

The Diversity-Ability Trade-Off in Scientific Problem Solving

Samuli Reijula and Jaakko Kuorikoski*†

According to the diversity-beats-ability theorem, groups of diverse problem solvers can outperform groups of high-ability problem solvers. We argue that the model introduced by Lu Hong and Scott Page is inadequate for exploring the trade-off between diversity and ability. This is because the model employs an impoverished implementation of the problem-solving task. We present a new version of the model that captures the role of ‘ability’ in a meaningful way, and we use it to explore the trade-offs between diversity and ability in scientific problem solving.

1. Introduction. Modern science is a deeply collaborative enterprise. Most genuinely important intellectual challenges cannot be tackled by a single scientific discipline, let alone by individual researchers. Science needs diversity—solving scientific research problems requires attaining specialized expertise and resources from a variety of perspectives.

Problem-solving groups in general are taken to benefit from diversity (Reagans and Zuckerman 2001; Mannix and Neale 2005; Jeppesen and Lakhani 2010; Steel et al. 2021). Among other important benefits, it is assumed that differences in how members of a group see a problem, in the cognitive resources they have at their disposal, and in the kind of heuristics they use make it more likely that the group as a whole has the resources to solve the problem. An important question, therefore, is whether the diversity of a group is in itself epistemically valuable, over and above the epistemic abilities of the group members.

*To contact the authors, please write to: PO Box 24, 00014 University of Helsinki, Finland; e-mail: Samuli Reijula, samuli.reijula@helsinki.fi; Jaakko Kuorikoski, jaakko.kuorikoski@helsinki.fi.

†We thank Kristina Rolin, Inkeri Koskinen, Renne Pesonen, and the other members of the TINT group (University of Helsinki), as well as the participants of the Diversity in Science workshop (Tampere University, May 5, 2019) and the poster sessions at EPSA 2019 (University of Geneva) for their helpful comments. Thanks to Kate Sotejeff-Wilson for editing the manuscript. This research was carried out as part of the Social and Cognitive Diversity in Science project funded by the Academy of Finland.

Philosophy of Science, 88 (December 2021) pp. 894–905. 0031-8248/2021/8805-0012\$10.00
Copyright 2021 by the Philosophy of Science Association. All rights reserved.

Apart from the empirical evidence cited above, a particularly influential argument in favor of diversity has been presented in the form of a mathematical theorem and an agent-based simulation. According to the diversity-beats-ability (DAB) theorem, groups of diverse problem solvers can outperform groups of high-ability problem solvers. This means that in assembling problem-solving teams, functional group diversity should sometimes be prioritized over selecting the most able individual members. Although they originate in computational social science, in management and organization studies, the DAB results have recently been discussed in the philosophy of science (Holman et al. 2018; Grim et al. 2019; Singer 2019).

We argue that the “can” (in “can outperform”) in the DAB theorem is ambiguous between several different modalities: in some of its uses, it is only a claim about conceptual possibility, whereas in its much advertised practical applications, it is clearly regarded as a more substantial possibility. This raises the question of when and under which exact conditions diversity really beats ability. We examine whether the original model by Hong and Page, and its further developments by Grim and associates, actually support the existence of the DAB phenomenon.

We show that because of their *impoverished task implementation*, these models cannot capture interesting trade-offs between functional diversity and individual ability: the problem-solving tasks portrayed in the models are too difficult (i.e., random noise) for ability to make any difference to the outcomes. We develop a new version of the model with an improved problem-solving task. The new task representation allows our model to capture the role of individual ability in problem solving. Only when both diversity and ability really affect the outcome can the trade-off between them be studied.

We start by briefly presenting the DAB theorem and the associated simulation models, focusing on the latter. In section 2, we highlight the “bait-and-switch” argumentative strategy used by Page to argue for DAB, showing that many of the modeling results supposed to support the theorem are problematic and do not replicate well. In sections 3 and 4, we present our main argument: the model template used by Hong and Page as well as Grim and colleagues is ill suited for exploring the trade-off between diversity and ability, because the problem-solving task is computationally implemented in a way that does not afford any advantage to individual ability or expertise. We introduce our version of the model, the stairway landscape, and demonstrate how it captures a substantial trade-off between diversity and ability. We draw two potentially interesting conclusions concerning the trade-off.

