<https://helda.helsinki.fi>

Value of cognitive diversity in science

Pöyhönen, Samuli

2017-11

Pöyhönen, S 2017, ' Value of cognitive diversity in science ', Synthese , vol. 194 , no. 11 , pp. 4519-4540 . <https://doi.org/10.1007/s11229-016-1147-4>

<http://hdl.handle.net/10138/230131>

<https://doi.org/10.1007/s11229-016-1147-4>

unspecified

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Value of cognitive diversity in science

Samuli Pöyhönen

Final corrected manuscript, 30 June 2016.

Abstract When should a scientific community be cognitively diverse? This article presents a model for studying how the heterogeneity of learning heuristics used by scientist agents affects the epistemic efficiency of a scientific community. By extending the epistemic landscapes modeling approach introduced by Weisberg and Muldoon, the article casts light on the micro-mechanisms mediating cognitive diversity, coordination, and problem-solving efficiency. The results suggest that social learning and cognitive diversity produce epistemic benefits only when the epistemic community is faced with problems of sufficient difficulty.

Keywords social epistemology · diversity · social learning · division of cognitive labor

1 Introduction: Why scientific communities should be diverse

The rationality of individual scientists is neither a sufficient nor necessary condition for achieving good collective outcomes in research. The literature on the social epistemology of science even suggests that having egoistic, stubborn, or otherwise epistemologically sullied agents as members of a scientific community can – under appropriate conditions – improve its epistemic efficiency (Weisberg 2010; Mayo-Wilson et al 2011). In such cases, the increase in the efficiency of knowledge production capacities of the community is often due to its increased diversity.

As especially feminist social epistemologists have emphasized, differences in how people see the world due to their different backgrounds, social identity,

Samuli Pöyhönen,
Social and Moral Philosophy / Department of Political and Economic studies,
University of Helsinki
E-mail: samuli.poyhonen@helsinki.fi

and gender are a precondition for effective critical discourse, and hence important for avoiding bias and producing objective scientific knowledge (Longino 1990, 2002).¹ Solving complex scientific problems often requires that they are attacked with a wide range of different research approaches (Solomon 2006; Page 2008), and in many cases, the connection between diversity in general and efficiency of collective problem-solving appears to be mediated by factors such as variation in background beliefs, concepts used, and reasoning styles of scientists, that is, *cognitive diversity*. The model presented in this article shows how a particular aspect of cognitive diversity – agents’ different learning heuristics – affects the epistemic division of labor within a community and thereby influences the epistemic performance of the community in different kinds of research domains.

Let us begin by considering a few examples. According to the popular perception of science, scientific discoveries arise from flashes of insight by exceptional individuals. For example, Nikola Tesla has long been thought of as a lone genius, whose numerous scientific and technological inventions appeared to arise mainly from his independent inquiries into the nature of electro-magnetic phenomena (Novak 2014). The lone genius model seems to also apply to Yitang Zhang, an unknown mathematician who in April 2013 proved a weaker variant of the twin prime conjecture, a great result in the history of number theory that had widely been regarded as too difficult to solve with the current resources of mathematics. In interviews after receiving the MacArthur award for his accomplishment, Zhang has attributed his success mainly to perseverance, refusing to switch topics even after long stretches of time with no progress on a problem (Klarreich 2013; MacArthur-Foundation 2014).

However, for the most part contemporary scientific research is far from a solitary endeavor. By talking to supervisors, colleagues, reading journals and going to conferences, academic researchers are constantly collecting new ideas from others, and on the lookout for ways to improve their research based on social feedback. We are often happy to align our research questions, methods, and theories with those of others in order to produce scientific results of at least moderate significance and impact. And as history of science shows, often even major breakthroughs in science have resulted not from solitary work but from *bricolage*, skillful and often lucky combination of ideas from a variety of sources (cf. Johnson 2011).

These examples illustrate what I mean by scientists having different learning heuristics. Since resources and time are always limited, each scientist repeatedly faces a decision of how to conduct her research: whether she should spend the day at the bench running experiments and analyzing her data, or whether to engage in social learning. It appears that successful communal knowledge production needs both individual and social learners. On the one hand, laying the foundations for new paradigms and scientific revolutions requires that a scientist or a group of scientists goes against the grain and

¹ Diversity in science obviously matters also beyond strictly epistemological concerns. For example, it can help fight epistemic injustice (Fricker 2007) and promote institutional epistemic virtues, e.g., by reducing prejudice and ethnocentrism (Anderson 1995).

develops new ways of thinking independently of the existing paradigm. On the other hand, efficient problem-solving in normal scientific research requires that a large share of research work is allocated to the currently most promising research approaches. But how much social learning should there be in a scientific community, and when should it occur?

To study the effects of the different research strategies on epistemic performance, I adopt a population-modeling approach to scientific problem-solving. This approach treats the scientific community in a research field as an epistemic system (cf. Goldman 2011). Scientific knowledge is not understood as residing primarily at the level of individuals, but instead, it is treated as a system property, determined both by the work done by individual scientist as well as the adequacy of their social coordination and division of cognitive labor (Polanyi 1962; Hull 1988; Longino 2002).

More precisely, the systems approach to the epistemology of science suggests the hypothesis that the efficiency of a scientific community is determined by at least three kinds of factors:

1. The distribution of the cognitive properties of individual agents in the community (cognitive diversity)
2. The organizational properties of the community (e.g., its communication structure, reward system)
3. The nature and difficulty of the problem-solving task faced by the community.

The model presented in this article examines the dependencies between these factors and the problem-solving capacity of a scientific community.

Since Philip Kitcher's (1990; 1993) seminal work on the topic, diversity and the social organization of science have been discussed within a variety of models and modeling frameworks (Strevens 2003; Weisberg and Muldoon 2009; Zollman 2010; De Langhe 2014; Muldoon 2013). Among these different modeling approaches, currently the most amenable framework for studying diverse learning heuristics in science is the epistemic landscapes model (EL model) by Michael Weisberg and Ryan Muldoon (2009), which represents scientific research as a population of scientist agents foraging on an epistemic landscape. The current article extends this modeling approach in two ways. First, I argue that the EL model (i) suffers from several interpretational problems, (ii) builds on problematic assumptions about the behavioral rules followed by scientist agents, and (iii) applies to an overly narrow set of research topics. Secondly, by introducing new assumptions regarding the implementation of the social-learning heuristics, measurement of epistemic performance, and the complexity of research topics, the broadcasting model (introduced in Section 3) aims to provide a more applicable account of the micro-mechanisms mediating cognitive diversity, division of cognitive labor, and epistemic success.

