

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The effect of diversity on group decision-making

Permalink

<https://escholarship.org/uc/item/6685n22z>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Karadzhov, Georgi Milev

Vlachos, Andreas

Stafford, Tom

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

The effect of diversity on group decision-making

Georgi Karadzhov (gmk34@cam.ac.uk)

Department of Computer Science and Technology, University of Cambridge
Cambridge, UK

Andreas Vlachos (av308@cam.ac.uk)

Department of Computer Science and Technology, University of Cambridge
Cambridge, UK

Tom Stafford (t.stafford@sheffield.ac.uk)

Department of Psychology, University of Sheffield
Sheffield, UK

Abstract

We explore different aspects of cognitive diversity and its effect on the success of group deliberation. To evaluate this, we use 500 dialogues from small, online groups discussing the Wason Card Selection task - the DeliData corpus. Leveraging the corpus, we perform quantitative analysis evaluating three different measures of cognitive diversity. First, we analyse the effect of group size as a proxy measure for diversity. Second, we evaluate the effect of the size of the initial idea pool. Finally, we look into the content of the discussion by analysing discussed solutions, discussion patterns, and how conversational probing can improve those characteristics.

Despite the reputation of groups for compounding bias, we show that small groups can, through dialogue, overcome intuitive biases and improve individual decision-making. Across a large sample and different operationalisations, we consistently find that greater cognitive diversity is associated with more successful group deliberation.

Code and data used for the analysis are available in the repository: <https://github.com/gkaradzhov/cognitive-diversity-groups-cogsci24>.

Keywords: group deliberation; decision-making; collaboration

Introduction

Individuals come together to discuss, analyse, and evaluate different points of view. Such group deliberation is considered an essential tool for tackling complex tasks. It is a cornerstone of decision-making in domains such as research, government and business. In academia, we have the peer review process (Nicholas et al., 2015; Marcoci et al., 2022), where scientists discuss and improve scientific reports and academic studies. In government, collective bargaining, resource allocation and community planning are driven by committees, task forces, and city councils, where groups of people try to come up with fair decisions and policies. In business, decisions are made again in small groups, such as boards of directors or management teams, responsible for running companies, budgeting, and hiring. The widespread deployment of group deliberation as a decision-making strategy makes it important to understand the forces which support or undermine good deliberation. The current work presents a data-driven approach to understanding one such key factor: cognitive diversity.

Group deliberation is the dialogue that facilitates the careful weighing of information, providing adequate speaking opportunities, and bridging the differences among participants'

diverse ways of speaking and knowing (Burkhalter, Gastil, & Kelshaw, 2002). Mercier and Sperber (2011) review evidence that, in the right conditions, groups can outperform even the most capable individual members in their decision-making. These conditions include common goals among group members and the exchange of arguments. As well as groups exceeding the capacity of any individual member, another relevant standard is the decision which would be reached by simple aggregation or averaging of independent individual judgements (the so-called “Wisdom of the Crowd” effect, Galton, 1907). In the study of Navajas, Niella, Garbulsky, Bahrami, and Sigman (2018) small groups outperform the wisdom of the crowd standard. Navajas et al. (2018) argue that small groups can be more effective than both individuals of superior ability, and large crowds. Further, this and other evidence (Mercier, 2016; Trouche, Sander, & Mercier, 2014) supports the idea that the exchange of arguments is essential to successful group decision-making.

That said, group decision-making is not universally successful. A tradition in social psychology has emphasised the potential for negative outcomes. A notable diagnosis is the perils of “Group Think” (Janis, 1972). Potential mechanisms of failures of group decision-making include social loafing, herding, deference to others’ preferences — real (Mojzisch & Schulz-Hardt, 2010), or imagined (Prentice & Miller, 1996) —, failure to share knowledge available to individual group members (Stasser & Titus, 1985), underweighting dissenting voices, and amplification of individual cognitive biases. We note that the mechanisms supporting the Wisdom of Crowds effect relies on individuals having independent access to some true information (signal), and uncorrelated errors (noise). To the extent that individuals share systematic error, such as a shared cognitive bias, aggregation of individual judgements will reinforce error, not eliminate it.

