

Tolerance for failure unleashes the benefits of cognitive diversity in collective problem solving

Justin Sulik (justin.sulik@gmail.com)

Cognition, Values & Behavior, Ludwig Maximilian University of Munich, Munich, Germany

Bahador Bahrami (bbahrami@gmail.com)

Department of Psychology, Ludwig Maximilian University of Munich, Munich, Germany

Department of Psychology, Royal Holloway, University of London, London, UK

Ophelia Deroy (ophelia.deroy@lrz.uni-muenchen.de)

Faculty of Philosophy & Munich Center for Neuroscience, Ludwig Maximilian University of Munich, Munich, Germany

Institute of Philosophy, School of Advanced Study, University of London, London UK

Abstract

Collective problem solving is supposed to benefit from cognitive diversity (e.g., when a team consists of individuals with different learning strategies). However, recent evidence for this claim fails to rule out an alternative explanation: that the benefit is due to moderate non-conformity, not diversity. We extend a previous agent-based simulation to distinguish these hypotheses, and demonstrate that diverse learning strategies alone do not yield the expected benefit. We extend the model further, based on an idea from the philosophy of science: Group-level benefits in complex problem solving often entail individual-level failures. Accordingly, we parameterize tolerance for failure, and show that there is an interaction between tolerance for failure and diversity. When tolerance for failure is zero, heterogeneous and homogeneous groups perform equally; when non-zero, diverse groups can outperform heterogeneous groups. Our agent-based simulations help clarify when cognitive diversity benefits collective problem solving.

Keywords: communication; diversity; problem solving; social learning; explore/exploit; epistemic landscape; agent-based simulation

Collectively, humans can solve problems of astonishing complexity. This success would be impossible without our communicative abilities, which allow us to share information efficiently. However, communication is a double-edged sword: communication is necessary to spread the information crucial to solving a problem as a group, but in doing so, it reduces the diversity of information in the group (Lazer & Friedman, 2007). This can lead to herding (Salganik, Dodds, & Watts, 2006) or premature convergence on suboptimal outcomes (Zollman, 2010). It can also reduce people's tendency to engage in further exploration (Yahosseini, Reijula, Molleman, & Moussaïd, 2018) or to share relevant private information (Gigone & Hastie, 1993).

While a complete lack of communication is clearly a barrier to collective problem solving, these findings imply that unconstrained communication is also not ideal. But how to strike a balance? Agent-based simulations show how group problem solving can be optimized by throttling communication to some extent. One approach focuses on social factors, such as the network structure that the agents are embedded

in, limiting who can talk to who (Zollman, 2010). A second approach focuses on cognitive factors, such as the strategies used by the agents to tackle the problem, where some strategies involve less communication than others (Pöyhönen, 2017; Thoma, 2015; Weisberg & Muldoon, 2009).

This second approach is particularly intriguing, as it is used to argue that cognitive diversity helps groups to solve complex problems. Building on these agent-based simulations, we demonstrate that, as things stand, this approach does not in fact make a clear-cut case for cognitive diversity. We then show how the inclusion of an additional factor — tolerance for failure — fills this gap, thus enabling cognitive diversity to yield a benefit.

A division of cognitive labor in epistemic landscapes

We draw from a modeling tradition rooted in the philosophy of science, not only because science is a prime example of the phenomenon of interest here — complex problem solving by human collectives (Muthukrishna & Henrich, 2016) — but also because it involves a tension between the best outcome for an individual and the best outcome for a group (Kitcher, 1990). To illustrate: imagine that two researchers are leading teams that develop vaccines, and each must decide which of two approaches (depending on the field, these may be methods, theories, frameworks, or research questions) to use during the upcoming year. Let's say approach A has a good chance of success while approach B has a poor chance of success.

From the point of view of each individual researcher, the optimal strategy is to choose approach A. However, given that there is a chance that approach A will fail and that approach B will succeed, from the point of view of the collective (i.e., medical science), the optimal strategy is for one researcher to pick approach A and the other to pick approach B. Science thus benefits from a 'division of cognitive labor' where individuals pursue different approaches, even if this means some of them pursue approaches that are more likely to fail (Kitcher, 1990). Two core ingredients for this division of labor, then, are: the pursuit of diverse approaches, and accep-




Figure 1: An epistemic landscape, with two hills of significant epistemic value topped by peaks of outstanding epistemic value.

tance of possible failure.