2 Research as foraging on an epistemic landscape

There are a number of reasons why epistemic landscape modeling is especially suitable for studying the effects of cognitive diversity on collective epistemic performance: First, unlike many other models, the framework is capable of representing genuine cooperation, not only competition, between agents (cf. D’Agostino 2009). Secondly, the agent-based approach allows for a natural representation of bounded rationality and cognitive diversity in terms of different learning heuristics employed by agents. Thirdly, unlike the otherwise elegant NK models used in some alternative approaches (Alexander et al 2015; Lazer and Friedman 2007; Page 2008), the three-dimensional representation of the fitness landscape allows easy manipulation of the epistemic structure of the studied research field (landscape topography). This is crucial for the experiments I run with my model. Let us begin by examining how scientific research, cognitive diversity, and division of cognitive labor can be represented in epistemic landscapes modeling.

2.1 The original EL model

As in most agent-based simulations, the model primitives in the EL model concern the attributes and behavior of agents, and the structure of their environment. Generally, the model builds on an analogy to fitness landscape models used in ecology: Collective search is portrayed as the movement of a population on a landscape, where the height parameter of a particular environment point represents its fitness value (Wright 1932).

Applied to the social epistemology of science, the model is interpreted as follows: A scientific research topic (e.g. synthetic biology, astrophysics, endocrinology) is represented as an n -dimensional space, where the dimensions up to $n-1$ constitute the different aspects of a *research approach*. For example, attempting to synthesize novel DNA nucleotides and studying the stability of these molecules by computational methods are independent but both necessary research approaches in synthetic biology (Weisberg and Muldoon 2009).

The last (n :th) dimension stands for the epistemic significance of the approach. Here epistemic significance is understood according to Philip Kitcher’s (1993, Ch. 4) analysis: Significant statements are the ones which answer significant questions. Significant questions, in turn, are ones that help us uncover the structure of the world, or at least organize our experience of it. Respectively, we can define the significance of a research approach as the significance of truth (or truths) that can be uncovered by using the approach. If we, furthermore, make the idealizing assumption that the scientists working in a field share the same judgments about significance, it is possible to represent the community as populating and perceiving a shared epistemic landscape.

For the sake of simplicity, in the model the large number of dimensions defining an approach are collapsed into two, and hence a three-dimensional landscape can be used to represent the epistemic structure of the research

topic at hand. Furthermore, the space is divided into discrete patches, where each patch represents a combination of (i) a research question being investigated, (ii) instruments and methods for gathering and analyzing data, and (iii) background theories used to interpret the data. The greater the elevation of a particular patch, the more significant truths the research done by using that approach discloses. The simulations that Weisberg and Muldoon report in their paper concern a smooth, mostly flat landscape with two Gaussian-shaped hills of positive epistemic significance.

An important difference to many other models in recent social epistemology is that in epistemic landscape modeling, agents are not portrayed as Bayesian conditionalizers or competent maximizers of expected utility. Instead, bounded rationality is implemented in the following way. Initially all of the agents working on the research topic are placed randomly on zero significance areas of the landscape, and at each turn of the simulation, they move at a velocity which is small compared to the size of the landscape. Their movement is guided by a satisficing search for increasing epistemic significance in their Moore neighborhood. This search-based implementation of learning embodies an assumption about the local nature of information available to scientists when deciding about how to proceed about their future research: Changing one's research approach is a gradual and costly process, and no individual agent has access to global information about how epistemic significance is distributed on the landscape, i.e. what the most effective research approaches are.

Although the notion of cognitive diversity could be seen to include a variety of factors ranging from background beliefs, training and education, perception of significance, to general intellectual style, and so on, the EL model actually focuses on one particular type of cognitive diversity, variation in the scientists' research heuristics in terms of individual and social learning.² While this is by no means the only important source of cognitive diversity, epistemic landscape models do show how it brings about and maintains another kind of diversity in the research field – division of cognitive labor, i.e., the distribution of agents over different research approaches (Alexander et al 2015, pp. 436–437).

In the EL model, the diversity in learning heuristics is implemented in terms of three kinds of agents (controls, followers, and mavericks) corresponding to different rules for engaging in individual and social learning. In this case, when conducting individual learning, agents search for epistemically significant results by interacting with nature, so to speak (e.g., by gradually improving their methods, experimenting with changes in instrumentation, and in their theoretical assumptions). Social learning, in turn, refers to improving one's research approach based on the exchange of information with other agents (e.g., by reading publications of other scientists). In the EL model, both individual and social learning are represented as variations of gradient-climbing on the

² This same aspect of cognitive diversity is also the topic of the subsequent epistemic landscape models introduced by Thoma (2015) and Alexander et al (2015). Zollman (2010), in contrast, deals with a different aspect of diversity. In the context of theory choice, he examines how (a) the flow of information in social networks and (b) the strength of agents' prior degrees of belief influence the emergence of consensus in theory choice.

landscape. Control agents ignore all social information, while followers prefer patches already visited by others, and mavericks avoid already examined approaches.

The most striking finding suggested by the model is that maverick scientists immensely improve the problem-solving efficiency of the scientific community. Intuitively, the EL model suggests that the researcher population benefits from the presence of explorer members who explicitly avoid methods and approaches employed by others. However, the simulations by Weisberg and Muldoon also imply a much stronger claim, according to which a homogeneous population of mavericks is more efficient than a diverse community consisting of both followers and mavericks. Both these results have been shown to be open to various kinds of criticism, and assumptions in the model regarding both the search rules of agents and the topography of the landscape have been challenged.

First, as shown by their critics, some of the central results of Weisberg and Muldoon's Netlogo simulation arise from implementation errors in the control and follower search rules (Alexander et al 2015). Moreover, as Thoma (2015) suggests, alternative forms of the search rules which appear to be just as compatible with actual scientists' behavior as those suggested by Weisberg and Muldoon, lead to clearly different outcomes. The general applicability of results from the EL model is further compromised by the fact that in the absence of robustness analysis (Grimm and Berger 2015), Weisberg and Muldoon's published findings are not sufficient for establishing that their results follow from the substantial assumptions of the model, and not from implementation-related auxiliary assumptions. On the contrary, the fact that the model structure is borrowed from another domain (ecology) raises the worry that at least some of the results might be artifacts produced by the imported modeling framework itself.

2.2 What does an epistemic landscape represent?

Another set of worries concerns the landscape and its interpretation. The choices that Weisberg and Muldoon make about landscape topography in their simulations are more controversial than they admit, as those choices stand for crucial assumptions about the epistemic structure of the research topic (Alexander et al 2015). Epistemic landscapes underlying real scientific research probably involve a greater number of interdependencies between the elements of approaches (questions, instruments, methods, theories) than the smooth two-dimensional landscape can represent. Consequently, results from the EL model are conditional on the choice of particular kinds of simple research topics, which might often not be the ones encountered in cutting-edge scientific research.

