Modeling cognitive diversity in group problem solving

Samuli Reijula (samuli.reijula@helsinki.fi)
Theoretical Philosophy, PO Box 24
00014 University of Helsinki, Finland

Jaakko Kuorikoski (jaakko.kuorikoski@helsinki.fi)
Practical Philosophy, PO Box 24
00014 University of Helsinki, Finland

Abstract

According to the diversity-beats-ability theorem, groups of diverse problem solvers can outperform groups of high-ability problem solvers (Hong and Page 2004). This striking claim about the power of cognitive diversity is highly influential within and outside academia, from democratic theory to management of teams in professional organizations. Our replication and analysis of the models used by Hong and Page suggests, however, that both the binary string model and its one-dimensional variant are inadequate for exploring the trade-off between cognitive diversity and ability. Diversity may sometimes beat ability, but the models fail to provide reliable evidence of if and when it does so. We suggest ways in which these important model templates can be improved.

Keywords: cognitive diversity; binary string model; distributed cognition; diversity-beats-ability theorem; simulation modeling

Introduction

Group problem solving benefits from cognitive diversity within the group. Differences in how members of the group see the problem, what kind of cognitive resources they have at their disposal, and what kind of heuristics they use, can all make it more probable that all the necessary ingredients for solving a complex problem are available to the group as a cognitive unit. According to a striking claim made by Lu Hong and Scott Page (2004), this benefit is so strong that *cognitively diverse groups may outperform groups consisting of more able individuals*, who are, due to shared expertise, inevitably cognitively less diverse. This claim has momentous practical implications: managers creating problem solving teams should favour group diversity even at the cost of individual expertise. No wonder the *diversity-beats-ability theorem* has had a large impact in a broad range of academic and practical fields (Reagans & Zuckerman, 2001; Mannix & Neale, 2005; Jeppesen & Lakhani, 2010; Steel, Fazelpour, Crewe, & Gillette, 2019; Aminpour et al., 2021).

The most important source of evidence for the diversity-beats-ability theorem is not to be found in data from controlled experiments, but from a set of mathematical and computational models (Hong & Page, 2001, 2004). In the absence of controlled experiments on the phenomenon, the use of formal modeling is an appropriate method of inquiry: The effects of the counteracting mechanisms of diversity and individual ability are hard to disentangle from purely observational data.[1] The counterfactual scenarios created by modeling can yield theoretical insight and assist the interpretation of observational data. Computational modeling, in particular, allows us to view group problem solving as the functioning of a distributed cognitive system in a task environment (cf. Hutchins, 1995; Sun, 2008), and the properties of both can be parameterized in the model.

We argue that despite their appeal, the models used by Hong and Page to derive their diversity theorem are fundamentally ill-suited to the task. To be clear, our findings in this paper do not call into question the importance of diversity, nor do they imply that diversity cannot beat individual ability. Generally we find it likely that cognitive diversity is beneficial for group problem solving, in various ways. Our findings show, however, that the models employed by Page and coauthors, despite their prominence, do not provide reliable support for the diversity theorem.

In the following section, we introduce the general modeling approach. Then we discuss the analytical results put forward by Hong and Page to support the diversity theorem. We then report findings from two simulation studies where we examined the model-based support for the diversity theorem. Our findings point to serious shortcomings in the original models employed by Hong and Page, but we show how an improved version of one of the models can be used to meaningfully study the trade-off between diversity and ability.

Group problem solving as heuristic search

The models used by Hong and Page join the tradition of modeling problem solving as heuristic search in a multi-dimensional solution space (Newell & Simon, 1972). Most of the conceptual work by Hong and Page is carried out with the help of a model introduced in Hong and Page (2001), where the problem task is represented as a binary string of finite length. Each bit in the string can be seen as portraying a yes/no decision regarding the solution to a particular sub-problem (Kauffman & Levin, 1987). The problem comprises of a sequence of *components* $S = \{s_1, s_2, \ldots, s_N\}$, where $s_i \in \{0,1\}$. The number of bits, N, represents the size of the problem and contributes to its difficulty. A *configuration*, or a possible solution to the problem, is a string $x_i = s_1, s_2, \ldots, s_N$. We denote the set of

[1] We do not suggest that insights from theoretical modeling could replace empirical evidence from observation and experiment. Instead, we see modeling as a tool for examining the validity of theoretical reasoning (Reijula & Kuorikoski, 2019).