A subsequent research tradition uses epistemic landscape models (Weisberg & Muldoon, 2009) to relate diverse approaches (features of the problem space) to cognitive differences (features of the individuals who explore that problem space). In these models, a population of agents explores an abstract landscape (Fig. 1). Each position on the landscape represents some approach in a scientific field, and the height the landscape at that point represents the epistemic value of that approach: how much useful knowledge can be gained by pursuing that theory or research question or method. In this example, there are two peaks, so two scientific theories of outstanding epistemic value. Generally, in such models, the agents explore the landscape with the aim of discovering the highest peaks.

Agents are imbued with various learning strategies, or ways of exploring the landscape. Two main types of learning strategy are common across models. An individual learning strategy involves an agent gathering information about the landscape directly: the agent moves on to new patches, and thereby learns the epistemic value of those patches. A social learning strategy involves the agents communicating with each other, thereby learning the value of the patches that the others are currently occupying. Thus, an agent can either learn where a peak is by visiting it, or by hearing from another agent who has discovered it.

A homogeneous population of mavericks (agents with ex-

treme individual strategies) is inefficient, as nobody benefits from sharing current knowledge about the location of potential hills, though such a group eventually succeeds in discovering both peaks. A homogeneous population of conformists (agents with extreme social learning strategies) typically does not succeed, as the group quickly discovers and converges on one hill, but commonly does not to discover the other. In a mixed population, consisting of some mavericks and some conformists, the mavericks explore new patches and reduce overall herding, preventing convergence on a single peak, while the conformists exploit the regions of known potential value. Thus, such models typically conclude that cognitive diversity leads to better collective problem solving (Pöyhönen, 2017; Thoma, 2015).

However, there are two reasons for caution regarding this conclusion, stemming from the fact that these models employ categorically different learning strategies (maverick vs. conformist, or some terminological equivalent). A practical concern is that human learning is not so categorical: there is stable variation in people's tendency to incorporate social information into their beliefs (Molleman, Kurvers, & van den Bos, 2019). This limits the relevance of such findings for human problem solving. A theoretical concern is that this qualitative distinction in learning strategies means that such models cannot distinguish between two possible hypotheses. For the sake of illustration, say that conformists do 0% individual learning while mavericks do 100% individual learning, and that it turns out — for a population of 10 agents on a given landscape — the optimal mix is 2 mavericks and 8 conformists. Is this due to the diversity of the group? Or could it just be due to the fact that the average strategy in the group is 20% individual learning? That is, perhaps a homogeneous population where everyone has the same 20% individual learning strategy would do just as well as a mixed population of 2 mavericks (all of whom are 100% individual learners) and 8 conformists (all 0% individual learners). In that case the model would support moderate non-conformity rather than cognitive diversity.

It is important to explore whether these two different compositions of collectives, that differ in diversity but not the mean of their learning strategy, could produce distinctive collective behaviors. To address this, we adapt an epistemic landscape model to include intermediate strategies, allowing us to test whether group performance is driven by the average level of social learning, or by its variance (i.e., cognitive diversity). Then, we further adapt the model to include a key element of a division of cognitive labor mentioned above: it involves variable risk of failure, so it is plausible that any benefits of diversity that derive from a division of cognitive labor should be sensitive to that.

The broadcasting model

We take the 'broadcasting model' by Pöyhönen (2017) as our starting point, as it offers two useful features for studying cognitive diversity in collective problem solving.

The first important innovation in Pöyhönen's model is how

epistemic success is evaluated. Rather than just exploring the landscape, agents also do some epistemic work to extract value. Concretely, this involves excavating the landscape: at every time point, each agent depletes the value of the patch that they are currently standing on by a small amount. This means that only so much value can be extracted from a patch before it becomes useless, analogous to how, in the real world, a given theory cannot be exploited indefinitely for epistemic gain. What this means, in terms of collective problem solving, is that it is not enough for one agent to discover one peak and for another agent to discover the other peak, because unless other agents come to collaborate with them in mining those scientific positions, it will take a long time for these two agents to fully excavate the hills of epistemic value. The relevance for our interest in communication and problem solving is that this makes collaboration crucial for group success. Otherwise, the claim that it is useful to throttle communication would be trivial.

