Philosophy of Science (2025), **00**, 1–25
doi:10.1017/psa.2025.20

ARTICLE

Challenging the Consensus: The Strategic Value of Homogeneous Groups in Collective Problem Solving

Sahar Heydari Fard

The Philosophy Department, The Ohio State University, Columbus, OH, USA
Email: heydarifard.1@osu.edu

(Received 29 August 2024; revised 18 December 2024; accepted 27 March 2025)

Abstract

As technology fosters connections among like-minded individuals, concerns about the effects of homogeneous clusters—often criticized as ideological bubbles and echo chambers—have intensified. While these clusters are commonly seen as obstacles to independent thought and progress, this paper argues that they can, under certain conditions, drive significant advancements. By revising computational models of collective problem-solving and examining historical cases, I demonstrate that clusters, particularly among minority groups with superior ideas, can overcome dominant resistance and promote progress. However, this clustering introduces trade-offs, including slower consensus formation.

1. Introduction

As technological advances make it easier to connect with like-minded others, there is a growing interest in, and anxiety about, the health of various epistemic communities. One significant concern is the generation of homogeneous clusters, such as ideological bubbles or echo chambers, which may limit independent thinking and prevent communities from growth. Such an environment is reasonably believed to hinder social and epistemic progress and contradict our collective recognition of the value of diversity. Consequently, many categorically deem these clusters harmful and advocate for their dissolution and prevention.

Yet, contrary to this widespread belief about homogeneity, numerous instances of scientific and societal progress reveal that the formation of small like-minded collectives can be essential for driving meaningful change. These clusters can be especially critical when a minority group possesses a valuable idea but struggles to gain traction because of prevailing social dynamics working against them. While not all changes are productive, some of the most significant examples of large-scale progress—those that shift consensus toward better solutions despite obstacles—begin with such homogeneous clustering. Therefore, explaining why and when these clusters can be beneficial is crucial for understanding progress in social collectives and identifying pathways for change.

The closest we get to understanding the potential mechanisms underlying the impact of homogeneity on progress, movement-building, and change comes from influential research in the philosophy of science and computational social sciences (Hong and Page 2004; Zollman 2010; Fazelpour and Steel 2022; O'Connor and Weatherall 2019; Weisberg and Muldoon 2009; Grim et al. 2019). Despite the variety of approaches in this literature, the common thread is the use of computational simulations to provide empirically testable claims and reliable mechanisms for testing causality across different contexts.

While these models involve artificial agents and simulated societies, their insights have had a tangible influence on public policy and institutional design. Unfortunately, however, these models say very little about the role of strategically homogeneous collectives in contexts marked by power asymmetries[1] and in response to entrenched problematic consensus. This gap is critical as many notable examples of scientific and social progress have emerged precisely from such outnumbered and homogeneous groups challenging the status quo.

The main goal of this paper is to address this gap by modifying existing models to explore conditions under which homogeneous clustering could facilitate change rather than hinder progress. To achieve this goal, in section 2 I situate my discussion within the existing and influential models of diversity to identify the necessary ingredients for addressing the gap. In section 3, drawing from historical examples of scientific progress, I provide a representative case where homogeneous clusters prove essential for generating change. I use this example to justify two key deviations from the existing modeling literature: one concerning asymmetry and the other regarding the conventional definition of progress. In section 4, I outline the design of a model that helps us explain the significance of homogeneous clusters for epistemic progress in a collective. And in section 5, I present and discuss the results.

The results of my model support the claim that homogeneous clustering can be essential for the progress of an epistemic collective—at least insofar as models of artificial societies and agents offer insights into real-world dynamics or serve as a basis for generating empirically testable hypotheses. When dealing with "complex problems," a minority group with superior answers can benefit significantly from clustering, as it enhances their ability to persuade a majority that has already made up its mind. If the minority group has inferior answers, however, their clustering does not shift the majority's views, indicating an asymmetrical relationship. The trade-off for this clustering is the longer time it takes, on average, to form a consensus. This finding suggests that homogeneous clusters of like-minded individuals are not inherently detrimental and should not be categorically avoided. More importantly, mandating their dissolution can inadvertently reinforce a problematic status quo.

Another significant finding of my model pertains to the network-level analysis of collective problem-solving and the impact of demographic or ideological composition on group performance. Prior research in this area (Fazelpour and Rubin 2022; Gomez and Lazer 2019; O'Connor and Weatherall 2019) has explored how factors such as homophily—individuals' tendency to associate with similar others—influences

[1] Some recent work, like Wu (2023), highlights the impact of power asymmetries.

outcomes, providing recommendations on effective strategies. My results indicate that clustering as a strategy has markedly different effects depending on a community's stage in the problem-solving process. In communities still searching for answers, clustering has a negligible impact, whereas in communities where a minority group clusters in response to a majority that has already formed an opinion, the effect is profound. This disparity suggests the need to rethink the common operationalization of progress in collective problem-solving models and to reconsider the application of recommendations from models that overlook this distinction.

Furthermore, my model demonstrates that factors such as size and conformity significantly hinder a minority group's ability to disseminate their superior solution, while trust plays a minimal role. It is also vital for the minority group to achieve internal consensus from the outset. These insights emphasize the importance of tailoring strategies when utilizing clustering to foster epistemic and social progress.

2. Modeling the impact of diversity and homogeneity

Can the clustering of similar people ever benefit the epistemic health of a community? Research across various fields highlights the value of the opposite: mixing different types of individuals and their ideas brings significant benefits, such as generating innovative solutions and improving our ability to correct errors. Studies show that diverse groups tend to produce more and better ideas in a variety of settings (Dezso and Ross 2012; Freeman and Huang 2014; Muldoon 2016; Hong and Page 2001; Gaus 2021), ranging from small deliberative collectives (Sommers 2006; Richard et al. 2003) to research clusters and citizen assemblies (Landemore and Page 2015; Mason and Watts 2012), which are often more creative and efficient because of their diversity.

Hong and Page (2004) provide an influential example of such studies. They demonstrate that groups with diverse members often outperform groups of similar but high-ability individuals. Diverse groups perform better because they bring a wider range of perspectives and approaches to exploring a *complex problem space* (for an extensive discussion of this model see Grim et al. 2019; Steel and Bolduc 2020; Singer 2019; Thompson 2014). Weisberg and Muldoon (2009) make a similar point about scientific communities by showing that a mix of diverse heuristic strategies, as opposed to individually virtuous strategies, leads to the most thorough exploration and the best outcomes for the collective. Their model suggests that diversity in scientific thinking and methodology prevents communities from getting stuck with suboptimal solutions and promotes better results. Consequently, these studies recommend increasing diversity within collectives to enhance their potential for progress and success.

Homogeneity, on the other hand, is often seen as a costly problem for a group's epistemic health (Philips and Apfelbaum 2012). As homogeneity increases, diversity and its epistemic benefits diminish. As a result, almost all recommendations for solving complex problems through collaboration argue against homogeneity. In fact, when solving complex problems, decreasing homogeneity and thereby increasing diversity are often treated as a golden solution for various pathologies of epistemic collectives, especially those believed to hinder progress (Gaus 2021; Muldoon 2016; Vallier and Muldoon 2021).

We have a "complex problem" when the problem at hand is much greater in scale or intricacy than individuals' capacity to solve it alone. Such problems often require exploring a vast array of potential solutions without a straightforward way to narrow down the options through a priori reasoning. This complexity can arise from interactions among many variables, as seen in the stock market, ecosystems, or cultures where components are not easily decomposable or isolatable.

Complex problems also feature many potential pitfalls (local optima) and only a few truly excellent solutions. As problems become more complex, the human mind's capacity to formulate and solve them alone or with preset strategies diminishes (Simon 1955; Kahneman 2003; Zollman 2010). Thus, we rely on epistemic collectives to explore the problem space, avoid pitfalls, and identify the best possible solutions. By dividing cognitive labor in diverse epistemic collectives, we enhance our epistemic capacity, broaden our exploration scope, and reduce the likelihood of blind spots (Kitcher 1990; Zollman 2010; Weisberg and Muldoon 2009).

Even outside of the modeling literature, the imperative to break homogeneous clusters has prominence. Daniel Kahneman's idea of "adversarial collaboration" among scientists is a notable example (Mellers, Hertwig, and Kahneman 2001; Kahneman 2012). By recognizing the complexity of problems scientists must manage, Kahneman argues that "adversarial collaboration" among them can resolve long-standing disagreements and inefficiencies that come with collective problem-solving. With a political spin, the recommendation to disincentivize ideological clusters grounds recent state-level regulations in the United States for increasing "ideological diversity" in academic spaces.[2] In the context of social scientific research, the suggestion has been to incentivize research teams composed of researchers with opposing ideologies to work on controversial issues such as trans medicine, the use of police force, affirmative action, etc. (Clark et al. 2022; Martel et al. 2024).