However, when Alexander and his coauthors treat the EL model as a special case of NK landscapes (2015, p.446), they risk falling prey to another misinterpretation of the landscape idea. In typical applications, NK landscapes are used to represent a search space for a problem, where hill-climbing towards

maxima points is interpreted as a search for better solutions to the problem, i.e., minimization of the error function (Lazer and Friedman 2007; Kauffman and Levin 1987). Although Alexander et al. purport to reinterpret the NK landscape as an epistemic landscape, in fact they still use their model to represent interdependencies between a set of propositions. Although it remains somewhat unclear what the *global performance* measured by their simulations refers to, their results against the usefulness of social learning seem to rely on observations about the effect of social learning on how quickly and how often the agents find the peaks on rugged landscapes.

As I argue below, such a measure tracks the wrong property – success on an epistemic landscape is more subtle. A landscape *does not represent a search space for a single problem* (where patches correspond to different states of the belief vector of an agent, as Alexander et al. suggest). Instead, the landscape stands for a distribution of epistemic significance over a set of different but complementary research approaches. The EL model is not (or at least it should not be) primarily used to study the performance of individual agents and whether they find the peaks on the landscape. Instead, the modeling results concern the *dependency between cognitive diversity and coordination*, where coordination refers to the distribution of agents on non-zero parts of the landscape. Consequently, the global maximum on an epistemic landscape should not be understood as the “correct solution to the problem.” There is simply not a single problem to solve. What the maximum (when unique) stands for is the most productive way to advance the inquiry on a research topic.

2.3 Measuring collective epistemic performance

Keeping in sight this coherent interpretation of the landscape suggests that both of the ways that Weisberg and Muldoon measure the epistemic success or efficiency of a scientist population are inadequate, or at least insufficient, as they track very particular aspects of success. First, Weisberg and Muldoon study (i) how often and (ii) how quickly peaks are reached. As argued above, (i) and (ii) do not directly measure community-level epistemic success: A situation where the whole population quickly lands on the global maximum looks good by these measures, but the coordination between the agents is very poor. Imagine, for example, a population of synthetic biologists all of whom converge on doing the same kind of computational modeling work, because at a certain time, that is what produces the most epistemic significance. In such a situation there would be no division of labor, and the success of such a strategy would surely prove short-lived.

The other measure, epistemic progress, defined as the proportion of explored non-zero patches on the landscape, captures the coordination aspect better. However, this measure disregards the epistemic significance level of the explored patches (as long as it is not zero). Consequently, beyond the simple landscape studied by Weisberg and Muldoon, epistemic progress can

be a misleading measure of success. To see why, imagine a situation where there is a certain number of *dead-end research approaches* within the research field, that is, research approaches of low epistemic significance surrounded by even worse alternatives. Due to factors such as badly formulated research questions and inappropriately chosen instrumentation and analysis methods, such low-significance research approaches where no incremental change can improve the approach are likely to crop up in all non-trivial scientific research fields. The resulting dead-end approaches can be represented as local maxima on the landscape (figure 1).

In general, there are often interdependencies between aspects of research approaches, which cannot be captured by Weisberg and Muldoon’s smooth landscape. For example, certain experimental approaches might be useful only when used with a particular data analysis method and background theory. If we assume that aspects of research approaches are mapped on each dimension of a landscape so that similarity with respect to a particular aspect is represented by distance along that axis, postulating a smooth landscape would amount to assuming that the interdependencies mentioned above do not arise in the particular research field. Hence, rugged landscapes characterized by several local maxima can more faithfully represent the complex problem-solving situations faced by scientific researchers (cf. Alexander et al 2015).³

What epistemic progress, as it is defined by Weisberg and Muldoon, fails to convey is whether the population has found areas of high epistemic significance, or whether all of the explored patches reside on low significance areas surrounding the dead-end approaches. Although these difficulties do not arise on smooth landscapes, on rugged ones epistemic progress is not an adequate way to measure the collective performance of an epistemic community.

2.4 Behavioral rules of agents

Finally, the model behavior resulting from adding dead-end approaches on the landscape also suggests that the search rules of agents implemented in the EL model are only applicable to the simplest epistemic landscapes. It is well known that gradient-climbing search is sensitive to local maxima and fails to find global maxima on rugged landscapes (Russell and Norvig 2003, Ch. 4.4.1). As all three behavioral rules of the agents in the EL model rely only on local-search based hill climbing, they cannot capture the research heuristics employed by scientists working on more complex research topics.

I conclude this section by reporting on a modeling experiment which suggests that the radical results obtained by Weisberg and Muldoon might be, at least partly, artifacts of the problematic aspects of their model discussed above.

³ For more general discussions of the use of rugged landscapes in problem-solving research, see, e.g., Page (2008, Ch.1), and Lazer and Friedman (2007).

Fig. 1 Landscape with dead-end research approaches

2.5 Are mavericks unbeatable?

Weisberg and Muldoon’s most striking finding is that a homogeneous population of maverick agents is epistemically superior to all mixed populations of controls, followers, and mavericks. Intuitively, this result follows almost trivially from the fact that epistemic progress is measured in terms of the proportion of explored patches on the landscape, and the maverick movement rule is tailored to maximize the efficiency of exploration. Let us, however, introduce a small cost of exploration into the model – for example a delay of ten time steps when an agent enters an unexplored patch. In the next section I will return to the question of how to measure epistemic success, but for now, let us, instead of epistemic progress, keep track of the *average epistemic significance* of the population (calculated as the mean of the elevations of individual agents) over time. It now turns out that the cost of exploration can make a mixed population of followers and mavericks more efficient than a pure maverick population (Figure 2). With this particular delay value ($t_d = 10$), a diverse population consisting of 50-50-mix of trailblazing mavericks and faster followers lands most quickly on high-significance areas.⁴

Furthermore, Figure 2 also shows that once the implementation errors in the search rule for control agents have been corrected (cf. Alexander et al 2015), controls perform even slightly better in the mixed population than maverick agents do. It appears that when exploration is more costly than

⁴ For a Python replication of Weisberg and Muldoon’s model, and the source code for the simulations in this article, see the repository at <https://github.com/samulipo/broadcasting/>. Simulations were conducted with $n=50$ for each data point. Error bars (when shown) stand for one standard deviation in sample.