The second important innovation by Pöyhönen is the eponymous ‘broadcasting’ implementation of learning strategies. An individual’s strategy is defined by their social-learning threshold: how much of an improvement in epistemic value they could accept if they moved to where another agent is (by analogy, adopting that other agent’s approach or theory), relative to the distance they would have to travel to get there. As this threshold amounts to the ratio between height climbed and distance travelled, it can be thought of as a slope m (Fig. 2). For a given threshold, an agent will consider learning socially from other agents that lie above that slope. In the context of a 3D landscape, all places above a given slope can be thought of as lying within a cone, centered on the agent in question. In that case, a maverick strategy involves a steep slope or narrow cone, whereas a conformist strategy involves a flat slope or wide cone.

In this model, a simulation run of t time steps begins with population of N agents with social learning thresholds $m_1 \dots m_N$. The agents are randomly placed on low-value positions in the landscape (i.e., not on the hills).

At the start of the run (t_0), every agent i learns the epistemic value of its starting position $e_{i,0}$. It then picks a random heading, and starts moving in that direction at a constant velocity. At each subsequent time point (t_j), it learns the epistemic value of its new position $e_{i,j}$ and compares this to the epistemic value of its previous position $e_{i,j-1}$. The difference in value at that time $\Delta e_{i,j} = e_{i,j} - e_{i,j-1}$ indicates whether the agent is going downhill ($\Delta e_{i,j} < 0$) or not ($\Delta e_{i,j} \geq 0$).

If an agent is not going downhill (i.e., is improving or maintaining its epistemic position), it simply continues along its original heading. However, if it is going downhill, it abandons its heading, and decides between an individual or social learning strategy on the basis of its social threshold m_i . Every agent polls all other agents who are above its social learning threshold. That is, each agent can only learn from those others who are close enough or high enough that they lie within its cone (Fig. 2).




Figure 2: The broadcast implementation of social learning (Pöyhönen, 2017). The focal agent has a social learning threshold m_a , and it is able to learn from other agents above that threshold. In this illustration, the epistemic value gained by traveling to target position B has slope m_t , and as $m_t > m_a$, the focal agent considers it worthwhile to travel to position B. Even though position A is within the focal agent’s cone, there is nobody there: the cone does not represent a view of the landscape, but rather a view of other agents. Even though position C is closer than position B, it falls below the focal agent’s threshold.

If there are no such others, the agent does individual learning: it explores its immediate vicinity and sets a new heading in the direction of the highest adjacent patch (or a random heading if all surrounding patches are lower than it). It thus engages in local hill-climbing. But if there are other agents above its social threshold, it does social learning: it identifies which other agent represents the greatest gain for the least travel, and sets its heading for that other agent. Regardless of whether the new heading derives from individual learning or social learning, it continues in that direction, checking its Δe at each time step, unless it goes downhill again, at which point the process repeats. Agents’ thresholds do not change, so a maverick agent is simply an agent with a very high threshold: it will be unlikely to ever encounter other agents it considers worth learning from. A conformist agent has a low threshold, and will consider any agent a higher point at any position on the landscape to be worth learning from.

As the final action of each time step, all agents do epistemic work, excavating some value w from their position, depleting the landscape by the equivalent amount. At the end of t time steps, the epistemic success of the group is the sum of all excavated values.

We stress that no agent has a god’s eye view of the landscape (which is why they have to explore), and they do not build a representation of the landscape while they explore. Each agent’s memory only retains the height of its current position and the height of its position at the immediately preceding time step. The only information it can access is the height and distance of other agents within its cone, and it can only access this information at a moment when it is going downhill

and deciding which learning strategy to employ next.

The DE-land model

Agents The first main departure in our Diversity in Epistemic Landscapes (DE-land) model is how the social learning thresholds are set. For each run, an aggregate threshold m is selected from the range of possible thresholds (see below for the parameter space), and the run is randomly assigned to one of two conditions (‘homogeneous’ vs. ‘heterogeneous’). In the homogeneous condition, N agents are all initialized with threshold m . In the heterogeneous condition, $m_1 \dots m_N$ thresholds are randomly sampled from a Beta distribution (α, β) that has mean $m = \frac{\alpha}{\alpha + \beta}$. Thus, both conditions have the same average threshold.