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

possible configurations by $X = \{x_1, \dots, x_{2^N}\}$. The configurations can be ordered according to their epistemic (or practical) value, captured by the value function $V : X \rightarrow \mathbb{R}$. We write $x_i \succeq x_j$ whenever x_i is (weakly) preferred to x_j. Given this notation, a problem is defined by the pair $(X, \succeq)$ (Marengo & Dosi, 2005).[2]

The agents portrayed in the model search for solutions of maximally high epistemic value by relying on hill-climbing search: If the epistemic value of a novel solution candidate is strictly higher than that of the current solution, $V(x_j) > V(x_i)$, it is adopted as the new solution. Otherwise the candidate is discarded.

The model in Hong and Page (2001) implements diversity in problem-solving heuristics as (group) agents possessing different *flipset heuristics*. A flipset can be thought of as a bit mask, where the bits set to '1' flip the state of the corresponding bit in the target string. For example, the flipset '001' applied to the target string '101' flips the rightmost bit, and results in the string '100'. Each distinct set of flipset heuristics results in a characteristic set of possible paths that an agent can follow in the search space.

In Hong and Page's model, each agent has a small set of such heuristics. A set of heuristics ϕ gives rise to a set of positions $\{x_j, \dots, x_k\}$ which can be reached from x_i by a single application of the agent's heuristics. We call this set the *neighborhood* of x_i.[3] Consequently, each distinct set of flipset heuristics results in a characteristic set of possible paths that the agent can follow through the search space. The motivating intuition underlying the diversity-beats-ability theorem is that, under appropriate conditions, the diversity supplied by the larger set of heuristics in a random group of problem-solvers is epistemically more beneficial than a less diverse set of individually high-performing expert heuristics.

The binary string model nicely captures many key intuitions about the benefits of cognitive diversity. Moreover, it can be used to model another, distinct aspect of the epistemic benefit of diversity in problem solving: the possibility of *division of cognitive labor*. Cognitive diversity is useful when the problem to be solved can be broken down into partially independent sub-problems and distributed to agents with differing competences. Organizational economists (Marengo & Dosi, 2005) have used this model to investigate the efficiency of decompositional heuristics, and our implementation of such a model can be found in Reijula et al. (n.d.). However, this division-of-labor mechanism is not the one grounding the diversity-beats-ability theorem. Moreover, although most of the intuitions about diversity in, for example Page (2008), are pumped using the binary string model described in this section, the analytic derivation and simulation results for the diversity theorem are, in fact, drawn from *a different model template*, to which we turn next.

The mathematical result

The analytical result proven in Hong and Page (2004) to support the diversity theorem is almost completely independent of specific assumptions concerning the solution space, i.e. the nature of the problem task. The result assumes a population of agents that satisfy the following – in this context, fairly weak – assumptions: (i) Agents are intelligent: given any starting point, an agent finds a weakly better solution, and the set of local optima can be enumerated. (ii) The problem is difficult in the sense that no agent can always find the optimal solution. (iii) Agents are diverse: for any potential solution that is not the optimum, there exists at least one agent who can find an improvement. (iv) The best agent is unique.

The derivation establishes that for such a population, with probability one, there exist positive integers N and N_1, $N > N_1$, such that the joint performance of the N_1 independently drawn problem solvers exceeds the joint performance of the N_1 individually best problem solvers among the group of N agents independently drawn from the population (according to any probability distribution with full support). The proof is based on two immediate consequences of the assumptions. First, as independently drawn agents are very unlikely to have common local optima, the probability of their having common local optima converges to zero as the number of agents in the group grows. Second, as the number of agents becomes large, the best problem solvers become more and more similar and therefore do not do better than the single best problem solver.

