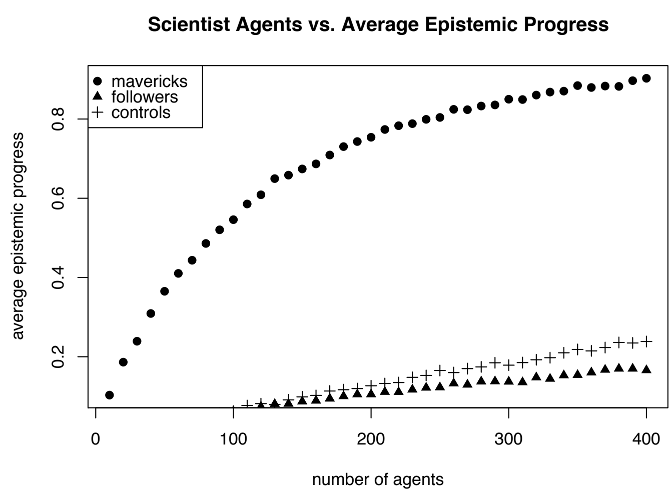

Figure 3: Comparison of the epistemic progress of controls, followers, and mavericks. Controls and mavericks measured after 200 cycles, followers after 1,000.

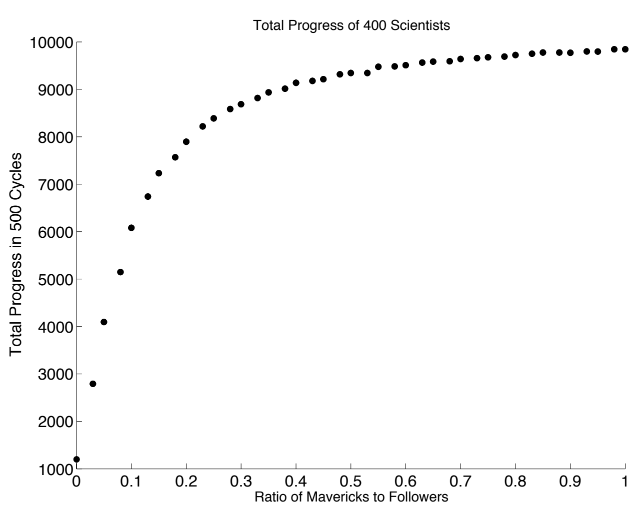

Figure 4:
Average number of approaches investigated by mixed populations of followers and mavericks after 500 model cycles. The landscape consists of 10,201 approaches total. The population size is held fixed at 400 scientists, but the ratio of mavericks to followers is varied from 0 mavericks to 400 mavericks.

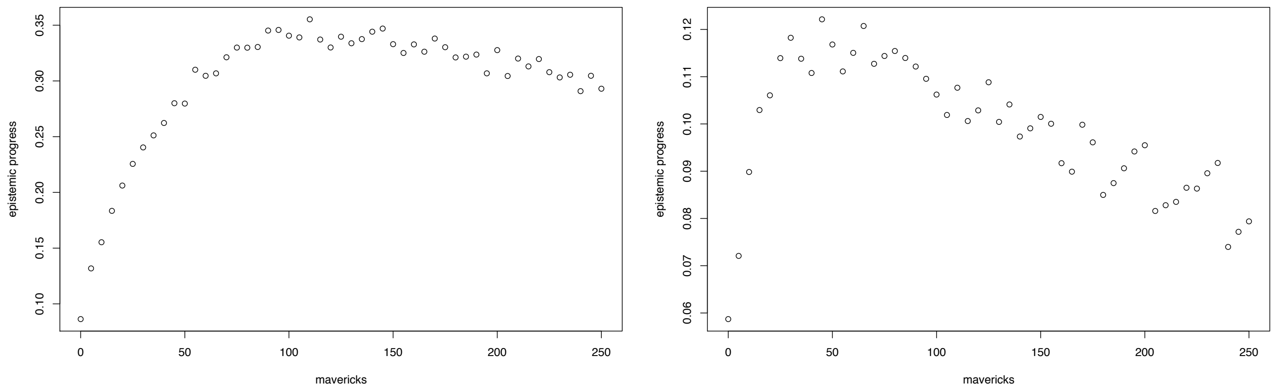

Figure 5: The effect of resource restrictions on epistemic progress. On the left, resources are restricted to a total of 75,000 agent-moves, where each maverick move costs ten times more than a follower move. On the right, only 25,000 agent-moves are allowed. In both cases, the populations have a total of 250 scientists, divided between followers and mavericks.