Fig. 2 EL model with cost of exploration. Progress measured as average epistemic significance of the population after 200 time steps.

exploitation, simple hill-climbing is a better strategy than a similar strategy biased toward avoiding already explored approaches. In sum, small modifications of the model compromise the generality of Weisberg and Muldoon’s result that a large proportion of maverick scientists in an epistemic community drastically improves its performance.

3 The broadcasting model

The broadcasting model differs from the EL model with respect to (i) how epistemic success is measured, (ii) which landscapes are studied, and (iii) what the search rules employed by the agents are. In doing so, it avoids the problems discussed in the previous section.

I begin with measurement and landscape topography. I suggest that a quantity, which can be called the *epistemic work* done by an individual agent (and aggregated into the epistemic work done by the community), is a meaningful and coherent measure of success. Given the characterization of epistemic significance in Section 2.1, it seems that a population of scientist agents can be said to be the more successful the more significant truths the agents can communally uncover, and the efficiency of the community is determined by how quickly the accumulation of truths is done. I define epistemic work w_i done by an agent i since the beginning of the simulation as the sum of the epistemic significances of the patches visited by the agent until time t , weighted by a constant time scale factor $\lambda \in [0, 1]$. Hence for the agent i , $w_{i,t} = \sum_{\tau=0}^t \lambda s_{i,\tau}$, where $s_{i,\tau}$ is the significance of the patch visited by agent i at time τ . The

epistemic work W_t done by the population at time t is simply the sum of work done by the individual agents.⁵

Measuring epistemic efficiency in terms of epistemic work over time is, alone, not sufficient for capturing the need for genuine coordination or division of labor between scientists. Like success in finding peaks, epistemic work could also be maximized by minimizing coordination and guiding the whole population to the global maximum. To capture the dynamics of the division of labor, adopting a research approach that has already been used by another scientist should be associated with diminishing marginal returns. In order to represent this assumption that research done using a particular approach decreases the payoff from further research with the same approach, let us introduce *landscape depletion*: When an agent receives the amount $\lambda s_{i,\tau}$ of epistemic payoff by visiting a patch, the same amount of “significance mass” is removed from the landscape. In the current simulation, this was done by lowering the elevation of the patch in question by the same amount. The coefficient λ determines how quickly depletion occurs: the smaller λ is, the more time it takes to deplete a patch.⁶

I have two main reasons for adopting this new measure of epistemic success. First, epistemic work avoids the rather serious conceptual problems with Weisberg and Muldoon’s two original measures. Secondly, it allows a more natural interpretation of what it means to visit a patch. By not attributing any epistemic payoff to revisiting a patch, Weisberg and Muldoon are drawn to suggest one of two implausible ideas: Either that there’s only one significant truth to be uncovered per research approach (2009, footnote 3), or that all of the significant truths from an approach are always uncovered within one time step of the simulation. In contrast, together epistemic work and landscape depletion allow for a plausible interpretation of collective search on the landscape: What matters for epistemic advancement is that an agent spends a unit of time on a patch of considerable epistemic significance – not only the fact that a non-zero patch has been visited. That is, collective problem-solving is advanced when a scientist spends some time applying a significant research approach to a meaningful problem. A core challenge of division of cognitive labor within a community can hence be formulated in the language of my model: How, as a community, should the scientists organize their joint work so that they can harvest as much significance mass from the landscape as quickly as possible?

The third major difference to the EL model concerns the learning heuristics used by the agents. As discussed above, all social learning in the EL model occurs by studying traces left by other agents in one’s Moore neighborhood. That is hardly a well-founded assumption. As specified by Weisberg and Mul-

⁵ Consequently, the *average epistemic significance* of the population at a particular time (used in Section 2.3) is the change in W over one time step, scaled by a constant.

⁶ Weisberg and Muldoon (2009, p.232), suggest implementing such interaction between agents and landscape as a possible extension of their model. Likewise, Thoma (2015, footnote 7) brings up the idea of a modified model where revisiting patches uncovers further epistemic significance, but she does not develop the idea further.

doon, a single epistemic landscape represents a rather constrained research topic (e.g., the study of opioid receptors in chemical biology, or critical phenomena in statistical physics), and so it can safely be assumed that researchers are aware of each others' work. There is no reason for social learning to be constrained only to agents who use approaches very similar to one's own (cf. Thoma 2015).

In the broadcasting model, the cognitive diversity in the population pertains to the agents' different *thresholds for social learning*. For several reasons (an agent's risk preference, her assumptions about the size and shape of the landscape and the length of her career, etc.), some agents prefer to collect immediate epistemic payoff from individual learning rather than investing in expected future gain from social learning, whereas others value long-term success more.

At each time step, every agent follows the same decision procedure. All agents on the landscape are potentially visible to each other as sources of social information (based on their ability to broadcast their findings through publications, conference presentations, etc.), and at each turn every agent gets to observe the differences in epistemic significance between her approach and those of others. She compares the expected payoff from social learning to the assumed gain from doing individual search in her local neighborhood, and decides which kind of research (social or individual learning) to conduct during the following time step.⁷ More specifically, the decision procedure goes as follows:

```

ASK: Is the epistemic significance of the current patch higher than
or equal to the one on previous time step?
  If YES: Do not alter heading.
  If NO:
    Calculate the highest expected payoff from social learning
    (for details, see below).
    ASK: Is that payoff higher than agent's threshold
     $\alpha_i$  for social learning?
      If YES: Take new heading  $h_i$  towards the most successful peer
      If NO: ASK: Are there any higher patches in the agent's
              Moore neighborhood?
                If YES: Take heading  $h_i$  towards a randomly
                    selected higher patch
                If NO: Set velocity  $v_i$  to 0
    Move with velocity  $v_i$  to heading  $h_i$ 

```

In this stepwise heuristic, an agent could calculate the expected payoffs from social and individual learning in various ways. In the current simulations we focus on the case where agents are motivated simply by gain in their epistemic significance level, i.e. trying to find as significant a research approach as possible. The behavioral rule used by the agents in the simulations reported in the next section is consistent with them conducting a simple cost-benefit analysis of whether to engage in social learning: The benefit from social learning ($\Delta s_{j,i}$) is the difference in epistemic significance level of the agent i and

⁷ For applications of similar social-learner-explorer strategies in the literature on cultural inheritance and evolution, see Enquist et al (2007) and Borenstein et al (2008).

