Temporary disagreements foster better solutions: How homophilic interactions in diverse teams can improve collective decision-making

Identity diversity in teams brings advantages for complex decision-making because it is associated with cognitive diversity among team members. At the same time, homophilic interactions along shared identity dimensions can hinder information exchange among dissimilar individuals and threaten successful exploitation of the team's cognitive diversity. We present an agent-based model to investigate how homophily impacts decision-making quality in diverse teams. Team members communicate information in a 'hidden profile' setting where some pieces of information are known only to single individuals while other pieces of information are known to subgroups with the same identity. While intuition may suggest that homophily impairs collective decision-making, our model reveals how homophilous environments lead to better collective decisions: homophily fosters temporary disagreements between dissimilar team members, which grant teams additional time to uncover crucial information that would not have been shared otherwise. Longer discussion time comes along with improvements in the quality of the final decision, indicating a trade-off between the time needed to deliberate and decision quality.

Keywords: homophily, collective decision-making, team diversity, hidden profile

Authors

Jonas Stein[+*] Vincenz Frey[*] Andreas Flache[*]

[+] j.d.stein@rug.nl

[*] University of Groningen
Faculty of Behavioral and Social Sciences
Department of Sociology
Grote Rozenstraat 31, 9712 TG Groningen, The Netherlands

Competing interests: The authors declare no competing interests.

Funding: This study is part of the research program Sustainable Cooperation – Roadmaps to Resilient Societies (SCOOP). The authors are grateful to the Netherlands Organization for Scientific Research (NWO) and the Dutch Ministry of Education, Culture and Science (OCW) for generously funding this research in the context of its 2017 Gravitation Program (grant number 024.003.025).

Supporting Materials: The full NetLogo code is stored on the CoMSES Computational Model Library under the following URL: https://www.comses.net/codebase-release/b4c24979-7e9c-42f3-b75e-c6f9befc9132/. The analyzed simulation output data and Stata syntax can be found on the Open Science Framework: https://osf.io/76hfm/

1. Introduction

A large literature on work teams documents that diverse teams have a greater pool of social, human, and cultural capital, translating into a higher potential for team performance. Yet, when diversity activates social identity processes (*1*), this potential may not be used. Homophily, the tendency to preferentially interact with those similar to oneself is a strong force in humans (*2*) and prevents team members from communicating with dissimilar others providing them with information needed to reach higher performance. With an agent-based model, we challenge here the intuition that homophily is detrimental to the performance of diverse teams. We demonstrate that homophily can improve team decision making, studying agent teams confronted with a hidden profile task (*3*) which requires team members to share information not known to others in order to collectively find the best solution to a problem. Our finding contrasts with the conjecture that homophily can hamper the performance of diverse teams (*4–6*) and highlights instead the benefits of limiting the flow of information between dissimilar team members.

Structuring collaboration processes and interactions patterns in diverse teams so that they enhance team decision-making has become an increasingly important issue, as a globalized division of labor, rising international migration, and increasingly diverse workforces have led to an ubiquity of heterogenous decision-making groups in organizations (*7*). Research investigating how this trend shapes collaborative work processes identifies both challenges and benefits (*8–10*). Successful integration of *cognitive diversity* referring to the wealth of perspectives, knowledge, and skills present in a team is found to have mostly beneficial outcomes for the quality of a task, especially when tasks are complex because a conjunction of diverse skills and perspectives is expected to enhance team creativity and foster innovative solutions (*11*).

At the same time, *identity diversity* in teams can provide a challenge to the successful integration of cognitive diversity. Identity diversity is sometimes also referred to as 'surface-level diversity' or 'demographic diversity' (*12*). Including potential demographic or 'surface-level' traits, we focus specifically on identities that are easily observable for others, salient during collaborative work processes, and plausibly correlated with cognitive traits. Even without assuming that identity diversity leads to intergroup conflict (*13*), stereotyping, and negative outgroup attitudes (*14–16*), it has been consistently documented that individuals tend to associate themselves with similar others and that similarities are usually recognized along common identities (*2*, *4*, *17*).

Simulation studies on opinion dynamics (*19*, *20*) and experimental studies (*21–23*) show that such *homophilous preferences* are sufficient to drive groups apart and induce polarization. Empirically, identity traits often correlate with cognitive traits (*16*) and theoretical as well as empirical studies have shown that such correlations tend to amplify polarization tendencies (*24–26*). If individuals socially influence each other but tend to interact with similar others, patterns can emerge where distinct sets of opinions revolve around similarities in other, seemingly unrelated dimensions.

Models of opinion dynamics highlight how opinion divergence in teams can disable consensus, but they do not clarify how opinion polarization links to the quality of decision-making in a team. In this paper we move beyond modeling the dynamics of opinions alone and develop expectations about how homophilous preferences and social influence affects decision-making *quality* in a team. Intuitively, one can expect this theoretical extension to show how homophily in diverse teams can negatively affect decision-making quality. For complex tasks, where diverse knowledge must be brought together to obtain an optimal decision, a lack of communication between team members with different identities hampers the conjunction of valuable information, leading to failure to realize good solutions. Second, even if some team members find the optimal solution to the task at hand, a lack of consensus endangers the possibility that this solution is adopted by the team. Finally, lacking communication slows down the deliberation process, making decision-making more costly.