To study group decision-making under conditions of bias, we use the Wason Card Selection Task (Wason, 1968), an iconic task used to study human reason (Evans, 2016). For our purposes, it has two important qualities. First, with high reliability, the majority of people get the task wrong, and they fail in a consistent way. This suggests the presence of a confirmation, or matching, bias which draws people to an

intuitive but incorrect response. This makes the task an attractive paradigm for investigating how group function may interact with individual biases. Second, the Wason task has a correct answer which can, once understood, be relatively easily explained and understood by others. This underlies the remarkable prior result that while individuals tend to fail on the Wason Task, small groups tend to succeed (discussed in Mercier & Sperber, 2011).

Cognitive diversity in group decisions

An established finding in studies of group decision-making is the potential benefits of diversity (Bang & Frith, 2017; Davis-Stober, Budescu, Dana, & Broomell, 2014; Page, 2008). Diversity can be in the form of differences in the backgrounds, knowledge or cognitive styles that individuals bring to a group, as well as the different roles they may adopt or be assigned within a group. In the limit, too extreme diversity can hamper group function — particularly where it facilitates intragroup conflict or hinders agreement on common goals — but otherwise, there are good reasons to expect greater diversity to enhance group decision-making.

Diversity can be improved in a number of ways. One approach is by differing demographic or professional characteristics (Wang, Fussell, & Cosley, 2011) of the participants. This can also be in part achieved by considering groups of larger sizes - i.e. by having a larger group, there is an increased chance that the group would have people from different backgrounds and experiences. An alternative way is to adopt procedures or structures which facilitate expressing and exchanging diverse ideas (Dalkey & Helmer, 1963). Our focus here is on *cognitive* diversity, which reflects differences between individuals in knowledge, processes or actions within the group context and is distinct from (but may stem from) demographic differences, such as gender or ethnic diversity.

The current work, hypotheses. In this work, we leverage a unique dataset of group decisions, the DeliData corpus (Karadzhov, Stafford, & Vlachos, 2023). In DeliData, 500 groups provided initial judgments and then discussed the Wason Card Selection Task, for which full transcripts were recorded. Individuals then submitted their updated answers. This dataset allows us to look, with a high degree of granularity, at the features of groups and the deliberative processes which generated successful and unsuccessful deliberation. Furthermore, DeliData contains data indicating participants’ solutions, arguments, and deliberative markers, which would be beneficial for rich and multifaceted analysis.

Concretely, we look into three aspects of cognitive diversity - (i) indirectly, via the proxy of group size, (ii) via the distribution of initial solutions possessed by the group, and (iii) directly, via the content of the discussions, the solutions considered by the group, as well as deliberation markers affecting said content.

Given previous research, as well as the three aspects of cognitive diversity, we will be testing the following hypotheses. First, that larger groups will demonstrate better performance.

# of dialogues	500
# of utterances	14003
AVG # of utterances / dialogue	28
AVG group size	3.16
AVG utterance length	8.59
# of Probing utterances	1739
Solo Performance	0.59
Group performance	0.72
Performance Gain	0.13

Table 1: Summary statistics for the DeliData corpus

This has been shown in other studies (Navajas et al., 2018; Hill, 1982), but not — to our knowledge — in the Wason Task, nor for other tasks which elicit strong individual biases (for which groups may function to aggregate rather than eliminate) as a characteristic that may affect group performance.

Second, we hypothesise that a greater diversity of initial ideas will support better group performance. Following Schulz-Hardt, Brodbeck, Mojzisch, Kerschreiter, and Frey (2006) we index this by inspecting the number of distinct individual solutions submitted by participants prior to entering the group discussion. Third and finally, that diversity in discussion content will support better group performance. This was shown for forecasting by Wintle et al. (2023). We index this based on the distinct solutions mentioned in the conversation, as well as the transitions between mentioned solutions. We further hypothesise, that conversational probing (i.e. asking questions that provoke further discussion (Karadzhov et al., 2023)) has an effect over the diversity of discussion content. The premise is that conversational probing can be considered as a possible intervention that would increase the content and cognitive diversity of a discussion.

Data & Methods

DeliData

The DeliData corpus is described in Karadzhov et al. (2023). It comprises 500 small group discussions of the Wason Task, which vary both in terms of group composition (2 to 5 participants) and performance (both high and low, and pre- (solo) and post-discussion (final), performance). Statistics of the DeliData corpus are presented in Table 1.