At their core, these recommendations assume that homogeneous clustering is purely problematic and must be eliminated without considering any potential benefits. They overlook the fact that some of these controversial issues, like trans medicine, are relatively new and often studied by small groups of scientists working against long-standing beliefs and opposition. The assumption is that regardless of the history or status of these topics, researchers' contributions will be fairly considered by promoting adversarial collaboration or increasing the intermixing of individuals with opposing viewpoints. Any discomfort or difficulty that arises from such collaborations can be justified because of the common goal of achieving the best possible answers as an epistemic collective. As some independently argue, such discomfort or distrust is an asset, not a problem, of diverse epistemic groups (Muldoon 2018; Fazelpour and Steel 2022).

The attempt to improve diversity by increasing the intermixing of individuals—such that homogeneous clusters would be dissolved and dispersed—is not groundless. The modeling literature on diversity and collective problem-solving supports this approach, with philosophical discussions explicitly using these models to advocate for what is called "perspectival diversity" (Vallier and Muldoon 2021; Gaus 2021; Muldoon

[2] For example, in May 2023, the Ohio Senate passed Senate Bill 83, mandating intellectual diversity in public colleges while banning mandatory diversity programs. Similar bills passed in Florida and are under consideration in Texas.

2016, 2018). In what follows, I elaborate on a prototype of such models that grounds these claims.

Consider Hong and Page's (2004) model as an example. They conceptualize solving complex problems as a group exploring a vast landscape with 2000 possible solutions of varying quality. Each solution is assigned a score between 1 and 100, with higher scores representing better solutions and higher peaks on this landscape. The group's goal is to find the best possible solution by navigating and exploring this landscape. If we imagine the landscape as a terrain, heuristics represent the steps agents can take to move to a new location/solution and evaluate its quality. The central problem that this model aims to answer is how individuals and groups can maximize their chances of finding the highest peak or the best solution given their limitations.

One of those limitations is that the agents in this model are myopic, meaning that they only learn about new solutions through experimentation. The assumption is that agents move together as a group and learn about the outcomes of their group members' experiments instantaneously, an assumption I problematize later. They move (find new solutions to experiment with) according to their heuristics. They stay in the new location if they find better answers, otherwise, they return to their starting point.

To identify which combination of individuals has the highest rate of success, Hong and Page compare the performance of groups with higher diversity (lower homogeneity) with groups whose members have previously outperformed others. They show that improving diversity (lowering the number of agents with the same heuristics in a group) is *always* the best strategy for enhancing the group's chance of finding the highest peak or best possible solution. This finding grounds their influential motto that "diversity trumps ability" in collective problem-solving.

The Hong and Page model supports recommendations like adversarial collaboration by showing that incorporating diverse problem-solving approaches leads to better outcomes than relying on our best judgments of individual performance in epistemic collectives. If enhancing diversity involves increasing the intermixing of individuals with different heuristics, integrating researchers with opposing ideologies is logical, as far as it effectively diversifies the range of heuristics. Conversely, intentionally forming homogeneous groups—for example academic departments with a single ideological leaning—should have the opposite effect.

The problem, however, is that the ideological diversity of an epistemic community does not necessarily change through adversarial collaboration. Forming research teams with opposing views only alters the local distribution of ideological diversity or the local intermixing that affects how often members interact or hold each other accountable. More importantly, for ideological diversity to have the intended effect, it must directly correspond to diversity in heuristics with no other social dynamics that could alter outcomes if a critical mass supports the project. However, as I elaborate shortly, one can model similar effects of ideological clustering without any impact on heuristics. In other words, we cannot assume that group-level diversity captures the problem with the clustering of like-minded individuals (without attending to diversity's local distribution), nor can we equate ideological diversity with diversity in heuristics.

3. Three steps toward an alternative model

If increasing diversity always enhances collective performance, how do we explain the potential benefits of homogeneous clustering? More importantly, how can we create a model that evaluates the various factors involved in this process? A helpful step is to identify cases where homogeneous clustering has been crucial. By understanding the mechanisms at play in these cases, we can see that the utility of homogeneous clusters can be compatible with the overall epistemic value of diversity, particularly when considering the structure of collaboration within a collective. This also highlights the need for other modifications to diversity models, such as accounting for power or numerical asymmetries and recognizing that the definition of progress can vary depending on the stage of exploration at which a collective finds itself. I summarize these points in two further steps.

3.1. Step 1: Identifying cases where homogeneous clustering has been crucial

A representative example demonstrating the value of strategic clustering is the process that led to the discovery of the smallpox vaccine.[3] After observing in Turkey the practice of variolation—intentionally infecting a person with material from smallpox sores to confer immunity—Lady Mary Wortley Montagu introduced it to the medical community in England in 1721.[4] The practice was met with skepticism and opposition from the scientific establishment for several reasons. Smallpox was a devastating disease that made the medical community wary of new and untested methods. Variolation also conflicted with the prevailing medical practices and understanding of disease prevention at the time. Additionally, Lady Montagu's status as a woman and outsider to the medical community, coupled with the foreign origin of her idea, did little to aid the acceptance of her proposal.

In response to this skepticism and rejection, Lady Montagu created a cluster of believers in the practice, primarily among other aristocratic women. She convinced them by performing variolation on her own children to prove its legitimacy. This kind of public demonstration garnered enough attention and support that variolation was eventually able to be tried and evaluated on its merits. The original support led to grassroots endorsement and ultimately to broader acceptance by both the general public and the medical community.[5] Widespread endorsement later set the stage for Edward Jenner's development of the smallpox vaccine in 1796, which employed a similar concept using the cowpox virus.

This example is one of many that shows how a group that is a minority, whether in power or number, and often distrusted by the majority, can overcome significant barriers through strategic clustering. As the literature on social movements also suggests, strategic clustering can amplify minority voices and enable them to challenge the status quo (Fraser 2021; Heydari Fard 2024b; Crutchfield 2018; Tilly 1998). Such clustering does not alter the overall diversity of a population but changes the way such diversity is locally distributed, affecting the level of intermixing, for

[3] This example is adapted from O'Connor and Weatherall (2019), who employ it in a similar context.

[4] This was not the first time that the British Loyal Society was informed of the practice, but the information was not taken seriously (see Barnes 2012).

[5] For a detailed discussion of this endorsement and Lady Mary's role in generating it, see Glynn and Glynn (2005); Willett (2021).

instance, of those who believed in variolation and those who did not. More importantly, this clustering was a response to the skepticism that arose *both* from being an outsider and from having beliefs incompatible with the majority consensus. Finally, this example hints at the struggle that is inherent in advocating for change against a consensus.

A comprehensive model that captures the value of clustering should incorporate four key elements. First, it should account for unchanged overall diversity with altered local distribution. Second, it should consider social dynamics that can trap good solutions and prevent the population from benefiting from them. Third, it should incorporate the possibility of asymmetry in these relations, such as when minority groups with superior insights struggle to spread their ideas due to distrust and pressure from the majority. Fourth, it should illustrate the potential difficulties of changing a consensus as opposed to participating in a collective exploration. In Step 2, I explain how the first three elements can be addressed, and in Step 3, I tackle the last element. For a model with these modifications, the hypothesis to be tested is that these minority groups can overcome significant barriers through strategic clustering, amplifying their voices, and challenging the status quo.

3.2. Step 2: Adding social structures

Following the work of Bala and Goyal (1998) in economics, Zollman (2007) shows that certain communication structures can enhance a group's overall performance by balancing the cost of exploration against the benefits of exploiting the knowledge produced by others. Exploitation, or copying the better answers one's neighbors have obtained, is crucial for the collective's ability to efficiently and collaboratively find the best possible solutions. However, it is also important for individuals not to halt their exploration prematurely and just copy the marginally superior answer of their neighbors. Zollman demonstrates that in highly connected communication structures, where individuals have numerous neighbors, they cease exploring more quickly and, on average, converge on lower-quality answers. In contrast, less connected networks search more thoroughly before converging on an answer. This impact of communication structures on collective performance is known as the "Zollman effect."

In addition to the impact of the overall structural features like connectivity level, there are other social dynamics in a network that impact the collective performance. For example, people often trust those who they perceive to be more similar to themselves over others (Gaither et al. 2018; Levine et al. 2002). They might also experience a greater social pressure to conform to those who belong to the same groups (Cialdini and Goldstein 2004). Thus, mixing individuals who belong to different groups, or what Fazelpour and Steel (2022) call increasing the "demographic diversity," alters the efficiency of their overall communication structure, which thereby changes the collective performance. According to their model, this diversity leads to slower communication in highly connected communication structures, which prevents individuals from prematurely converging on suboptimal answers and tempers the Zollman effect.