Epistemological studies on decision-making in diverse groups, on the other hand, oppose the expectations suggested by the opinion dynamics literature and argue that bounded communication between different individuals is instead beneficial to decision-making quality (*27*). Restrictions on social influence will prevent individuals from prematurely adopting others' solutions, which can result in deciding for superior options on a group level. In other words, boundaries in communication hinder the rapid dissemination of inferior knowledge, ultimately ensuring that individuals explore the full spectrum of possible decisions before exploiting suboptimal knowledge (*27*–*29*). Similar to the opinion dynamics literature, these works also suggest that 'transient diversity' will lengthen the deliberation process but point out instead that added discussion time gives room to ensure that knowledge is optimally explored and disseminated.

Studies of transient diversity emphasize possible advantages of boundaries to communication in team decision making, but they leave us in the dark as to how homophily interacts with diversity in teams. Scholars of this canon are primarily concerned with communication processes in science and argue that skepticism and sparse communication can be induced by adapting macro-level incentive structures such as changing funding policies in research. Yet, we would expect that meso-level social processes such as homophily induced by identity diversity can also bound excessive communication by limiting exchange between dissimilar members, even if macro-level structural boundaries are absent. This notion is similar to the 'value in diversity' hypothesis (*30*), arguing that salient markers of diversity can be beneficial to decision-making quality even when they are unrelated to cognitive traits (*11*, *12*, *31*, *32*). Easily observable diversity in identities can help groups to apply healthy skepticism, prevent the placement of undue trust, and foster constructive discussion. It follows from both the transient diversity literature and the value in diversity literature that tendencies to associate with similar over dissimilar individuals improve collective decision-making quality by helping groups to examine information critically instead of converging around early, suboptimal consensus.

Research on opinion dynamics in diverse teams and studies of transient diversity lead to competing intuitions: preferential interactions among similar over dissimilar team members are either beneficial or detrimental to a team's performance. The present paper uses an agent-based model to theorize how homophilous interaction preferences shape decision-making quality in diverse teams. The model combines central aspects that have not been studied in tandem before: first, it evaluates the quality of the decision that is made, which opinion dynamics models have paid little attention to so far. Second, it considers that background traits shape interaction preferences between individuals without having to assume exogenous incentive structures as outlined by the transient diversity literature.

Our model builds on hidden profile tasks, an established paradigm that has been widely used in experimental research to study decision-making in groups (*3*, *33*–*37*). In a hidden profile, a team of decision-makers is equipped with a set of information pieces and instructed to deliberate before choosing one of several available decision alternatives, which differ in quality. Individuals are given different pieces of information at the onset of the deliberation task. Pieces of knowledge that point towards inferior options are 'common information', i.e., known to everyone in the group. Common information anchors decision-makers to initially prefer inferior options. Anchoring effects, a tendency to share knowledge supporting one's own views, and social validation from others with similar decision preferences can prevent members from sharing or accepting dissenting information. This makes hidden profiles difficult to solve (*33*), which resonates with the conjecture that diversity has a much greater impact on the outcome of a task when it is complex and challenging (*38*). Information supporting the optimal option is 'unique', in that it is known to not more than one individual. However, a conjunction of multiple pieces of unique information will reveal the optimal option, which captures the well-studied phenomenon of cognitive diversity that bringing together knowledge from different individuals will produce better solutions (*39*). In addition, hidden profiles are a suitable paradigm for the purpose of the present research because they incorporate a number of features that are often hard to observe or hold constant in natural settings: Hidden profiles allow the experimenter to control the distribution of knowledge, as a predefined set of decision options can be perfectly ranked according to their quality, and all communication processes and their outcomes can be observed and subsequently analyzed.

In addition, our model extends the traditional hidden profile framework by assigning identity traits to individuals, making them either similar or dissimilar to each other. We capture the aspect that identity is associated with the kind of knowledge individuals possess by distributing a separate set of common information to any group of individuals with a common identity. We further condition interaction preferences on identity traits in such a way that higher homophily levels reflect how much identity-similarity increases the chances of communication between individuals. Our key interest is in assessing the effects of homophily on decision-making quality. Thus, as further outlined in the 'setup of simulation experiments' section, we measure how likely teams were to obtain optimal consensus given different levels of homophily, and how long it took them to reach a decision. In the results section, we show how homophily affects consensus outcomes, uncovering the underlying mechanism and investigating how other discussion features such as deliberation length and belief changes among team members are affected.

2. Model description

Our model develops a formal representation of a diverse work team facing a decision problem as implemented in the experimental setup of the hidden-profile paradigm. We implement a setup where a team seeks to identify the best out of a set of possible decision options. Individuals are equipped with different pieces of information that need to be combined to identify the best option. To this end, we assume a team of N agents. Each agent belongs to one of M groups where each group consists of agents who share a common identity. Identities could represent, for example, different branches in an organization, or different academic disciplines in an interdisciplinary project. For simplicity, we assume throughout that groups are assumed to be of equal size, i.e., we do not consider unbalanced group sizes.

The virtual teams in our model face a decision problem, in that the best option o_{max} out of a set of J discrete options needs to be identified. Every team member forms her own belief about which decision option is best but is open to influence by other team members. Influence is implemented as a sequence of communication events. Agents take turns in sharing an argument with an interaction partner. Every time an argument is emitted, the recipient updates her beliefs and tells her team what option she currently believes to be best. This influence process continues until all agents prefer the same option. This option is the team´s decision. Alternatively, if no consensus is reached after a large number of interaction events (5,000 interactions), the simulation is stopped.