In this article, we are only concerned with the purely instrumental value of cognitive diversity; we are not arguing against the DAB phenomenon as such. We only ask whether the particular models we discuss are an informative and reliable way of exploring the possible trade-off and provide what we regard as a better alternative way for doing so.

2. The Diversity-Ability Trade-Off in Group Problem Solving. Consider design tasks, such as designing an automobile, a space shuttle, or a piece of software, or scientific tasks such as measuring the mass of an elementary particle or discovering the structure of a macromolecule. Heterogeneous cognitive and material resources need to be applied to solve all these problems, and as the set of solution candidates is not known beforehand, a search for solutions is needed. Simon (1989) suggested viewing the scientific research process through the lens of heuristic search. For instance, scientists search for formulations of problems, experimental designs, patterns in data, mechanisms behind data, and implications of their theories. On some occasions, these multi-dimensional search trajectories result in beneficial epistemic design; in other cases, they yield research approaches of little cognitive value. Importantly, most scientific problems worth solving lie beyond the capacities of a single knower, and scientific progress relies on a successful division of labor and collaboration between researchers and research groups and sometimes even between scientific disciplines. Hence, scientific research should be understood as a socially distributed problem-solving process.

Such a picture of a collective search immediately suggests a possible trade-off. On the one hand, as Newell and Simon (1972) suggested, expert performance often relies on highly specific search heuristics. On the other hand, more diversity in the group's cognitive resources is beneficial, all other things being equal, as more varied resources provide access to larger portions of the solution space. Diversity may, however, conflict with individual ability. Experts are often more alike (in the relevant respects) than nonexperts. Herein lies the trade-off: individual ability and group diversity both contribute to group performance, but, at least in some circumstances, the two factors may be in conflict.

Explicit modeling of the epistemic benefits of diversity in collective problem solving is needed, because the phenomenon involves multiple group-level mechanisms as well as possible interactions between different epistemic, processual, and social factors. Therefore, purely verbal and conceptual theorizing is not a reliable tool for drawing out the implications of theoretical assumptions, and empirical (experimental or case-based) evidence does not usually unambiguously discriminate between alternative mechanistic explanations for why, in any particular case, diversity may or may not facilitate successful problem solving. Group problem solving has proved challenging to model, however. The computational implementation of the problem (task), cognitive resources (and differences therein), problem-solving behavior and cognition, and interaction between the group members all present difficult methodological and theoretical choices for the modeler, easily resulting in complex and intractable models with too many methodological degrees of freedom. Such models yield results that are hard to interpret. We believe that the heuristic-search paradigm proposed by Newell and Simon (1972) still provides the most promising approach for addressing these

modeling challenges (see also Kauffman and Levin 1987; March 1991; Darden 1997). The models discussed and developed in this article join this tradition.

In a series of articles and books, Lu Hong and Scott Page have provided model-based evidence for the existence of the diversity-ability trade-off (Hong and Page 2001, 2004; Page 2008). They, in fact, use two distinct models to investigate diversity. The first model, introduced in Hong and Page (2001) and described at length by Page (2008) in the context of the diversity theorem, represents the problem to be solved as a binary string of finite length, where each bit could be seen as portraying a yes-no decision regarding a solution to a particular subproblem (Kauffman and Levin 1987). A group of problem solvers of limited ability attempts to maximize a value function defined over the possible states of this string (potential solutions to the problem). Diversity is represented in the model by each agent having a different set of possible ways of flipping the bits (“flipset heuristics”) of the candidate solution string shared between the group members. Measures of problem difficulty can be assigned to alternative value functions (see Page 1996), and so the model can be used to represent a range of problems of different difficulty and complexity. This model template therefore corresponds well to pretheoretic intuitions about how cognitive diversity can facilitate collective problem solving.

It is therefore rather surprising that the influential DAB results are not derived from this model. Our replication of the model in Hong and Page (2001) did not provide evidence to support the DAB phenomenon (see fig. 1).¹ As the figure illustrates, no systematic difference emerges between groups of high-ability problem solvers and groups of randomly selected problem solvers. A more careful look at Page’s (2008) argument reveals that it is based on evidence for the diversity theorem from an altogether different model introduced in Hong and Page (2004). We refer to this simplified model as the *ringworld model*. In sum, the substantial intuitions about diversity and ability in collective problem solving are first formalized in one model, but the results are derived from a different model based on assumptions that do not correspond as neatly to the original intuitions. We find such a “bait-and-switch” argumentative strategy confusing and not appropriate for transparent and epistemically sustainable use of theoretical models.

