Epistemic landscapes, optimal search and the division of cognitive labor

Prof. J. McKenzie Alexander[a], J. Himmelreich[a], and Dr. C. Thompson[b]

[a]London School of Economics and Political Science
[b]University of Groningen

November 30, 2014

This paper examines two questions about scientists' search for knowledge. First, which search strategies generate discoveries effectively? Second, is it advantageous to diversify search strategies? We argue pace Weisberg and Muldoon (2009) that, on the first question, a search strategy that deliberately seeks novel research approaches need not be optimal. On the second question, we argue they have not shown epistemic reasons exist for the division of cognitive labor, identifying the errors that led to their conclusions. Furthermore, we generalize the epistemic landscape model, showing that one should be skeptical about the benefits of social learning in epistemically complex environments.

---

1. Introduction. A well-known example of the benefits conferred by the division of labor appears in Adam Smith's *The Wealth of Nations*. In his discussion of the pin factory, Smith noted that the efficiency gains derived from specialization could yield an increase of productivity between 240- and 4,800-fold from that of a single individual. Whereas a single worker might endeavor to produce twenty pins in a day, a group of ten in a single factory had been seen to produce upwards of twelve pounds.

Does the same hold true for the division of *cognitive* labor? Would there be more discoveries or would discoveries come faster if scientists divided their labor? For a number of reasons, the answer is obviously yes: undivided cognitive labor would lead to unnecessary repetition, scientists would fail to benefit from the unequal distribution of skill and talent and, finally, complex projects would become unmanageable given only a single worker. However, some have argued for a positive answer for other reasons. In particular, Weisberg and Muldoon (2009)[1] suggest that diversity of research strategies may "*stimulate* [...] greater levels of epistemic production" (225) and contend that even small steps towards a more diverse community of scientists "massively boosts the *productivity* of that population" (246–47, italics inserted).

The argument Weisberg and Muldoon provide for this claim utilizes a formal model of search strategies on an "epistemic landscape", a natural reinterpretation of the idea of a fitness landscape from evolutionary biology. In what follows, we show that, contrary to what they report, a careful examination of their formal model does not actually support many of the conclusions they attempt to draw regarding the division of cognitive labor. There are three main reasons. First, for the *particular* epistemic landscape they consider, the purported benefits of cognitive diversity are exaggerated due to a failure to consider a broad enough comparison class of search strategies. We provide several examples of homogeneous populations which prove surprisingly efficient at searching the space and identifying the points of epistemic interest. Second, the apparent benefits of cognitive diversity reported largely derive from implementation errors in two of the three search strategies they discuss. And third, if the model of epistemic landscapes is generalized to more rugged, higher dimensional landscapes whose overall topography is not discernible by the individuals,[2] social learning and the division of cognitive labor only helps in particular circumstances. The upshot is that, although there clearly are real benefits from the division of cognitive labor, the reasons have nothing to do with the epistemic reasons suggested by Weisberg and Muldoon's formal model.

The overall structure of the paper is as follows. In section 2 we briefly revisit the original epistemic landscape model. In section 3 we derive results providing an upper bound for efficient search strategies on that landscape; essentially, no rational scientist should perform worse than this value. We then show that some of the search strategies investigated by Weisberg and Muldoon fare far worse than this constraint. Section 4 shows why the search strategies considered by Weisberg and Muldoon performed so badly. Section 5 considers two key hypotheses which they claim to have substantiated, and we show that our re-examination of the model invalidates both hypotheses, effectively undermining their attempts to provide epistemic reasons for the division of cognitive labor. Section 6 demonstrates that homogeneous populations can do even better than heterogeneous populations, in some cases. And, finally, in section 7 we generalize the epistemic landscape

1. References given as page numbers only refer to this article.
2. We use the general framework of *NK*-fitness landscapes of Kauffman and Levin (1987).

model and argue that whether the division of cognitive labor is advantageous or not depends upon features of the landscape which may well be unknowable.

2. Epistemic Landscapes. The basic idea of an epistemic landscape derives from Sewall Wright's (1932) insight in population biology that one can represent the fitness values of genotypes in terms of an abstract landscape. There, a particular genotype corresponds to a point in a highly multidimensional landscape, with the fitness value of that genotype as its "height", and "nearby" points on the landscape being genotypes accessible via point mutation, recombination and so on.[3] Analogously, we can think of an "epistemic landscape," where a point of the landscape represents a particular research approach to investigating a topic of inquiry. A research approach consists of the composite set of research questions, instruments, techniques and methods used, as well as background theories a scientist or group of scientists rely on.

Not all research approaches to a topic of inquiry are equally fruitful, though. Some research approaches bring better results, or more publications, or more useful applications, than others. Following Weisberg and Muldoon, we can treat each approach as having a significance between 0 and 1. This generates an abstract landscape over the various research approaches, with the height of each approach corresponding to its significance. The entire landscape itself represents a *single* topic of inquiry.[4]

How do they proceed to model this? First, they fix the form of the epistemic landscape. For simplicity, they work with a discrete 101×101 lattice, wrapping at the edges to form a torus, with two peaks as illustrated in figure 1. Second, they must operationalize the concept of finding points of epistemic significance on this landscape. This can be broken down into two components:

Epistemic Success. The time required to visit the two peaks.

Epistemic Progress. The percentage of the significant regions explored after a given time.

Given these two aims, how should scientists search the space? Since a point in two dimensions represents a research approach, this becomes the question of how

3. Of course, this description ignores the fact that frequently one cannot meaningfully speak of the fitness of a genotype separate from the distribution of genotypes/phenotypes in a population. For the purposes of investigating their model, we share this idealization with Weisberg and Muldoon. However, we agree that this assumption is highly dubious and future work should investigate the consequences of dropping it.

4. This can be seen from the fact that the same set of research methods may have very different degrees of fruitfulness over different topics of inquiry: randomized controlled trials are highly fruitful for determining the efficacy of various drugs, less so for purposes of literary theory.

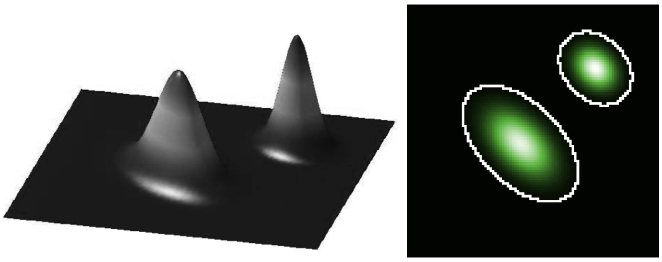

Figure 1: 3D and 2D projections of the epistemic landscape. On the right, significance hills are indicated by shaded gradient, with the perimeter highlighted.

a scientist should move about the landscape in light of the information available to her.[5] A variety of different kinds of information exist which one might use. There is *epistemic information*, such as the significance of one's current approach. There is also *social information*, such as how often a certain approach has been tried before, or whether it has been tried at all. Finally there is the possibility of using *metric information*, such as how far away is the nearest scientist.[6]

Weisberg and Muldoon consider three search strategies: a "control", and two others which they call "mavericks" and "followers". The control strategy, also referred to as the "HE rule" (short for "hill-climbing with experimentation") only uses epistemic information, whereas mavericks and followers use both epistemic information and social information. Essentially, the HE rule instructs a scientist to try to climb a hill towards the peak (without being able to detect the gradient), and otherwise to follow a straight line on the landscape with, occasionally, a random change in direction. Followers, as the name suggests, are intended to favor explored approaches which have been previously considered, and only take significance into consideration as a secondary consideration. The maverick strategy, on the other hand, is sensitive to research approaches which have been explored previously and deliberately seeks out unexplored approaches at random.

5. In turn, this requires specifying just exactly what information is, in fact, available to a scientist. This purportedly simple model has a lot of detail which needs to be posited before one can begin to make any headway with the two questions.

6. Strictly speaking, the epistemic landscape provides a *topological* model rather than a metric model, since the spatial "positions" are supposed to be abstract representations of variations in some research approach. (Consider, by way of comparison, the concept of "distance" between two genomes identical except at one base. If the differing base was, say, adenine instead of thymine, does that make the second genome closer or farther than it would have been if the base had been guanine or cytosine? It's hard to make sense of this question.) Nonetheless, one could impose a metric onto the epistemic space by simply imposing a Euclidean metric onto the landscape. Whether this would mean anything is, of course, unclear.