2.1 Decision options and arguments

To create a decision problem as implemented in the hidden profile paradigm, we assume that there is a set of I arguments $A = \{a_1, a_2, \ldots, a_I\}$ pertaining to a predefined set of J decision options $O = \{o_1, o_2, \ldots, o_J\}$. An example of a set of arguments and decision options is presented in Figure 1A. The set of arguments is fixed, implementing a setting in which agents cannot invent new arguments during the deliberation process. Each argument contains J weights reflecting the degree to which the argument supports the different decision options, i.e. $a_i = \{w_{i,1}, w_{i,2}, \ldots, w_{i,J}\}$. Following the standard assumption of the hidden profile literature, we assume that each argument has a positive weight for only one of the options and a weight of zero for the other options. We further assume that there is an equal number of arguments with a positive weight for each decision option. For a given simulated team, weights are randomly drawn from a uniform distribution so that $w_{i,j} \in [0,1]$. The sum of the weights associated with a decision option determines its true quality $Q_j = \sum_1^i w_{i,j}$, and the decision option j with the highest quality is the optimal option, o_{max}.

Figure 1. Arguments, decision options and option quality. **Panel A:** example set of available arguments and decision options. **Panel B:** argument distribution across groups and agents. Highlighted arguments represent common arguments within groups. Quality scores in red indicate the optimal option in panel A and agents' beliefs about what option is optimal in panel B.

(A) Argument availability

arguments	decision options o_1	o_2	o_3
a_1	.76	0	0
a_2	.17	0	…*
a_3		.63	
a_4		.39	
a_5			.64
a_6			.55
Q_j	.93	1.02	1.19

(B) Argument distribution in the team

Group g_1

agent 1	o_1	o_2	o_3
a_1	.76		
a_4		.39	
$q_{x,j}$	.76	.39	0

agent 2	o_1	o_2	o_3
a_1	.76		
a_5			.64
$q_{x,j}$	.76	0	.64

Group g_2

agent 3	o_1	o_2	o_3
a_3		.63	
a_2	.17		
$q_{x,j}$	.17	.63	0

agent 4	o_1	o_2	o_3
a_3		.63	
a_6			.55
$q_{x,j}$	0	.63	.55

* Empty cells represent weights with a value of zero.

2.2 Unique and common arguments

The last ingredient needed to implement a hidden profile task is concerned with the initial distribution of arguments across team members. In hidden profile experiments, participants are provided with so-called 'common' and 'unique' information. An argument is common when all participants know it already at the outset of the deliberation process. Unique information, in contrast, is provided only to a single participant. This principle is modified for the study of diverse teams: We added that a given piece of common information is provided only to the members of one group (Figure 1B). In this sense, common arguments can reflect certain disciplinary basics that everyone of the same profession was trained with, or specific skills shared by everyone working in the same organizational branch. We distribute common arguments in such a way that each agent of group g_m receives the same set of arguments in favor of option o_j. Agents in group g_1 receive C arguments favoring option o_1, agents in group g_2 receive C arguments favoring option o_2, and so on. Like scientists of one discipline thinking that their approach is superior to others or employees being convinced that the 'way things are done' within their organizational branch is best, the distribution of common arguments biases agents of one group to initially favor a specific, but not necessarily optimal option over others. We assume that there are more decision options than groups ($J > M$) so that the optimal option can reside outside of those options supported by common arguments.

Once common arguments are assigned to groups, we consider all arguments that remain and let agents take turns at randomly drawing from those arguments without replacement until all arguments are distributed. By doing so, these arguments represent 'unique information' that is held by single agents but not by groups. In line with the hidden profile paradigm, we select for our simulation experiments those tasks where none of the common arguments in either of the groups favor the optimal option. Altogether, our initialization procedure creates the situation of most theoretical interest to us in this study: common arguments ensure cognitive diversity on the group level and unique arguments provide cognitive diversity at the individual level. The latter are dispersed over agents across groups and must be brought together to outweigh common stocks of arguments pointing to inferior solutions.

2.3 Argument processing and communication

Similar to how objective quality scores Q_j are computed, agents form a perceived quality score for each decision option, $q_{x,j}$, by summing over the weights of the arguments they possess. Agents always believe the decision option to be best that has the highest perceived quality to them and communicate this belief publicly to everyone at the onset of the simulation.

Over the course of the simulation, agents share arguments and update beliefs, thereby deliberating which option is best. Agents are activated sequentially and according to their identifier. The first task of the active agent is to select a decision option o_j^* she wants to support. Psychological research suggests that individuals are most inclined to advocate options they deem most preferable themselves (*38*). For this reason, we assume that the agent is more likely to choose options with higher perceived quality scores relative to the quality score of other options. The probability to choose option *j* at a given moment is formalized by equation 1.

$$p_{o_j} = \frac{e^{\beta * q_{x,j}}}{\sum_{j=1}^{J} e^{\beta * q_{x,j}}} \quad (1)$$

The parameter β reflects agents' adherence to choosing an option of higher versus lower perceived quality. When $\beta \to \infty$, the probability of choosing the option with highest perceived quality approximates one while probabilities of choosing other options are zero. When β is one, an option is chosen with a probability proportional to its perceived quality. When β is zero, all options are chosen with equal probability, regardless of their perceived quality.

After an agent has decided which option to support, the second task of the agent is to determine which argument to communicate. Here, an agent regards all arguments she holds but only considers those weights w_{i,j^*} that contain information on her chosen option o_j^*. She picks one of her arguments with the probability given by Equation 2:

$$p_{a_i} = \frac{e^{\beta * w_{i,j^*}}}{\sum_{i=1}^{I} e^{\beta * w_{i,j^*}}} \quad (2)$$

Again, the β parameter determines agents' adherence to choosing stronger versus weaker arguments pertaining to her chosen option. As long as β does not approach large numbers, the discrete choice equation assigns all arguments a positive probability of being chosen, including those with a weight of zero. For simplicity, we assume that the value of beta in Equation 1 and Equation 2 is the same, representing a general tendency to select arguments that most strongly support the alternative an agent believes to be best, given the information she possesses.