Fig. 3 The cognitive diversity in the population can be visualized as agents having "cones of vision" of different breadths. Given α_A , agent A would consider B, but not C, as a potential source of social information.

another agent j . The distance between the two agents' approaches determines the time Δt it takes for i to adopt j 's approach, and hence it represents the foregone possibilities for individual exploration. Each agent is characterized by a particular value of the social-learning threshold α_i . An agent engages in social learning if

$$\alpha_i < \frac{\Delta s_{j,i}}{\Delta t} \quad (1)$$

In other words, if at a particular time in the simulation the expected gain per timestep from adopting a fellow scientist's approach exceeds the agent's social-learning threshold α_i , she chooses social learning. However, if there are currently no agents around following whom would exceed the threshold, an agent defaults to individual learning, which is implemented as a satisficing form of gradient climbing (similar to the control rule in the EL model). Hence, α_i can be understood as the agent's expectation about the average increase in significance level that she can obtain from individual learning.⁸

Geometrically, the cognitive diversity in the population can be represented as each agent having a "cone of vision" of particular breadth, within which she agrees to pursue a more successful approach of a peer (Figure 3). Individualist scientists – like Tesla, at least when portrayed as in the example in section 1 – could be seen as having a narrow cone of vision, whereas less ambitious members of the research workforce are happy to settle for social-learning opportunities of smaller expected payoff, and hence have broader cones.

⁸ Simulation experiments were also run with agents who, instead of aiming for maximum significance level, aim to maximize the epistemic work done over a future time period, and yet more sophisticated ones who try to take landscape depletion into account by exponentially discounting for distant payoffs based on λ . Such decision rules result in a field of vision delimited by a surface of revolution drawn by a non-linear function. Careful analysis of such situations must be left as a task for future work, but in initial experimentation the differences in the shape of the cone did not lead to qualitative changes in the results.

To summarize, all the agents in the population share the structurally same decision procedure, and the cognitive diversity concerns their different tendencies to engage in social learning. Unlike in the EL model where diversity is represented in terms of discrete agent categories, the model allows continuous variation, and social learning thresholds can be drawn freely from a continuous distribution. This way of representing cognitive diversity avoids the problematic artifacts resulting from the EL model’s all-or-nothing implementation of the different learning profiles, and crucially, the new implementation of the learning heuristics captures the idea that social learning can help avoid myopia and motivate long-term projects involving movement downhill through areas of lower significance.

Apart from these modifications, other structural assumptions in the model are the same as in the EL model.

4 Results

In the simulations, the behavior of a population consisting of 50 agents was studied on a 101x101 toroidal landscape. The central question of interest concerned how the different tendencies for social learning present in the scientist population affect its epistemic efficiency on different kinds of landscapes. Variations in environment structure were introduced by varying the smoothness of the landscape and the levels of the time scale parameter λ .

4.1 Smooth landscapes

On a smooth landscape resembling that of the original EL model, social learning does not increase the epistemic efficiency of a population of agents. Instead, exploration of the parameter space (summarized in Table 1; more extensive data available in the supplementary materials, documents 1 and 2) shows that for all values of the time scale parameter λ , the epistemic work done by the community slightly increases as values of α in the population increase, that is, when agents become less eager to engage in social learning. Similarly, diverse populations with lambdas drawn from a uniform distribution $U(1, 100)$ slightly outperform pure populations of social learners.⁹

In addition to epistemic work, the performance of the different kinds of populations of agents was also measured by keeping track of *epistemic progress**, a modified version of the measure used by Weisberg and Muldoon. Epistemic progress* is defined as the proportion of visited patches among all significant (elevation >100 units) patches, and it reflects how exhaustively the patches on the two hills of significance have been visited at a particular time. For

⁹ However, as the value of λ becomes smaller, harvesting epistemic significance from a patch becomes slower, and movement on the landscape becomes relatively less costly. Consequently, the choice of a search heuristic becomes less critical. Real research topics with small lambdas would be ones where changing one’s approach is relatively quick compared to the time it takes to produce results by using a chosen approach.

Table 1 Epistemic work W in diverse and non-diverse populations on a smooth landscape. Progress measured at $t=400$. Italics indicate the most efficient population type.

Depletion rate	individual learners $\alpha \sim N(100, 1)$	social learners $\alpha \sim N(1, 1)$	diverse population $\alpha \sim U(1, 100)$
$\lambda = 0.001$	$7.14 * 10^{-2}$	$7.00 * 10^{-2}$	$7.25 * 10^{-2}$
$\lambda = 0.01$	0.405	0.401	0.411
$\lambda = 0.1$	0.934	0.897	0.939

smooth landscapes, epistemic progress* leads to very similar conclusions as those suggested by epistemic work: Social learning does not improve epistemic progress*, and the differences between the three kinds of populations are rather insignificant (see supplementary materials, documents 1 and 2).¹⁰

The categories represented in the columns of the table feature also in the subsequent simulations discussed below. These particular categories were chosen for presenting the results for two reasons. First, they capture the range of relevant variation in the social-learning thresholds. Experimentation with the model shows that most of the variation in the social-learning thresholds that influences epistemic work occurs when values of α fall between 1 and 100.¹¹

Secondly, the categories make it easy to draw qualitative conclusions from the model. A social-learning threshold of 1 represents an agent who only needs to expect a gain of one unit of epistemic significance per time period from social learning in order to engage in it. In contrast, an alpha of 100 represents a strongly individualist learning profile, where an agent on the zero-significance plane only decides to follow a peer on the global peak (at 1000 units) when the distance between them is less than 10 units. Henceforth, I refer to these kinds of agents as *social learners* and *individual learners*, respectively. Finally, cognitively diverse populations are modeled either by increasing the variance of the distribution from which alphas are drawn or by constructing mixed populations by merging sub-populations of social and individual learners.

These results seem to conflict with those obtained by Weisberg and Muldoon. According to my model, on the smooth landscape examined also in the EL model, individual learning ($\sim$ control agent rule) is more efficient than the use of social information, and diversity in the population produces only very small effects. This is hardly surprising, however. Visual inspection of the simulation runs reveals that both individual and social learners find the two peaks on the landscape, albeit in different ways. Whereas a population of individual learners gets evenly distributed on both hills, groups of agents with low social-learning thresholds harvest the epistemic significance mass in a sequential manner: As one of the agents happens to find the higher peak, the whole

¹⁰ Thanks to the anonymous referee for insisting on the use of alternative measures of epistemic success.

¹¹ Due to the nature of the model (see the discussion in Section 5), the qualitative results from the model should not depend on the choice of particular point values for the parameters. The scaling of the parameter space was chosen mainly for convenience and for its continuity with the EL model.

group flocks to that hill. Only after the whole hill is depleted does the population start to search for another. Once the other hill is found, the agents repeat the same flocking procedure.¹² What makes individual learners slightly more effective is the robustness of their strategy: Both hills of epistemic significance are investigated in parallel, whereas social learners often end up spending a long time searching for the second peak. As is seen below, this difference in collective search strategy turns out to be quite important on more complex landscapes.