We note that samples from Beta distributions are in the range $(0, 1)$, so even our most maverick agent will choose to learn from another agent whose position offers a climb of 1 for a travel distance of 1. Thus, our mavericks are not as extreme individualists as in previous models, where mavericks are typically unwilling to learn from anybody, no matter how large of improvement in epistemic value they offer. However, such extreme maverick attitudes are unlikely to reflect the target phenomenon (doing science, yet being unwilling to be influenced by anybody else, no matter how valuable their approach or theory). Further, our results show a ceiling effect around this maximum. Finally, previous work has shown that the optimal group composition is mostly conformists with a minority of mavericks (Pöyhönen, 2017), so the most interesting region of the parameter space is lower, rather than higher social learning thresholds.

A further departure is that we parameterize the agents’ tolerance for failure. Previously, an agent abandoned its heading and decided on a learning strategy every time it went downhill ($\Delta e < 0$), which amounts to zero tolerance for failure. Now, this decision kicks in when $\Delta e < f$, with f representing the tolerance for failure. Thus, with $f = 0.2$ an agent will tolerate a decrease in epistemic value of up to 0.2 in one time step, continuing its original heading rather than making a decision about which learning strategy to employ next.

Landscape The landscape is toroidal (wrapping around on both sides), and has 2 peaks of equal height, situated maximally far apart and generated from multivariate normal distributions. To this is added a layer of Perlin noise, with variable amplitude a : a higher amplitude means the base landscape is noisier, with local hills and troughs (though far smaller than the main 2 peaks). This allows us to parameterize problem difficulty, as when a is low, the landscape is smooth, so social information about positive epistemic value reliably indicates the location of the 2 hills. With some noise, however, one agent may end up following another agent who is higher than it, but is not near the hills. At the beginning of the run, the total amount of positive epistemic value on the landscape is calculated as its total epistemic mass. A group’s problem solving success is the proportion of this total mass excavated during the run.

Simulation 1: Does cognitive diversity yield a group-level benefit?

A division of labor cognitive based on learning strategies is one way to improve collective problem solving outcomes. However, we raised the concern that this may be due to the average level of conformity, rather than cognitive diversity, in previous epistemic landscape models. For our first simulation, we compare homogeneous and diverse groups when, as in previous models, there is no tolerance for failure ($f = 0$). We ran 3000 simulations each of 400 time steps with 40 agents on a 40×40 landscape, and all possible combinations of the following parameters.

- Perlin noise amplitude a in: [1, 6, 10]
- Social learning threshold m from Beta(α, β) in: [(1, 9), (3, 7), (5, 5), (7, 3), (9, 1)]
- Condition in: [‘homogeneous’, ‘heterogeneous’]

Thus, for instance, in a heterogeneous run where the distribution parameters are (1,9), agents’ social learning thresholds were sampled randomly from this distribution, whereas in a homogeneous run, agents’ thresholds were set to the mean of this distribution ($m = \frac{1}{1+9} = 0.1$).

Fig. 3a shows the main effect of social learning threshold. Having a higher social learning threshold produces better group outcomes, but this benefit plateaus in the middle-to-upper end of our range (Fig. 3a). A balanced strategy ($m = 0.5$) does quite well overall, and a moderately non-conformist strategy ($0.5 \leq m \leq 0.9$) only slightly better. Although all strategies start off similarly, the difference emerges well into the run. All populations discover at least one hill early on. But then conformist populations, having converged on that hill, fail to explore much more and are unlikely to discover the second hill (consistent with findings outlined in the Introduction). In contrast, balanced and moderately non-conformist populations are quite likely to discover the second hill, which is how they discover substantially more than 50% of the landscape’s epistemic mass.

Fig. 3b shows the main effect of noise. Performance for moderate noise ($a = 6$) was better than for low noise ($a = 1$). This occurs because, in low-noise landscapes, as soon as any agent starts climbing one of the hills, it represents an unambiguous signal as to where a peak is. The other agents converge, abandoning further exploration. In contrast, in noisier landscapes, even if one agent is starting to climb the hill, some agents may not converge on it because they are distracted by other agents who are at higher positions, even if those positions are only higher due to random noise. This is enough to prevent premature convergence, and ultimately, both peaks are discovered.