To create the dataset, the Wason Card Selection Task was presented in its traditional form: it is explained that cards have a number on one side and a letter on the other. Four cards are shown displaying an even number (‘E’), odd number (‘O’), consonant (‘C’) and vowel (‘V’). The task is to answer the question “Which card(s) should you turn to test the rule: **All cards with vowels on one side, have an even number on the other**”. The intuitive but wrong answer (henceforth the Lure) is to turn the even number and vowel (‘EV’). The correct answer (henceforth the Target) is to turn over the odd number and the vowel (‘OV’).

Karadzhov et al. (2023) report that 64% of the groups performed better following the discussion compared to their initial

performances. In 44% of the groups who had at least one correct answer as their final solution, none of the participants had begun with the solution. This demonstrates the benefits of group discussion and suggests that more is occurring than mere propagation of the best answers among the groups.

Moreover, Karadzhov et al. (2023) provides utterance-level annotations that mark deliberative and argumentation cues. The full details of the annotation scheme and corpus are presented in Karadzhov et al. (2023), however in the context of this work, we are interested in the phenomenon of *conversational probing*. Conversational probing is defined as any utterance that provokes discussion, deliberation or argumentation. This is further split into three categories - Moderation (i.e. modelling group dynamics), Reasoning (probing further argumentation), and Solution (i.e. enquiring about a specific solution).

Analytic strategy

The Wason Card Selection task has one correct answer, allowing for a direct and objective measure of performance. In this work, we leverage the fine-grained scoring from Karadzhov et al. (2023), where 0.25 points each are given for i) selecting either card in the Target (O and V) and (ii) for **not** turning unnecessary cards (E and C). Therefore, if the participant submitted a correct solution, their score would be 1, if they are off by one card - 0.75 and so on. If participants selected the Lure (EV), their fine-grained score would be 0.5. Similarly, selecting all cards (CEOV) or none, will have a fine-grained score of 0.5.

Performance gain is calculated by subtracting the **task performance** of the solo solution from the task performance in the final submission. For example, if a participant’s solo submission was 0.5 and improved to 0.75 after the discussion, the performance gain would be $0.75 - 0.5 = 0.25$. In this work, we investigate **group performance gain** which is the averaged performance gain of each of the participants.

Different measures of diversity First, we investigate group size as a proxy measure of diversity, and can be compared to a baseline of pre-discussion, solo performance (groups of size 1). The premise is that more members participating in a dialogue would naturally increase cognitive diversity. Following, we consider the diversity of initial ideas, which is calculated by counting the number of unique solo solutions submitted by the group’s participants. Finally, we analyse the factors measuring the diversity of discussed solutions and their contribution towards the performance gain. To this end, we evaluate the following characteristics of a diverse discussion.

- **dialogue length** measured by the number of utterances. The premise is that longer conversations will naturally contain more mentions of solutions, thus contributing towards a more diverse discussion:
- Number of **unique solutions** that are mentioned in a discussion. We hypothesise, that conversations with positive

performance gain will contain more mentions of unique solutions, indicating that participants were actively discussing a diverse set of possible solutions to the Wason task.

- We analyse how participants discuss solutions. To measure this, we extract the solutions participants talk about in every utterance by keeping track of Wason card mentions ¹. Then, we calculate the occurrences of all transitions between solutions in the form of a from-through-to transition triple. For example, if the conversation starts with the lure (EV), then transitions to a solution that turns every card, and then to the correct solution, the transition triplet would be “EV-CEOV-OV”. Having these transition triplets, we measure the following three quantities:
 - **Number of unique transitions**
 - **Number of stuck transitions** - where participants recalled the same solution within a triplet, for example, EV-EV-EV.
 - **Number of circular transitions** - where participants mention one solution, then an interim solution and finally return to the first solution. In our notation, this would be represented as: OV-CEOV-OV, if participants first discuss odd number + vowel, then discuss turning in all cards, and then they return back to odd + vowel. This is a special case of stuck solutions, where participants may try to “break the cycle” but ultimately return to a previous solution.