How can we use structures to see whether homogeneity can ever be beneficial? Structural models of collective problem-solving distinguish group-level diversity from its local distribution. For example, it is perfectly compatible with Zollman's model that local homogeneity, where connections are concentrated among individuals with the same solution, can create pockets of diversity that persist longer, thereby reducing the risk of premature consensus. These models also allow us to incorporate identity-based dynamics that can alter the spread of good or bad information within a network (Fazelpour and Steel 2022). These dynamics can also coexist with the clustering of individuals who share the same solution or belief. For instance, members of a minority group can happen to have a distinct ideology or solution from the majority while this membership also determines various aspects of their social dynamics. These dynamics can occur between groups of unequal size or with asymmetrical relations of trust or conformity.

In sum, structural models provide empirically testable ways to measure how local clustering of like-minded individuals can impact collective performance while preserving group diversity. If such clusters ever enhance collective performance, we must reconsider the utility of adversarial collaboration or blanket statements about intermixing as universal solutions to social epistemic problems, without denying clustering's benefits in certain conditions. More importantly, we can identify conditions under which these local clusters serve as engines for progress rather than problems to be dissolved. Figure 1 represents an example of a structural model that distinguishes global diversity from its local distribution.

Moreover, structural models offer alternative ways to conceptualize diversity tailored to various explanatory goals. For example, while Hong and Page's model emphasizes diversity in heuristics—the strategies individuals use to explore problem spaces—Zollman focuses on temporary variation in solutions themselves, rather than on the methods used to find them. Fazelpour and Steel's demographic diversity highlights the impact of interaction between group dynamics and differing opinions within communication structures.

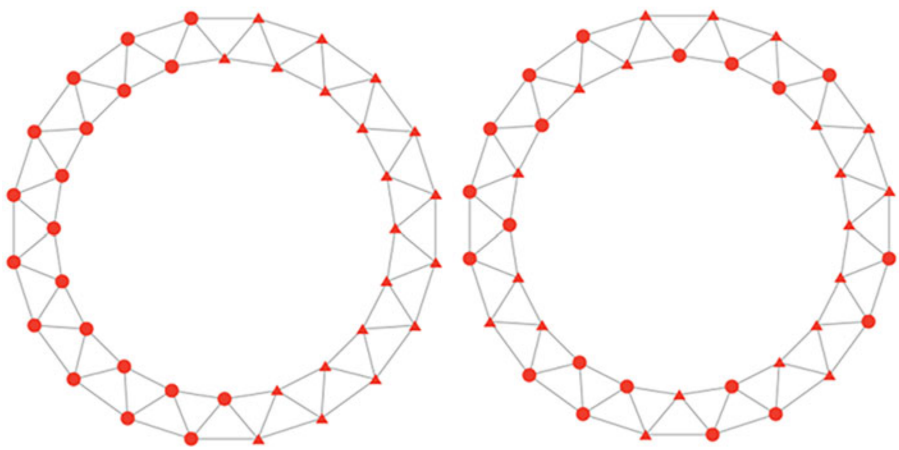

Figure 1. The figure shows two networks with identical global diversity in terms of variety and proportion. In the graph on the left, the two types of agents (circles and triangles) are clustered, while in the graph on the right, they are randomly distributed and maximally intermixed.

The way diversity is defined significantly influences the outcomes of models (Singer 2019). For example, ideological diversity might describe fixed heuristics—stable, unchanging rules followed by individuals or groups. Alternatively, it can reflect differences in adaptable solutions that are temporarily stabilized by group dynamics. This perspective aligns more closely with intuitions about ideology and with real-world phenomena like variolation. I use this dynamic view of ideological diversity to build my model and to evaluate their interpretations and outcomes.

3.3. Step 3: Revisiting the operationalization of progress

The final step is to challenge the conventional view of progress in the modeling literature, which largely confines it to the "exploration period"—the phase before consensus is formed, where diverse opinions, strategies, and group dynamics help a population converge on the best solution. These models overlook the possibility that progress does not stop when consensus is reached. In fact, some of the most significant advances occur when a consensus shifts from a lower-quality solution to a better one. In what follows, I identify three assumptions underlying this oversight and propose modifications to challenge and refine them.

First, these models assume a closed population with no meaningful inflow of individuals or information. In such a scenario, once a consensus is reached and everyone is satisfied with the perceived best answer, this consensus becomes a stable equilibrium. There is no incentive for anyone to move away from it or explore alternatives, leaving no means to disrupt the consensus. However, real societies are open systems with a constant inflow of new information and even new individuals. Therefore, any consensus in such societies can be subject to *internal* change.

Another reason is the assumption that individuals choose from a finite and known number of solutions. For example, instead of exploring a landscape, Zollman's model (2007, 2010) reimagines the process of discovery and progress in scientific communities as a two-armed bandit problem.[6] Each "arm" of his model represents a different scientific theory or treatment. Scientists must decide which theory or treatment to pursue, each with varying likelihoods of success, which they can only determine by experimenting with one of the options. For instance, they might test a treatment on patients or exchange information with other scientists, reflecting the exploration–exploitation dynamics previously discussed. However, in the end, there are only two clear candidates to choose from, with no other alternatives. Therefore, once experiments on both options are complete and consensus forms around the superior answer, there is no reason for change. However, in real-life scenarios we rarely know whether we have exhausted the problem space or whether a better solution is possible.

A third reason for dismissing the possibility of changing consensus is the assumption that we are merely in another exploration period whenever a consensus is challenged. Thus, while it is technically possible to form a consensus and then change it, the processes before and after consensus formation are considered identical (see, for example, Zollman 2010). Consequently, any recommendations

[6] A two-armed bandit problem involves choosing between two slot machines, each with an unknown payout probability. The goal is to maximize total rewards by balancing between exploiting the machine you believe has the better payout and exploring the other to gather more information.

applicable to the exploration period before consensus would also apply afterward. However, even intuitively, contributing to collective problem-solving when everyone is actively exploring solutions differs from doing so when people resist changing their minds.

To capture the fact that we live in open societies, we can model a society that begins with a majority consensus and then consider the impact of an outsider group or influence generating a small subgroup that adheres to an alternative solution (corresponding to the first reason). Additionally, we can systematically incorporate the landscape metaphor, instead of a two- or many-armed bandit model, to reflect the reality of our uncertainty about whether we have exhausted all solution spaces and whether a better alternative is possible (corresponding to the second reason). By doing so, we can test the legitimacy of the third reason: whether interventions that are effective during the exploration period have the same effect after a consensus is formed. As I elaborate shortly, my model demonstrates that they do not have the same effect and that the relationship is not symmetrical.

4. Simulation setup

My model builds on the work of Gomez and Lazer (2019), who use complex epistemic landscapes (NK landscapes) alongside social structures to explore whether homogeneous clusters can ever be beneficial. Their model is a structural version of that by Hong and Page (2004) with significant improvements, which I explain below. I expand their model by incorporating aspects of demographic diversity, including variations in social dynamics such as trust and conformity pressure based on group membership, as described by Fazelpour and Steel (2022). This approach allows for the regeneration of ideological or solution-based clustering with overlapping group-based dynamics. Additionally, I introduce various asymmetries to measure their impact on collective performance and consider progress as something that can occur at any stage of a collective's exploration.

4.1. NK landscape

As I mentioned before, complex problems can be visualized as rugged terrains where agents (like climbers) seek the highest peak (best solution). A powerful way to study such problems is to formalize these terrains based on two factors: N, representing the number of involving variables, and K representing the number of interactions among them. Each point in the NK space represents a combination of the N variables, captured by a sequence of 1s and 0s. This sequence is like a "basket of ideas," where each "1" indicates the presence of an idea and each "0" its absence. Every point of this landscape, then, has an address represented by N variables that can be either 1 or 0 and a score between 0 and 1 that represents the quality of this point.

K is the level of interdependencies between the variables and represents the complexity of the problem at hand. For example, a healthy lifestyle involves many variables: a good diet, physical activity, emotional and financial security, and many other factors. Arguably, these variables are interdependent, meaning that a high-calorie diet can work well for a farming community that involves manual labor and many levels of interdependence and support to provide emotional and financial

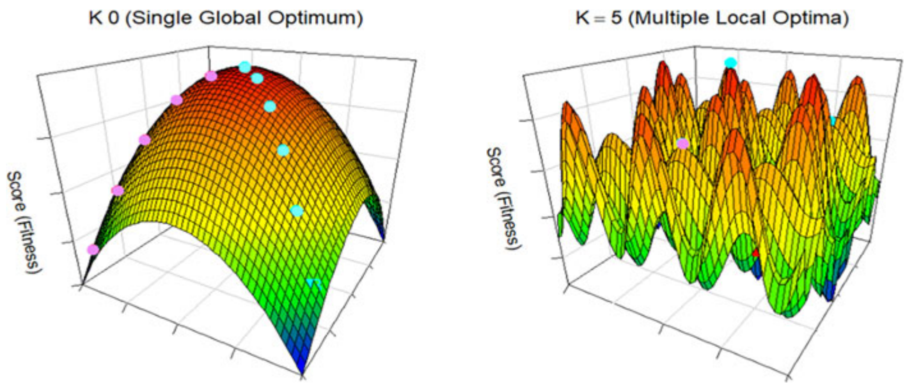

Figure 2. An NK problem space comparing the topography of two landscapes with different levels of interdependency.

security. But the same diet can be a problem for a culture with a sedentary lifestyle in which food becomes a way to soothe feelings of loneliness and despair.