4.2 Rugged landscapes

As was argued in Section 2, the collective problem-solving tasks in real scientific research can hardly be represented by smooth landscapes. It also turns out that the most interesting dynamics of the broadcasting model occur on more complex landscapes.

To study research topics with dead-end research approaches, noise consisting of small amplitude two-dimensional Gaussian bumps (Gaussian kernels) was added onto the landscape (see Fig 1). Figure 4 shows that as the amount of ruggedness increases (the number of bumps denoted by β), the search problem obviously becomes harder for all kinds of agents. Notice, however, that this is where the power of social learning starts to show: Because the individual learners easily get trapped on local maxima, their performance drops more than that of agents with lower social-learning thresholds.¹³

Table 2 reveals that the difference in performance between social and individual learners is sensitive to the time-scale parameter, the advantage of social learners being larger for small values of λ . At large values of λ ($\lambda \geq 0.1$), the effect is again reversed and a population of individual learners can be more effective than social learners.¹⁴ In the following, I focus on the case where λ is 0.01, because it nicely captures the conditions in a research field where coordination and division of cognitive labor are needed: When $\lambda = 0.01$, it takes one agent roughly 500 time steps to deplete a patch to a 1/100 of its original height, i.e. to produce nearly all of the significant results available by using that approach. With a population size of 50, this leads to a situation where during a 1000-round simulation, the population can deplete both hills of epistemic significance only with coordinated effort. Large values of λ represent less interesting cases, where the problem faced by the community is easy enough so that only a handful of agents who happen

¹² Animations of individual simulation runs with different kinds of populations can be found in the online supplementary materials.

¹³ As the ruggedness of the landscape increases even more, local search heuristics generally become less and less useful, and even the diverse communities make little progress. Such landscapes can be seen to represent research problems beyond the cognitive capacities of the scientist agents, where attaining significant results becomes increasingly a matter of luck.

¹⁴ As document 1 in the supplementary materials shows, these results hold remarkably well across β values ranging from 50 to 300, and across different times of measurement.

Fig. 4 Example of epistemic work on rugged landscapes. ($t = 400, \lambda = 0.01$). See supplementary materials document 3 for a more comprehensive presentation of the results.

Table 2 Epistemic work on rugged landscapes ($\beta = 200$)

Depletion rate	individual learners $\alpha \sim N(100, 1)$	social learners $\alpha \sim N(1, 1)$	diverse population $\alpha \sim U(1, 100)$
$\lambda = 0.001$	$9.49 * 10^{-3}$	$3.79 * 10^{-2}$	$2.10 * 10^{-2}$
$\lambda = 0.01$	$7.36 * 10^{-2}$	0.146	0.128
$\lambda = 0.1$	0.411	0.340	0.445

to find the hills of epistemic significance can exhaust them on their own, and there is no need for successful community-wide learning and coordination.

Hence, the first result of interest on rugged landscapes is the *necessity of social learning to overcome getting stuck on local maxima*. The mechanism underlying the success of social learners is the following. In a population of social learners, a large group of scientists can take advantage of one of their peers landing on a successful research approach. Commitment to adopting the approach of a peer provides scientists a rationale to accept temporary losses in the epistemic significance level due to ruggedness of landscape, when promised longer-term large rewards.

However, although a population consisting of social learners is significantly more efficient than a population of individual learners, it suffers from a problem of its own – herding.¹⁵ When all agents are sensitive to social information, the whole population often ends up on one of the two hills. As already described above, once the hill is depleted, members of the population have no

¹⁵ By herding, I refer to undesirable behavior where agents do what others do even in situations where they should be relying on their own information (Banerjee 1992).

information about the location (or even existence) of the other hill on the landscape. Because ruggedness makes rediscovery difficult, often such populations fail to discover the other hill during 1000 simulation rounds.

This suggests that perhaps adding a few individual agents into an otherwise social learner population might improve its exploratory capacities. As Figure 5 suggests, this is indeed what happens. In a population of 50, replacing 10 agents, who have low social learning thresholds, with agents strongly preferring individual learning significantly increases the population’s efficiency. Especially at later stages of search (after 500 rounds), a mixed population performs clearly better. This is because, with high probability, in the mixed population at least one of the individualist agents remains on the hill ignored by social learners. Once the first hill is sufficiently depleted (around 500 rounds), the individualist agent(s) can guide the majority to the new source of significance.¹⁶

Adding 10-20% individual learners both increases the epistemic work done by the community and improves the reliability of the population in finding both hills on the landscape. However, having more than 20% of individual learners in the population is counterproductive. As more individual learners are added, more agents get stuck on local maxima, and the collective efficiency decreases.¹⁷ Hence, at least *when faced with moderately challenging research topics (in terms of β and λ), a mixed population consisting of a few individual learners together with a large majority of social learners achieves the best epistemic outcomes.*

4.3 Summary

The three most important above results can be summarized as follows:

1. On smooth landscapes, no social learning or cognitive diversity is needed for efficient epistemic work. In fact, populations of agents following social-learning heuristics suffer slightly from herding, whereas individual learning leads to effective exploration and division of labor between agents.
2. On rugged landscapes, individual learners easily get stuck on local maxima, and populations of social learners achieve significantly better outcomes.
3. The power of cognitive diversity shows on landscapes corresponding to moderately challenging research topics. Mixed populations consisting mostly

¹⁶ Measuring epistemic progress* (see supplementary material, documents 3 and 4) suggests a further advantage of diversity. Unlike populations of social learners, given enough time, diverse populations and populations of individual learners do not leave behind unexamined significant patches on the landscape. This is due to their local search strategy: Agents conducting individual learning only leave a neighborhood of patches once all its patches have been depleted. Hence, in this way, diversity can make the population of scientists more pedantic in its work.

¹⁷ As full results reported in document 4 in the supplementary materials confirm, the ordering in Figure 5 is stable across the whole range of examined β values. However, for the higher ruggedness values ($\beta \in [200, 300]$), adding some extra individual learners can lead to a small additional payoff in last stages of the 1000-round simulation run.

Fig. 5 Homogeneous and diverse populations on rugged landscapes ($\beta = 100$, $\lambda = 0.01$)

of social learners with a minority (in the current simulations, 10-20% percent) of individualistic agents who tend to ignore social information are more effective than a homogeneous population of social learners.