We conclude this section by also analysing the interaction of the measures above and the probing utterances, as annotated in DeliData (Karadzhov et al., 2023)

Results

Effects of group size on diversity




Figure 1: Average performance gain against group size

As a first step of this analysis, we are hypothesising that group size is a factor contributing to better cognitive diversity and ul-

¹And the two special cases when participants mentioned “none” or “all” where we consider that their solution is to select none or all of the cards

Groups	p-value
2 & 3	0.17
3 & 4	0.75
4 & 5	0.06
2 & 3+	0.03

Table 2: Comparison between the effect of group size on performance gain. Significant values in **bold**.

timely improved decision-making performance (as measured by performance gain).

Considering only the size of the group and the performance gain, we show that the performance gain of a group increases consistently with group size, as shown in Figure 1. To confirm the significance of these results, we perform a 1-sample t-test against the baseline performance gain of 0. This yields p-values of 2.5^{-8} , 3.5^{-15} , 4.1^{-12} , and 2.8^{-10} for groups of 2, 3, 4, 5 respectively. This indicates that for all group sizes, the performance gain was significantly above 0. Next, we performed an ANOVA (Girden, 1992) comparison between groups of sizes 2 to 5. This showed a significant result ($p=0.0251$), confirming that group size affects performance gain. Following this, we did a post-hoc analysis, performing a t-test between different group sizes. The results (also presented in Table 2) of the test show that there is no significant difference between each of the groups. This is also consistent and confirms the findings of (Hill, 1982). That said, Hill (1982) argues that groups (3+ people) perform better than dyads (dialogues). To this end, we also confirm this claim on the 500 dialogues of DeliData, showing that groups (size ≥ 3) significantly outperform dyads (size=2) with a p-value of 0.03.




Figure 2: Comparison between conversational statistics of two-party dialogues (dyads) and group dialogues (groups)

Overall, we find that group size has an effect on performance gain. However, group size being a proxy measure, we cannot directly test whether the increase in performance gain is due to increased cognitive diversity or other factors.

That said, in this work, we further investigate two characteristics that measure cognitive diversity directly. First, larger groups may benefit from an increased idea pool, which we investigate in detail in the next section - **Diversity of initial ideas**. Secondly, larger groups may engage in more diverse conversational strategies, thus incorporating more ideas and arguments during the group discussion. Figure 2 compares basic conversational statistics between dyads and groups. Dialogues between two interlocutors have mostly between 10 and 25 utterances, while group dialogues have a wider range of lengths, with a long tail of dialogues longer than 50 utterances. Likewise, groups in these dialogues tend to use a larger vocabulary, as shown on the violin plot of the unique tokens. We further evaluate conversational patterns and characteristics in section **Diversity of discussed solutions**.

Diversity of initial ideas




Figure 3: Performance gain of groups with high (purple) and low (pink) diversity of initial submissions. Groups are also split by whether any initial submissions contained a correct answer (right two) or did not contain any correct answers (left two).

Herein, we investigate how the diversity of the initial pool of ideas (indicated by submitted solutions in participants' solo phase) affects performance gain. We hypothesise that more diverse initial ideas will provide a basis for more discussion and a more balanced starting point, ultimately leading to an improved conversation performance (i.e. increase in performance gain). Figure 3 shows how initial ideas, in terms of both diversity and the presence of a correct solution among group members, affect performance gain. We compare two factors - (i) whether the initial idea pool contained a correct answer and (ii) the level of diversity (i.e. number of unique solutions) of the initial idea pool. We show that regardless of whether the initial idea pool contained a correct answer, high diversity is always beneficial for the group's performance gain. Further, we analyse the effects of these two factors by fitting a linear regression, where performance gain is predicted based on the level of diversity and whether the initial idea pool contained a correct answer. The *diversity* component has a statistically significant contribution with a p-value of 0.001. On the other hand, the p-value of the component indicating a correct initial solution is 0.079.

	Correlation	PCR coef.
dialogue length	0.10	0.01
unique solutions	0.28	0.02
unique transitions	0.22	0.02
stuck transitions	0.02	-0.02
circular transitions	0.19	0.01

Table 3: Results of correlation and PCR analyses between different factors representing the diversity of solutions and performance gain. In bold, we show the statistically significant values with p-value < 0.05.