The higher the value of *K*, or the interdependencies between variables, the harder it is for individuals to strike a balance among their needs and find an optimal solution by themselves. Higher values of *K* change the topography of the problem space and create many local optima, which makes it harder to find the best possible solution and easier to get stuck with mediocre answers. Figure 2 compares such a topography for two hypothetical landscapes with two levels of interdependency. I also assume that agents are myopic, meaning that there is no way for them to have a view of the whole landscape and use it to pick the best solutions. In fact, the number of available solutions, local or global, is unknown to anyone. Agents only see the consequences of their own and their neighbors' choices.

Since my focus is to see how progress happens after consensus is made, I take the starting position for the majority group to be consensus on a locally optimal solution. But unbeknownst to the agents who occupy this position, there are better peaks available. We can further assume the presence of a subgroup within the population that is aware of a significantly better solution, though not necessarily the best possible one. Progress, in this context, would be for the entire population to reach a consensus on this better solution. To provide a consistent point of comparison, I hold the landscape constant, as opposed to generating a new one for each simulation as done by Gomez and Lazer (2019).[7] I also identify two local optima on this landscape as the starting position for groups: one significantly worse than average and one significantly better.

4.2. Network structure

I use a torus network of 40 agents, each connected to four others, as shown in Figure 1. This setup is informed by the research of Lazer and Friedman (2007), which points to a

[7] Doing so does not affect the results, as I have tested over 20 different randomly generated landscapes with the same *N* and *K*.

trade-off between efficiency (the speed at which a solution is found) and accuracy (the quality of the solution).[8] Their findings suggests that a moderate level of connectivity —having four connections per agent in my case—strikes an optimal balance between efficiency and accuracy. This makes the torus network with four neighbors an ideal choice for exploring problem-solving dynamics. Keeping the structure constant also ensures that the impact of clustering is not conflated with variations in connectivity levels.

Dividing the population into two subgroups enables the exploration of group dynamics. We can also use the two network structures, shown in Figure 1, to distribute the members of these groups: one with group clustering and one without. When there is no clustering, the members of the two groups are randomly distributed to represent maximal intermixing between them. We can further manipulate the proportion of the groups as well as their in-group/out-group trust or conformity.

4.3. The experimental setup

As in most models I have described so far, in each iteration, agents aim to enhance their solution by using one of two options: *exploitation* and, if this fails to improve their solution, *exploration*. First, they look around to see whether any of their immediate four neighbors has a better solution than theirs. If they find a better answer, they simply copy that answer. If they cannot find a better solution, they explore by randomly changing one of the N variables from 0 to 1 or vice versa, thereby incrementally changing their position on the landscape. They observe the outcome, and if their incremental change leads to a solution better than their starting point, they adopt it and move to the new location. If it doesn't, they stick with the old solution. They repeat until either there is no better answer to be found or everyone has the same solution, which marks the end of their exploration period.

The dependent variable in models of collective problem-solving is usually the fraction of times agents converge on the best possible solution.[9] To capture the essence of exploration periods, the starting position for each agent is unique and randomly chosen. In Zollman's two-armed bandit models, for instance, agents start with a random distribution of credence or subjective utility for the two solutions at hand, and their success is measured by the fraction of times they converge on the best one. In Gomez and Lazer's (2019) model and similar ones, agents start with a random distribution of solutions or positions on the landscape. However, to explore the possibility of an evolving consensus, two changes are necessary. First, instead of starting from random, agents should be able to start from a semi-consensus, where each of the two groups has its own consensus that is not shared with the other group. Second, instead of defining progress as finding the best possible solution, we should consider converging on the better of the two consensus points as an instance of progress.

The focus here is not on diversity in heuristics or exploration methods, as with Hong and Page (2004); instead, it is on a mix of demographic diversity and diversity in knowledge. In other words, we are looking at cases where a solution, or a starting position, is shared by members of a group and this membership also influences other

[8] Other network structures may be more realistic, but the torus is well-established in the literature and demonstrates that clustering can be beneficial.

[9] While modeling approaches vary, as Huang (2024) shows, these differences are often translatable.

social dynamics like trust and conformity. The correlation is particularly strong when both groups start with a consensus. However, we can also examine scenarios where only one group starts with a consensus, or where neither group does. Exploring this correlation enables us to investigate cases like ideological or information-based clustering, where shared group identity or similarity of opinion/like-mindedness coalesce.

With this setup, we can now determine whether clustering among those who share a common starting point helps or hinders their chances of success, given other factors. One key factor is size proportion. The dynamics between two equally sized groups, which can resemble polarization, are vastly different from those between a significant majority and a small minority (Heydari Fard 2024a; Axelrod, Daymude, and Forrest 2021). Clustering during periods of polarization might be a primary issue that needs addressing, but in the latter case, it could be the only way for the minority to make its voice heard.

Another potentially relevant factor is the dynamics of trust. As Fazelpour and Steel (2022) note and incorporate in their model, individuals often condition their trust based on group membership. This can mean that when an agent observes their four neighbors achieving a superior solution, they will adopt their answers with an ω% likelihood if the neighbor belongs to an out-group, and with a 100% likelihood if the neighbor is part of their in-group. We can hold ω constant or let it change with the degree of similarity in solutions. However, for the sake of simplicity I assume that ω is constant in each simulation. We can manipulate ω to examine how change in trust affects the necessity of homogeneous clusters for progress.

Conformity pressure is another crucial factor to consider when evaluating the impact of homogeneous clustering on progress. Conformity can be sensitive to the number of people who endorse a particular set of beliefs. For instance, if we suppose x out of a person's four neighbors believe that variolation is effective, we can then consider the conformity incentive as a function of $x/4$ and an additional term to the objective utility of variolation. Therefore, even if no one believes that variolation is effective, there remains a baseline utility determined by the objective and stable factors. However, the additional conformity boost can make the NK landscape as perceived by agents dynamic, with its peaks changing value depending on the number of agents occupying them. This dependency can vary depending on whether the neighboring agents belong to one's identity group, and it can be asymmetrical. For example, the minority group may feel conformity pressure from the majority, but not vice versa.

The following equation simplifies Fazelpour and Steel's formalization (2022) of the conformity's impact on individuals' perceived utility of a solution. In this equation, U_j^i represents the perceived utility for agent i and solution j, while V_j denotes the objective utility of that solution. The tendency to conform is represented by κ, which I assume to be constant and equal to 0.5 following Fazelpour and Steel (2022). This means that agents assign equal value to both the objective utility and the conformity pressure. As previously mentioned, x represents the number of agents who share one's solution and belong to the same group. Since the number of neighbors is constant and equal to four for all agents, the conformity factor is a function of $x/4$.[10]

[10] For discussions about the impact of variation in K (kappa) please check the appendix.

$$U_j^i = (1 - \kappa) \times V_j + \kappa \times {}^{x}\!/_{4}$$

5. Results

With the setup I explained above, I ran 10,000 simulations, all using the same NK problem space (with $N = 10$ and $K = 5$) and a torus network occupied by two groups and where each agent has four neighbors. In what follows, I discuss two major results. First, for an absolute minority group (20% of the population) with a high-quality solution (though not the best possible one), I demonstrate that homogeneous clustering is strongly correlated with significant improvement in collective performance. I assume that both the majority and the minority groups start with a consensus. Second, I examine factors that make homogeneous clusters a strategically advantageous choice for minority groups, including their size, trust, conformity, and starting position. I also show that the impact and necessity of clustering differ between periods of exploration and after a consensus is formed.[11]

5.1. The value of clustering for the minority

Figure 3 summarizes the results for the clustering effect of the minority groups. In this figure, the x-axis represents the final score on which the whole population converges, while the y-axis indicates the fraction of simulation runs that ended up with a particular score. On the x-axis, the lower score, 0.290, represents the majority's inferior starting position and locally optimal solution, while 0.619 represents the minority's superior starting position, which is also a local optimum. The fill color in this figure distinguishes clustering networks from those with random distribution. As shown, the odds of converging on the superior solution with clustering—indicating progress—are 88%; the chance of converging on the lower-quality solution is only 12%. Without clustering, the odds of progress—the whole population converging on the better solution—drop to 25%, and there is a 75% chance that the majority will convince everyone, including the minority groups, to stick with the inferior solution.