Together these results suggest the more general observation that no learning strategy is *per se* more rational than others, but instead the efficiency of individual learning, social learning, and the usefulness of cognitive diversity all depend on the task faced by the community. In order to achieve good outcomes, a correct mix of agents employing different learning rules must be applied in the right context.¹⁸

Compared to the EL model, these results provide a more fine-grained, and as I have argued, more realistic, view of the conditions and processes related to cognitive diversity. Looking at individual simulation runs confirms that a mixed population performs best on rugged landscapes because it combines effective exploitation of socially learned information with exploration of new patches conducted by more individualistic agents: The success of a search conducted by the social learners is sensitive to their initial distribution on the landscape, and adding some individual learners in the population increases the probability that both peaks are found. Therefore, although social learners can be fast, individual learners improve the reliability of the search and reduce variance in the amount of epistemic work conducted over several trials.

¹⁸ The interaction effects between task difficulty and the distribution of social learning thresholds shows in Table 2, where no column dominates the others. For example, at different values of λ , pure social learning can be the best, second, or even the worst learning heuristic among the three.

5 Discussion

Before concluding, I address some concerns regarding the reliability of these results, their interpretation, and their implications for understanding real science and research policy.

First, the generality and relevance of the findings from the broadcasting model might raise concerns. Compared to analytical models, agent-based models usually come with more parameters, and the results often only hold in limited parts of the parameter space. However, as Marchi and Page (2014) point out, this should not be seen as a general shortcoming of agent-based modeling. Instead, the expansion of the parameter space is a result of having to actually run the simulation on a computer. In computational modeling even seemingly trivial modeling choices (regarding, e.g., timing, learning, interaction) must be made explicit. In analytical modeling it is easier for similar assumptions to go unnoticed, but this does not reduce the modeler’s responsibility to justify such choices.

One possible strategy for meeting the challenge of relevance for real-world science would be to calibrate the model with empirical data. However, data about the difficulty of scientific problems, learning strategies employed by scientists, and dependencies between aspects of research approaches is not readily available – and often it is not even clear how relevant evidence should be obtained. So while empirical calibration is a laudable aim, it remains outside the scope of my current endeavor. The model presented in this paper serves a different purpose. Rather than being a high-fidelity model of a particular target system in the world, it could be called a *how-possibly* model or an intuition engine (Marchi and Page 2014) aimed to reveal qualitatively described dependencies between the components of the model – the broadcasting model is designed to illuminate the micro-mechanisms mediating cognitive diversity, coordination, and problem-solving efficiency.

One should always be cautious about drawing conclusions about real science based on results from a simple theoretical model. I share Alexander and his coauthors’ worry that, since the true nature of epistemic landscapes in real science is beyond our knowledge, an epistemic landscape model cannot directly be used to argue for the desirability of cognitive diversity in a particular scientific field. I believe, however, that epistemic landscape modeling can legitimately serve a more modest role: Agent-based models which have not been calibrated with empirical data can be conceived as computational thought experiments, where modeling assumptions are seen as premises and results as conclusions from extended arguments (Beisbart 2012). Hence, the modeling results are of a conditional nature. They answer questions on what would happen to epistemic efficiency (and why), given certain hypothetical ranges of parameter values, learning heuristics, and the rest of the model structure. From this perspective, agent-based models in philosophy can be seen as argumentative devices, and the added value of agent-based modeling resides in their usefulness in deriving simple conclusions from a massive number of

premises (regarding, for example, distributions of the cognitive properties of agents).

Furthermore, Ylikoski and Aydinonat (2014) have recently argued that the epistemic contribution of abstract theoretical models can often only be understood in the context of a cluster of models of the same phenomenon. This suggests that we should not study the explanatory contributions of different epistemic landscape models in isolation, but instead see these computational thought experiments as variations on a theme, as a family of models with similar explananda and explanantia. Exploration of the parameter space of a single model provides information about the robustness of modeling results within the scope of modeling assumptions embodied in that particular model. In like manner, cross-model comparisons allow more general robustness assessments in light of the variation of modeling assumptions within the model cluster. Therefore, general results to be drawn from the model cluster should typically be ones which can be shown to be immune to changes in the auxiliary assumptions.

Insofar as the epistemic landscape models put forward by Weisberg and Muldoon, Thoma, Alexander and his coauthors, and myself all share the same basic model structure, they definitely form a cluster within which comparisons across models can increase our understanding of the target phenomenon. The main differences between these models concern changes in the behavioral rules of the agents and the way epistemic performance is measured: Alexander and his coauthors' examination of the swarm rule is basically a replication of the original EL model with the addition of a new homogeneous population engaging in flocking behavior. Similarly, Thoma slightly changes the way collective epistemic performance is measured, suggests new implementations of the maverick and follower rules, and examines the usefulness of diversity when the assumption of strict locality of search and movement is relaxed.

Like Thoma's model, also the broadcasting model extends social learning beyond the immediate Moore neighborhood. However, unlike Thoma's model, it does not allow non-local movement (jumping), and it derives the choice between individual and social learning from rudimentary cost-benefit analysis done by the scientist agents. Together with the implemented agent-environment interaction (landscape depletion) and examination of rugged landscapes, these changes make the broadcasting model perhaps the clearest departure from the original EL model. However, as was pointed out above (see footnote 6), similar modifications were envisioned already by Weisberg and Muldoon as worthwhile extensions of their work.

As this article's starting point and its target of critical appraisal has been the original EL model, more detailed comparisons between the broadcasting model and those put forward by Thoma and Alexander et al. must be left for future work. It should, however, be noted that results from the broadcasting model differ in interesting ways from findings in the other models. In contrast to Thoma's findings, my model suggests that on smooth landscapes there is no noticeable benefit from diversity. However, when compared to Alexander and his coauthors' results, the broadcasting model paints a more positive picture of

the usefulness of social learning and diversity in more complex problem-solving situations. By showing how the usefulness of social learning and cognitive diversity are related to the difficulty of the problem, the broadcasting model goes some way towards explaining the *prima facie* contradictory results obtained in the earlier models.

The cautious conclusion to be drawn from these differences is that, in its entirety, the relationship between diversity and epistemic performance is likely to be more complex than can be captured by any simple model. Nonetheless, the model cluster approach to the epistemology of theoretical modeling suggests that instead of trying to determine which model provides the correct picture of cognitive diversity, each model could be seen as a candidate for illuminating some possible scenarios and processes related to cognitive diversity. By tracing differences in outcomes to differences in modeling assumptions, the different models can together be seen to lead to a clearer picture of the potential and correct interpretation of epistemic landscape modeling – and more generally, to a better understanding of the possible mechanisms through which cognitive diversity influences the conduct of scientific research.