The argument in Hong and Page (2004) has a two-pronged structure: the basic assumptions of the ringworld model are used to derive an analytical proof intended to provide support for the theorem. However, as argued by Thompson (2014), the implications of the proof are unclear: even after technical corrections, the theorem only provides a highly abstract proof of possibility,

1. For details about the bit string model, see Hong and Page (2001). All program code for the simulations and the generated data sets are available for download at <https://osf.io/a6f5e/>.




Figure 1. Score differential between high-ability and random groups in the bit string model.

and its implications for a nontechnical interpretation of diversity are difficult to judge. Although we agree with Singer (2019) that the proof does rely on diversity and not merely on randomness (see Thompson 2014), it still remains the case that as such, the proof tells us little about the conditions under which the trade-off between diversity and ability can be expected to be significant. Mere logical possibility is not enough for the far-reaching practical implications suggested by Hong and Page. Their more persuasive evidence for DAB and its relevance for real-world group problem solving are derived from their agent-based simulation of the ringworld model. It is to this simulation that we now turn.

3. Problems in the Ringworld. The “computational experiment” used by Hong and Page to demonstrate DAB portrays a group of agents collectively searching for optimal solutions in a one-dimensional landscape. The discrete landscape consists of positions $1 \dots n$ on the number line, wrapped as a circle.² Value function V defined over the set of positions assigns to each position a payoff value drawn from the uniform distribution $[0, 100]$. The agents’ goal is to find the largest possible values on this landscape. To do so, each agent employs a heuristic ϕ . A heuristic is defined as consisting of k different jumps of length $1 \dots l$ (e.g., $[1, 5, 11]$ and $[3, 4, 12]$ are two examples of heuristics

2. It turns out that the circular topology of the landscape does not make a difference to the results, as the distance explored by the individual agents (and groups) typically does not exceed 20 steps along the 2,000-step circle.

with parameters $k = 3$, $l = 12$). Starting from its current position, an agent sequentially applies these jumps along the landscape and moves to a new position along the circle if the payoff associated with that position is strictly larger than the current one. When no further improvement is possible, the agent stops. The performance of an agent is defined as the expected payoff of the stopping points over the different starting positions of the landscape and over a set of landscapes.

Hong and Page implement group problem-solving behavior as a sequential, iterative search. First, one agent initiates the search. As its local maximum is found, the second agent in the group takes the baton and applies the jumps included in its heuristic as long as they lead to improvements. After all group members have taken their turn, a new round begins. The collective search stops when no agent can make further progress. Group performance is defined as the expected value of the position at which the group search stops.

In order to compare groups of high-ability problem solvers to more diverse ones, an exhaustive set of agents (with respect to possible heuristics) is first ranked according to their individual performance on a set of landscapes. A high-ability group of size g is constructed from the g highest performers in such a tournament, whereas the diverse group consists of g agents sampled randomly from the population.

In their model analysis, Hong and Page (2004) report results for various sets of parameter values. For example, for $l = 12$, $k = 3$, and $n = 2,000$ they find that the best individual agents scored 87.3, whereas the worst agent's score was 84.3. For groups of 10, the high-ability group scored 92.56, and the random group, 94.53. This difference in favor of the random group is the diversity effect discovered in the simulation. Similar results were found by Grim and associates (2019), and we were also able to replicate the findings.

Hong and Page suggest that there are reasons to believe that the random group scored higher because of its diversity. An alternative way to express this finding is in terms of *effective group size*. In our replication, we noticed that the difference in performance ('performance differential') between the random and the high-ability group was strongly correlated (.65) with the difference in effective group size between the two groups, where effective group size was defined as the size of the group heuristic from which overlapping elements had been removed. In other words, the similarity between the members of a high-ability group results in the group being functionally smaller (from the perspective of the problem-solving task). Because the performance of a group generally increases as its effective group size gets larger, it is not surprising that smaller effective group size leads to worse performance.