These results show that the diversity of initial ideas is a significant component of a constructive conversation. However, a reasonable question can be asked - *how does the size of the group confound the diversity of the initial ideal pool with respect to performance gain?* Intuitively, larger groups would have more ideas, thus the chance of encountering a diverse pool of ideas is higher. To evaluate this interaction between group size and the diversity of initial ideas, we predict performance gain using a linear regression model fitted with the least squares method. We consider three linear models with the following setups - we predict performance gain based on (i) group size, (ii) diversity of initial ideas, and (iii) a combination of (i) and (ii), i.e. group size + diversity of initial ideas. We then compare these models using ANOVA.

First, we compare models (i) and (iii), testing whether including diversity of ideas to group size gives us extra predictive power. The ANOVA comparison estimates that these two models are significantly different, with a p-value of 0.0006. Following, we compare the inverse - models (ii) and (iii), to test whether including group size to diversity of ideas gives us extra predictive power. This time, the difference is not significant, with a p-value of 0.81.

This analysis shows that most of the predictive power comes from the diversity of initial solutions, and including the number of participants does not bring additional predictive power. Notably, another perspective would be that having a higher diversity of initial solutions already captures the information of having a larger group size.

Diversity of discussed solutions

Here we use the five indicators outlined in the **Analytic strategy** subsection - dialogue length, number of unique solutions, number of unique transitions, stuck transitions and circular transitions. In order to evaluate these five indicators and their relationship to performance gain, we perform (i) a correlation study using ("Pearson's Correlation Coefficient", 2008) and (ii) a regression analysis using Principal Component Regression (PCR) (Massy, 1965). For the latter, we transform the input feature space using principal component analysis. This can handle adequately the multicollinearity present among the features in the dataset. We implement it by fitting a linear regression over the transformed feature space. We then transform the coefficients of the regression back to the original feature space to obtain the final coefficients.

The values of the correlation analysis and the PCR analysis are presented in Table 3. If we consider the correlation statistic, out of the five factors, four are statistically significant in predicting the performance gain - the dialogue length, the number of unique discussed solutions, the number of unique transitions, and the number of circular transitions. Further, if we consider the correlation coefficients, the number of unique solutions (0.28) is considered the most important feature, followed by the number of unique transitions (0.22), and the circular transitions (0.19). Similarly, if we analyse the results of the PCR analysis, dialogue length, number of unique solutions, and circular transitions are significant. Similarly to the correlation analysis, the number of unique solutions has the highest importance, based on the coefficients from the principal component regression.

These results confirm the hypothesis that the diversity of discussed solutions correlates with group performance gain. Similarly to the diversity of the initial ideas, the number of unique solutions discussed is highly indicative of well-performing groups. Following similar intuition, we argue that this is both a measure of how diverse the discussion is, as well as it is a proxy for how open the group is towards new ideas.

In the previous subsection (**Diversity of initial ideas**), we showed that groups with a high diversity of initial solutions achieve better performance gain. Likewise, in this section, we showed that the number of discussed solutions is also associated with increased performance gain. However, not all groups start with a diverse set of initial ideas. A reasonable research question would be - *what happens in conversations where groups with low initial diversity achieve a high number of discussed ideas.* To evaluate this, we test the hypothesis that conversational probing affects dialogues with a low diversity of initial ideas. The premise is that if participants *probe* the conversation more often, this will push the group to discuss novel ideas. To this end, we calculate a Pearson-R correlation of 0.34 between the number of conversational probes and the diversity of discussed solutions.

To further analyse this effect, we take advantage of the more granular *probing* annotations in DeliData, namely whether the probing was related to **Moderation**, **Reasoning** behind a solution, or whether a specific **Solution** should be picked. To that end, we perform correlation and PCR studies to investigate the relationship between different kinds of probing and the diversity of discussed solutions. Results of this analysis are shown in Table 4. We see that both probing for reasons and solutions are a significant contributor towards predicting the diversity of discussed solutions. The correlation coefficients show that probing for reasons (0.41) has a higher association with the diversity of discussed solutions than simply enquiring about solutions (PearsonR = 0.26). Likewise, the PCR coefficient for probing for reasons is 0.83, and probing for solutions is 0.43. This indicates that regardless of the type of analysis (correlational and predictive), we show the importance of probing for reasons, followed by probing for solutions. This is an interesting result, showing the importance of reasoning

and also confirming the findings of (Mercier & Sperber, 2011) about the value of deliberation. Moreover, this gives a basis for future work on interventional analysis and applications, where conversational probing can increase cognitive diversity and improve conversational outcomes.