6 Conclusion

Diversity clearly makes a difference in the research community. The broadcasting model suggests that the picture of the micro-mechanisms mediating cognitive diversity and epistemic efficiency is roughly the following: Cognitive diversity produces efficient division of labor between scientists by maintaining a beneficial mix of exploration and exploitation in the population. In simple research domains where areas of high epistemic significance are easily identifiable, no social learning between agents is needed for good coordination. However, on rugged landscapes which better capture the complexity of scientific problem-solving, the presence of social learners is crucial for enabling the population to converge to research approaches of high epistemic significance. However, high occurrence of social learning increases the unreliability of the research process due to herding behavior.

My results suggest that for moderately demanding research topics, the best mix of problem-solving speed and reliability is reached by a cognitively diverse population of scientists, where most of the agents are eager to engage in social learning, counterbalanced by a minority of individualistic researchers who mostly conduct explorative research on their own. In such a population, most of the epistemic work is done by the majority of agents highly sensitive to social information about the currently most efficient research approaches. However, maverick scientists are indispensable for providing an alternative to the consensus, ensuring that the community does not lose sight of valuable research approaches currently ignored by the conformist majority.

Acknowledgements I would like to thank the anonymous referees for their useful comments and suggestions on earlier drafts of the paper. I am also thankful to Manuela Fer-

nández Pinto, Marion Godman, Jaakko Kuorikoski, Otto Lappi, Caterina Marchionni, Carlo Martini, and Petri Ylikoski for helpful discussions about the paper and about epistemic landscapes modeling in general. The paper also benefited from comments by the participants at the Agent-based Models in Philosophy conference at LMU Munich, TINT brown bag seminar, and the cognitive science research seminar at University of Helsinki. This research has been financially supported by the Academy of Finland and the University of Helsinki.

References

- Alexander JM, Himmelreich J, Thompson C (2015) Epistemic Landscapes, Optimal Search, and the Division of Cognitive Labor. *Philosophy of Science* 82(3):424–453
- Anderson E (1995) Knowledge, human interests, and objectivity in feminist epistemology. *Philosophical Topics* 23(2):27–58
- Banerjee AV (1992) A Simple Model of Herd Behavior. *The Quarterly Journal of Economics* 107(3):797–817, 05045
- Beisbart C (2012) How can computer simulations produce new knowledge? *European Journal for Philosophy of Science* 2(3):395–434
- Borenstein E, Feldman MW, Aoki K (2008) Evolution of Learning in Fluctuating Environments: When Selection Favors Both Social and Exploratory Individual Learning. *Evolution* 62(3):586–602
- D’Agostino F (2009) From the organization to the division of cognitive labor. *Politics, Philosophy & Economics* 8(1):101–129
- De Langhe R (2014) A unified model of the division of cognitive labor. *Philosophy of Science* 81(3):444–459
- Enquist M, Eriksson K, Ghirlanda S (2007) Critical Social Learning: A Solution to Rogers’s Paradox of Nonadaptive Culture. *American Anthropologist* 109(4):727–734
- Fricker M (2007) *Epistemic Injustice: Power and the ethics of knowing*. Oxford University Press, Oxford
- Goldman A (2011) A guide to social epistemology. In: Goldman A, Whitecomb D (eds) *Social Epistemology: Essential Readings*, Oxford University Press, Oxford, pp 11–37
- Grimm V, Berger U (2015) Robustness analysis: Deconstructing computational models for ecological theory and applications. *Ecological Modelling* p forthcoming
- Hull DL (1988) *Science as a process: an evolutionary account of the social and conceptual development of science*. Science and its conceptual foundations, University of Chicago Press, Chicago ;London
- Johnson S (2011) *Where Good Ideas Come From: The seven patterns of innovation*. Penguin, London
- Kauffman S, Levin S (1987) Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology* 128(1):11–45
- Kitcher P (1990) The Division of Cognitive Labor. *Journal of Philosophy* 87(1):5–22
- Kitcher P (1993) *The advancement of science: science without legend, objectivity without illusions*. Oxford University Press, New York ;Oxford
- Klarreich E (2013) Unheralded Mathematician Bridges the Prime Gap. *Quanta Magazine* URL <https://www.quantamagazine.org/20130519-unheralded-mathematician-bridges-the-prime-gap/>, 00005
- Lazer D, Friedman A (2007) The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly* 52(4):667–694
- Longino HE (1990) *Science as Social Knowledge: Values and objectivity in scientific inquiry*. Princeton University Press, Princeton, N.J.
- Longino HE (2002) *The Fate of Knowledge*. Princeton University Press, Princeton, N.J. ;Oxford
- MacArthur-Foundation (2014) MacArthur Fellows: Yitang Zhang. URL <https://www.macfound.org/fellows/927/>, 00000
- Marchi Sd, Page SE (2014) Agent-Based Models. *Annual Review of Political Science* 17(1):1–20

- Mayo-Wilson C, Zollman KJS, Danks D (2011) The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science* 78(4):653–677
- Muldoon R (2013) Diversity and the Division of Cognitive Labor. *Philosophy Compass* 8(2):117–125
- Novak M (2014) Tesla and the lone inventor myth. URL <http://www.bbc.com/future/story/20130322-tesla-and-the-lone-inventor-myth>
- Page SE (2008) *The difference: how the power of diversity creates better groups, firms, schools, and societies*, paperback edn. Princeton University Press, Princeton, N.J.; Woodstock
- Polanyi M (1962) The Republic of Science. *Minerva* 1(1):54–73
- Russell SJSJ, Norvig P (2003) *Artificial intelligence: a modern approach*, 2nd edn. Prentice Hall series in artificial intelligence, Prentice Hall/Pearson Education, Upper Saddle River, N.J.; London
- Solomon M (2006) Groupthink Versus The Wisdom of Crowds. *Southern Journal of Philosophy* 44(Supplement):28–42
- Strevens M (2003) The Role of the Priority Rule in Science. *Journal of Philosophy* 100(2):55–79
- Thoma J (2015) The Epistemic Division of Labor Revisited. *Philosophy of Science* 82(3):454–472
- Weisberg M (2010) New Approaches to the Division of Cognitive Labor. In: Magnus PD, Busch J (eds) *New Waves in Philosophy of Science*, Palgrave Macmillan
- Weisberg M, Muldoon R (2009) Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science* 76(2):225–252
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress on Genetics* pp 355–366
- Ylikoski P, Aydinonat NE (2014) Understanding with theoretical models. *Journal of Economic Methodology* 21(1):19–36
- Zollman KJS (2010) The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72(1):17–35