These results indicate that clustering is an extremely effective strategy for minority groups to help the majority transition to a better solution. Conversely, dissolving such clusters can be an effective way to maintain the status quo and prevent change. One might argue, however, that while homogeneous clustering by the minority groups with superior solutions leads to improvement, the potential drawbacks of minority clusters with inferior solutions outweigh the benefits. For instance, flat-earthers or conspiracy theorists may only cause trouble or irreversible harm. Given our inability to assess the quality of solutions a minority group possesses, it might seem prudent to dissolve any clustering indiscriminately. While this perspective might hold for many factors not considered here, the setup of my model suggests otherwise. Clustering by a minority group that possesses an inferior solution does not have the same effect. As seen in Figure 4, clustering does not increase the likelihood of a minority group convincing the population to adopt their inferior solution. They almost always converge on the majority's superior alternative.

When minority members are maximally dispersed, their initial score is 0.619, with no additional conformity benefits since it is likely that all four of their neighbors belong

[11] The code for this model can be found at https://github.com/Saharhf/challengingconsensus.git.

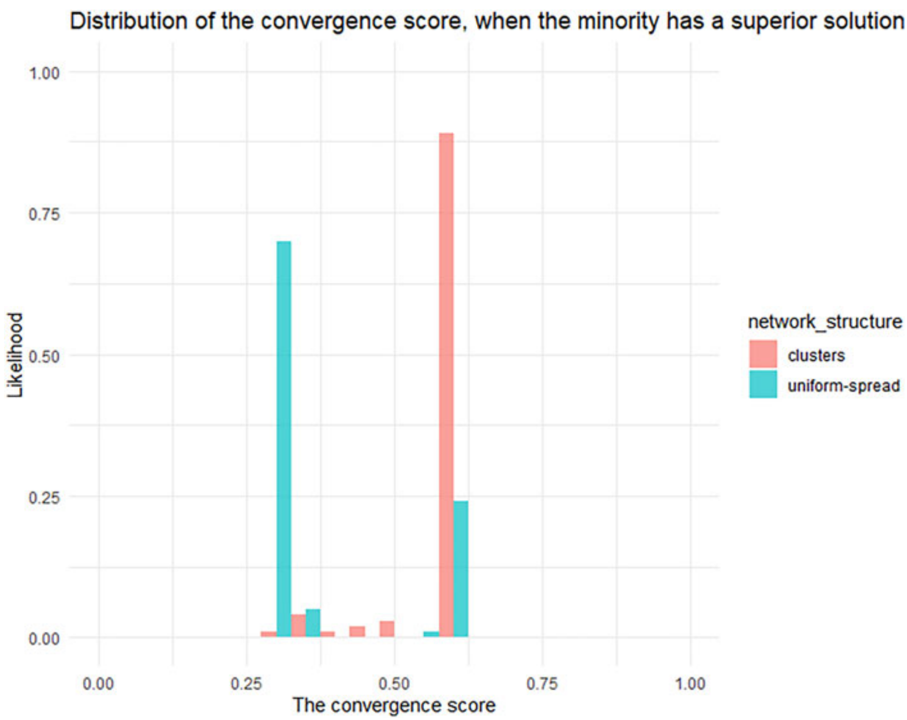

Figure 3. This graph represents the odds of the whole population converging on either solution, depending on the clustering of the minority group.

to the majority. In contrast, the majority members surrounding them start with an inferior score of 0.290. However, because three out of four of their neighbors share their solution, they receive a conformity boost. This results in the dispersed minority having a perceived score of $0.5 \times (0.619 + 0) = 0.31$, while their majority neighbors have a utility of $0.5 \times \left(0.290 + \frac{3}{4}\right) = 0.52$. Consequently, it is rational for the minority members to adopt the majority's solution to improve their score. There is no incentive for majority members to consider the minority's position since, overall, their own solution appears more advantageous.

However, if the minority members form a cluster instead of being maximally dispersed, some will have only minority neighbors. In this case, their perceived starting utility is 0.809, while a majority member surrounded entirely by majority neighbors has a utility of 0.645. More importantly, minority members at the edge of their cluster often have higher scores than their majority neighbors. For example, a minority member with three minority neighbors and one majority neighbor has a perceived utility of 0.684, which exceeds their majority neighbor's highest possible score of 0.52. Even with two majority neighbors, a minority member's score of 0.559 is still higher than that of a majority member with one minority neighbor.

With maximal intermixing, conformity pressure prevents the better idea of minority members from spreading, making it only a matter of time before they conform to the majority. In contrast, clustering allows conformity pressure to echo

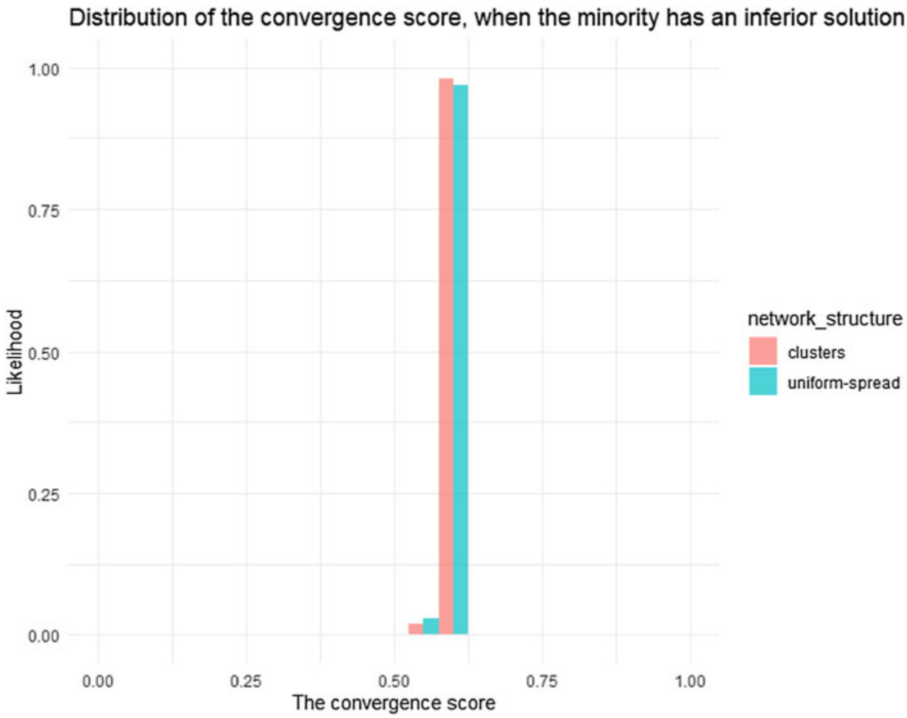

Figure 4. This figure represents the likelihood of the whole population landing on the better solution. It allows us to compare the effect of clustering, which, as is clear in the graph, is insubstantial.

within the minority cluster, amplifying the utility of the better solution and helping it endure long enough to capture the attention of the broader population. Essentially, clustering creates a protective shield for the minority, preventing them from succumbing to the majority's pressure. The impact of this protective effect for the minority, even when it has an inferior solution, highly depends on the difference in solution scores and the tendency of individuals to conform to others. However, even with a lower score difference and a higher tendency to conform, the effect is not symmetrical.

Clustering helps sustain minority perspectives but comes at a cost: as shown in Figure 5, it significantly slows the population's convergence on a solution, creating a trade-off between accuracy and efficiency. This trade-off, however, is not absolute. In contexts where preserving diverse viewpoints is vital for innovation or fairness, the slower pace may be worth the cost. Conversely, in situations requiring quick coordination, the drawbacks of clustering might outweigh its benefits. This underscores the importance of assessing clustering's value based on the specific goals and constraints of each situation.