3. Controls, followers, mavericks, and the efficiency of search. Before getting into the details of the Weisberg and Muldoon model, let us first establish some clear upper bounds on the search efficiencies one might expect to find. Recall that the control strategy attempts to hill climb, if in an area of epistemic significance, and otherwise follows a straight line, with occasional random changes in direction. If we dispense with the requirements of hill climbing and the occasional random reorientation when in areas of zero significance, we get a search strategy which can be proven to almost always visit both peaks and exhaustively search the entire space. How so? Recall that, since the two-dimensional landscape wraps at the edges, it is topologically equivalent to a torus. The Kronecker foliation of the torus is obtained by projecting a straight line in the real plane with slope θ onto the surface of the torus. If the slope θ is rational, the projected line forms a closed loop; if the slope is irrational, the line will be dense in the surface of the torus.[7] Figure 2 illustrates both types of foliations. Since the epistemic landscape under consideration is divided into discrete cells, the fact that the Kronecker foliation is dense in the surface of the torus guarantees that a single agent who simply follows a straight line will, in finite time, search the entire space.[8] Call an agent employing such a strategy a "foliator".[9] A population of foliators will manage to explore the entire space more quickly than a single agent. And since almost all of the possible slopes θ which an agent may follow are irrational, a population of foliators use a *simpler* search strategy than Weisberg and Muldoon's control agents, but can be proven to almost always succeed in achieving both epistemic aims.

Let us estimate the efficiency of this search rule by simulation. Figure 3 illustrates one result from a simulation containing 10 foliators beginning at random locations in the area of zero significance, which we call the "desert". Notice that the foliator strategy can be quite quick: within 500 steps, one peak of the landscape had already been found and 36% of the entire landscape explored. Out of 5,000 simulations run with ten foliators and random initial conditions, 4,988

7. Strictly speaking, the slope of the line needs to be classified as rational or irrational relative to both the major and minor radii of the torus. Let the major radius have a length of $\frac{1}{2}\sqrt{3}$ and the minor radius a length of $\frac{1}{2}$. The surface of the torus is equivalent to a rectangular lattice of height 1 and width $\sqrt{3}$ which wraps at the edge. A line in the plane with slope of $\frac{1}{\frac{1}{3}\sqrt{3}}$ (and hence irrational) will, when mapped onto the torus, self-intersect after three loops. However, if we take 'irrational' to mean 'irrational with respect to the major and minor radii of the torus', then the claim holds.

8. Although the lattice is divided into discrete cells, the heading of an agent may vary continuously.

9. Note that for the purpose of establishing a theoretical upper bound, strictly speaking, we may not require the existence of a real-world analog to this strategy. This strategy requires to impose a Euclidian metric onto the landscape, which is a plausible assumption when real-world constraints are set aside (see our remarks in footnote 6). We discuss the real-world applicability of the epistemic landscape model at length in section 7.

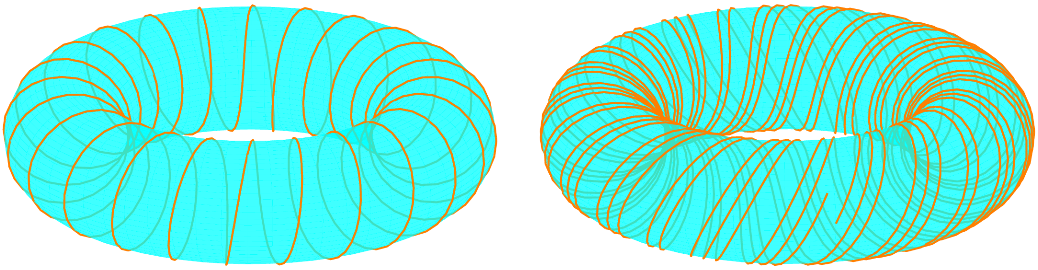

(a) Partial foliation of the torus by a line with a rational slope. Notice that the line self-intersects.

(b) Partial foliation of the torus by the initial segment of a line with an irrational slope. Notice that the line fails to self-intersect and will densely cover the torus over repeated windings.

Figure 2: Foliations of the torus.

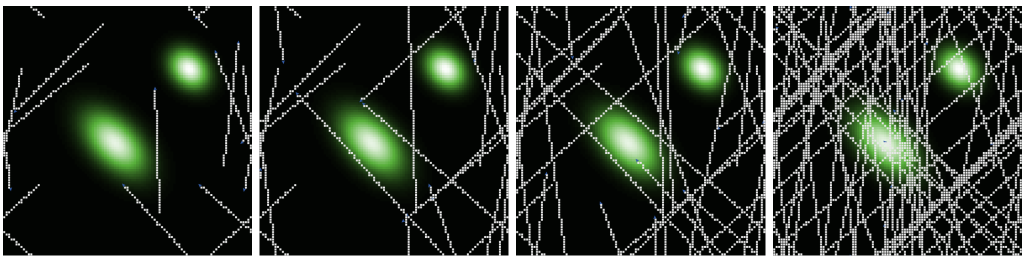

(a) 50 iterations (b) 100 iterations (c) 200 iterations (d) 500 iterations

Figure 3: Epistemic search by 10 foliators. The squares are 'breadcrumbs' showing approaches that have been visited.

managed to find both peaks with 50,000 steps. The mean time required to find both peaks was 1,855 steps, with minimum and maximum times of 43 and 32,167 steps, respectively, and a median time of 1,430 steps.

Now compare these results with the results reported by Weisberg and Muldoon (see figure 6). The discrepancy reveals something rather curious. In a simulation with 100 repetitions of populations of 10 HE rule agents, only 95 populations found both peaks within the time allowed (50,000 steps). Of these 95 populations:

> "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." (236)

In short, foliators — who never attempt to hill climb, and who never pick a new direction of travel — are both more *effective* at finding both peaks (with a success rate of 99.7% instead of 95%) and *faster* (a mean time of 1,855 steps, as opposed to 6,075). Let us bracket this observation, for the moment, and return to it at the end of this section.

Now consider, as an alternative search strategy, the case of a simple random walk. Assume that at points of zero significance (the "desert") the agent randomly moves to one of its eight nearest neighbors, with all transition probabilities equal; at points with positive significance, the agent follows the gradient. This means that at points of zero significance, an agent's movement at a time is independent from her movement at all previous times and that once the agent enters a region with positive epistemic significance, she never leaves. Thus we can treat the two hills as a single absorbing states and model the movement of the agent in the desert as a Markov process.

Let $k_{n,m}$ denote the expected time to absorption for an agent starting at location (n,m). By construction, $k_{i,j} = 0$ for all points (i,j) having positive epistemic significance. Furthermore, for all points (n,m) having zero epistemic significance, we know that

$$
\begin{aligned}
k_{n,m} = 1 + \frac{1}{8}\Big(k_{n-1,m+1} + k_{n,m+1} + k_{n+1,m+1} \\
+ k_{n-1,m} + k_{n+1,m} \\
+ k_{n-1,m-1} + k_{n,m-1} + k_{n+1,m-1}\Big).
\end{aligned}
$$

Figure 4(a) illustrates the local transition diagram for a point in the desert bordering the perimeter of a hill. This gives a system of 10,201 simultaneous linear equations.

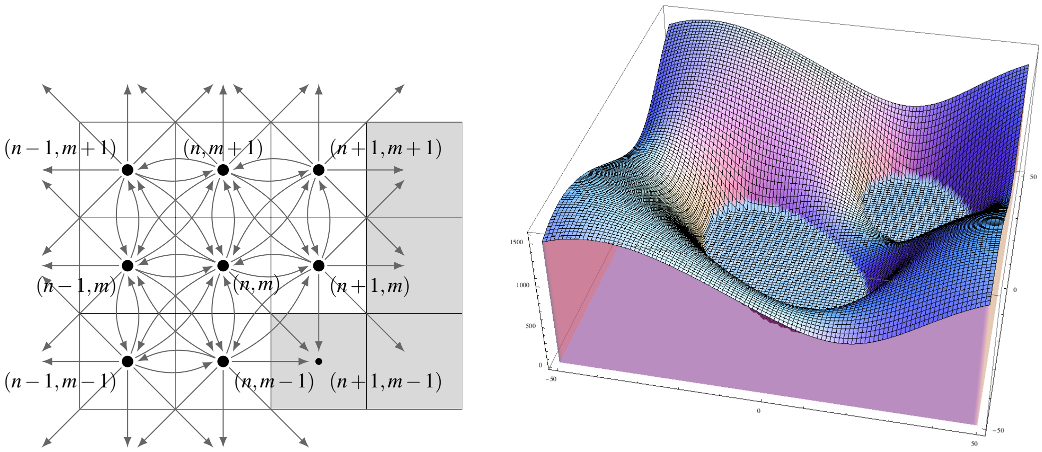

(a) A portion of the Markov chain transition diagram for the random walk search strategy. Gray areas represent absorbing states. For simplicity, only the state transitions exiting the 3×4 block of points have been shown.

(b) A plot of the expected number of steps required for an agent to encounter a region of epistemic significance via a random walk (the height of each column represents the expected number of steps).

Figure 4: Analytically solving for the expected waiting time of the random walk search strategy.