Fig. 3c shows the comparison between heterogeneous and homogeneous populations, for varying levels of noise and social learning threshold. There is a clear interaction between noise and social learning threshold: for low levels of noise,

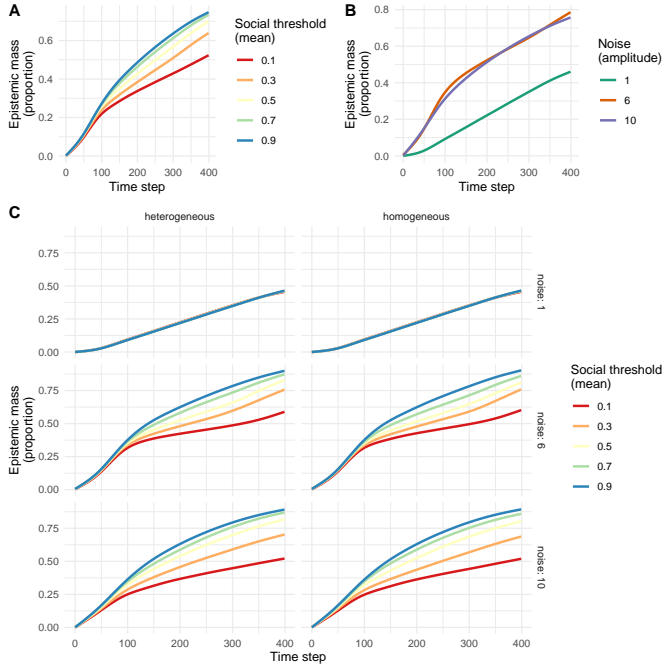


Figure 3: Results of simulation 1, showing aggregates across simulation runs with Loess splines. (A) The main effect of mean social learning threshold. A higher threshold implies more maverick behavior; (B) The main effect of noise (C); The interactions between threshold, noise and condition.

the threshold does not matter (the lines in the top row overlap). No matter whether an agent is more maverick or more conformist, if they are on a fairly flat plain, and some other agent starts climbing a hill, everyone converges on that point. The previously described effect of social learning threshold emerges for higher levels of noise. Thus, moderate non-conformity becomes valuable in noisier environments.

However, the most notable feature of Fig. 3c is the lack of any substantial difference between homogeneous and heterogeneous populations. Thus, consistent with the concern we raised earlier, even if epistemic landscape models claim to show a benefit of cognitive diversity, in fact they really show a benefit of a population that is, on average, moderately non-conformist. Moderate non-conformity can occur when the entire population is moderately non-conformist, or when a population has diverse strategies, but their average threshold is moderately non-conformist.

Simulation 2: Does tolerance for failure contribute to a benefit of diversity?

Variable chance of failure was one of the two elements of the division of cognitive labor outlined above (Kitcher, 1990), and the degree of failure that agents are willing to accept is one factor that has not been implemented in previous epistemic landscape models. Converging evidence from social epistemology (Thoma, 2015) and cultural evolution (Boyd &

Richerson, 1985) shows that it is costly to explore individually, without the benefit of social learning. While some mavericks may encounter epistemic rewards, many do not, and consequently do worse than the average conformist (Thoma, 2015). Nonetheless, their presence helps the group, as when those rare mavericks discover a hill, this allows conformists to follow them.

As occasional failure is an unavoidable part of the maverick strategy, perhaps tolerance for failure is what is missing, for a mixed population of mavericks and conformists to do well. To test this, we ran 9000 simulations as described in Simulation 1, additionally setting tolerance for failure to f in $[0, 0.2, 0.4]$.

To facilitate presentation of results with the inclusion of this additional variable, we focus on the epistemic mass excavated at the end of the run (at $t = 400$). Fig. 4 shows this final standing. Here, the points where tolerance for failure is zero simply recapitulate the end-points of the Loess splines at $t = 400$ in Fig. 3c), where there is no substantial difference between heterogeneous and homogeneous groups. What is striking, though, is that values of $f > 0$ make a difference for heterogeneous populations, but not for homogeneous ones.

For low levels of noise ($a = 1$), increasing tolerance for failure decreases the overall performance of heterogeneous groups (though there is still no effect of social learning threshold for low-noise landscapes). For higher levels of noise, the outcome is more complex. The aforementioned decrease in performance remains (if weakly) for more maverick diverse groups, but performance improves with increasing tolerance for more conformist groups. In particular, compare the red lines (for $m = 0.1$) between heterogeneous and homogeneous groups. Whereas these more conformist populations do consistently poorly when they are homogeneous, tolerance drives a marked improvement when they are heterogeneous. Thus, tolerance for failure interacts with learning strategy to produce an effect of diversity, though the direction and strength of this effect depends on landscape noisiness. When social learning thresholds make a difference at all (i.e., when noise $a > 1$), there is a major boost to performance for more conformist populations who are more tolerant to failure, but a minor dip in performance for more maverick populations.