5.2. Other factors

It is crucial to expand the simulations to examine how factors such as size asymmetry, trust, conformity, and starting position influence the impact of clustering for a

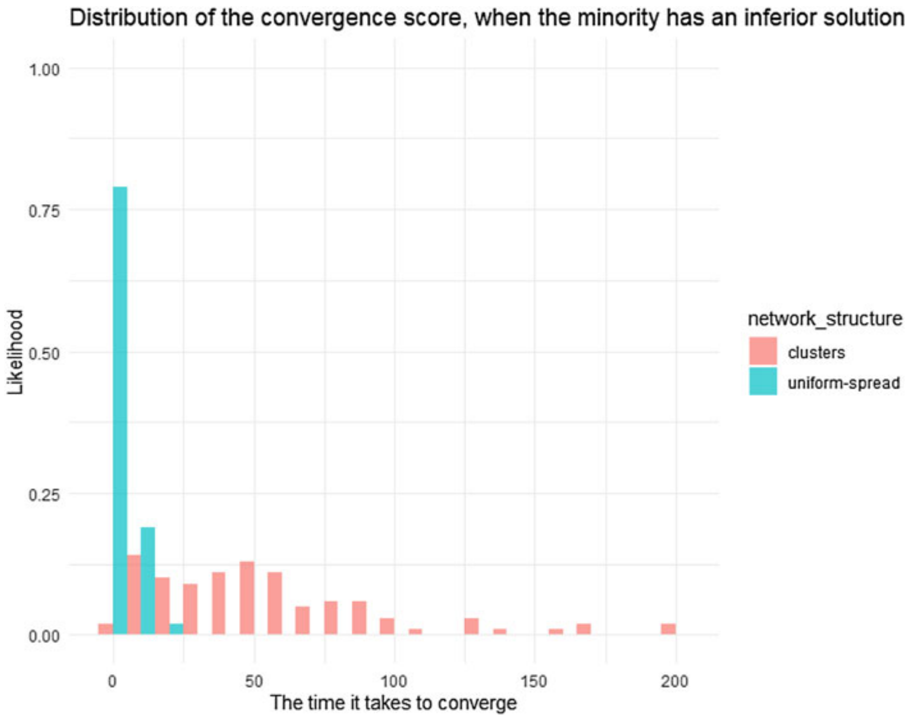

Figure 5. The time it takes to converge is significantly lower with a uniform spread than it is with clustering.

minority group, especially since the results observed so far assume these variables to be fixed.[12] If my analysis is right, size asymmetry and conformity should have the most significant effect. Moreover, if the observed impact results from the combination of shared solution and group dynamics, as stated in discussions about ideological diversity and the variolation example, then consensus among both the majority and the minority groups is necessary for clustering to be effective. Finally, it is important to compare the effect of clustering during periods of exploration with its effect after a population has reached a consensus. As I demonstrate shortly, there is a substantial difference in the efficacy of clustering between scenarios where the majority group starts from a consensus and those where it does not.

5.2.1. Size asymmetry and trust

Figures 6 and 7 represent the variation in collective performance of a population regarding two factors: (i) different size proportions between the minority and the majority (between 0.1 to 0.5) and (ii) different levels of intergroup trust (between 0.1 to 1), as defined in section 4.3. Here, both the minority and the majority still start with

[12] These variables are discussed here because they were fixed in the previous results section and, as noted in section 5.1, are central to the argument that clustering like-minded agents can aid epistemic progress. Additional variables, such as solution quality, and the difference between solutions, are analyzed in the appendix to streamline the main discussion.

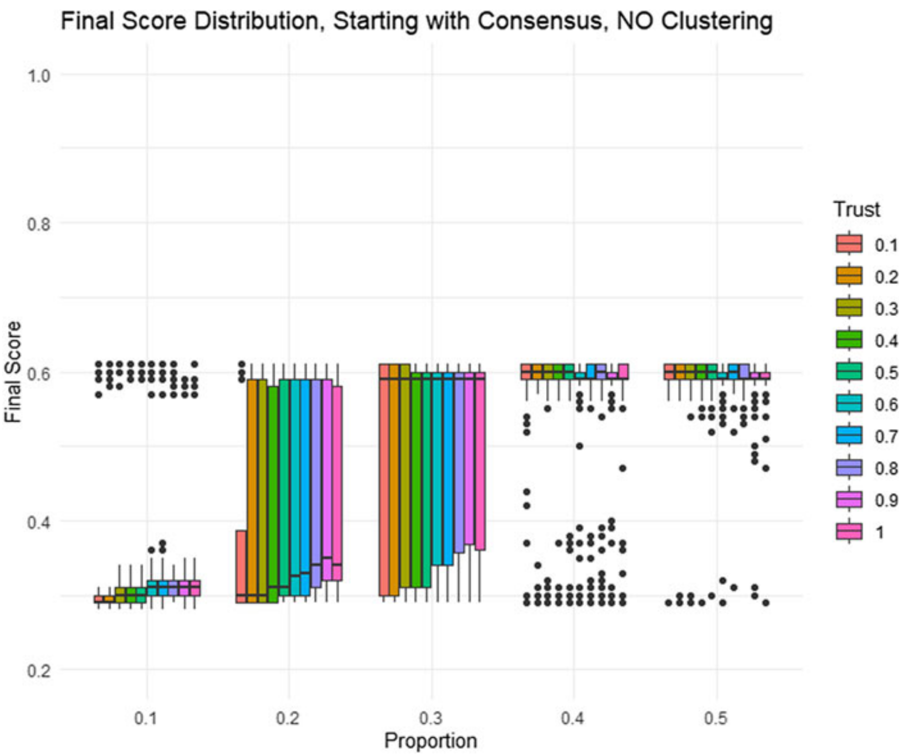

Figure 6. The uniform distribution of two groups, when the smaller group has a much better answer than the larger one. It is evident that the point of convergence changes with size (proportion) and level of trust.

consensus. Figure 6 depicts the collective performance without clustering while Figure 7 represents the results for groups that form homogeneous clusters. Comparing the two results helps to identify conditions that make clustering necessary.

As you can see in Figure 6, without clustering (with a uniform spread of the two groups or with maximal intermixing), when the minority group composes 0.1 or 0.2 of the population, the chance of progress—converging on the better of the two solutions—is quite meager. When the proportion is 0.1, landing on the better solution is almost an anomaly, even though it can happen occasionally. The higher the trust, the better the average score of the population, but the variation is quite small. With the proportion of 0.2, still on average the odds of converging on the better solution are not high, but it is higher with clustering than without. With the proportion of 0.3, the mean convergence score is almost the same as the superior solution at 0.619, even though it is quite likely that the population ends up with the worse of the two outcomes. As the proportion grows, the likelihood of converging on the better answer grows as well.

When we allow agents to cluster, however, the outcomes are quite different, especially with greater asymmetry in size between the groups. As Figure 7 indicates,

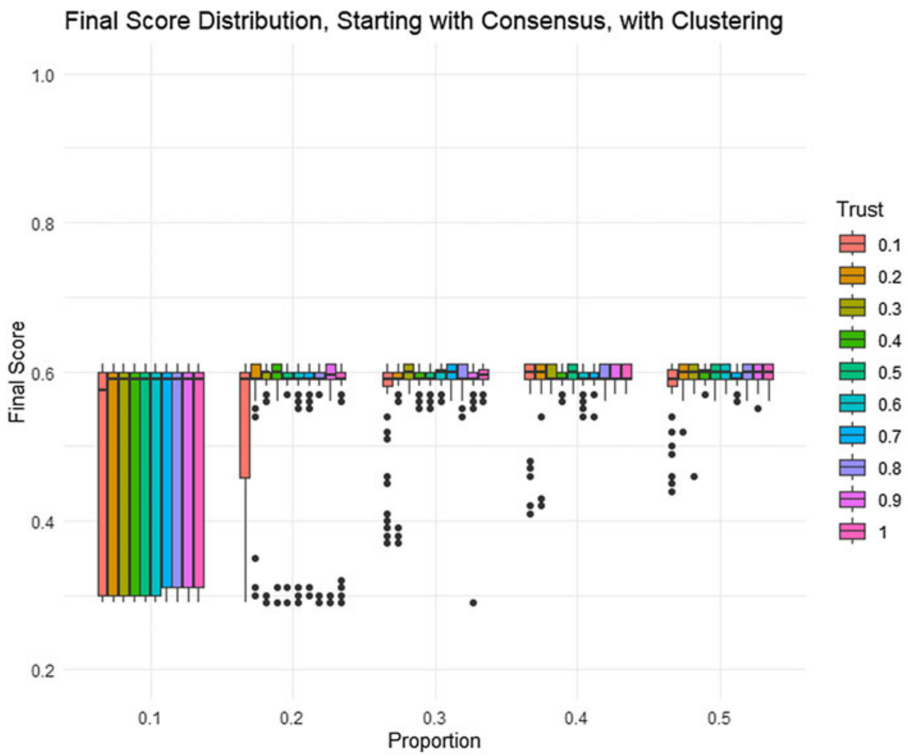

Figure 7. With homogeneous clustering, the convergence point of the population moves to the better solution even for smaller proportions like 0.1 and 0.2.

clustering can completely flip the outcomes for the smallest proportion, 0.1, and change the average convergence score from 0.290 to 0.619. The same drastic change comes with clustering when the proportion grows to 0.2, with some variation in *trust*. In other words, clustering counteracts the extent of asymmetry in size and gives the minority group a chance. When the proportion is 0.3, the effect of homogeneous clustering is still salient, even though not as strong. As the proportion grows, this effect becomes less significant but is still in the positive direction.

5.2.2. Conformity and the starting position

Figure 8 represents the same results as Figure 6 but without conformity pressure. Unsurprisingly, both the problem and the need for clustering is far less salient. The population has a reasonable chance of converging on the better solution without the aid of clustering. This result confirms that conformity is the most significant force restricting the spread of useful information, which clustering can effectively resolve.