2.4 Homophilous interactions

After an agent chose which argument to communicate, she decides whom to share it with. Because we are interested in the effects of homophilic interaction patterns, we assume that agents share arguments in dyadic encounters in which they preferentially interact with those of identical group membership. Interactions are regulated through a homophily parameter *h* which ranges from zero to one. The greater *h*, the more likely agents are to interact with team members from their own group. Choosing an interaction partner is operationalized as follows: whenever a sending agent *x* becomes active, we define all remaining team members as potential receiving agents $Y = \{y_1, y_2, \ldots, y_{N-1}\}$. Each agent y_k is assigned a similarity value s_k, which takes on the value of $h / 2 + 0.5$ if sending and receiving agent share the same identity and $1 - (h / 2 + 0.5)$ otherwise. Exactly one of the other members of the team is chosen as recipient, where the probability of choosing agent y_k as the receiving agent is given by equation 3.

$$p_{y_k} = \frac{s_k}{\sum_{k=1}^{N-1} s_k} \quad (3)$$

When h = 1, homophily is maximal and the sending agent will always choose a member of her own group. When h = 0, no homophily is present and all agents are chosen with equal probability. Once the receiving agent has been determined, the sending agent shares the argument she has picked before, and the receiving agent updates her set of arguments and beliefs. If the new argument changed the receiving agent's belief about which option is best, she communicates this change immediately with the team. Receiving agents do not forget arguments or value them differently according to recency, frequency of receival, or group membership and beliefs of the sender. The process of activating an agent, sharing an argument, and updating beliefs of the receiver represents one iteration t and is repeated until the team obtains consensus and all agents agree on which option is best. If the team does not reach consensus after a large number of iterations (t = 5000), the simulation is stopped.

3. Setup of simulation experiments

To investigate how homophily affects decision-making quality, we vary homophily while holding all other parameters constant. For each homophily level, we simulate 5,000 teams and observe how often they obtain optimal consensus and how long it takes them to reach consensus. Homophily is varied from h = 0 to 0.9 in steps of 0.1. Between h = 0.9 and 0.98 we vary homophily in steps of 0.02 because at such high levels of homophily we found particularly strong effects on discussion time. We do not consider teams in which no arguments are exchanged across groups at all (h = 1) because they would represent two separate discussions instead of one.

We simulate teams with the following properties. Our teams have N = 6 members, which is a common size among decision-making teams in real-world contexts and hidden profile experiments alike. They are split into M = 2 groups, reflecting a setting where relevant identities proxy a binary. A total of J = 3 decision options is chosen and, given the distribution of common arguments, each group initially supports one of the two inferior options. The optimal option, on the other hand, is only supported by unique arguments. A total of I = 18 arguments is available in a team (i.e., six arguments per decision option given J = 3) and each group starts off with C = 3 common arguments. There are thus 18 arguments – 2 groups * 3 common arguments = 12 unique arguments, which are distributed across the six agents. Hence, each agent initially holds three common and two unique arguments. This distribution of arguments implies a hidden profile: Common arguments point towards inferior options and agents initially believe such options to be optimal. Throughout the main analyses, we examine teams where agents probabilistically select preferred decision options and arguments according to adherence values of β = 3.5. This was found to be a reasonable value for agents to select strong arguments supporting options with higher levels of perceived quality while still allowing for small probabilities of stochastic deviation.

Apart from the analyses reported in the main text, we conduct extensive robustness analyses in which we vary all variable parameters of the model – i.e., the number of I arguments, J decision options, C common arguments, N team members, M groups and adherence values β. An overview of these analyses is accessible at https://osf.io/76hfm/ and reveals remarkably robust effects.

4. Results

We start off by comparing how often simulated teams reach consensus on each of the three decision options of different quality, given different homophily levels. As Figure 2A shows, more teams form consensus on the optimal option as homophily levels increase and agents tend to interact less with team members with a different identity. The percentage of teams reaching consensus on the second-best option, on the other hand, sinks symmetrically with rising fractions of teams reaching optimal consensus. The share of teams which form consensus on the worst decision option is always below 5 percent, irrespective of the level of homophily. Considering this, and the fact that one group initially favors the second-best and the other group favors the worst option, we conclude that suboptimal consensus is most often made when one group convinces the other of the second-best option. To prevent suboptimal consensus on the second-best option, higher homophily levels could thus be helpful against the diffusion of arguments favoring this alternative from one group to the other.

Figure 2: Option choice, belief changes and discussion time by level of homophily

