Beyond the Echo Chamber: Modelling Open-Mindedness in Citizens' Assemblies

JAAMAS Track

Jake Barrett
University of Edinburgh, U.K.
j.t.l.barrett@sms.ed.ac.uk

Kobi Gal
University of Edinburgh, U.K.
Ben-Gurion University, Israel
kgal@ed.ac.uk

Loizos Michael
Open University of Cyprus, Cyprus
loizos@ouc.ac.cy

Dan Vilenchik
Ben-Gurion University, Israel
vilenchi@bgu.ac.il

ABSTRACT

A Citizens' Assembly (CA) is a democratic innovation tool where a randomly selected group of citizens deliberate a topic over multiple rounds to generate, and then vote upon, policy recommendations. Despite growing popularity, little work exists on understanding how CA inputs, such as the expert selection process and the mixing method used for discussion groups, affect results, and therefore on how to systematically set such parameters to optimize the process.

In this work, we model CA deliberation and opinion change as a Multi-Agent Systems problem. We introduce and formalise a set of criteria for evaluating successful CAs using insight from previous CA trials and theoretical results. Although real-world trials meet these criteria, we show that finding a model that does so is non-trivial; through simulations and theoretical arguments, we show that established opinion change models fail at least one of these criteria. This is an extended abstract of a JAAMAS article [2].

KEYWORDS

Opinion Change, Social Influence, Participatory Democracy, Citizens' Assemblies

ACM Reference Format:
Jake Barrett, Kobi Gal, Loizos Michael, and Dan Vilenchik. 2025. Beyond the Echo Chamber: Modelling Open-Mindedness in Citizens' Assemblies: JAAMAS Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

A Citizens' Assembly (CA) is a democratic innovation tool used to inform political decision-making by assigning political power to a random panel of demographically representative participants. Over several weekends, participants hear from a range of expert speakers on topics related to a main theme, such as climate change,[1] gender equality,[2] and constitutional design.[3] These span from the municipal level — such as filling a $1.2 billion funding gap in Melbourne[4] — to the Global Assembly convened to address the climate crisis.[5] Participants deliberate in smaller groups after hearing each expert, combining their own lived experiences with expert information to generate a series of recommendations. An overview of the process is shown in Figure 1.

CAs aim to reach agreement based on a sample of the informed views of the population — the views of the population if everyone had the time and resources to research the issue [1, 7]. CAs thus differ from referenda, which seek to sample the current views of the population [8], as participants are encouraged to, and invariably do, change their opinions. In both political theory and real-world CA trials, participants engage with the process, change their views, and reach sophisticated policy recommendations [5, 9].

Despite their relevance to group decision-making and multi-agent deliberation, CAs are understudied computationally. Other computational works [3, 4, 6] focus on either existing algorithmic input or applying theoretical models. There is an accompanying paucity of quantitative data, with the exception of the Citizens' Assembly of Scotland (CA of Scotland) [5], whose data from participant surveys we utilise in this work. This quantitative gap leaves open several questions. Why do participants engage meaningfully with different views, when we might expect them to stick to pre-existing opinions and only engage with those inside their own echo-chamber [10]? Can the process be manipulated to obtain a pre-ordained result, as some critics fear?[6] These open questions are a limiting factor in the continued expansion of CAs as a tool for supporting participatory democracy.

In this work we attempt to address these questions by motivating, formalising, and optimising a model structure that usefully captures CA dynamics. The individual opinion dynamics of CA agents, combined with the algorithmic structure of their interactions, leads naturally to a Multi-Agent System (MAS). Based on theoretical works and real-world CA implementations, we motivate

[1] e.g. https://www.climateassembly.uk/ (2020) among many others.

[2] https://citizensassembly.ie/ (2017).

[3] https:www.gov.scot/publications/research-report-citizens-assembly-scotland/ (2019-20).

[4] https://involve.org.uk/resources/case-studies/melbourne-peoples-panel (2014).

[5] https://globalassembly.org/ (2021).

[6] For example, see https://politicalreform.ie/2016/05/16/citizens-assemblies-are-open-to-manipulation/ https://www.counterfire.org/article/the-problem-with-citizens-assemblies/.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowe (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

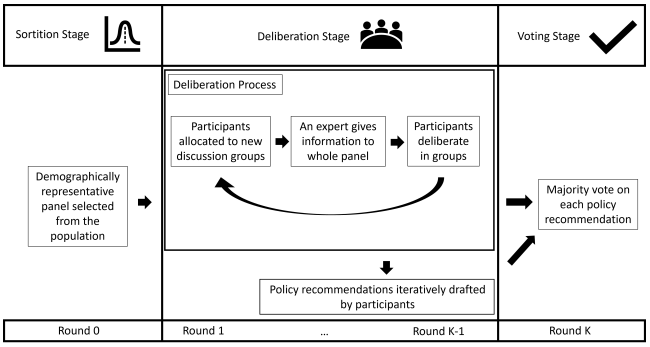

Figure 1: Graphical description of the CA process with a total number of rounds K

four key criteria of a sufficient CA model: agents (i) move their opinions closer to experts; (ii) maintain some variance in opinions; (iii) converge to a unique opinion; and (iv) abandon extremist views.

Our work makes three contributions. The first is that we have formalised the CA process as a MAS problem. Second, we formalise criteria for a sufficient model of CAs based on both results from the literature and real-world CA experiments. We collaborated with CA facilitators in our model design. Finally, we validate our model using survey data from the CA of Scotland.

Although we are not the first to model CA deliberation from a computational viewpoint, our specific formulation of a latent 'open-mindedness' variable is the only one that is rooted in theory, consistent with empirical results, and improves the performance of deliberative models in CAs.