If we only consider equations generated from points having zero significance, and the perimeter of the two epistemically significant regions, the system reduces to 8,273 simultaneous linear equations. Figure 4(b) illustrates the expected hitting time of the absorbing state, for each point in the landscape.[10] The maximum expected hitting time is a fraction over 1,600 steps. The average hitting time, for a single agent starting in the desert, is 881.9 steps.

Since it is not feasible to obtain an analytic solution for a population of ten agents engaged in independent random walks, we examined this via simulation. Out of 5,000 populations of 10 agents, all starting from random initial conditions, 4,944 found both peaks. In all remaining 56 cases, all 10 agents had found the same peak. (With the random walk search strategy, when an agent finds a peak it stays there.) The mean time required to find both peaks was 249 steps, with a minimum and maximum time of 14 and 4,241 steps respectively, and a median of 141 steps. After 500 iterations on average 15.6% of the total landscape had been explored.

In summary, we have shown the following: if we are interested in both effectiveness (were both peaks found?) and efficiency (how long did it take to find

10. The system of equations was programmatically generated and then solved using *Mathematica*. Since the system of equations is both linear and described with a sparse coefficient matrix, a solution is found quite quickly.

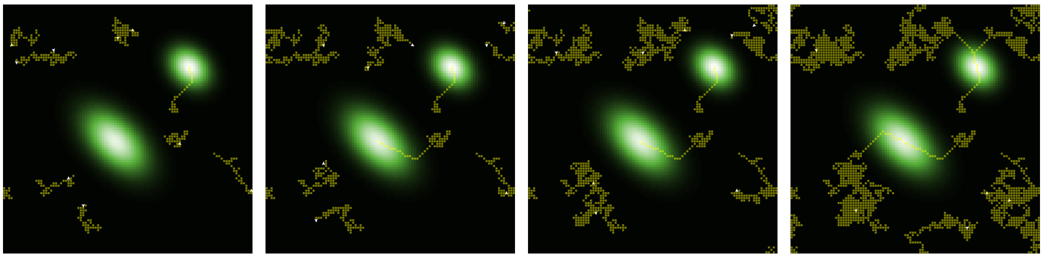

(a) 50 iterations (b) 100 iterations (c) 200 iterations (d) 500 iterations

Figure 5: Epistemic search by 10 agents using the random walk strategy. One peak has been found after less than 50 steps. Within 100 steps both peaks have been found. Within 500 steps, 5 agents have found peaks and roughly 17% of the landscape have been explored.

	Epistemic success					
Strategy	**P**	**Mean**	**SD**	**Min**	**Max**	**Median**
HE rule	0.95	6,075	8,518	553	43,004	2,553
Random Walk	0.9888	249	327	14	4,241	141
Foliator	0.9976	1,855	1,755	43	32,167	1,430

Figure 6: A comparison of simulation results reported in Weisberg and Muldoon on 10 agents using the HE rule with results of agents using the random walk and Kronecker foliation search strategies. *P* denotes the proportion of populations that found both peaks.

both peaks?), the HE rule did *worse* than pure populations of foliators or random walkers. (Figure 6 has a side-by-side comparison of all three strategies.) In particular the HE rule — despite taking epistemic information into account by hill-climbing — did worse on average by a factor of three when compared to the foliator strategy, which simply ignored this information. Furthermore, the HE rule did worse than the random walk strategy by a factor of roughly 24. Admittedly, our random walk strategy would be expected to do better, in regions of positive significance, because it simply followed the gradient, whereas the HE rule used a probe-and-adjust method to hill climb. But that surely doesn't explain *everything* about why it did over 24 times worse.

4. A closer look. Why did the control scientists perform so badly? Let us examine the exact statement of the HE rule as described in the original paper. It is as follows:

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.

Scope ambiguities regarding the nested conditionals in step 2 make the rule, as stated, open to multiple interpretations. The pseudocode representation provided in figure 7(a) makes precise the scope relations between the else-clauses of the conditionals under one interpretation. Here, step 1 of the published version of the HE rule corresponds to the `forward()` command in line 2. Step 2 corresponds to the nested conditionals and commands in lines 3–20. There is nothing corresponding to the "Begin again at Step 1" instruction because we assume each agent calls `HE_rule()` at the start of each iteration.

From this, it is clear that our foliator rule approximates the HE rule quite well: when exploring areas of zero significance, the test on line 3 fails and the test on line 7 succeeds. As the test on line 8 succeeds only 2% of the time, the remaining 98% of the time the HE rule will not reset its heading. Thus, when the `forward()` command on line 2 is invoked at the start of the next iteration, the agent continues to move forward according to its previous heading, i.e., in a straight line. Furthermore, this shows that the `stay_put()` command on line 13 is not, strictly speaking, necessary, since the agent has already taken one step forward during the current iteration. (Indeed, this serves to highlight a small error in the Weisberg-Muldoon statement of the HE rule: when in areas of increasing significance, the combination of the the `forward()` command on line 2 and the successful test on line 3 ensures that the agent will step forward *twice* in the same iteration.)

Inspection of the original code used in the Weisberg-Muldoon simulation, which they made available, revealed that their actual implementation was as in figure 7(b). There are several things to note. First, the absence of a `forward()` command at the beginning means that agents are not guaranteed to move at least once each iteration. Second, the test condition at line 2 contains the >= operator. When exploring regions of zero significance, this means the test at line 2 will *succeed*, dropping us immediately into the test at line 3, which will also succeed. A control agent, given lines 4–6, moves forward at a random heading with a 2% probability and otherwise *remains stationary* 98% of the time.

In other words, whereas the description of the HE rule in the paper suggests that agents ought to behave rather like foliators, their actual implementation has those agents behaving like lethargic random walkers. This explains the difference between our baseline results and those reported by Weisberg and Muldoon. However, it also calls into question the *meaningfulness* of the comparison between control agents who use the HE rule and other search strategies. Since control agents *do nothing* 98% of the time, unless a similar delay is incorporated into the definition of any compared search strategies, we are comparing rules which operate on fundamentally different time-scales.

```
1  HE_rule() {
2    forward();
3    if (curr_sig > prev_sig) {
4      forward();
5    }
6    else {
7      if (curr_sig == prev_sig) {
8        if (random() < 0.02) {
9          random_heading();
10         forward();
11       }
12       else {
13         stay_put();
14       }
15     }
16     else {
17       backward();
18       random_heading();
19     }
20   }
21 }
```

(a) As in the paper.

```
1  HE_rule() {
2    if (curr_sig >= prev_sig) {
3      if (curr_sig == prev_sig) {
4        if (random() < 0.02) {
5          random_heading();
6          forward();
7        }
8        else {
9          stay_put();
10       }
11     }
12     else {
13       forward();
14     }
15   }
16   else {
17     backward();
18     random_heading();
19   }
20 }
```

(b) As implemented.

Figure 7: Two implementations of the HE rule search strategy. The variables `curr_sig` and `prev_sig` refer to the significance of the site currently occupied by the agent and the significance of the site previously occupied by the agent, respectively.

Let us now turn to the follower strategy. Here is the definition, as in the original paper:

Follower rule:

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.

The following pseudocode disambiguates the scope of the nested conditionals:

```
1  Follower_rule() {
2    if (any Moore neighbors visited?) {
3      let n = random visited neighbor with max significance;
4      if (significance(n) > current_significance) {
5        go_to(n);
6      }
7      else {
8        if (any unvisited neighbors?) {
9          let m = random unvisited neighbor;
10         go_to(m);
11       }
12       else {
13         stay_put();
14       }
15     }
16   }
17   else {
18     let m = random unvisited neighbor;
19     go_to(m);
20   }
21 }
```

In this version, there is no need to specify the tie-breaking rule explicitly because we have already selected, at line 3, a random visited neighbor with maximum significance.

As interpreted, the follower rule performs a biased random walk which can get stuck. In the presence of sites which have been previously visited, and which are of greater significance, the rule moves to one of those sites at random. When no visited sites are of greater significance, it moves to a random unvisited site. However, when entirely surrounded by visited sites of *equal* significance — as can happen in the desert — the follower rule will always end up selecting `stay_put()` at line 13.

How does this interpretation compare with the results Weisberg and Muldoon report? They write:

> "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 [...] 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 [...] of 56 cycles, which suggests that the randomly placed agents were near the boundary of significance at the beginning of the simulation." (240)
>
> Weisberg and Muldoon (2009, 240)

Even though the biased random walk performed by the followers can get stuck in the desert, it seems strange that a population of 200 followers only found one peak 60% of the time, and both peaks only 12% of the time. Further cause for concern should arise when one reads that when followers did find both peaks it happened "very rapidly." How can a search strategy perform so badly at search, generally, yet succeed so rapidly when it does?

Inspection of the code used by Weisberg and Muldoon showed that their implementation was functionally equivalent to our interpretation above *except for the test at line 4*, which in their version was:

```
if (significance(n) >= current_significance) {
```

The use of the >= operator instead of the > operator is a serious error. Given a sparse distribution of followers, where each agent is at least three squares distant from every other in regions of zero significance, the agents get stuck in a loop.