We do not question the validity of the proof nor the claim that the result really depends on the diversity assumption (cf. Thompson, 2014; Singer, 2019). What is problematic is that the conclusion is, in the end, a relatively weak existence claim based on asymptotic reasoning concerning a large number of agents. The choice of such an approach is not surprising, as the assumptions used are very general. There is, however, no responsible way to infer from such a conceptual possibility, proven at the limit, to any finite real-life case in which expertise and diversity trade-off against each other, and choices about optimal group composition have to be made. What is needed for such inferences is a model template which allows quantitative modeling of the diversity-ability trade-off, and the factors on which it depends. What is arguably more persuasive is Hong and Page's agent-based "computational experiment" which is supposed to demonstrate that diversity can actually beat ability.

Simulation study 1: Ringworld model

The simulation model in Hong and Page (2004) can be seen as a simplified version of the binary string model summarized above. We call the simplified model the *ringworld model*. It differs from the original mainly in its problem structure. In the ringworld model, agents conduct the search for optimal solutions in a one-dimensional landscape consisting of positions $1 \dots n$ on the number line, wrapped as a circle, with a value

[2] For a comprehensive formal specification of the model, see Reijula, Kuorikoski, and MacLeod (n.d.), available at `https://osf.io/hw7zj/?view_only=b6081af9b43e459c9d47c491c451f380`

[3] To be precise, this is the 1-neighborhood created by ϕ around x_i. 2-neighborhood is the set of positions that can be reached by two applications of heuristics from ϕ, and so on.

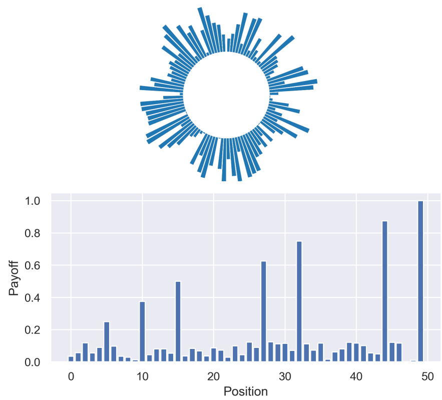

Figure 1: Above, an illustration of the original search space in the ringworld model, wrapped as a circle, and below a section of a similar landscape with the superimposed stairway sequence (=stairway landscape, step set [5,12]).

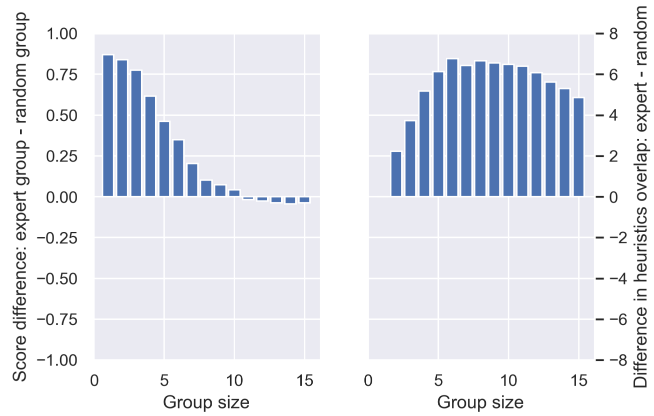

Figure 2: High-ability vs. random groups on a stairway landscape, step size 3. ($k = 3, l = 12, n = 2000$; 100 repetitions over 100 landscapes)

function V assigning a value drawn from the uniform distribution $[0, 100]$ to each of the 2000 positions (see upper part of figure 1). Each agent employs a heuristic ϕ, now defined as consisting of k different jumps of length $1 \dots l$ along the circle. (For example, [1,5,11] and [3,4,12] are heuristics with parameters $k = 3, l = 12$.) An agent applies these jumps sequentially along the landscape and moves to a new position if the payoff associated with the new position is strictly larger than the current one. Group problem solving is implemented as a sequence of individual searches by the agents in the group, in which the next agent begins where the previous stopped. Group performance is defined as the expected value of the position at which the group search stops (group's local maximum). The diversity-beats-ability result is demonstrated by comparing the performance of a high-ability group, produced by selecting the g individually best-performing agents in a tournament, to a randomly drawn group of size g. With some parameter values for k, l and g, a modest difference in favor of the random group emerges. For $(n = 2000, l = 12, k = 3, g = 10)$, the high-ability group scored 92.56 and the random group 94.53.