	Correlation	PCR Coef.
Moderation	-0.017	-0.06
Reasoning	0.41	0.83
Solution	0.26	0.43

Table 4: Results of correlation and PCR analyses of different types of probing and predicting the number of unique solutions discussed. In bold, we show the statistically significant values.

Discussion

Across three different operationalisations, we show that greater diversity is associated with enhanced benefits of discussion in small groups.

On the chosen task, which has a high risk of individual bias and error, online discussion groups outperform individuals, and groups of 3 and above outperform dyads.

Groups incorporating more diverse individuals inherently have access to greater cognitive diversity. We find support for this idea more directly by showing that the diversity of initial solutions was a stronger predictor of group performance gain than the fact of any individual group member possessing the right answer.

Finally, at the greatest level of analytic granularity we deploy, we inspect the content of group discussion and show that groups which discuss a wider range of potential solutions are also more successful. The circularity and number of discussed solutions transitions were also predictive of an increase in performance gain, while stuck transitions were insignificant. Moreover, we show that conversational probing can increase the diversity of discussed solutions, even in groups which had a limited idea pool.

Contributions

This analysis of the DeliData corpus shows that it replicates the finding that group discussion of the Wason Task can reduce, rather than reinforce, the individual susceptibility to error on the task. It extends this finding to the increasingly relevant domain of online discussion among strangers. The finding is all the more remarkable for the low-commitment² and the brief nature of the discussions (typically only 20 - 40 utterances).

By focusing on diversity, we confirm prior findings that the initial pool of ideas is an important factor in successful group deliberation (Schulz-Hardt et al., 2006), as is the exchange of ideas during the discussion (Mercier, 2016). We show that the association of diversity with successful deliberation extends beyond the mere content and volume of discussion to diversity

²In DeliData the participants were recruited from the crowd-sourcing website MechanicalTurk

in structure. Circularity in discussion is not necessarily negative — it could represent focus and persistence which we also show is important predictor of performance improvement.

Additionally, we show how analysis of features of groups and group dialogue, captured at different levels of description, can be applied to the study of group decision-making. These approaches are applicable to other domains where dialogue supporting group decision-making is captured.

Limitations

Since our study is observational, not experimental, in nature, it is only possible to note associations, not causation. Ultimately, we suspect that successful groups generate diverse discussions, as well as vice versa. Our analysis supports the idea that diversity of membership, ideas and discussion will promote successful group decision-making, but decisive proof of this idea awaits testing by intervention studies.

Further, in this study, we investigated the effect of cognitive diversity on performance gain. By definition, performance gain is affected by the baseline performance of the group - i.e. if a group has a low starting performance, the performance of this group can grow more compared to groups that start with high baseline performance. The vice-versa is also true - if a group starts with high (or even perfect) solo performance, during the discussion they can only decrease their performance leading to negative performance gain. We recognise that this is a limitation of the current study, as we only look at how much a conversation improved, disregarding the starting points. We performed a small-scale control experiment to evaluate the effect of this limitation. First, we consider the dialogues with high baseline performance, which we define as dialogues where more than 50% of the participants had the correct answer in their solo submissions. In DeliData, there are 9 dialogues out of 500 (1.8%) where this is the case, suggesting that this affects a limited portion of the data. Second, we re-ran all of the analyses in this work without those 9 dialogues and confirmed that the removal of those does not affect the statistical significance tests, nor does it change any of the analyses performed as part of this work.

Conclusion

Despite the reputation of groups for compounding bias, these findings show that small groups can, through dialogue, enhance individual decision-making. Further, we offer a worked example of the potential for fine-grained analysis of dialogue to illuminate exactly how the benefits of deliberation come about. Across a large sample and different operationalisations, we consistently find that greater diversity is associated with more successful group deliberation.

Acknowledgements

Georgi Karadzhov's PhD was supported by an EPSRC doctoral training scholarship. This research was additionally supported by a donation from Google.