Figure 8 illustrates how this happens in detail. In the initial state shown up left, the `if`-test at line 2 of the follower rule fails, but the `if`-test at line 8 succeeds, resulting in the agent moving to an adjacent site selected at random. But then, in a configuration like that shown up right, the `if`-test at line 2 now succeeds, and because the Weisberg-Muldoon implementation of the follower rule has the >= operator at line 4, the agent simply moves back to its previous position.[11] From this

11. This happens because there is exactly one previously visited neighbor, and so it is guaranteed to be chosen when we select "a previously visited neighbor with maximum significance".

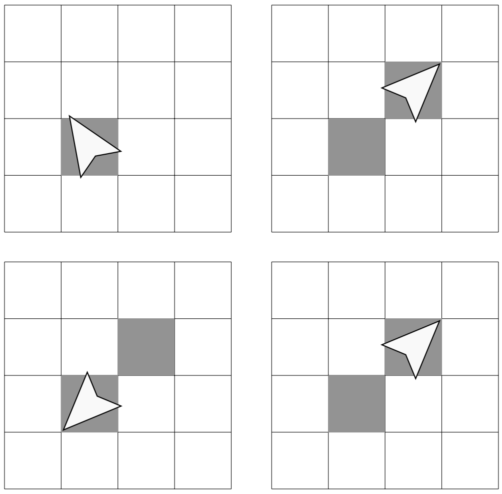

Figure 8: The implementation of the follower rule typically results in cycles of length two. (All shown sites have zero significance.) The agent moves from the initial configuration (up left) to a randomly selected unvisited square (up right). It then returns to a previously visited site whose significance is greater than or equal to the significance of the current site (bottom left) and is stuck in a loop (bottom right).

point on, the follower oscillates between the two visited sites. Instead of following others, the agent 'chases his own tail'.

Thus we see why only 3% of the 100 simulations of 10 followers managed to find a single peak: most of the time they were trapped in regions of zero significance, with at least one visited site in their neighborhood. It also explains why larger populations of followers, when they managed to find both peaks, found both peaks so quickly: the random initialization positioned a few followers in a site of zero significance which was adjacent to a region of positive significance. If the follower happened to randomly move into an area of positive significance, it will then proceed to climb up the hill via a random walk. With this in mind, consider the one graphic from the original Weisberg and Muldoon paper which showed the paths traveled by a population of 300 followers, reproduced in figure 9. Here we see that only agents positioned near the region of epistemic significance eventually managed to climb towards the top of the peak. Furthermore, the vast majority of paths traveled by followers consist of mere oscillation between two adjacent squares, exactly as one would expect given the code analysis.

5. On the division of cognitive labor. Weisberg and Muldoon note that "modern science requires the division of cognitive labor" (225) and claim that their simula-

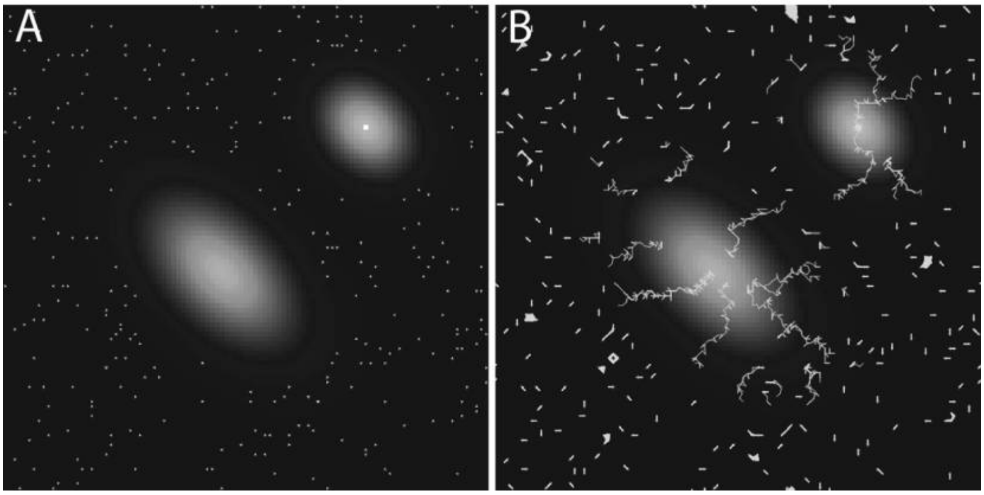

Figure 9: A figure from the Weisberg and Muldoon paper showing the exploration of the epistemic landscape by 300 followers before (A) and after (B) movement began. The tails show a plot of the path they followed. Note that the vast majority of followers are trapped in cycles of length 2.

tion results illustrate the epistemic benefit so conferred. But what, exactly, do we mean by the "division of cognitive labor"? Consider the following disambiguation. On one hand we have the phenomenon that scientists choose different approaches to investigate a research topic. In the epistemic landscape (which represents the research topic), the agents can occupy different points (which represent research approaches). Scientists specialize in different approaches. This is one meaning of division of cognitive labor. It describes a phenomenon of coordination. On the other hand, we have the phenomenon that scientists use different strategies to choose their research approach. The agents move across the epistemic landscape according to different strategies. Some scientists employ methods simply because they are trending (these would be the followers) whereas others favour approaches because they are exotic or unusual (these would be the mavericks). This is a second meaning of division of cognitive labor. It describes a phenomenon of diversity. Let's call the phenomenon where different people work on different projects *epistemic coordination*, and the latter, where different people have different reasons to work on different projects, *cognitive diversity*.

When Kitcher (1993) and Strevens (2003) use the Marginal Contribution and Reward model to explain why scientists pursue different research approaches, they seek to explain the phenomenon of epistemic coordination. They show that research behavior of scientists is coordinated when scientists are sensitive to social and epistemic information. Very roughly, scientists seek to maximize the rewards that accrue from scientific discoveries. They consider both the likelihood that an avenue will generate results (epistemic information) and the number of other scientists working on that avenue (social information). This yields the phenomenon that the scientific community spreads out across different possible avenues for research. However, because Kitcher and Strevens assume that all scientists choose among possible avenues for research in essentially the same way, they do not address cognitive diversity.

In contrast, an example of research on cognitive diversity can be found under the headings of "swarm intelligence" or "wisdom of the crowds".[12] For example, similar to Weisberg and Muldoon, Hong, Page, and Baumol (2004) develop a model of agents searching for local maxima in a space. The agents are cognitively diverse in two ways. Not only do they use different strategies to explore the epistemic space, they also represent the epistemic space differently; each individual has its own language to describe the points in the space. This is an example of a model of agents that are cognitively diverse.[13] More generally, models of cognitive diversity can be found in areas ranging from complex systems research and theoretical biology (Bonabeu, Dorigo, and Theraulaz 1999; Krause et al. 2011), management and organization studies (Thomas and Ely 1996; Polzer, Milton, and Swarm 2002; Jackson, Joshi, and Erhardt 2003) to psychology (Kerr and Tindale 2004), computer science (Clearwater, Huberman, and Hogg 1991) and economics (Hong, Page, and Baumol 2004; Arrow et al. 2008).

The epistemic landscape of Weisberg and Muldoon models division of cognitive labor in both senses. Agents take different approaches on a research topic (epistemic coordination) and they use different strategies in choosing those approaches (cognitive diversity). The central observation of "general trends about the division of cognitive labor" (249) made by Weisberg and Muldoon is that *cognitive diversity* gives a scientific community an epistemic advantage. An increase in epistemic performance ensues if a community of researchers uses different epistemic search strategies. They argue that "to be maximally effective, scientists need to really divide their cognitive labor" (227) and that a "healthy number of followers with a small number of mavericks" would provide an "optimal way" (251) to do so. They write that this is because there is a "very significant indirect affect that mavericks have on the research progress via their ability to stimulate the followers." (249)

12. However, a very different research project also uses the same label. It surrounds the phenomenon that large numbers of people are better at solving epistemic tasks. "Wisdom of the crowds" in this sense is related to the Condorcet Jury Theorem and not so much to cognitive diversity.

13. Hong, Page, and Baumol (2004) use the term "functional diversity".

That sounds like a fine example of epistemic benefits of cognitive diversity.[14] Is it true?

Weisberg and Muldoon observe that the epistemic performance of a population of followers increases when mavericks are added to it (247-8).[15] Does this vindicate the thesis that cognitive diversity improves epistemic performance? It does not if the improvement rests on a defective implementation of the search strategies. And it does not vindicate the thesis if the improvement is only due to the epistemic performance of the agents which have been added. It seems that both is the case here.