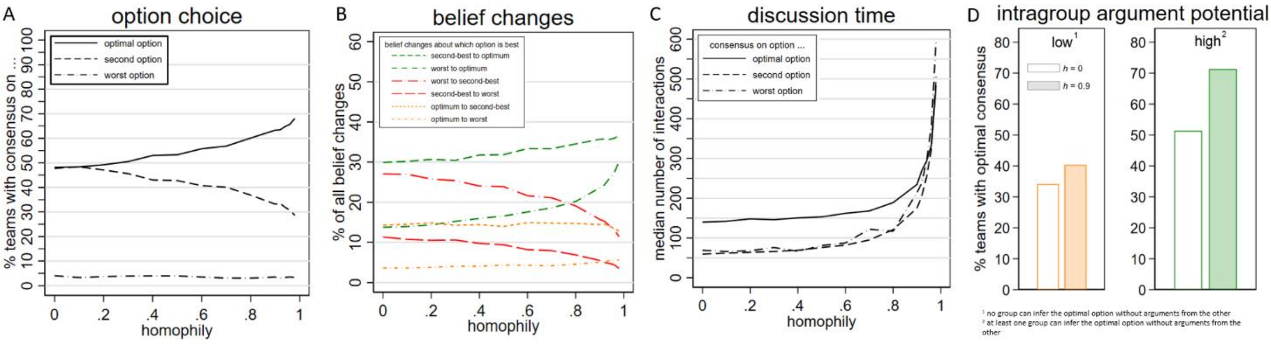

The explanation provided here implies that as homophily increases and interactions between members of different groups become less likely, fewer agents should change their belief towards finding the second-best option optimal. Figure 2B supports this conjecture, showing how the overall proportion of belief changes from the worst option to the second-best option decreases from 28 % to 12 % between $h = 0$ and $h = 0.98$. Similarly, since homophily limits communication between groups both ways, fewer arguments supporting the worst option are being passed over to the group favoring the second-best option, and fewer belief changes towards the worst option occur. Conversely, belief changes from both the second-best and the worst option towards the optimal option become increasingly frequent in higher homophily levels. While changes towards this option are obviously necessary to obtain optimal consensus, they are not easily explained. If homophily hinders the exchange of arguments across groups, including those that point towards the optimal option, why do more team members change their belief towards the optimal option?

4.1 Longer deliberation uncovers optimal arguments

An explanation to this is that because homophily limits argument exchange between groups, disagreement in beliefs across groups is preserved, and neither group can convince the other of their initially preferred option. Hence, discussion continues. Figure 2C illustrates this, showing how higher homophily levels result in higher median discussion time. Prolonged discussions, in turn, enable arguments favoring the optimal option (called 'optimal arguments' hereafter) to be revealed and spread within a respective group. This is so because optimal arguments are unique arguments, which need more time than common arguments to be selected. In comparison to common arguments, unique arguments face a sampling disadvantage and are initially disfavored by agents' argument selection procedure. But because this procedure is stochastic, small probabilities of choosing optimal arguments remain. When exchange between groups is limited and premature consensus kept at bay, optimal arguments are selected and spread within a group, and agents' perceived quality of the optimal option rises. Once all members in one of the groups realized what the optimal option is, they are unlikely to be swayed: optimal arguments tend to have the highest weights and are difficult to surpass by other arguments. Hence, as soon as one group has discovered the optimal option, the danger of a suboptimal consensus is minute, which gives this group ample time to still convince the other group.

This explanation implies that homophily has a much more powerful effect when at least one of the groups has sufficiently strong optimal arguments to identify the best option by themselves and without the help of the other group. Figure 2D supports this conjecture, showing how homophily increases the share of teams with optimal consensus to a great extent (i.e., from 51 % with $h = 0$ to 71 % with $h = 0.9$) when the arguments initially provided to one group are sufficient to infer the optimal option. However, when neither group can infer the optimal option without the other, homophilous interactions still increase the chances of making optimal consensus, but only slightly (34 % to 40 %). The smaller effect is explained by the fact that homophily still prolongs discussion time, making it more likely that optimal unique arguments spread within groups by chance. Altogether, it follows that homophily improves consensus quality mostly because it grants one group with the time to uncover unique arguments and arrive at the optimal option. At the same time, it hinders another group from quickly convincing the team to prefer a suboptimal option through the dissemination of inferior arguments.

4.2 Homophily is crucial when tasks are especially difficult

In Figure 2, we have shown that homophily improves decision-making quality because it prevents teams from prematurely adopting a second-best decision option. To further test this proposition, we investigate if optimal consensus is less likely when the second-best and the optimal option are close to each other in quality, and therefore hard to distinguish, and if the worst option has much lower quality than the second-best option and is therefore likely to be neglected quickly. For each simulation run, we compute a distance score reflecting the difference in quality between the optimal and the second-best option, and the second-best and the worst option. If the mechanism works as we have suggested, a *smaller* quality difference between the second-best and the optimal option should make it easier for the group initially supporting the worst option to get persuaded into the second-best option, thus reducing the chances of reaching optimal consensus. Similarly, a *large* difference in quality between the second-best and the worst option should make it easier for the second-best group to convince the worst group of the second-best option. In both cases, homophily should have a bigger effect because it is needed more to prevent suboptimal consensus on the second-best option.

In Figure 3A, we present the percentage of homophilous versus non-homophilous teams with optimal consensus by their differences in quality between the optimal and the second-best option, and the differences between the worst and the second-best option. As suspected, a smaller difference between the optimal and the second-best option increases the chances of finding optimal consensus. Greater differences between the second-best and the worst option, on the other hand, lead to lower proportions of optimal consensus. Comparing the fraction of populations with optimal consensus under high homophily versus no homophily, the previous finding persists that higher homophily levels render more populations with optimal consensus. This supports our explanation and shows that the positive effects of homophily we observe generalize to a wide set of different combinations of option quality.

However, Figure 3B also indicates that the positive effect of homophily varies among problems with different quality combinations across options. In line with our proposed mechanism, homophily appears to matter especially for those problems where chances to obtain optimal consensus are low to begin with. Here, homophily provides the crucial barrier to the team-wide adoption of second-best arguments that are dangerous precisely because they are either almost as strong as arguments supporting the optimal option, or because arguments pertaining to the worst option are weak in comparison. As becomes evident from the figure, the increase in teams obtaining optimal consensus under high homophily levels is largest when differences between the second-best and the optimal option are small, and differences between the worst and the second-best option are large.