References

- Bang, D., & Frith, C. D. (2017). Making better decisions in groups. *Royal Society open science*, 4(8), 170193.
- Burkhalter, S., Gastil, J., & Kelshaw, T. (2002). A conceptual definition and theoretical model of public deliberation in small face-to-face groups. *Communication theory*, 12(4), 398–422.
- Dalkey, N., & Helmer, O. (1963). An experimental application of the delphi method to the use of experts. *Management science*, 9(3), 458–467.
- Davis-Stober, C. P., Budescu, D. V., Dana, J., & Broomell, S. B. (2014). When is a crowd wise? *Decision*, 1(2), 79.
- Evans, J. S. B. (2016). Reasoning, biases and dual processes: The lasting impact of wason (1960). *Quarterly journal of experimental psychology*, 69(10), 2076–2092.
- Galton, F. (1907). Vox populi. *Nature*, 75(1949), 450–451.
- Girden, E. R. (1992). *Anova: Repeated measures* (No. 84). Sage.
- Hill, G. W. (1982). Group versus individual performance: are n + 1 heads better than one? psychological bulletin.. Retrieved from <https://api.semanticscholar.org/CorpusID:32389401>
- Janis, I. L. (1972). Victims of groupthink.. Retrieved from <https://api.semanticscholar.org/CorpusID:169716095>
- Karadzhov, G., Stafford, T., & Vlachos, A. (2023). Delidata: A dataset for deliberation in multi-party problem solving. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2), 1–25.
- Marcoci, A., Vercammen, A., Bush, M., Hamilton, D. G., Hanea, A., Hemming, V., ... Fidler, F. (2022). Reimagining peer review as an expert elicitation process. *BMC Research Notes*, 15(1), 127.
- Massy, W. F. (1965). Principal components regression in exploratory statistical research. *Journal of the American Statistical Association*, 60(309), 234–256.
- Mercier, H. (2016). The argumentative theory: Predictions and empirical evidence. *Trends in Cognitive Sciences*, 20(9), 689–700.
- Mercier, H., & Sperber, D. (2011). Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences*, 34(2), 57–74.
- Mojzisch, A., & Schulz-Hardt, S. (2010). Knowing others' preferences degrades the quality of group decisions. *Journal of personality and social psychology*, 98(5), 794.
- Navajas, J., Niella, T., Garbulsky, G., Bahrami, B., & Sigman, M. (2018). Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour*, 2(2), 126–132.
- Nicholas, D., Watkinson, A., Jamali, H. R., Herman, E., Tenopir, C., Volentine, R., ... Levine, K. (2015). Peer review: still king in the digital age. *Learned Publishing*, 28(1), 15–21. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1087/20150104> doi: <https://doi.org/10.1087/20150104>
- Page, S. (2008). *The difference: How the power of diversity creates better groups, firms, schools, and societies-new edition*. Princeton University Press.
- Pearson's correlation coefficient. (2008). In W. Kirch (Ed.), *Encyclopedia of public health* (pp. 1090–1091). Dordrecht: Springer Netherlands. Retrieved from https://doi.org/10.1007/978-1-4020-5614-7_2569 doi: 10.1007/978-1-4020-5614-7_2569
- Prentice, D. A., & Miller, D. T. (1996). Pluralistic ignorance and the perpetuation of social norms by unwitting actors. In *Advances in experimental social psychology* (Vol. 28, pp. 161–209). Elsevier.
- Schulz-Hardt, S., Brodbeck, F. C., Mojzisch, A., Kerschreiter, R., & Frey, D. (2006). Group decision making in hidden profile situations: dissent as a facilitator for decision quality. *Journal of personality and social psychology*, 91(6), 1080.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48, 1467–1478. Retrieved from <https://api.semanticscholar.org/CorpusID:34000088>
- Trouche, E., Sander, E., & Mercier, H. (2014). Arguments, more than confidence, explain the good performance of reasoning groups. *Journal of Experimental Psychology: General*, 143(5), 1958.
- Wang, H.-C., Fussell, S. R., & Cosley, D. (2011). From diversity to creativity: Stimulating group brainstorming with cultural differences and conversationally-retrieved pictures. In *Proceedings of the acm 2011 conference on computer supported cooperative work* (pp. 265–274).
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly journal of experimental psychology*, 20(3), 273–281.
- Wintle, B. C., Smith, E. T., Bush, M., Mody, F., Wilkinson, D. P., Hanea, A. M., ... others (2023). Predicting and reasoning about replicability using structured groups. *Royal Society Open Science*, 10(6), 221553.