Furthermore, Weisberg and Muldoon describe that the improvement in epistemic performance is due to the following "indirect affect". They write "Mavericks help many of the followers to get unstuck, and to explore more fruitful areas of the epistemic landscape." (247) Maverick scientists *may* help follower scientists getting unstuck, but the followers should not have been stuck in the first place. As shown above, the search strategy of follower scientists suffered from a defective implementation such that they ended up chasing their own tail. The beneficial "indirect affect" that Weisberg and Muldoon describe requires a population of follower scientists that has hardly left the place where they started. If a search strategy performs so direly, the result that a complementary strategy improves overall performance is hardly surprising. It does not vindicate the thesis that there is an epistemic reason for cognitive diversity in any interesting sense.

If the follower strategy worked properly, in that it did not get stuck almost immediately in a cycle, would there still be an indirect affect to vindicate the thesis that there is an epistemic reason for cognitive diversity? We show that this is not the case. It turns out that the improvement in epistemic performance is exclusively due to the performance of the mavericks that are added to the population of followers. It should not be surprising that the epistemic performance of a population increases when agents are added. In particular, when these agents are mavericks, who have been shown to perform quite well.

We set up a simulation to record the epistemic progress of followers and mavericks separately as the mavericks are added to a population of followers. In these simulations, we used a correct implementation of the follower search strategy. This separate bookkeeping enables us to identify which sub-population caused the increase in the total epistemic progress of the mixed population. We ran simulations with 100 repetitions for each condition, for populations of followers

14. Assuming that each stimulated follower would do better than a maverick would in her place. Otherwise, why should we ideally not have only maverick scientists to begin with?

15. More precisely, Weisberg and Muldoon consider two settings. In one setting the size of the mixed population remains fixed and merely the proportion of followers to mavericks changes. In the second setting mavericks are added to a population of followers. We focus on the second because it is more instructive. Our findings apply to both.

with an initial size of 100, 200, 300 and 400. We observed how the epistemic progress of this population changes as maverick scientists are added. Replicating the experiment from Weisberg and Muldoon we added 10, 20, 30, 40 and 50 maverick scientists. We measured the epistemic progress after 1,000 iterations recording the total epistemic progress of the mixed population, and the progress of each sub-population.

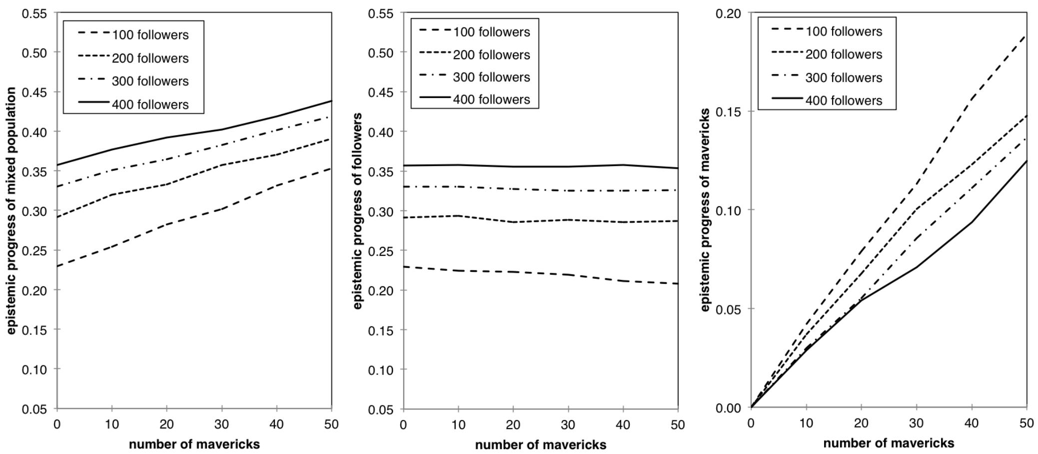

(a) mixed population (b) follower sub-population (c) maverick sub-population

Figure 10: Epistemic progress of a population as mavericks are added. The increase of the mixed population (a) is almost entirely due to the maverick sub-population (c). The epistemic progress of the follower sub-population remains unchanged while mavericks are added (b).

For all initial sizes of follower populations, the epistemic progress of the followers remains virtually unaffected by the presence of mavericks. In the case of the initial population size of 100 followers, the epistemic progress of the follower sub-population even decreased. While a pure population of 100 followers managed to explore 23% of the significant points after 1,000 steps, when mavericks are added to the population this number goes slightly downwards. The epistemic progress of the follower sub-population in the presence of 50 mavericks reaches only 21%.

Thus the increase in the epistemic progress of the mixed population can be solely accounted for by the epistemic progress of the maverick sub-population. Indeed, if anything, the followers seems to get in the way of the mavericks. The mavericks are doing particularly well when only few followers are present. Notice that the epistemic progress of a maverick sub-population of the size of 50 is 19% in the presence of a follower sub-population of a size of 100. The epistemic progress of the maverick sub-population decreases as the size of the follower sub-population increases. The maverick sub-population of the same size achieves an epistemic progress of 15% when the follower sub-population has a size of 200, 14% when it has a size of 300, and 12% when it has a size of 400.

6. Efficient search in homogeneous populations. The Weisberg and Muldoon epistemic landscape model features three different kinds of information: epistemic, social, and metric. Although control agents use epistemic, and followers and mavericks use both epistemic and social information, none of the three strategies considered utilize metric information. To an extent, using metric information (how far away other nearby scientists are and where they are going) can be interpreted as a "strategic follower" strategy: it pays attention to where other scientists are going rather than seeking parts of the landscape where they have been. Since there is good reason to suspect that some scientists behave in roughly similar ways (in that they consciously align their "research brand" with what is trending), let us consider how such a search strategy performs. Let us call this strategy the "swarm". Stated informally, a swarm scientist receives information about what others in her area are working on via journals and conferences and adjusts her own approach such that it is always similar but yet distinct to the approaches pursued by others in her area. Furthermore, when she observes that many of her colleagues incorporate a certain turn into their approaches, she will try to imitate this change.

There are at least four interesting parallels between the epistemic search of a scientific community and the foraging behavior of animal groups. First, individuals in a school of fish or a flock of birds need to coordinate their behavior to avoid occupying the same space at the same time. In scientific research, this is the problem of epistemic coordination: we don't want everyone to be attempting to do the same thing at the same time, as such redundancy would often be a waste of effort. Second, the animal group often has a common goal, such as finding food or traveling to a nesting site. Analogously, the scientific community has epistemic aims, such as determining a high-yield, cost effective way of manufacturing graphene. Third, information is distributed differently among individuals in a group: only a small subset of individual animals in a group may have information about the location of particular food site. Analogously, only a small subset of researchers has experience with a particular approach to the research topic. Finally, just like how some herd-based animals follow others who take the lead, a similar behavior may be found in the scientific community.[16]

Theoretical biology has a rich literature on collective behavior (see Sumpter 2010, ch.5, for an overview). The swarm search strategy we use is a simplified

16. An instructive example on the last two points are honeybees (see Seeley 2010). Once a decision about a nest site has been made, "[u]p to around 10,000 bees of which only 2 or 3% are informed of the location of the nest site fly as a single swarm to the site" (Sumpter 2010, 123). See also (Ward, Krause, and Sumpter 2012).

version of one by Couzin et al. (2005) and similar to the Boids model (Reynolds 1987). Roughly, if another agent gets too close, then the agent swerves to avoid collision. Otherwise, the agent aligns its direction of travel with the other agents around it. More precisely, the space surrounding the agent is divided into two different "zones", as shown in figure 11, a zone of repulsion and a zone of orientation. If there is another agent in the front half of the agent's zone of repulsion, then the agent changes its direction to the right if the closest individual is ahead and to the left, and to the left if it is ahead and to the right. If there are individuals in an agent's zone of orientation but not its zone of repulsion, then the agent adjusts its direction by the mean of the differences between its own direction and the directions of the individuals in the zone of orientation.

Although figure 11 assumes that the radius of the zone of repulsion is smaller than that of the zone of orientation, this need not be the case. If the radius of the zone of repulsion is greater or equal to the radius of the zone of orientation, the resulting swarm has considerably different properties. To distinguish between these two cases, let us introduce some terminology: call the latter case a "repulsing swarm", and the former case a "flocking swarm". The flocking swarm can be thought of as being composed of "strategic followers," with the repulsing swarm being composed of "strategic mavericks".

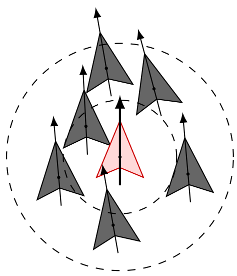

Figure 11: The focal agent of a flocking swarm with its zone of orientation (outer circle) and repulsion (inner circle).