Figure 3: Optimal consensus by quality difference in options* and homophily level

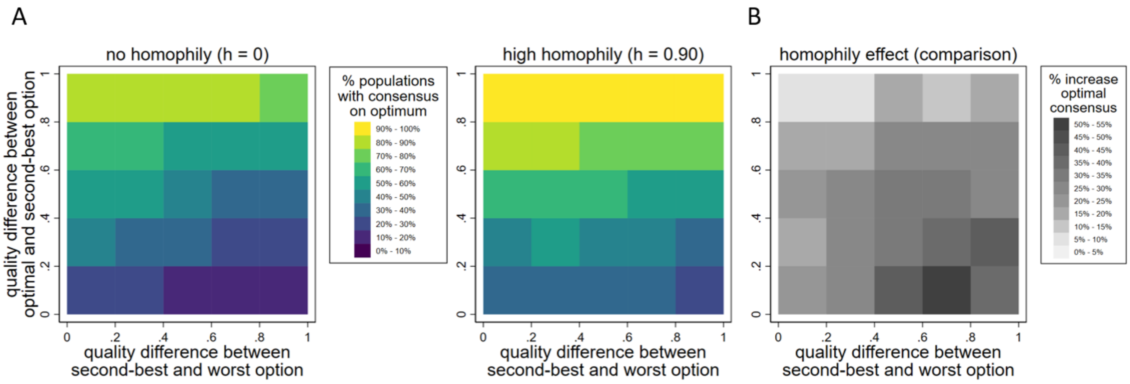

[*] Data grouped by quintiles of the distribution of distances scores to ensure that each cell represents at least 200 observations.

4.3 Homophily facilitates optimal intragroup consensus

The results from Figure 2 suggest that homophily fosters optimal consensus by granting one of the groups the crucial time to uncover their optimal arguments and, in a subsequent step, convince the rest of the team. To further elucidate this mechanism, Figure 4 shows for teams with high versus no homophily the fraction of teams that have reached optimal consensus after a given number of interactions in the whole team (dashed lines) and the first occurrence of optimal consensus within any of the two groups (dash-dotted lines). If the mechanism works as described, we should find that optimal consensus within either group is more frequent under high homophily and occurs much sooner than in teams with no homophily.

A comparison of teams with and without homophily reveals that at any timepoint more teams reach consensus on the optimal option in either group when homophily is high (dash-dotted lines in Figure 4). This finding aligns with previous results that homophily facilitates optimal consensus within a group without influence from the other group. However, teams with high homophily also need more time from the first occurrence of optimal consensus within a group until team-wide optimal consensus is established. The reason behind this is that homophily also slows down the sharing of arguments from the group with consensus on the optimal option to the other group. This contributes to increased discussion time under homophily. For both teams with high homophily and no homophily, most teams that reach optimal consensus in a group also establish optimal consensus on a team level. Some teams, however, fail to achieve optimal consensus despite having reached optimal consensus in a group before. This becomes apparent from the vertical difference between the dash-dotted and the dashed lines towards the right end of the figure. Suboptimal team consensus despite previous optimal intragroup consensus can occur because unless all optimal arguments have been uncovered already, belief changes within team members from the optimal to a suboptimal option can still occur. In sum, however, this tendency is insufficient to offset the mechanism that optimal consensus in a group predates and fosters optimal consensus in the team, both of which occurring more often under homophily.

Figure 4: Discussion time until the first occurrence of optimal consensus in any group and until optimal consensus in the whole team by homophily level

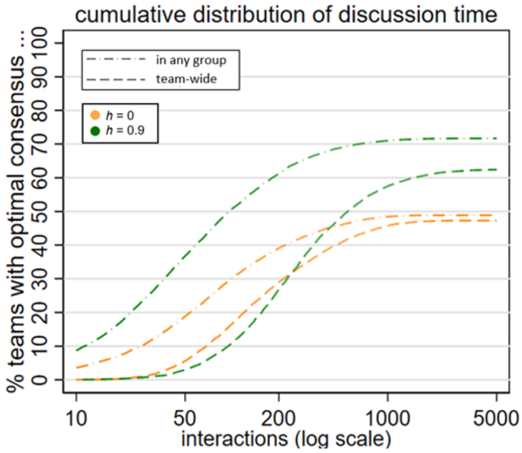

Altogether, our analyses reveal that in hidden profiles, homophilous interactions along common group identities substantially improve collective decision-making quality. Additional robustness analyses reported at https://osf.io/76hfm/ reveal this result to be remarkably robust.

5. Discussion

Our model identifies homophilic interactions in diverse teams as a key factor to the quality of a team consensus in difficult decision-making tasks. This finding resonates with the notion of 'transient diversity' (*27*). However, our investigation extends this notion to the new realm of hidden profile problems and revealed a mechanism that had not been considered by the transient diversity literature before. While transient diversity models find that unbounded communication leads to the insufficient *generation* of diverse information, our hidden profile task showed that even when all information had been created prior to the task, its *distribution* could lead to unfavorable situations in which non-homophilous interactions resulted in suboptimal decisions. Our results thus run contrary to expectations suggested by the opinion dynamics literature – namely, that homophily will undermine team functioning.

6. Future research

Integrating aspects of prominent opinion dynamics models (*19*, *20*, *25*), a possible extension to this model is to condition homophilic encounters on endogenously changing interaction preferences based on beliefs. While this may result in unresolvable disagreements, the necessity to study alternative means for making decisions arises. A common approach in real-world teams is to rely on voting procedures and other aggregation rules when failure to obtain consensus is immanent (*40*). Hence, studying decision quality while assuming voting procedures within the context of this model, or an adapted version thereof, provides a promising avenue for future research.