In our replication of the ringworld model we aimed at a better understanding of the process driving the observed outcomes (Reijula & Kuorikoski, 2021). Our analysis indicates that with the parameter values studied by Hong and Page (above), in nearly half of the cases, the random group ends up with a combined heuristic consisting of all the possible jumps, and only 13% of the cases the random group was missing more than two of the twelve possible heuristics. Consequently, even if the set of heuristics employed by the high-ability group was highly adaptive for the task at hand, the exhaustive search conducted by the random group would still outperform it.

Furthermore, as Grim et al. (2019) also pointed out, the purely random problem spaces studied by Hong and Page are not hospitable to anything that could be meaningfully interpreted as “ability” or “expertise”. Heuristic search makes sense only if the task has some structure or redundancy that the heuristic can exploit (Kauffman & Levin, 1987; Kahneman & Klein, 2009). Hence, aggregated over several random landscapes and starting positions, no significant performance differences are to be expected between different heuristics, expert or not. Hence, the ”ability mechanism” simply does not get any traction on the landscapes studied by Hong and Page, and we conclude that the model cannot adequately capture *the trade-off* between diversity and ability.

In Reijula & Kuorikoski, 2021, we modified the problem structure in the ringworld model by superimposing an increasing payoff sequence on the random value function V employed by Hong and Page (see lower part of figure 1).[4] Depending on its complexity this *stairway* sequence can be climbed by different combinations of heuristics available to agents, and hence the difficulty of the problem can be parametrically manipulated. Figure 2 summarizes our findings. The left panel presents the difference between the performance of high-ability and random groups (positive values standing for high-ability group advantage, and negative values for random group advantage). The results indicate that with these parameter values, stairway landscapes generally favor the high-ability groups.

The right panel provides a possible explanation for why the performance difference shrinks as group size increases. By the overlap between heuristics, or heuristic redundancy, we refer to the number of duplicate jumps in an (group) agent's heuristic (e.g., [2,2,7] has redundancy 1). The figure presents the difference between the redundancy in the high-ability and the random group (value 0 meaning that the overlap of heuristics in both groups is the same). As group size grows from 1, the redundancy in the high-ability group increases more than in the random group. This suggests that when the group size is larger, random groups again begin to approach the full heuristic, and can resort to exhaustive search. For this reason, at group sizes larger than 10, random groups catch up, and no significant performance difference is observed between high-ability and random groups (left panel).

We argue that this tension between the ”ability mechanism” and the ”diversity mechanism” captures the trade-off

[4] In the stairway model, we utilize V_n, value function scaled to the unit interval [0,1].

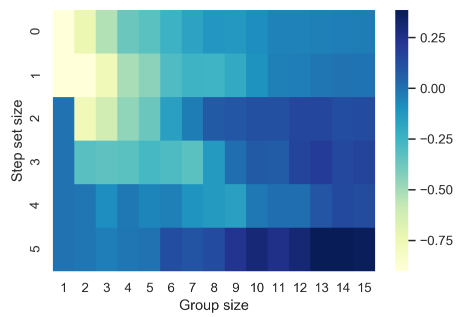

Figure 3: High-ability-vs-random group performance differential on stairway landscapes (50 repetitions, each over 50 landscapes).

addressed by the diversity theorem. What happens, however, when the level of ability or expertise required by the task changes? Figure 3 illustrates what happens in tasks of different problem complexity, where complexity is measured by *step set size*, the number of heuristic jumps needed for climbing the stairway sequence. In the figure, group size is represented on the horizontal axis, and step set size on the vertical axis. The color represents the performance differential between the high-ability group and the random group; lighter shades standing for high-ability group advantage. A genuine trade-off between diversity and ability can be seen. Observe the contrast between the upper-left quadrant, where ability dominates, and the lower-right, where random groups have a slight advantage over the high-ability groups; ability dominates when group size and step set size (i.e. complexity) are small, whereas diversity leads to better performance when the group size and step set size are larger.