Note that swarm scientists also use epistemic information, but in a manner somewhat different from the HE rule. Since some scientists have the ability to intuit the correct way to develop a theory — think of Newton, Einstein, von Neumann, and Feynmann — we incorporate this into the model by assigning to each agent a probability of being "clairvoyant." That is, each agent who happens to be in a region of positive significance has a probability of guessing the direction to the top of the hill.[17] When this happens, the clairvoyant agent adjusts its heading to point in the direction of its insight. Clairvoyance does not last, though, and so the initial

17. The probability is the same for all agents.

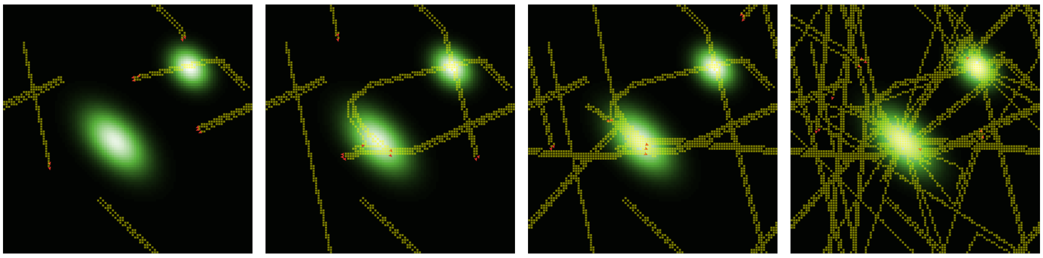

(a) 50 iterations (b) 100 iterations (c) 200 iterations (d) 500 iterations

Figure 12: Epistemic search by 10 swarm agents in the flocking configuration. Notice that one peak has been found within 50 steps. Within 100 steps both peaks have been found and around 14% of significant points and 8% of the total landscape have been explored. Within 500 steps, 47% of significant points and 32% of the total landscape have been explored.

flash of insight might disappear as the agent further adjusts its behavior to the rest of the surrounding swarm. Brilliant ideas may go unrecognized.

Compare this with the complexity of the behavior of the other agents. The HE rule, maverick and follower agents each change their behavior when they enter an area of positive significance. Each of them uses epistemic information in each move to climb towards the top of the hill. The swarm strategy exhibits no richer behavior or greater complexity than these strategies.

In short, in each iteration a swarm scientist does the following. With probability p, and only if the agent is on a point of positive significance, it has a one-off moment of clairvoyance, which causes it to changes its direction towards the closest peak, taking one step forward. Otherwise, it adjusts its direction to align with all the other agents in its zone of orientation, or it swerves away from the closest nearby agent in its zone of repulsion, if there is any. After aligning or swerving, the agent moves one step forward.

We ran simulations starting with populations of 10 agents, increasing it by 10 to up to 400 agents with 100 repetitions each.[18] We investigated the flocking swarm strategy and the repulsing swarm strategy.[19] The probability of clairvoyance was 0.03 for the flocking swarm and 0.015 for the repulsing swarm. We compare our

18. Let the number of agents be n, the radius of the zone of repulsion r, and the number of groups s. The agents were placed in $s = 3 + n/10$ groups that were randomly located in the desert. The radius for each group, in which each agent was randomly placed is $\sqrt{rs}$. To start the simulation without a pre-run to form stable swarms, all members of a group were given the same random heading.

19. The radii were set to 3 and 1 respectively. Note that when the zone of repulsion is greater than the zone of orientation, the agents never align their direction with that of the agents in their vicinity.

		Epistemic success						*Epist. progress*	
n	**Strategy**	**P**	**Mean**	**SD**	**Min**	**Max**	**Median**	**200**	**500**
10	Maverick	0.75	121	58	52	299	109	7.74	8.79
	Follower	0.79	123	63	50	509	110	6.76	7.24
	Control	1	107	71	37	668	93	19.32	27.41
	Swarm-F	1	167	110	34	602	133	17.94	39.51
	Swarm-R	1	151	88	17	488	131	17.80	39.42
100	Maverick	1	54	8	37	75	54	36.95	36.99
	Follower	1	52	7	38	73	52	22.92	22.92
	Control	1	42	10	23	71	41	67.35	74.25
	Swarm-F	1	64	29	21	154	57	80.85	97.54
	Swarm-R	1	41	18	12	111	37	86.49	99.32
200	Maverick	1	49	6	37	66	48	51.23	51.23
	Follower	1	43	5	31	59	43	29.10	29.10
	Control	1	33	8	20	54	32	84.24	88.70
	Swarm-F	1	45	20	20	135	40	93.87	99.78
	Swarm-R	1	32	11	13	66	31	98.09	99.99

Figure 13: Results comparing epistemic success and epistemic progress of different strategies. 'Swarm-F' and 'Swarm-R' are respectively the flocking and the repulsing configuration of swarm scientists. The Follower and Control strategies used here are as intended: Followers do not necessarily get stuck in a cycle of length two when isolated (although they might get stuck eventually) and Controls do not have an artificial time delay (nor do they move twice per iteration in regions of increasing significance). Regarding the column headings, *n* is the number of agents, *P* the proportion of populations that found both peaks with the time allotted, and *200* and *500* is percentage of significant landscape that has been explored after 200 and 500 iterations respectively. (Simulations were stopped after any number of iterations equal to 0 modulo 500 if no epistemic progress had been made in the last 500 iterations.)

results with similar simulations using maverick and the follower scientists.[20] The results for 10, 200 and 400 agents can be found in figure 13.[21]

On epistemic success we found that mavericks and followers perform better than the swarm strategy only if the populations are sufficiently small and the population of maverick or follower agents actually manages to find both peaks. We found that as the size of the population increases beyond 30, this result no longer holds. Then the swarm strategy performs better than the maverick and the follower strategy. The repulsing swarm configuration performs particularly well. Consider the median time to find both peaks for populations of size 100. Mavericks, followers, and the flocking swarm configuration have a median time between 50 and 60 iterations to find both peaks, whereas half of all the repulsing swarm populations find both peaks already after 37 iterations.

On epistemic progress, as previously shown, we found that the maverick scientists have a greater epistemic performance than follower agents. However, swarm scientists have an even greater epistemic performance than the maverick scientists; the repulsive swarm configuration seems to do slightly better than the flocking swarm configuration. The take-home message is the following: cognitively homogeneous populations of agents can do very well.

7. A generalized epistemic landscape model. One further concern with the Weisberg and Muldoon model derives from the simplicity of the epistemic landscape considered. Although we are generally ignorant about the shape of the epistemic landscapes underlying real scientific research, it is clear that they have at least two properties which are largely absent from the Weisberg and Muldoon landscape. First, on real epistemic landscapes, it is much easier to get trapped at a local optimum and much harder to identify the global optimum. And second, when we consider the "epistemic fitness" conferred by a combination of scientific methods, theories, techniques, and so on, there is a much greater degree of interdependency than a two-dimensional landscape would allow. If we consider more "realistic" epistemic landscapes, what, if anything, can we infer about the benefits of cognitive diversity?

20. We used a repaired implementation of the follower strategy, not the defective implementation that was originally used, as discussed above.

21. Comparisons between the two swarm strategies and the correct implementation of the Weisberg and Muldoon strategies regarding epistemic progress are not likely meaningful because controls, followers, and mavericks remain at epistemic peaks, once found. That said, it is worth noting that the control strategy — when implemented as intended — performs very well in terms of epistemic progress compared with pure populations of mavericks or followers (thereby again undermining the results of the original paper). Comparisons between all strategies regarding epistemic success, though, are meaningful.

$$\cdots\ b_{i-3}\ \underbrace{b_{i-2}\ \overbrace{b_{i-1}\ b_i\ b_{i+1}}^{K=2}\ b_{i+2}}_{K=4}\ b_{i+3}\ \cdots$$

Figure 14: Two different regions of epistasis for the bit b_i.

We can begin to consider this question by reinterpreting the NK-landscape model of Kauffman and Levin (1987) and Kauffman and Weinberger (1989) as an epistemic landscape. The idea is straightforward: suppose we have a set of N scientific propositions, where these propositions may consist of both abstract, general statements of high theory as well as specific statements of particular laboratory technique. The belief state of an individual scientist can be represented by a vector $\vec{b} = \langle b_1, \ldots, b_N \rangle$, where $b_i = 0$ if the scientist does not believe the i^{th} proposition, and $b_i = 1$ if the scientist does believe the i^{th} proposition.

The reason why NK-landscapes are useful for thinking about epistemic landscapes is that they allow one to model *interdependencies* between the various propositions believed (or not believed) by a scientist. That is, the fitness contribution of b_i may depend on not just the value of b_i (0 or 1), but on the value of several other entries in the scientist's overall belief state. The fitness function, in a word, may have varying degrees of *epistasis*. Let $0 \leq K \leq N$ denote the number of interdependencies contributing to the fitness contribution of b_i. (See figure 14 for an illustration of two different epistatic regions.) One can think of the amount of epistasis in a fitness function for an epistemic landscape as a formal model of the Quinean web of belief.