A feature that is inherent to hidden profiles is that individuals ultimately share the same goal and are likely to agree on one option to be best when faced with complete evidence. While this is applicable to many real-world situations, it abstracts from the possibility that team members of different identities may have group-based interests that make them attach different values to decision options, or even attach value to maintain disagreement with other groups. Such a case would make it necessary to redefine what an 'optimal' solution is and poses interesting distributive and ethical questions. While outside the scope of this paper, an extension of this model could be used as a starting point to investigate whether homophily is helpful in reaching decisions that maximize welfare for the team as a whole versus solutions that optimize payoffs for some groups of team members at the expense of reduced team performance. In a similar vein, future research may consider status distinctions or different group sizes that would introduce inequality in the influence that one group has over another. Here, homophilous interactions could again provide a crucial mechanism to improve deliberation tasks that would otherwise have been dominated by the group with the greatest influence.

For the purpose of this paper, probabilistic encounters between team members represented homophilic interaction *preferences*. However, the same encounters can also be seen as a manifestation of underlying *social foci* that structure team deliberation (*41*). Translating the insights of this paper to such a perspective implies that better decisions will be made in settings where team members are structurally guided to interact with similar over dissimilar others more frequently. This resonates with suggestions made by the transient diversity literature (*27*). Simultaneously, studying the effects of interactions in structurally embedded environments calls for a possible extension of our model in which encounters are not probabilistic but occur along a network that specifies who exchanges information with whom. If the mechanism proposed here holds, networks in which members of different groups are increasingly kept apart should also feature better decisions.

While the simulation results reported here convincingly show how homophily fosters the quality of the team decision, it is important to note that homophilous interactions may have other, unintended consequences. Limiting interactions between members of different identities may amplify social identity processes that can lead to negative outgroup attitudes, lower levels of trust, and less cooperative behavior in general (*8*, *13*). This raises the question whether improved decision-making can be reached through alternative means. Within the context of our model, such means could involve increased skepticism towards information coming from dissimilar members. However, this would imply that spillover effects from increased skepticism resulting in negative outgroup attitudes had to be tempered all the same. Similar to the extension suggested in the paragraph above and in line with simulation research suggesting that the timing of outgroup contacts matters crucially for multigroup discussions (*42*), an alternative to achieve improved consensus while minimizing negative affective consequences is to structurally embed conversations. For example, deliberation could be broken up into phases where groups are first kept apart and given enough time to uncover crucial information without influence from other groups, and only then brought together to find a consensual solution. The effectiveness of such an intervention could easily be tested in an experimental setting where teams with a structurally embedded deliberation procedure likely made better decisions than those without.

Lastly, the finding that keeping groups apart has positive effects on decision quality relies to some degree on at least one of the groups having sufficient information to infer an optimal solution by themselves. This provides a scope condition for the mechanism found by this paper, but also raises the important question whether homophily can be helpful in cases where only a specific conjunction of arguments from different groups can reveal the best solution. While exceeding the scope of this paper, this calls for a promising model extension in which this is addressed more explicitly – namely, a model where a complex underlying function enables certain argument combinations to have nonlinear impacts on team members' beliefs.

7. Conclusion

The work presented here provides a novel insight on how to better shape interactions in diverse teams with regards to their decision-making abilities. When tasks are difficult, unbounded communication among team members can be a liability. In such cases, homophilous interactions improve decision quality because they keep individuals from convincing dissimilar others of their suboptimal ideas too quickly. However, homophily also resulted in increased discussion time, which points to a trade-off between decision quality and efficiency. This is an important finding to consider when finite resources have to be weighed against convex returns to optimal over inferior solutions.

References

1. H. Tajfel, *Social Identity and Intergroup Relations* (Cambridge University Press, 2010).
2. M. McPherson, L. Smith-Lovin, J. M. Cook, Birds of a feather: Homophily in social networks. *Annual review of sociology*. **27**, 415–444 (2001).
3. G. Stasser, W. Titus, Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*. **48**, 1467–1478 (1985).
4. G. Ertug, J. Brennecke, B. Kovács, T. Zou, What Does Homophily Do? A Review of the Consequences of Homophily. *ANNALS*. **16**, 38–69 (2022).
5. R. Reagans, Demographic diversity as network connections: Homophily and the diversity-performance debate. *The Oxford handbook of diversity and work*, 192–206 (2013).
6. C. P. Estévez-Mujica, A. Acero, W. Jiménez-Leal, C. García-Díaz, The Influence of Homophilous Interactions on Diversity Effects in Group Problem-Solving. *Nonlinear Dynamics, Psychology & Life Sciences*. **22** (2018).
7. S. E. Jackson, K. E. May, K. Whitney, R. A. Guzzo, E. Salas, Understanding the dynamics of diversity in decision-making teams. *Team effectiveness and decision making in organizations*. **204**, 261 (1995).
8. A. B. Carter, K. W. Phillips, The double-edged sword of diversity: Toward a dual pathway model. *Social and Personality Psychology Compass*. **11**, e12313 (2017).
9. F. J. Milliken, L. L. Martins, Searching for Common Threads: Understanding the Multiple Effects of Diversity in Organizational Groups. *AMR*. **21**, 402–433 (1996).
10. K. W. Phillips, C. O'Reilly, "Demography and Diversity in Organizations: A Review of 40 Years of Research" in *Research in Organizational Behavior* (1998), vol. 20, pp. 77–140.
11. S. Page, *The Diversity Bonus: How Great Teams Pay Off in the Knowledge Economy* (Princeton University Press, 2019; http://www.degruyter.com/document/doi/10.1515/9780691193823/html).
12. U. Peters, Hidden figures: epistemic costs and benefits of detecting (invisible) diversity in science. *Euro Jnl Phil Sci*. **11**, 33 (2021).
13. D. C. Lau, J. K. Murnighan, Demographic diversity and faultlines: The compositional dynamics of organizational groups. *Academy of management review*. **23**, 325–340 (1998).
14. D. A. Harrison, K. H. Price, M. P. Bell, Beyond Relational Demography: Time and the Effects of Surface- and Deep-Level Diversity on Work Group Cohesion. *AMJ*. **41**, 96–107 (1998).
15. G. B. Northcraft, J. T. Polzer, M. A. Neale, R. M. Kramer, "Diversity, social identity, and performance: Emergent social dynamics in cross-functional teams" in *Diversity in work teams: Research paradigms for a changing workplace* (American Psychological Association, Washington, DC, US, 1995), pp. 69–96.
16. K. W. Phillips, The Effects of Categorically Based Expectations on Minority Influence: The Importance of Congruence. *Pers Soc Psychol Bull*. **29**, 3–13 (2003).
17. W. A. Brechwald, M. J. Prinstein, Beyond Homophily: A Decade of Advances in Understanding Peer Influence Processes. *Journal of Research on Adolescence*. **21**, 166–179 (2011).
18. D. Z. Levin, E. M. Whitener, R. Cross, Perceived trustworthiness of knowledge sources: The moderating impact of relationship length. *Journal of Applied Psychology*. **91**, 1163–1171 (2006).
19. A. Flache, M. Mäs, T. Feliciani, E. Chattoe-Brown, G. Deffuant, S. Huet, J. Lorenz, Models of Social Influence: Towards the Next Frontiers. *JASSS*. **20**, 2 (2017).
20. R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation*. **5** (2002).
21. C. A. Bail, Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*. **113**, 11823–11828 (2016).
22. S. Balietti, L. Getoor, D. G. Goldstein, D. J. Watts, Reducing opinion polarization: Effects of exposure to similar people with differing political views. *Proceedings of the National Academy of Sciences*. **118**, e2112552118 (2021).