Simulation study 2: Binary string model

Despite Hong and Page's model-switching, we also engaged in an attempt to computationally replicate and expand their findings by using the original binary string model (see section *Group problem solving as heuristic search* above). Hong and Page (2001) do not provide a computational implementation of the binary string model, and in model construction, we utilized, *mutantis mudandis*, the assumptions made in Hong and Page (2004): The payoff values of configurations are drawn from $U(0,100)$, the coordination of group problem solving is implemented as in the ringworld model, and the number of heuristics possessed by each agent, $k \in \{3,4\}$ (results appear robust across parameter values).[5]

The replicability of two results were of particular interest. The first one, originally proved analytically in Hong and Page (2001), reads as follows:

Arbitrary marginal product thesis. *For a randomly selected group of problem solvers, for group sizes larger than two, the marginal added value (increase in epistemic payoff) of an additional problem solver is not necessarily diminishing.*[6]

This is an intuitive result which states that the order in which diverse heuristics are applied matters to their contribution to the overall performance of search, and that increasing diversity *can* be more beneficial than simply adding agents with high individual ability. From an economists' perspective, this is a highly relevant result, as it distinguishes knowledge intensive work from standard economic assumptions about labor in general; it suggests that cognitive labor should not be conceptualized as an ordinary factor of production. However, as Hong and Page acknowledge, such a possibility proof fails to tell us anything about the conditions or the probability of occurrence of an increasing marginal product of an added problem solver.

Table 1: The probability of increasing marginal returns to an added problem solver

2	3	4	5	6	7	8
0.16	0.23	0.19	0.17	0.15	0.13	0.12

We studied whether this result can be observed in our computational implementation of the binary string model. Table 1 presents the probability of the marginal product of an added problem solver ($k = 1$) being larger than the marginal product of previous problem solver in the group, i.e. increasing marginal returns occurring. Across group sizes 2-8, the probabilities are not trivially small. This suggests that the increasing marginal returns phenomenon could be more than just a distant conceptual possibility. However, as group size increases, the probability of increasing marginal returns decreases. The arbitrary marginal product thesis is clearly important for the practical implications of the diversity-beats-ability claim, because the counter-intuitive power of the diversity thesis rests on the idea that it can be a good idea in team building to increase diversity even at the cost of individual ability *when no further information about the relative value and mutual interaction between individual sets of cognitive abilities is available*: If the manager already knew which additional heuristics would be beneficial (and which not) in solving the problem at hand, then increasing diversity *per se* would clearly be a sub-optimal strategy. Yet our aggregate results suggest that such a general recommendation cannot be directly drawn from the model, as the marginal return of diversity does diminish on average and that an investment on individual ability may thus always remain a better option.

Our second target is the famous diversity-beats-ability result already discussed in sections above. Applied to the binary string model, we formulate the thesis as follows:

Diversity theorem (binary string model). *For a given problem, randomly generated groups of agents outperform more*

[5] The source code used to generate the findings in this section can be viewed at https://osf.io/eypcq/?view_only=84bd3febfeca4010a6055e372197618b

[6] A detailed and rigorous presentation of this highly interesting proof is omitted here for reasons of space. See Hong & Page, 2001.