Going back to the original DAB theorem, however, the explanation above seems to capture only one side of the diversity-ability trade-off. Although the correlation between effective group size and performance is an indication of

the functioning of the “diversity mechanism,” it is still unclear why that effect is stronger than the influence of the “ability mechanism,” that is, the fact that some heuristics should lead to higher performance than others and that those high-performing heuristics should be more common in high-ability groups. A closer inspection of the model provides a solution to this puzzle.

Unlike Hong and Page, we regard the effect sizes from the simulation as remarkably small, given that they originate from theoretical modeling in which the modeler is free to explore a broad range of hypothetical scenarios. One would expect a purely theoretical model, purpose-built to examine and demonstrate a specific mechanism using heavy idealizations, to reveal relatively unambiguous effects of the modeled mechanisms. As a matter of methodological principle, we believe that conclusions drawn from agent-based modeling would be strengthened by showing how the effect size can be manipulated by changing model parameters. In other words, being able to “turn the dials” and observe how changes in model inputs result in systematic changes in the modeled effect suggests that we have reached an understanding about the dependencies between model inputs and outputs (see Woodward 2003; Aydinonat, Reijula, and Ylikoski 2020). Regarding the ringworld model, we argue there are two reasons to believe that the results reported by Hong and Page do not provide genuine insight into the diversity-ability trade-off.

First, with the parameter values studied by Hong and Page, in nearly half of the cases, the random group ends up with a *full heuristic*, that is, a heuristic consisting of all possible jumps [1, . . . , 12]. Furthermore, only 13% of the random groups have an effective group size smaller than 11. Hence, even if the agents in the high-ability group can make the jumps leading to high performance, it is highly likely that the same jumps will also be included in the heuristic of the random group—there is simply no way the high-ability group could systematically outperform the random one.

Second, as Grim and his colleagues (2019) also noted, the purely random landscapes studied by Hong and Page are simply not hospitable to anything that could be meaningfully interpreted as “ability” or “expertise.” For a heuristic search to be applicable, the task needs to have some structure or redundancy that the heuristic can exploit (Kauffman and Levin 1987; Kahneman and Klein 2009). Hence, aggregated over several random landscapes, no significant performance differences emerge between the different heuristics. This is seen in the very small performance differences between the best and worst performing individual agents (see above) in Hong and Page’s simulations: the “ability mechanism” does not get any traction on the landscapes they studied. Therefore, we argue that the model does not appropriately capture the trade-off between diversity and ability.

Grim and his coauthors (2019) propose to remedy this problem by partially smoothing out the random landscape (by adding interpolated values between randomly generated values). They argue that such a task representation can

better capture ability, because on smoothed out landscapes individual performance is more transportable to other landscapes of similar smoothness. Yet a closer numerical examination of the results of this remedy again reveals only small differences between diversity and ability. Even on smoothed random landscapes, the expected performance difference between best performing and random individuals is minute. This suggests that these landscapes simply do not represent a problem that is suitably complex for exploring trade-offs between ability and diversity.

4. Modeling the Diversity-Ability Trade-Off on Stairway Landscapes.

In order to better understand the tension between diversity and ability, we need to portray scenarios in which ability plays a role as well. In our own simulations, we introduce a type of problem in which high ability—either at the individual or group level—leads to noticeably increased performance. In science, having the right methodology for the problem at hand often sharply increases the epistemic payoff. Our stairway model differs from the Hong and Page ringworld model only in problem structure. The specifications of agent and group behavior remain the same as in the ringworld model. In generating problem landscapes, we start from the uniform noise distribution employed by Hong and Page. On top of those landscapes, however, we superimpose an increasing sequence of values, where the positions of the values are separated by intervals drawn from a finite set of integers in $1 \dots l$ (see fig. 2). We call this set the *step set*.




Figure 2. Stairway landscape with step set $\{5, 12\}$ and, therefore, step set size 2.

For an agent to climb the increasing subsequence, the *stairway sequence*, it must possess the heuristic jumps corresponding to the steps used to generate the sequence (e.g., {5, 12} in fig. 2). This strongly favors some heuristics over others: whereas an agent who does not possess the full step set is bound to remain in the noise region of the landscape, a “high-ability” agent that has the necessary heuristic can climb through the whole sequence (and even reach the maximum payoff on the landscape, normalized to 1.0).