Figure 15 illustrates how the fitness of a belief vector is calculated for a bitstring of length 8 and epistasis 2. The fitness function f is defined in terms of eight other functions $f_1, \ldots, f_8$, where function f_i is used to determine the fitness contribution of bit b_i. Since the degree of epistasis is 2, the fitness of bit b_i also depends on the values of b_{i-1} and b_{i+1} (where, at the end of the bit string, we wrap around the ends to avoid edge effects). The individual fitness functions $f1, \ldots, f_8$ are defined using the lookup table in figure 15(a). In general, if N is the length of the bitstring $\vec{b}$, then

$$f(\vec{b}) = \sum_{i=1}^{N} f_i \left(b_{i-\frac{K}{2}} \cdots b_i \cdots b_{i+\frac{K}{2}} \right).$$

Although we consider, simply for reasons of simplicity, only two possible values for each b_i (full belief or full denial), there is no reason why we could not allow more finely-grained credal states. If we denote the number of credal states

$N = 8$ $\overbrace{\hphantom{000\ 001\ 010\ 011\ 100\ 101\ 110\ 111}}^{K=2}$

	000	001	010	011	100	101	110	111
f_1	1	0	0	0	1	0	1	-1
f_2	1	1	1	-1	2	0	1	-1
f_3	1	1	-1	0	0	1	-1	0
f_4	1	1	-1	-1	1	1	2	1
f_5	2	-1	-1	2	2	-1	1	1
f_6	0	2	0	-1	-1	0	-1	-1
f_7	2	0	-1	0	-1	2	0	0
f_8	1	1	0	1	0	-1	2	-1

(a) The lookup table for the fitness function f.

$$
\begin{aligned}
f(01001101) = \\
&f_1(101) + f_2(010) \\
&+ f_3(100) + f_4(001) \\
&+ f_5(011) + f_6(110) \\
&+ f_7(101) + f_8(010) \\
&= 0 + 1 + 0 + 1 \\
&+ 2 - 1 + 2 + 0 \\
&= 5.
\end{aligned}
$$

(b) Applying the fitness function f to the bitstring 01001101.

Figure 15: Applying a fitness function defined via a lookup table to a bitstring.

by A, then we see that the Weisberg-Muldoon epistemic landscape model is simply an NK-landscape with $N = 2$, $K = 2$ and $A = 101$, and a particular fitness function. (See appendix A for further details regarding the calculation of fitness functions on NK-landscapes.)

In order to see whether social learning and cognitive diversity help people reach the peak of greatest epistemic fitness on NK-landscapes, let us consider — as a baseline result — how a single independent agent would fare. Assume that the agent searches via *probe and adjust* as follows. The agent starts with a randomly selected belief vector of length N. In each iteration, the agent probes one randomly selected belief by considering its alternate value (0 if 1, and 1 if 0). If changing that belief yields an overall increase in fitness, the agent keeps the change; if changing that belief decreased the fitness, the change is rejected. Figure 16 illustrates the simulation results for a range of values of N and K. For each set of values of N and K, one hundred simulations were performed, running for 1,000 iterations. Each simulation used a randomly generated uncorrelated fitness function.[22] The mean fitness over all simulations (and the standard deviation) are shown.

		$N=8$	$N=16$	$N=24$	$N=48$	$N=96$
No social learning	$K=0$	0.66 (0.08)	0.67 (0.05)	0.66 (0.05)	0.67 (0.04)	0.66 (0.03)
	$K=2$	0.70 (0.07)	0.70 (0.05)	0.70 (0.04)	0.70 (0.03)	0.71 (0.02)
	$K=4$	0.69 (0.06)	0.70 (0.04)	0.70 (0.03)	0.70 (0.03)	0.70 (0.02)
	$K=8$	0.67 (0.06)	0.68 (0.04)	0.68 (0.03)	0.68 (0.02)	0.68 (0.02)
	$K=16$		0.64 (0.03)	0.65 (0.03)	0.66 (0.02)	0.66 (0.02)
	$K=24$			0.62 (0.03)	0.64 (0.02)	0.64 (0.02)
	$K=48$				0.60 (0.02)	0.61 (0.02)
	$K=96$					0.58 (0.01)
Social learning	$K=0$	0.66 (0.10)	0.67 (0.06)	0.67 (0.05)	0.67 (0.04)	0.67 (0.03)
	$K=2$	0.73 (0.07)	0.72 (0.04)	0.72 (0.04)	0.71 (0.02)	0.71 (0.02)
	$K=4$	0.70 (0.06)	0.72 (0.04)	0.72 (0.03)	0.71 (0.02)	0.72 (0.02)
	$K=8$	0.66 (0.05)	0.68 (0.04)	0.70 (0.03)	0.70 (0.02)	0.70 (0.02)
	$K=16$		0.64 (0.04)	0.65 (0.03)	0.66 (0.02)	0.66 (0.02)
	$K=24$			0.62 (0.03)	0.64 (0.02)	0.65 (0.02)
	$K=48$				0.60 (0.02)	0.61 (0.01)
	$K=96$					0.58 (0.02)

Figure 16: Social learning makes no difference to the global performance given uncorrelated fitness functions. (When $K=N$, the fitness function f_i for b_i depends upon the entire belief state of the agent.)

Now consider the possibility of social learning. Suppose we have a population consisting of some fixed number of agents. At the end of each iteration, each individual polls every other. If agent M changed b_i such that it yielded greater fitness to M, then all other agents incorporate that change.[23] As figure 16 shows, social

22. According to Kauffman and Weinberger (1989), a fitness function is said to be *uncorrelated* if "the fitness of 1-mutant neighbors [is] assigned at random from some fixed underlying distribution." For the purpose of this paper, an uncorrelated fitness function assigns the fitness to 1-mutant neighbors at random using the uniform distribution over $[0,1]$. More precisely, the fitness function f_i, specifying the fitness contribution for bit b_i and its surrounding epistatic region, is defined on the substring $b_{i-\frac{K}{2}} \cdots b_i \cdots b_{i+\frac{K}{2}}$ of $\vec{b}$, which has length $K+1$. The fitness contribution of each of the possible 2^{K+1} arguments to f_i are set to a randomly chosen value in $[0,1]$, drawn from the uniform distribution. With such an uncorrelated fitness function, knowledge of the values of the entries in the region of epistasis around some b_i gives no information as to whether the 1-mutant neighbor will have greater or lower fitness than the present belief vector.

23. Think of this as a model where each agent publishes the result of each experiment, and people always trust each other's results.

$$\cdots \boxed{1}\overbrace{\boxed{0}\boxed{1}\boxed{1}\underset{b_i}{\boxed{0}}\boxed{1}\boxed{0}\boxed{0}}^{K=6}\boxed{1} \cdots$$

$$w(b_i) = \tfrac{4}{7}$$

Figure 17: The correlated fitness function. The local fitness of any bit b_i is the relative frequency of the current value of $b_i = 1$ in its epistatic region.

learning *makes absolutely no difference* in cases of uncorrelated fitness functions. Furthermore, social learning may actually be *detrimental* for the performance of the agents, given epistasis.[24]

The crucial difference between the NK-model just described and the Weisberg and Muldoon epistemic landscape is this. Expressed as an NK-model, Weisberg and Muldoon assume fitness functions are highly *correlated*. Let us suppose, for the sake of argument, that the fitness function used is similarly highly correlated. In particular, assume that the fitness of b_i is simply the relative frequency of the current value of b_i in its epistatic region. (See figure 17 for an illustration.) The results for simulations where the fitness functions are correlated in this way, as in figure 18, show that, in this case, there is indeed a positive effect of social learning on the rate of epistemic progress.

Why does this matter? It matters because the Weisberg and Muldoon model builds into the basic topology of the epistemic landscape correlations which make social learning advantageous. As such, we should not be surprised to find, in the case they consider, that cognitive diversity and social interactions between agents can be beneficial. But, as the generalisation to NK-landscapes shows, social learning is not always beneficial. Whether social learning is beneficial or harmful depends on the topology of the epistemic landscape, a point of which we know very little. In some cases we might well suspect that social learning will be advantageous (e.g., does the problem decompose into sub-problems? Does the problem require a diverse set of skills possessed by no single individual?), but we cannot be sure that social learning will be advantageous because we won't know the epistemic landscape's topology.

We stress this point because, although Weisberg and Muldoon acknowledge that "[l]andscapes can be made more rugged, they can contain more information"

24. Suppose agent M finds a fitness-enhancing change for b_i. The reason why it may be detrimental for other agents to adopt this change is because the fitness increase, for M, depends on the particular value of the other beliefs in the region of epistasis around b_i. If other agents in the population don't have the same K other beliefs that agent M has, there's no guarantee that changing the value of b_i for them will have the same effect.