As the results of my model show, having consensus for the minority groups is as important as clustering. In other words, clustering of like-minded people, not just clustering of those with the same social identity, is the key to progress. Figure 9 represents this finding by showing what happens to a clustered network that resist majority consensus to start from a random distribution of solutions instead of a

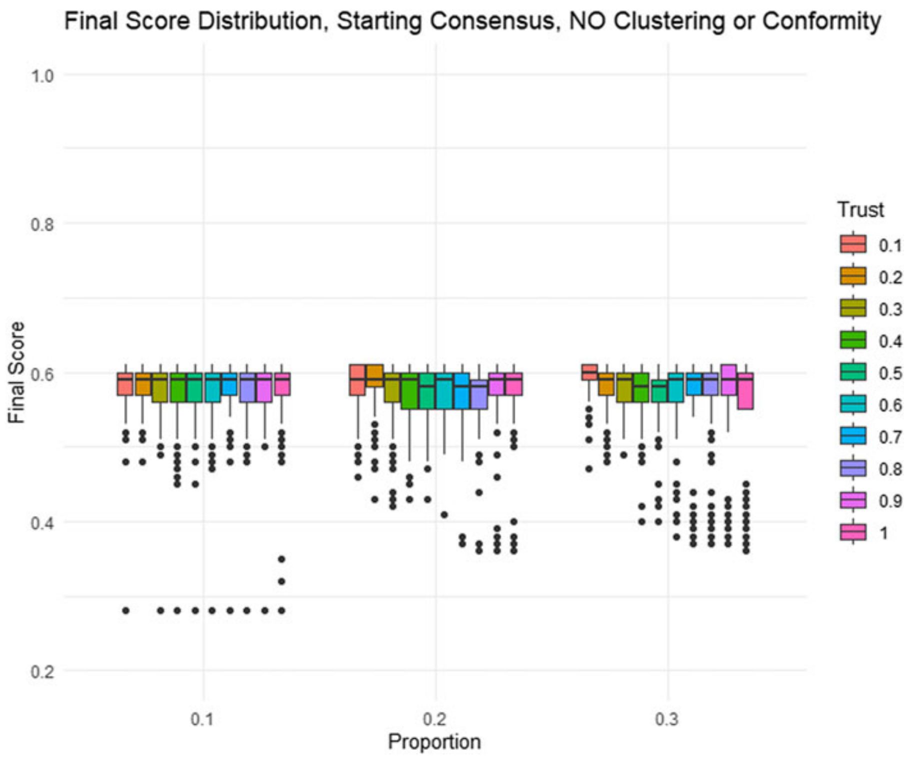

Figure 8. No clustering and no conformity pressure.

consensus. That might explain the extra conformity pressure some minority groups feel when they find themselves in opposition to a majority and its consensus.

5.2.3. The exploration stage

Finally, is there a difference between generating consensus and changing it? More specifically, would our recommendations based on modeling periods of exploration be useful for altering consensus after it has formed? To address these questions, we can compare the effects of clustering on a population that has already reached a consensus with one still exploring possibilities. In Figure 10, the positive impact of clustering on a majority with consensus can be seen by comparing graphs (a) and (b). The difference is striking, showing a reliable and drastic improvement in the likelihood of progress. Conversely, a comparison between graphs (c) and (d) shows that clustering during periods of exploration, or when the majority is still considering possible solutions, has a negligible impact. Therefore, generating consensus and changing it are not the same, or at least clustering as a strategic intervention does not have the same effect in these two situations. Moreover, the insights from the models regarding the value of interventions such as clustering during exploration periods do not necessarily apply once a consensus has been formed.

Figure 9. Even with clustering, when the minority lacks consensus, the odds for progress are low.

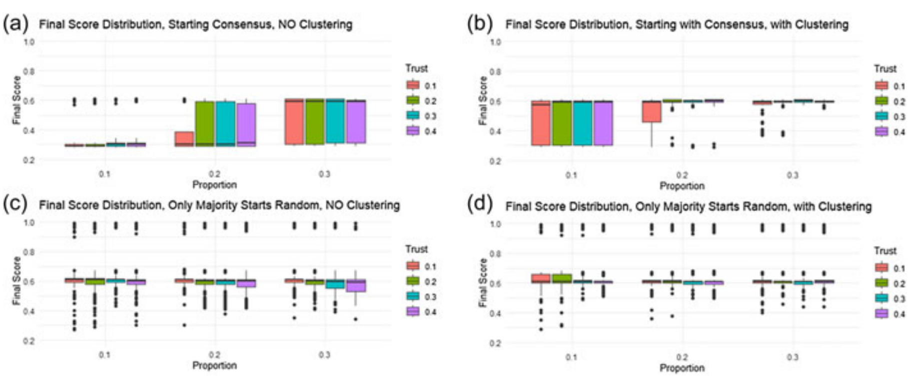

Figure 10. This compares the effect of clustering in two scenarios: when the majority has already decided (graphs (a) and (b)) and when the majority is still exploring (graphs (c) and (d)). The comparison of (a) and (b) highlights clustering's positive impact on progress with majority consensus. In contrast, (c) and (d) show that clustering is unnecessary and ineffective during the exploration period.

6. Conclusion

Is it ever beneficial for the epistemic health of a community for similar types of people to form clusters? A network-level or structural analysis allows us to examine this question closely. Recent research on homophily—the tendency of individuals to associate with similar others—provides valuable insights. For instance, Rubin and O'Connor (2018) investigate the sparsity created by homophily in interaction networks in bargaining games. Fazelpour and Rubin (2022) extend this idea to collective problem-solving, revealing that homophilic networks are beneficial in some dimensions but not others. Additionally, Gomez and Lazer (2019) explore clustering without altering network sparsity, showing that "clustering knowledge and dispersing abilities enhance collective problem-solving in a network" (p. 1).

However, more needs to be said about how individuals can and do strategically leverage their social connections to alter the outcome of their group's collective performance. Recent literature on adversarial collaboration and its policy implications for "ideological diversity" represents one instance of such strategic intervention in social networks. Another example is the more organic, grassroots clustering of like-minded individuals, as seen in the variolation example I discussed in this paper. The common thread in these examples is the clustering of those who share some beliefs/solutions in conjunction, which also defines the boundaries of group dynamics. While adversarial collaboration aims to dissolve ideologically homogeneous clusters, the variolation example shows how homogeneity can be advantageous.

Some argue that ideological or "perspectival diversity" functions exactly as Hong and Page (2004) describe diversity in heuristics or abilities. If this is so, dissolving homogeneous clusters would indeed be the best course of action to increase a collective's chances of progress. However, I have shown that similar dynamics can be modeled differently. To provide an alternative model, I deviated from the literature in two respects: first, by introducing asymmetry, highlighting that context and the stage of a problem matter, to show that dissolving homogeneous clusters is not always the right intervention; second, by revising our concept of progress and the methods by which we operationalize it.

According to my revised model, creating homogeneous clusters is sometimes necessary for a population to have a chance at improvement. Systematically limiting this capacity, on the other hand, can be a tool to maintain the status quo. In other words, adversarial collaboration can be effective when both sides have equal power and influence within the scientific community or general public. Without this balance, or when the opposition challenges well-established beliefs with millennia-old histories, compulsory collaboration among those with opposing views may merely ensure that change does not occur.

Acknowledgment. I am deeply grateful to Kevin Zollman, Daniel Singer, and Matthew Willis for their invaluable feedback and constructive suggestions on earlier drafts of this paper. I also wish to thank the participants of the Computational Social Philosophy Seminars (CSPS), including Dunja Šešelja, Hannah Rubin, Sina Fazelpour, Rory Smead, and Bert Baumgaertner, for their insightful discussions and comments. Additionally, I extend my gratitude to the Buffalo PPE seminar participants and organizers—Ryan Muldoon, Justin Bruner, Kirun Sankaran—and all the graduate students for their thoughtful engagement and perspectives that helped refine my ideas. Finally, I am grateful to the anonymous reviewers whose rigorous and thoughtful comments greatly contributed to improving the paper.

Funding and Declarations. None to declare.

References

Axelrod, Robert, Joshua J. Daymude, and Stephanie Forrest. 2021. "Preventing Extreme Polarization of Political Attitudes." *Proceedings of the National Academy of Science* 118 (50): e2102139118. https://doi.org/10.1073/pnas.2102139118.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbors." *Review of Economics Studies* 65 (3):595–621. https://doi.org/10.1111/1467-937X.00059.

Barnes, Diana. 2012. "The Public Life of a Woman of Wit and Quality: Lady Mary Wortley Montagu and the Vogue for Smallpox Inoculation." *Feminist Studies* 38 (2):330–62.

Cialdini, Robert B., and Noah J. Goldstein. 2004. "Social Influence: Compliance and Conformity." *Annual Review of Psychology* 55:591–621. https://doi.org/10.1146/annurev.psych.55.090902.142015.