Figure 3 illustrates outcomes from our model with parameter values corresponding to those studied by Hong and Page (2004) and by Grim and his colleagues (2019). Figure 3a 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 always favor high-ability groups. Especially when the group size is small, because it is made up of high-performing individuals (who typically possess valuable elements of the step set) the high-ability group performs significantly better than the random group. Figure 3b presents the difference between the redundancy of heuristics between the high-ability and random group (value 0 means that the overlap of heuristics in both groups is the same). As suggested by findings by Hong and Page (2004), random groups tend to have comparatively lower levels of overlap in their heuristics. As group size increases, the redundancy in




Figure 3. High-ability versus random groups on a stairway landscape, step size 3 ($k = 3$, $l = 12$, $n = 2,000$; 100 repetitions over 100 landscapes).

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, which obviously is sufficient for climbing the stairway sequence. 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 (fig. 3a).

We argue that this tension between the “ability mechanism” and the “diversity mechanism” captures the trade-off addressed by the DAB theorem. What happens, however, when the level of ability or expertise required by the task changes? Different levels of task difficulty can be represented by stairway landscapes with different step set sizes. For example, landscapes with step set sizes up to 3 lie within the abilities of the individual agents studied in the simulation ($k = 3$). Climbing the stairway for step sizes larger than 3 requires pooling heuristics from several agents.

Figure 4 summarizes tentative findings from our studies with landscapes of varying difficulty. In the figure, group size is represented on the horizontal axis, and step set size (complexity of the problem) is on the vertical axis. The color represents the performance differential between the high-ability group and the random group; lighter shades stand for high-ability group advantage. A genuine trade-off between diversity and ability can be seen.

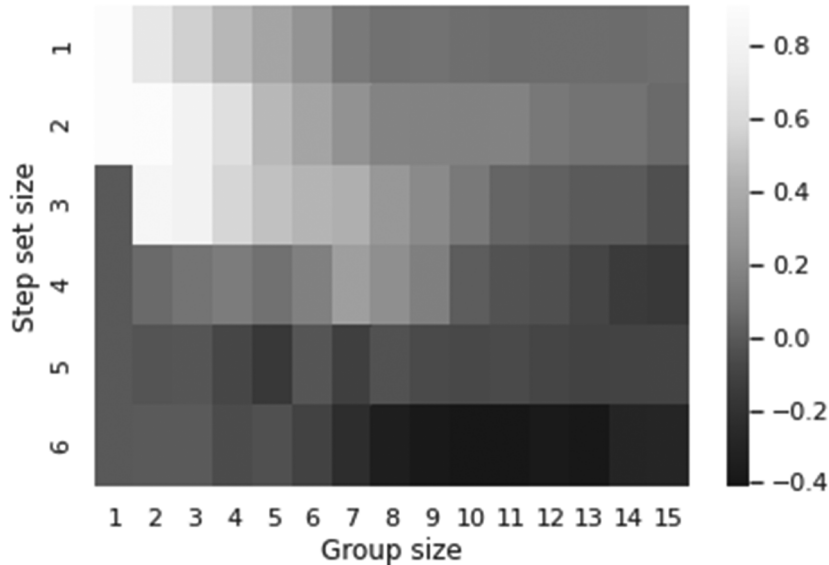


Figure 4. High-ability versus random group performance differential on stairway landscapes (50 repetitions, each over 50 landscapes). Color version available as an on-line enhancement.

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 are small, whereas diversity leads to better performance when the group size and step set size are larger.

Finally, our results suggest a conceptual distinction between the complexity and difficulty of a problem: perhaps not surprisingly, ability dominates when the problem is simple in the specific sense that multiple cognitive resources do not need to be combined to solve it. Note that if the problem is simple in this sense, this does not necessarily mean that it is easy to solve. When the problem becomes complex, requiring efficient division of cognitive labor, the diversity effect begins to dominate over individual abilities. The results demonstrate how diversity and group size begin to outdo individual ability only when the problem complexity exceeds the cognitive resources of any single individual.