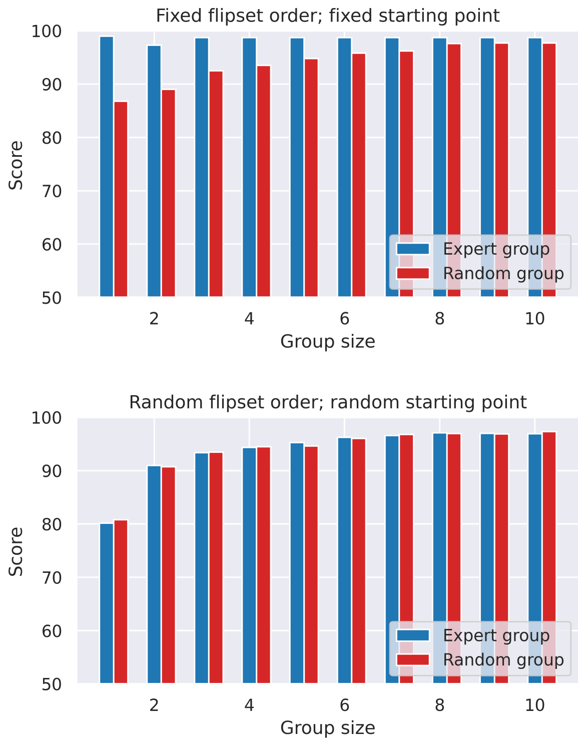

Figure 4: Random and high-ability groups in the binary string model ($N = 20, k = 3$, 50 repetitions over 50 different value functions.)

homogeneous groups of agents characterized by the individually best performing flipset heuristics.

Figure 4 summarizes our findings from a setting analogous to the one in the ringworld model. The lower panel shows the performance of high-ability and random groups of different group sizes under the assumption that each trial begins from a random point in the problem space, and the flipset heuristics are not always applied in the same order. Under such conditions, no systematic difference emerges between high-ability and randomly recruited groups. For comparison, the upper panel indicates what happens when group search begins from the same position in the landscape as individual search, and heuristics are always applied in the same order. Under these conditions, highly effective individual heuristics retain their effectiveness in the high-ability group, and hence the effects of a strong "ability mechanism" are observed. The fact that the pattern disappears when initial position and order of heuristics are randomized reflects the fact that on random landscapes with no redundant patterns, expertise fails to generalize to new situations.

In sum, our simulations of the original binary string model provide *no evidence for the diversity theorem*. The reasons for this failure turn out be interesting. As in the ringworld, the randomly generated multi-dimensional solution space of the binary string model fails to provide any signal which the heuristic search could detect; the problem task does not have any structure or redundancy that expertise could exploit. Therefore, the expertise embodied in particular heuristics fails go

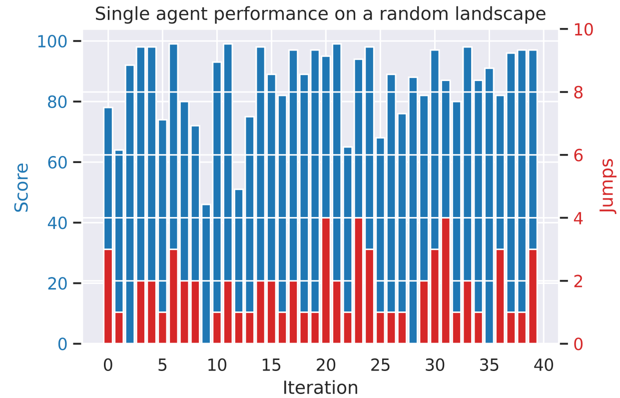

Figure 5: Single agent performance in the binary string model. Score mean: 85.9 ($N = 12, k = 4$)

generalize. This is an issue that Hong and Page (2001, pp.149–151) acknowledge, but it leads to two further questions: What kinds of problem spaces (e.g., analogous to stairway landscapes in the ringworld model) would be more amenable to heuristic search? And if the "ability mechanism" in fact does not function in the binary string model, why do we not see the diverse groups outperforming the more homogeneous expert groups in the lower panel of figure 4?