2 FOUNDATIONS OF A CA MODEL

The CA process can be distilled to three main stages, as visualized in Figure 1. In the first stage, a demographically representative panel of participants is drawn from the wider population. We treat this process as exogenous to our work assuming the panel contains representative demographics.

The second stage, which we focus on in this paper and is represented by the central box in Figure 1, is the deliberation stage. Participants are assigned to small discussion groups, typically of around 8-10 people, before an expert witness provides evidence on the topic [5]. Participants then deliberate in a moderated table discussion, adding insight from their own lived experiences. This process iterates over the course of the assembly as participants interact with an increasing number of other participants, due to seat reshuffling in each iteration, and selected expert witnesses.

The deliberative process naturally lends itself to an algorithmic formulation. As we are interested in creating an agent-based simulated environment, we henceforth refer to CA participants as 'agents', with the terms to be considered interchangeable. Similarly, we refer to 'deliberation' as the entire deliberative process, including group allocation, group discussions, and expert input, while the closely-related 'opinion change' refers specifically to the (latent) updating of agents' opinions.

In the third stage, agents iteratively design policy recommendations based on what they have learned, synchronous with the latter parts of the second stage, and then vote on each policy design by majority vote. As the recommendation design process is complex [5], and can vary from assembly to assembly, we leave this aspect open to future work on voting mechanisms. However, it is worth noting that the final majority vote reflects the agents' final opinions at the end of the process, and is therefore closely linked to opinion changes driven by the deliberative stage.

3 MAIN RESULTS

In our results, we first demonstrate that existing opinion change models fail these criteria and therefore do not accurately capture the opinion change dynamics in real CAs. We then demonstrate that an opinion change model that incorporates a latent open-mindedness variable passes the criteria, by fitting to surveyed opinion data from the CA of Scotland and then simulating agents' behaviour across a range of participant and expert types.

Having derived a sufficient model with fixed global inputs, we then simulate what would happen if CA facilitators were to change these inputs. In particular, our simulations show that our model is robust to arbitrary orderings of expert speakers, with the criteria satisfied whether experts are chosen randomly, given extremist views, or alternate between partisan extremes. The model is also robust to participant types, whether agents are chosen randomly, feature bipartisan subgroups, or contain up to 20% committed extremists. In simulated results with a large sample size (around 100 agents), the model also indicates robustness to different types of group allocation algorithm: random allocation gives almost identical results to optimised allocation.

4 CONCLUSION

In this work, we first developed a model of CA deliberation that sufficiently explains observed features of CAs. Based on this model, we simulated results to explore how certain levers affect deliberation in the model space, in the spirit that general trends in the model space may give some counterfactual insight to real-world settings. Our simulated results indicate that, controlling for group size, adequate run-time, and optimised group allocation schedules, CAs reach robust results irrespective of participant typologies, expert witness selection, and expert witness ordering.

5 ACKNOWLEDGEMENTS

We would like to thank Dr Oliver Escobar, whose knowledge and experience of Citizens' Assemblies at all levels significantly informed our aims and model design.

REFERENCES

[1] Arash Abizadeh. 2021. Representation, Bicameralism, Political Equality, and Sortition: Reconstituting the Second Chamber as a Randomly Selected Assembly. *Perspectives on Politics* 19, 3 (2021), 791–806. https://doi.org/10.1017/S1537592719004626

[2] Jake Barrett, Kobi Gal, Loizos Michael, and Dan Vilenchik. 2024. Beyond the echo chamber: modelling open-mindedness in citizens' assemblies. *Autonomous Agents and Multi-Agent Systems* 38, 2 (2024), 30.

[3] Jake Barrett, Ya'akov Kobi Gal, Paul Gölz, Rose Hong, and Ariel D. Procaccia. 2023. Now we're talking: Better deliberation groups through submodular optimization. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*. 5490–5498. https://doi.org/10.1609/aaai.v37i5.25682

[4] George Butler, Gabriella Pigozzi, and Juliette Rouchier. 2020. An Opinion Diffusion Model with Vigilant Agents and Deliberation. In *Multi-Agent-Based Simulation XX*. 81–99. https://doi.org/10.1007/978-3-030-60843-9{_}7

[5] S Elstub, O Escobar, A Henderson, T Thorne, N Bland, and E Bowes. 2022. *Research Report on the Citizens' Assembly of Scotland.* Technical Report. Scottish Government. https://www.gov.scot/isbn/9781802018943

[6] Bailey Flanigan, Paul Gölz, Anupam Gupta, and Ariel D. Procaccia. 2020. Neutralizing self-selection bias in sampling for sortition. *Advances in Neural Information Processing Systems* 34 (2020). https://doi.org/10.48550/arXiv.2006.10498

[7] John Gastil and Erik Olin Wright. 2018. Legislature by lot: Envisioning sortition within a bicameral system. *Politics and Society* 46, 3 (2018), 303–330. https://doi.org/10.1177/0032329218789886

[8] Hannu Nurmi. 1998. Voting paradoxes and referenda. *Social Choice and Welfare* 15, 3 (1998), 333–350. https://doi.org/10.1007/s003550050109

[9] Eoin O'Malley, David M. Farrell, and Jane Suiter. 2020. Does talking matter? A quasi-experiment assessing the impact of deliberation and information on opinion change. *International Political Science Review* 41, 3 (2020), 321–334. https://doi.org/10.1177/0192512118824459

[10] Joshua Weinert. 2022. Radical Revisionism : A principle to escape echo chambers. *Studies in Philosophy and Economics* 4, 1 (2022), 22–27.