One could object to our stairway model on seemingly similar grounds to those on which we based our criticism of the original ringworld model. We questioned the DAB results on the basis that the model was built to favor diversity over ability. Why would our model fare any better, as it was clearly built to favor ability over diversity? This objection misses our point, however. Our argument is that the original model cannot be used to model the trade-off between diversity and ability, because it cannot be used to represent the gains from ability. Of course, we fully admit that the stairway landscape is built to favor ability, but the model nevertheless also retains the gains from diversity. Stairway landscapes give both ability and diversity their due and, therefore, can illuminate the trade-off between them. This, we argue, was the original and interesting interpretation of the DAB results to begin with.

5. Conclusion. The original results by Hong and Page do not provide reliable evidence for the DAB theorem because the ringworld model, especially its task implementation, does not allow for the ability to adequately influence individual or group performance. This one-sidedness implies that their model cannot be used to explore the possible trade-offs between diversity and ability in problem-solving groups. Our exploration of stairway landscapes illustrates how Hong and Page (2004) rely on a problematic task structure to get their results.

Stairway landscapes provide a better model for “medium-hard” problems that require specialized abilities and true division of cognitive labor. Such landscapes can be used to model the relevant interplay between diversity and ability and its effects on the division of cognitive labor in science.

Our tentative modeling results suggest a trade-off between diversity and ability. Ability is favored when the problem is moderately difficult, requiring only a few different expert heuristics, and when groups are small. Diversity is favored when the problem is complex, requiring multiple component

solutions, and when the groups are large. A further qualitative effect can be observed at the point where problem complexity increases beyond the capacity of a single agent and necessitates division of cognitive labor: simple problems solvable by individuals favor ability regardless of group size.

REFERENCES

- Aydinonat, N. E., S. Reijula, and P. K. Ylikoski. 2020. "Argumentative Landscapes: The Function of Models in Social Epistemology." *Synthese*, forthcoming.
- Darden, L. 1997. "Recent Work in Computational Scientific Discovery." In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, 161–66. Mahwah, NJ: Erlbaum.
- Grim, P., D. J. Singer, A. Bramson, B. Holman, S. McGeehan, and W. J. Berger. 2019. "Diversity, Ability, and Expertise in Epistemic Communities." *Philosophy of Science* 86 (1): 98–123.
- Holman, B., W. J. Berger, D. J. Singer, P. Grim, and A. Bramson. 2018. "Diversity and Democracy: Agent-Based Modeling in Political Philosophy." *Historical Social Research* 43 (1): 259–84.
- Hong, L., and S. E. Page. 2001. "Problem Solving by Heterogeneous Agents." *Journal of Economic Theory* 97 (1): 123–63.
- . 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences of the USA* 101 (46): 16385–89.
- Jeppesen, L. B., and K. R. Lakhani. 2010. "Marginality and Problem-Solving Effectiveness in Broadcast Search." *Organization Science* 21 (5): 1016–33.
- Kahneman, D., and G. Klein. 2009. "Conditions for Intuitive Expertise: A Failure to Disagree." *American Psychologist* 64 (6): 515–26.
- Kauffman, S., and S. Levin. 1987. "Towards a General Theory of Adaptive Walks on Rugged Landscapes." *Journal of Theoretical Biology* 128 (1): 11–45.
- Mannix, E., and M. A. Neale. 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.
- March, J. G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2 (1): 71–87.
- Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Page, S. E. 1996. "Two Measures of Difficulty." *Economic Theory* 8 (2): 321–46.
- . 2008. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press.
- Reagans, R., and E. W. Zuckerman. 2001. "Networks, Diversity, and Productivity: The Social Capital of Corporate R&D Teams." *Organization Science* 12 (4): 502–17.
- Simon, H. A. 1989. "The Scientist as Problem Solver." In *Complex Information Processing: The Impact of Herbert A. Simon*, ed. David Klahr and Kenneth Kotovsky, 375–98. Hillsdale, NJ: Erlbaum.
- Singer, D. J. 2019. "Diversity, Not Randomness, Trumps Ability." *Philosophy of Science* 86 (1): 178–91.
- Steel, D., S. Fazelpour, B. Crewe, and K. Gillette. 2021. "Information Elaboration and Epistemic Effects of Diversity." *Synthese* 198:1287–307.
- Thompson, A. 2014. "Does Diversity Trump Ability?" *Notices of the AMS* 61 (9): 1024–30.
- Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford Studies in Philosophy of Science. New York: Oxford University Press.