An answer to both these questions is suggested by an important disanalogy between the ringworld and the binary string model. The portrayal of the core model features common to both (see section *Group problem solving as heuristic search*) gives rise to an intuition of a small set of heuristics resulting in meandering search paths in the solution space. In the ringworld model, this can in fact happen. Yet a moment's reflection reveals that it does not occur at all in the binary string model. Due to the binary nature of the problem space and the implementation of heuristics as (xor) *masks*, the components of heuristic ϕ do not in fact mutually interact in a way which could create true novelty in the search, i.e., open up genuinely new areas of the solution space. For example, the iterative application of heuristic jumps A, B, and A results in the agent ending up in the point reached by a single application B alone (as the second application of A simply cancels its original effect). It is thus easy to show that the total size of the neighbourhood reachable from a point x_i with k distinct heuristics is simply the number of all possible subsets of a set with k elements, namely 2^k. Likewise, the maximum path length (for paths not ending up in the starting position) in the solution space is $2k - 1$.

For example, independently from the dimensionality of the problem and cardinality of the set of solutions in the space, an agent A with 4 heuristics can reach 16 distinct positions, and the length of its longest search path is 7 steps. Flipset-based search thus limits a (group) agent to a very small portion of the total solution space of size 2^N.

This perhaps unexpected short reach of flipset-based search can be observed in figure 5. The vertical axis on the right side of the figure indicates the number of successful applications of a flipset heuristic in the course of one iteration of search: the maximum number of jumps is 4. The blue bars represent the final payoff from search, and with these parameter values, the mean score is roughly 86 points. In other words, on average, modest search distance ($\leq$ 4 jumps) already leads to relatively good payoffs. This is to be expected. Bearing in mind that the epistemic values of positions in the solution space are drawn from $Z \sim U(0,100)$, a useful baseline for search performance can be obtained by calculating the expected value of the maximum value from k independent draws from such a distribution.

$$E[\max(Z_1,\ldots,Z_k)] = 100 - \frac{100}{k+1}$$

According this baseline, an exhaustive search of the 16 positions reachable from agent A's position gives an expected epistemic value of 94.4. The fact that the average performance of the agent in figure 5 falls somewhat short of this baseline is due to it occasionally getting stuck in a local maxima *within its own* 16-solution search space. This explains why the diversity effect does not lead to a noticeable performance increase in the binary string model: Increasing diversity by having an additional non-duplicate flipset heuristic available does not unlock whole new sections of the search space. Instead, it has only a marginal effect on the number of solutions reachable from a given location. This suggests that while it may still be possible to model cognitive diversity using the binary string approach, the flipset-heuristics implementation must be replaced by a less problematic alternative (see, for example, Marengo & Dosi, 2005).

Discussion

The pioneering research done by Hong and Page has given rise to both theoretically and practically important discussions on the role of cognitive diversity in group problem solving. The phenomenon of group problem solving has proved notoriously difficulty to model, and the binary string model satisfies many of the desiderata of such a modeling approach. That said, our exploration of the two model templates employed by Hong and Page revealed significant representational shortcomings and interesting differences between the two models.

Although intuitively less appealing as a model of cognitive diversity, the one-dimensional ringworld model does capture the power of diversity in that the search paths of groups with multiple different heuristics are longer and reach a broader neighborhood than more homogeneous expert groups. The main problem with the ringworld model is that the random landscape does not provide any signal for the experts to exploit, and thus does not adequately represent the ability side of the trade-off. This problem and a solution we described in Reijula and Kuorikoski (2021).

The binary string model, which appeared to nicely capture many salient features of cognitive diversity and group problem solving, in fact fails to adequately represent both diversity as well as ability: as in the one-dimensional case, the random landscape fails to provide any structure that heuristic search could utilize. Furthermore, the implementation of heuristic search through flipset heuristics leads to subtle but, in our view, insurmountable challenges to the whole model in its current form.

It should be noted, however, that the challenges faced by Hong and Page's approach do not call into question all models employing the multidimensional binary search space approach. For example, the problem of stunted path lengths does not arise in models combining NK-landscapes with local 'one-bit-mask' search (e.g., Lazer & Friedman, 2007; Yahosseini & Moussaïd, 2020).