Discussion

Under the rubric of a division of cognitive labor, epistemic landscape models are used to argue that cognitive diversity benefits collective problem solving. However, we have shown that this conclusion is premature, as the same findings can be produced by homogeneous populations that are moderately non-conformist. However, individual failure is another factor entailed by the division of cognitive labor, and mavericks are on average more likely to fail. By parameterizing tolerance for failure (which was effectively at 0 in previous models), we have shown that epistemic landscape models do in fact support an effect of diversity, though this interacts with landscape noisiness and average social learning threshold.

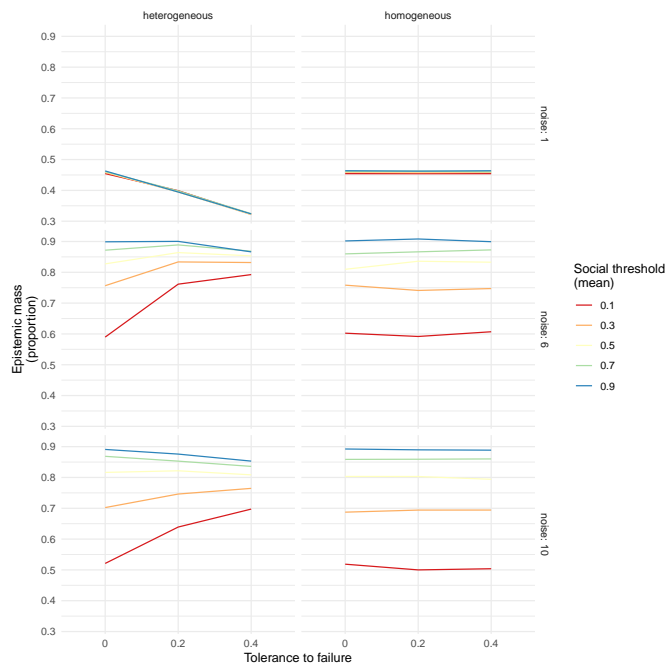


Figure 4: Results of Simulation 2, showing the final epistemic value excavated at $t = 400$, for varying levels of landscape noise (in 3 rows), mean social learning threshold (in color), and tolerance for failure (along the x -axis). The lines for social threshold overlap substantially for noise $a = 1$.

We also replicate previous findings, from entirely different paradigms (e.g., models with payoffs derived from two-arm bandits and with agents embedded in social networks, Zollman, 2010), that some disruption to communication can in fact benefit group epistemic outcomes. Here, this is accomplished by adding noise to the landscape, so that signals of positive epistemic value are unreliable guides to the state of the world.

Why is useful to know that the benefits of diversity are clearest for more conformist groups (e.g., the red and orange lines in Fig. 4)? These results are consistent with previous findings that a minority of mavericks in a general population of conformists yields the best group outcomes (Pöyhönen, 2017). More importantly, though, humans are a social species and our greatest accomplishments depend heavily on cumulative culture (Muthukrishna & Henrich, 2016), and many of our decisions weight conformity heavily (Cialdini & Goldstein, 2004). If there is to be any benefit of diversity in learning strategies for humans, this is the area of the parameter space where that we should expect one.

Does tolerance for failure simply mean that agents have fewer chances to engage in social learning, as this only occurs once they fail to improve, and as we vary what counts as failure? If so, the results would be less interesting, as they would simply recapitulate what Simulation 1 showed about the benefits of moderate non-conformity. However, if this worry were true, we would expect the same to hold

for homogeneous groups, and Fig. 4 shows that it does not. This is not just individual learning by another name. Further, the effect of tolerance is sometimes negative (as in low-noise situations), so this is also not simply a matter of over-parameterizing the model. Rather, it is specifically a matter of leveraging diversity in cases where social information is a flawed guide to the state of the world.