Clark, Cory J., Thomas Costello, Gregory Mitchell, and Philip E. Tetlock. 2022. "Keep Your Enemies Close: Adversarial Collaborations Will Improve Behavioral Science." *Journal of Applied Research in Memory and Cognition* 11 (1):1–18. https://doi.org/10.1037/mac0000004.

Crutchfield, Leslie. 2018. *How Change Happens: Why Some Social Movements Succeed While Others Don't*. Chichester: Wiley.

Dezso, Christian L, and David Gaddis Ross. 2012. "Does Female Representation in Top Management Improve Firm Performance? A Panel Data Investigation." *Strategic Management Journal* 33 (9):1072–89. https://doi.org/10.1002/smj.1955.

Fazelpour, Sina, and Hannah Rubin. 2022. "Diversity and Homophily in Social Networks." In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

Fazelpour, Sina, and Daniel Steel. 2022. "Diversity, Trust, and Conformity: A Simulation Study." *Philosophy of Science* 89 (2):209–31. https://doi.org/10.1017/psa.2021.25.

Fraser, Nancy. 2021. "Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy." In *Public Space Reader*, edited by Miodrag Mitrašinović and Vikas Mehta, 34–41. Abingdon: Routledge.

Freeman, Richard B., and Wei Huang. 2014. "Collaboration: Strength in Diversity." *Nature* 513:305. https://doi.org/10.1038/513305a.

Gaither, Sarah E., Evan P. Apfelbaum, Hillary J. Birnbaum, Laura G. Babbitt, and Samuel R. Sommers. 2018. "Mere Membership in Racially Diverse Groups Reduces Conformity." *Social Psychology and Personality Science* 9 (4):402–10. https://doi.org/10.1177/1948550617708013.

Gaus, Gerald. 2021. *The Open Society and Its Complexities*. Oxford: Oxford University Press.

Glynn, Ian, and Jenifer Glynn. 2005. *The Life and Death of Smallpox*. London: Profile Books.

Gomez, Charles J., and David M. Lazer. 2019. "Clustering Knowledge and Dispersing Abilities Enhances Collective Problem Solving in a Network." *Nature Communications* 10:5146.

Grim, Patrick, Daniel J. Singer, Aaron Bramson, Bennet Holman, Sean McGeehan, and William J. Berger. 2019. "Diversity, Ability, and Expertise in Epistemic Communities." *Philosophy of Science* 86 (1):98–123. https://doi.org/10.1086/701070.

Heydari Fard, Sahar. 2024a. "Diversity, Polarization, and Dynamic Structures: A Structural Turn in Social Contract Theory." In *New Approaches to Social Contract Theory: Liberty, Equality, and the Open Society*, edited by John Thrasher and Michael Moehler, 101. Oxford: Oxford University Press.

Heydari Fard, Sahar. 2024b. "The Transformative Power of Social Movements." *Philosophy Compass* 19 (1): e12951. https://doi.org/10.1111/phc3.12951.

Hong, Lu, and Scott E. Page. 2001. "Problem Solving by Heterogeneous Agents." *Journal of Economic Theory* 97 (1):123–63. https://doi.org/10.1006/jeth.2000.2709.

Hong, Lu, and Scott E. Page. 2004. "Groups of Diverse Problem Solvers can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Science* 101 (46):16385–89. https://doi.org/10.1073/pnas.0403723101.

Huang, Alice C. 2024. "Landscapes and Bandits: A Unified Model of Functional and Demographic Diversity." *Philosophy of Science* 91 (3):579–94.

Kahneman, Daniel. 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review* 93 (5):1449–75. https://doi.org/10.1257/000282803322655392.

Kahneman, Daniel. 2012. "The Human Side of Decision Making." *Journal of Investment Consulting* 13 (1):5–14.

Kitcher, Philip. 1990. "The Division of Cognitive Labor." *The Journal of Philosophy* 87 (1):5–22. https://doi.org/10.2307/2026796.

Landemore, Helen, and Scott E. Page. 2015. "Deliberation and Disagreement: Problem Solving Prediction, and Positive Dissensus." *Politics, Philosophy & Economics* 14 (3):229–54. https://doi.org/10.1177/1470594X14544284.

Lazer, David, and Allan Friedman. 2007. "The Network Structure of Exploration and Exploitation." *Administrative Science Quarterly* 52 (4):667–94. https://doi.org/10.2189/asqu.52.4.667.

Levine, Mark, Clare Cassidy, Gemma Brazier, and Stephen Reicher. 2002. "Self-Categorization and Bystander Non-intervention: Two Experimental Studies." *Journal of Applied Social Psychology* 32 (7):1452–63. https://doi.org/10.1111/j.1559-1816.2002.tb01446.x.

Martel, Cameron, Steve Rathje, Cory J. Clark, and Gordon Pennycook. 2024. "On the Efficacy of Accuracy Prompts Across Partisan Lines: An Adversarial Collaboration." *Psychological Science* 35 (4):435–50.

Mason, Winter, and Duncan J. Watts. 2012. "Collaborative Learning in Networks." *Proceedings of the National Academy of Science* 109 (3):764–69. https://doi.org/10.1073/pnas.1110069108.

Mellers, Barbara, Ralph Hertwig, and Daniel Kahneman. 2001. "Do Frequency Representations Eliminate Conjunction Effects? An Exercise in Adversarial Collaboration." *Psychological Science* 12 (4):269–75. https://doi.org/10.1111/1467-9280.00350.

Muldoon, Ryan. 2016. *Social Contract Theory for a Diverse World: Beyond Tolerance*. Abingdon: Routledge.

Muldoon, Ryan. 2018. "The Paradox of Diversity." *Georgetown Journal of Law & Public Policy* 16:807.

O'Connor, Cailin, and James Owen Weatherall. 2019. *The Misinformation Age: How False Beliefs Spread*. New Haven CT: Yale University Press.

Philips, Katherine W., and Evan P. Apfelbaum. 2012. "Delusions of Homogeneity? Reinterpreting the Effects of Group Diversity." In *Looking Back, Moving Forward: A Review of Group and Team-Based*, edited by Margaret A. Neale and Elizabeth A. Mannix, 185–207. Bingley: Emerald Publishing Limited.

Richard, Orlando, Alanson McMillan, Ken Chadwick, and Sean Dwyer. 2003. "Employing an Innovation Strategy in Racially Diverse Workforces: Effects on FirmPperformance." *Group Organizational Management* 28 (1):107–26. https://doi.org/10.1177/1059601102250022.

Rubin, Hannah, and Cailin O'Connor. 2018. "Discrimination and Collaboration in Science." *Philosophy of Science* 85 (3):380–402. https://doi.org/10.1086/697744.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69 (1):99–118. https://doi.org/10.2307/1884852.

Singer, Daniel J. 2019. "Diversity, Not Randomness, Trumps Ability." *Philosophy of Science* 86 (1):178–91.

Sommers, Samuel R. 2006. "On Racial Diversity and Group Decision Making: Identifying Multiple Effects of Racial Composition on Jury Deliberations." *Journal of Personality and Social Psychology* 90 (4):597–612. https://doi.org/10.1037/0022-3514.90.4.597.

Steel, Daniel, and Naseeb Bolduc. 2020. "A Closer Look at the Business Case for Diversity: The Tangled Web of Equity and Epistemic Benefits." *Philosophy of The Social Sciences* 50 (5):418–43. https://doi.org/10.1177/0048393120911130.

Thompson, Abigail. 2014. "Does Diversity Trump Ability?" *Notices of the American Mathematical Society* 61 (9):1024–30.

Tilly, Charles. 1998. "Social Movements and (All Sorts of) Other Political Interactions—Local, National, and International—Including Identities." *Theory and Society* 27:453–80.

Vallier, Kevin, and Ryan Muldoon. 2021. "In Public Reason, Diversity Trumps Coherence." *Journal of Political Philosophy* 29 (2):211–30. https://doi.org/10.1111/jopp.12227.

Weisberg, Michael, and Ryan Muldoon. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76 (2):225–52. https://doi.org/10.1086/644786.

Willett, Jo. 2021. *The Pioneering Life of Mary Wortley Montagu: Scientist and Feminist*. Barnsley: Pen & Sword.

Wu, Jingyi. 2023. "Epistemic Advantage on the Margin: A Network Standpoint Epistemology." *Philosophy and Phenomenological Research* 106 (3):755–77. https://doi.org/10.1111/phpr.12895.

Zollman, Kevin J. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5):574–87. https://doi.org/10.1086/525605.

Zollman, Kevin J. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72 (1):17–35. https://doi.org/10.1007/s10670-009-9194-6.

Cite this article: Heydari Fard, Sahar. 2025. "Challenging the Consensus: The Strategic Value of Homogeneous Groups in Collective Problem Solving." *Philosophy of Science*. https://doi.org/10.1017/psa.2025.20