Our main result concerns the diversity-beats-ability theorem. This striking claim about the power of cognitive diversity is highly influential within and outside academia, from democratic theory to management of teams in professional organizations. Our results show that the model-based arguments by Hong and Page for this claim are unsound. The analytical theorems establish only highly abstract proofs of possibility, from which nothing about the probability or conditions for diversity actually beating ability in the real world can be responsibly inferred. The computational experiment in Hong and Page (2004) showing a minute advantage in favor of diversity is demonstrated by using a model template which cannot adequately represent individual expertise. Moreover, the template is different from the original multi-dimensional model which is used to tie the relatively abstract models to salient intuitions about distributed and heuristic problem solving. When the question of diversity-vs-ability is computationally formulated in the original model, no systematic effect is observed. Diversity may well beat ability in many contexts of collaborative problem solving, but establishing this requires better models and preferably empirical experiments.

References

Aminpour, P., Gray, S. A., Singer, A., Scyphers, S. B., Jetter, A. J., Jordan, R., … Grabowski, J. H. (2021). The diversity bonus in pooling local knowledge about complex problems. *Proceedings of the National Academy of Sciences*, *118*(5).

Grim, P., Singer, D. J., Bramson, A., Holman, B., McGeehan, S., & Berger, W. J. (2019). Diversity, ability, and expertise in epistemic communities. *Philosophy of Science*, *86*(1), 98–123.

Hong, L., & Page, S. E. (2001). Problem solving by heterogeneous agents. *Journal of economic theory*, *97*(1), 123–163.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(46), 16385–16389.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Jeppesen, L. B., & Lakhani, K. R. (2010). Marginality and problem-solving effectiveness in broadcast search. *Organization science*, *21*(5), 1016–1033.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, *64*(6), 515–526.

Kauffman, S., & Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, *128*(1), 11–45.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694.

Mannix, E., & Neale, M. A. (2005). What differences make a difference? The promise and reality of diverse teams in organizations. *Psychological science in the public interest*, *6*(2), 31–55.

Marengo, L., & Dosi, G. (2005). Division of labor, organizational coordination and market mechanisms in collective problem-solving. *Journal of Economic Behavior & Organization*, *58*(2), 303–326.

Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 104) (No. 9). Prentice-Hall Englewood Cliffs, NJ.

Page, S. E. (2008). *The difference: how the power of diversity creates better groups, firms, schools, and societies* (Paperback ed.). Princeton, N.J. ;Woodstock: Princeton University Press,.

Reagans, R., & Zuckerman, E. W. (2001). Networks, diversity, and productivity: The social capital of corporate R&D teams. *Organization science*, *12*(4), 502–517.

Reijula, S., & Kuorikoski, J. (2019). Modeling epistemic communities. In M. Fricker, P. Graham, D. Henderson, & N. Pedersen (Eds.), *The Routledge Handbook of Social Epistemology* (pp. 240–249). Routledge.

Reijula, S., & Kuorikoski, J. (2021). The Diversity-Ability Trade-Off in Scientific Problem Solving. *Philosophy of Science*, *88*(5), 894–905.

Reijula, S., Kuorikoski, J., & MacLeod, M. (n.d.). Division of cognitive labor: The costs and benefits of interdiscipinarity. *under review*.

Singer, D. J. (2019). Diversity, not randomness, trumps ability. *Philosophy of Science*, *86*(1), 178–191.

Steel, D., Fazelpour, S., Crewe, B., & Gillette, K. (2019). Information elaboration and epistemic effects of diversity. *Synthese*, forthcoming.

Sun, R. (2008). Cognitive social simulation. In *The Cambridge handbook of computational psychology* (pp. 530–548). Cambridge University Press.

Thompson, A. (2014). Does diversity trump ability? *Notices of the AMS*, *61*(9), 1–024.

Yahosseini, K. S., & Moussaïd, M. (2020). Comparing groups of independent solvers and transmission chains as methods for collective problem-solving. *Scientific reports*, *10*(1), 1–9.