In future work, a necessarily step will be a deeper and more detailed exploration of the parameter space, including depletion rates or the size of the landscape (affecting the sparsity of hills). This includes alternative distribution families for social learning: Here we have used the Beta distribution, but other continuous non-negative distributions such as the Gamma would allow for more extreme maverick behavior. In addition to tweaking or redefining parameters in this way, another goal will be to allow social learning thresholds and tolerance for failure to covary: Here, tolerance for failure is uniform across the population, but it will be interesting to learn if this benefits everyone, or if it has better effects for specific types of learning strategies (Thoma, 2015). Moreover, the model can be extended to include parameters (such as flexibility or range of movement, Thoma, 2015) that contribute to group performance. Overall, a substantial challenge will be to integrate two quite different traditions — cultural evolution and the philosophy of science — that study individual vs. social learning and individual-level vs. group-level benefits.

Moving beyond these models, it will be useful to understand the analogues of these parameters in real-world problem solving. Whereas our range of noise parameters was chosen because it yielded interesting variation (e.g., no effect of learning threshold for noise = 1 in Fig. 3c, with this effect maxing out by noise = 10), it would be useful to know what concrete measures of noise or complexity in real world problem solving show the strongest benefit of diversity (Sulik, Bahrami, & Deroy, in press). Similarly, whereas our parameter for tolerance of failure is an abstract decision threshold, failure in the real world may be a matter both of individual psychology (e.g., how people feel when they fail, or what motivates them to persevere) and social context (e.g., how incentives are set up in societies or institutions in such a way that encourages an optimal amount of risky exploration, Mann & Helbing, 2017). Further, problem solving here focuses on discovery (e.g., discovering the peaks of high epistemic value), whereas doing science in the real world requires a balance between discovering new things and checking that they can replicate, and as this latter issue is another important dimension in the study of diversity (Devezer, Nardin, Baumgaertner, & Buzbas, 2019), it would be useful to see whether these dimensions (diversity of learning strategies and diversity of roles within science) are orthogonal or related.

In any case, do we really need more evidence for the benefits of cognitive diversity? There is already evidence pointing to such benefits (e.g., Hofstra et al., 2020). However, the pressing issue is that there is also evidence pointing in

the opposite direction (e.g., Almaatouq, Yin, & Watts, 2020). Thus, the question is not ‘Does cognitive diversity benefit group epistemic outcomes?’ but rather ‘When does cognitive diversity benefit group epistemic outcomes?’ (Sulik et al., in press). Identifying (and then expanding on) the limits of computational simulations is one way to do just that.

References

- Almaatouq, A., Yin, M., & Watts, D. J. (2020). Collective problem-solving of groups across tasks of varying complexity. *PsyArXiv*, January 29. doi: 10.31234/osf.io/ra9qy
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
- Devezer, B., Nardin, L. G., Baumgaertner, B., & Buzbas, E. O. (2019). Scientific discovery in a model-centric framework: Reproducibility, innovation, and epistemic diversity. *PLoS One*, 14(5), e0216125.
- Gigone, D., & Hastie, R. (1993). The common knowledge effect: Information sharing and group judgment. *Journal of Personality and Social Psychology*, 65(5), 959–974.
- Hofstra, B., Kulkarni, V. V., Galvez, S. M.-N., He, B., Jurafsky, D., & McFarland, D. A. (2020). The diversity–innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17), 9284–9291.
- Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, 87(1), 5–22.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
- Mann, R. P., & Helbing, D. (2017). Optimal incentives for collective intelligence. *Proceedings of the National Academy of Sciences*, 114(20), 5077–5082.
- Molleman, L., Kurvers, R., & van den Bos, W. (2019). Unleashing the BEAST: A brief measure of human social information use. *Evolution and Human Behavior*, 40(5), 492–499.
- Muthukrishna, M., & Henrich, J. (2016). Innovation in the collective brain. *Phil. Trans. R. Soc. B*, 371(1690), 20150192.
- Pöyhönen, S. (2017). Value of cognitive diversity in science. *Synthese*, 194(11), 4519–4540.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854–856.
- Sulik, J., Bahrami, B., & Deroy, O. (in press). The diversity gap: when diversity matters for knowledge. *Perspectives on Psychological Science*, <https://psyarxiv.com/hx7ze>.
- Thoma, J. (2015). The epistemic division of labor revisited. *Philosophy of Science*, 82(3), 454–472.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2), 225–252.
- Yahosseini, K. S., Reijula, S., Molleman, L., & Moussaïd, M. (2018). Social information can undermine individual performance in exploration-exploitation tasks. In C. Kalish, M. Rau, J. Zhu, & T. Rogers (Eds.), *CogSci 2018* (pp. 2473–2478).
- Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17–35.