23. M. Mäs, A. Flache, Differentiation without distancing. Explaining bi-polarization of opinions without negative influence. *PloS one*. **8**, e74516 (2013).
24. D. DellaPosta, Y. Shi, M. Macy, Why do liberals drink lattes? *American Journal of Sociology*. **120**, 1473–1511 (2015).
25. M. Mäs, A. Flache, K. Takács, K. A. Jehn, In the short term we divide, in the long term we unite: Demographic crisscrossing and the effects of faultlines on subgroup polarization. *Organization science*. **24**, 716–736 (2013).
26. T. H. Stark, A. Flache, The Double Edge of Common Interest: Ethnic Segregation as an Unintended Byproduct of Opinion Homophily. *Sociol Educ*. **85**, 179–199 (2012).
27. K. J. Zollman, The epistemic benefit of transient diversity. *Erkenntnis*. **72**, 17–35 (2010).
28. D. Frey, D. Šešelja, Robustness and Idealizations in Agent-Based Models of Scientific Interaction. *The British Journal for the Philosophy of Science*. **71**, 1411–1437 (2020).
29. J. Wu, C. O'Connor, How Should We Promote Transient Diversity in Science? (2021), (available at http://philsci-archive.pitt.edu/19428/).
30. P. L. McLeod, S. A. Lobel, T. H. Cox, Ethnic Diversity and Creativity in Small Groups. *Small Group Research*. **27**, 248–264 (1996).
31. S. S. Levine, E. P. Apfelbaum, M. Bernard, V. L. Bartelt, E. J. Zajac, D. Stark, Ethnic diversity deflates price bubbles. *Proceedings of the National Academy of Sciences*. **111**, 18524–18529 (2014).
32. K. W. Phillips, D. L. Loyd, When surface and deep-level diversity collide: The effects on dissenting group members. *Organizational Behavior and Human Decision Processes*. **99**, 143–160 (2006).
33. L. Lu, Y. C. Yuan, P. L. McLeod, Twenty-Five Years of Hidden Profiles in Group Decision Making: A Meta-Analysis. *Pers Soc Psychol Rev*. **16**, 54–75 (2012).
34. S. Schulz-Hardt, A. Mojzisch, How to achieve synergy in group decision making: Lessons to be learned from the hidden profile paradigm. *European Review of Social Psychology*. **23**, 305–343 (2012).
35. S. G. Sohrab, M. J. Waller, S. Kaplan, Exploring the Hidden-Profile Paradigm: A Literature Review and Analysis. *Small Group Research*. **46**, 489–535 (2015).
36. G. Stasser, D. Stewart, Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality and Social Psychology*. **63**, 426–434 (1992).
37. G. Stasser, W. Titus, Hidden Profiles: A Brief History. *Psychological Inquiry*. **14**, 304–313 (2003).
38. G. M. Wittenbaum, A. B. Hollingshead, I. C. Botero, From cooperative to motivated information sharing in groups: moving beyond the hidden profile paradigm. *Communication Monographs*. **71**, 286–310 (2004).
39. L. Hong, S. E. Page, Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*. **101**, 16385–16389 (2004).
40. G. Levy, Decision making in committees: Transparency, reputation, and voting rules. *American economic review*. **97**, 150–168 (2007).
41. S. L. Feld, The Focused Organization of Social Ties. *American Journal of Sociology*. **86**, 1015–1035 (1981).
42. A. Flache, M. Mäs, How to get the timing right. A computational model of the effects of the timing of contacts on team cohesion in demographically diverse teams. *Computational and Mathematical Organization Theory*. **14**, 23–51 (2008).