	$N = 96$			
	No social learning		Social learning	
	Mean (SD)	Mean steps	Mean (SD)	Mean steps
$K = 2$	0.92 (0.01)	994.43	0.94 (0.02)	508.14
$K = 4$	0.93 (0.01)	997.74	0.95 (0.02)	277.66
$K = 8$	0.90 (0.01)	994.77	0.93 (0.03)	253.35
$K = 16$	0.90 (0.02)	989.29	0.92 (0.04)	215.37
$K = 24$	0.90 (0.02)	927.92	0.92 (0.06)	197.83
$K = 48$	0.93 (0.04)	916.71	0.97 (0.08)	183.61
$K = 96$	1.00 (0.00)	600.94	1.00 (0.00)	185.58

Figure 18: When correlation exists for the fitness function, social learning makes a positive difference in the rate at which epistemic progress occurs.

and so on, they suggest that some of their findings do, in fact, generalise. They state that "[e]ven with our current models and current landscape, we have observed a number of very interesting general trends about the division of cognitive labour." What are some of these general trends? For one, that

> "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." (249)

And also:

> "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." (250)

And, finally (italics added):

> "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." (250)

Each of these claims would be uninteresting if "the [epistemic] landscape" only referred to the epistemic landscape modelled in the paper. The reason why these claims are interesting is that they gesture towards general properties of scientific practice and suggest fruitful ways of organising scientific research. Yet it only makes sense to say that one *should* employ a population of mavericks in cases where the epistemic landscape is such that the maverick strategy would be beneficial, and it is far from obvious that the maverick strategy will prove to be beneficial on an arbitrary epistemic landscape, or when the model is adjusted to allow for greater realism by, say, incorporating observation error.[25] So, although the general trends have a certain degree of intuitive appeal, it is unclear to what extent these claims are, in general, justified by the Weisberg and Muldoon epistemic landscape model.

8. Conclusion. In recent years, the Weisberg-Muldoon model has received a considerable amount of attention regarding its purported claim to show that there are epistemic reasons for the division of cognitive labor. In particular, Weisberg and Muldoon alleged to show that the "maverick" research strategy is far better than its competitors,[26] and one of the "general trends" (249) they observed is that "to be maximally effective, scientists need to really divide their cognitive labor" (227). We have argued that these two claims are not true. Maverick scientists do not perform far better than their competitors, such as the HE rule, once the implementation errors which handicapped the other search types have been corrected. By proper bookkeeping, we have shown that the increase in the performance of the mixed population is only due to the performance of the added mavericks. As for the benefits of cognitive diversity, we have constructed at least one other search strategy, the "swarm scientist", which, in some cases, outperforms the maverick scientists.

In saying this, we do not wish to be understood as arguing that there are no epistemic reasons for cognitive diversity. We are simply pointing out that, despite its intuitive appeal, the Weisberg and Muldoon model does not succeed in showing

25. One way this could be done would be to incorporate a probability that an agent incorrectly perceives the true epistemic fitness of the point they currently occupy on the landscape. Suppose agents "publish" their findings, and this incorrect report will be the epistemic fitness ascribed to that point on the landscape by other agents. (If another agent visits that same point on the landscape, they have the chance to perceive correctly the "true" value, of course.) If the change of experimental error is sufficiently high, the maverick strategy, in these cases, could well prove disadvantageous: mavericks, by avoiding parts of the landscape already explored, could miss finding peaks for a very long time if that peak was explored previously, but incorrectly identified.

26. "[I]f 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." (250)

that there are epistemic reasons for cognitive diversity. Furthermore, since so much in their methodology turns on assumptions regarding the specific nature of the epistemic landscape — something whose very nature is beyond our ken — we are skeptical as to whether their particular method of arguing for the epistemic benefits of cognitive diversity can ever succeed.

A. Generating uncorrelated fitness functions on large NK-landscapes. Following Kauffman and Levin 1987, fix integers N and K such that $0 \leq K \leq N$. The total fitness of a bit string $b_1 b_2 \cdots b_N$ is calculated using fitness functions $f_1, \ldots, f_N$, where f_i is applied to bit b_i and the surrounding region of epistasis consisting of the K bits flanking b_i on the left and right. (To prevent edge effects, we assume the bit string wraps at the edges to ensure that all bits have regions of epistasis the same size.) For small N and K, the fitness function can be defined using a lookup table specifying all 2^{K+1} values for each of the N fitness functions, as shown in figure 15.

For large N and K, it is not feasible to define fitness functions using a lookup table. If $N = 100$ and $K = 60$, an uncorrelated fitness function would have $100 \cdot 2^{61}$ different values. However, it is possible to *procedurally* define an uncorrelated fitness function using a trick similar to that used by programmers of the 1984 video game *Elite*. (That game had to fit 8 different universes, each containing 256 planets with unique properties, into 16 kilobytes of memory.) We used a common 19937 Mersenne Twister RNG to procedurally generate the fitness functions.

Given fixed values of N, K with $K < N$, let $S = (s_1, \ldots, s_N)$ be a list of N different *salt values*. Denote the region of epistasis around bit i at time t by $\eta_i^t = \{\eta_{i,1}^t, \ldots, \eta_{i,K+1}^t\}$. The fitness contribution of b_i at t is determined by seeding the RNG with the bits of $s_i \oplus \eta_i^t$ (zeroing out the remaining bits in the state space of the RNG) and generating a random real between 0 and 1. As long as the *length* of $s_i \oplus \eta_i^t \leq 19{,}968$, we obtain a unique seed. This is because the state of the 19937 Mersenne Twister is 624 words, each with 32 bits, so there are $624 \cdot 32 = 19{,}968$ bits in the state. If the length of $s_i \oplus \eta_i^t$ is sufficiently less than the length of the state space, the random value will be unique (and reproducible). This yields a completely uncorrelated fitness function.

References.

Arrow, Kenneth J., Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, and Forrest D. Nelson. 2008. "The promise of prediction markets." *Science* 320 (5878): 877.

Bonabeu, Eric, Marco Dorigo, and Guy Theraulaz. 1999. *Swarm Intelligence : From Natural to Artificial Systems: From Natural to Artificial Systems.* Oxford University Press.

Clearwater, Scott H., Bernardo A. Huberman, and Tad Hogg. 1991. "Cooperative Solution of Constraint Satisfaction Problems." *Science* 254 (5035).

Couzin, Iain D., Jens Krause, Nigel R. Franks, and Simon A. Levin. 2005. "Effective leadership and decision-making in animal groups on the move." *Nature* 433 (7025): 513–516.

Hong, Lu, Scott E. Page, and William J. Baumol. 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences of the United States of America* 101 (46): 16385–16389.

Jackson, Susan E., Aparna Joshi, and Niclas L. Erhardt. 2003. "Recent Research on Team and Organizational Diversity: SWOT Analysis and Implications." *Journal of Management* 29 (6): 801–830.

Kauffman, Stuart A., and Edward D. Weinberger. 1989. "The NK model of rugged fitness landscapes and its application to maturation of the immune response." *Journal of Theoretical Biology* 141 (2): 211–245.

Kauffman, Stuart, and Simon Levin. 1987. "Towards a general theory of adaptive walks on rugged landscapes." *Journal of Theoretical Biology* 128 (1): 11–45.

Kerr, Norbert L., and R. Scott Tindale. 2004. "Group Performance and Decision Making." PMID: 14744229, *Annual Review of Psychology* 55 (1): 623–655.

Kitcher, Philip. 1993. *The Advancement of Science: Science Without Legend, Objectivity Without Illusions.* Oxford University Press.

Krause, Stefan, Richard James, Jolyon J. Faria, Graeme D. Ruxton, and Jens Krause. 2011. "Swarm intelligence in humans: diversity can trump ability." *Animal Behaviour* 81 (5): 941–948.

Polzer, Jeffrey T., Laurie P. Milton, and William B. Swarm. 2002. "Capitalizing on Diversity: Interpersonal Congruence in Small Work Groups." *Administrative Science Quarterly* 47 (2): 296–324.

Reynolds, Craig W. 1987. “Flocks, herds and schools: A distributed behavioral model.” In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques,* 25–34. SIGGRAPH ’87. New York, NY, USA: ACM.

Seeley, Thomas D. 2010. *Honeybee Democracy.* Princeton University Press.

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

Sumpter, David J. T. 2010. *Collective Animal Behavior.* Princeton University Press.

Thomas, David A., and Robin J. Ely. 1996. “Making differences matter.” *Harvard business review* 74 (5): 79–90.

Ward, Ashley J. W., Jens Krause, and David J. T. Sumpter. 2012. “Quorum Decision-Making in Foraging Fish Shoals.” *PLoS ONE* 7 (3): e32411.

Weisberg, Michael, and Ryan Muldoon. 2009. “Epistemic Landscapes and the Division of Cognitive Labor.” *Philosophy of Science* 76 (2): 225–252.

Wright, Sewall. 1932. “The roles of mutation, inbreeding, crossbreeding, and selection in evolution.” In *Proceedings of the Sixth International Congress on Genetics,* 355–366.