Epistemic Selection of Costly Alternatives: The Case of Participatory Budgeting (Extended Abstract)

JAAMAS Track

Simon Rey
ILLC, University of Amsterdam
The Netherlands

Ulle Endriss
ILLC, University of Amsterdam
The Netherlands

ABSTRACT

We initiate the study of participatory budgeting using the epistemic approach, where one interprets votes as noisy estimates of some ground truth regarding the objectively best set of projects to fund.

KEYWORDS

Participatory Budgeting; Epistemic Social Choice; Voting Theory

ACM Reference Format:
Simon Rey and Ulle Endriss. 2025. Epistemic Selection of Costly Alternatives: The Case of Participatory Budgeting (Extended Abstract): 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

The term *participatory budgeting* (PB) covers a range of mechanisms that involve citizens in public spending decisions [3]. People vote on grassroots projects (e.g., building a playground), and then the most popular projects—that fit a given budget constraint—get funded. PB is flourishing around the world, and it also received significant attention in the literature on computational social choice [1, 12].

Given people's votes, it is not always obvious which projects to fund. That is, there are many different voting rules one could use. The dominant approach to choosing a voting rule is the *normative approach*, where we evaluate rules in terms of how well they respect the subjective preferences of individuals. Under the *epistemic approach*, we instead assume that there is an objectively best budget allocation and we evaluate voting rules in terms of their ability to recover this *ground truth* given noisy estimates (votes). While the epistemic approach is a basic methodological staple in computational social choice more generally [5–7, 10, 13, 15], our recent paper on the topic [11] is the first attempt to systematically apply it to PB. Here we briefly outline our main findings.

But first we should address the elephant in the room: Selecting projects for PB seems to be an inherently subjective choice. So why does it nonetheless make sense to speak of an objective "ground truth" in this context? To appreciate this, consider the fact that whether a given project is a success will often become clear only some time *after* it has been implemented: Will people really use the compost bins? Will the new speed camera reduce accidents? So we might interpret a citizen's vote as an imperfect estimate of this objective quality of a given project. At the time of voting, citizens do not know what the best set of projects is, but each one of them is more likely to vote for a good rather than a bad set of projects.

Of course, this perspective captures the essence of PB only imperfectly, so in the context of political decision making the epistemic approach can and should merely supplement the normative approach, not replace it. But there are other scenarios of collective decision making that are mathematically equivalent to PB where the epistemic approach is less controversial. For instance, referees assessing grant proposals are often asked to predict the longterm impact of a proposed research project. A further example is the online EteRNA platform [14], where users can submit suggestions for folding a given protein. A subset of the proposed configurations is then synthesised in a laboratory to find the most stable ones. We can think of this as a PB process: the projects are the protein foldings; their cost is that of synthesising them; the budget limit is the amount of money allocated to this process; and the protein foldings submitted by a given user constitute their vote. Mathematically speaking, this is thus a well-defined PB process with a clear ground truth: a set of objectively most stable protein configurations.

2 PARTICIPATORY BUDGETING

A PB instance $I = \langle \mathcal{P}, c, b \rangle$ consists of a set $\mathcal{P}$ of *projects*, a *cost function* $c : \mathcal{P} \to \mathbb{N}$ mapping any given project p to its cost $c(p)$, and a *budget limit* $b \in \mathbb{N}$. If we ask several *voters* to each submit an *approval ballot* $A \subseteq \mathcal{P}$, we obtain a *profile*, i.e., a vector $\boldsymbol{A}$ of ballots, one for each voter. Based on this information, we need to select a *budget allocation* $\pi \subseteq \mathcal{P}$, consisting of projects to implement. It needs to be *feasible*, i.e., its total cost $c(\pi) = \sum_{p \in \pi} c(p)$ should not exceed b. If π cannot be enlarged without violating the budget limit, i.e., if $c(\pi \cup \{p\}) > b$ for all $p \in \mathcal{P} \setminus \pi$, we call π *exhaustive*.

Some of our results apply only to *unit-cost instances* $I = \langle \mathcal{P}, c, b \rangle$, where all projects p have the same cost $c(p) = \ell$ and b is divisible by this ℓ. As is well known, unit-cost instances are equivalent to instances of approval-based multi-winner voting [8].

A *voting rule* for PB is a function F that takes as input an instance I and a profile $\boldsymbol{A}$, and that returns a feasible budget allocation (or, in case of a tie, several such budget allocations). The most widely used voting rule in real-life PB elections is the *Greedy Rule*, where we accept projects in order of the number of approvals they received, skipping over any projects that would require us to exceed the budget limit. The rule favoured by most social choice theorists working in the area is probably the *Method of Equal Shares* [2, 9],[1] which—vey roughly speaking—works by endowing each voter with an equal amount of virtual currency and simulating the process of

[1]Strictly speaking, the Method of Equal Shares is an entire family of voting rules. Please refer to the full paper for more precise statements of definition and results [11].

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

groups of voters banding together to buy projects they approve of. A further example are *welfare-maximising rules*. Suppose we are willing to make the assumption that a voter with approval ballot $A \subseteq \mathcal{P}$ derives a level of satisfaction from budget allocation $\pi \subseteq \mathcal{P}$ that is equal to $c(\pi \cap A)$, the total cost of the selected projects she approves of.[2] Then we can define voting rules that maximise *utilitarian social welfare* (sum of individual levels of satisfaction) or *Nash social welfare* (product of individual levels of satisfaction).

3 THE EPISTEMIC APPROACH

Fix an instance I. We assume that there exists an objectively best feasible budget allocation for it, the *ground truth* $\pi^\star$, and we want every reasonable voting rule to select $\pi^\star$. The ground truth is not known, neither to the voters nor to the decision maker.

A *noise model* $\mathcal{M}$ is a generative model that produces random approval ballots for any given ground truth $\pi^\star$. Given $\pi^\star$ (which we cannot observe directly), every ballot A (and thus also every profile $\boldsymbol{A}$) has a certain likelihood of being generated by $\mathcal{M}$. Then any function F from profiles $\boldsymbol{A}$ to feasible allocations π that maximises the likelihood that π is the ground truth $\pi^\star$ that generated $\boldsymbol{A}$ is called a *maximum likelihood estimator* (MLE) for $\mathcal{M}$. We can think of F as voting rule. If $\mathcal{M}$ were to accurately model how voters form their preferences under a given ground truth, F would be a good voting rule that, for any given profile, returns the budget allocation most likely to be the objectively best feasible budget allocation.

To give a concrete example, when asked to generate a ballot A given ground truth $\pi^\star$, a natural noise model $\mathcal{M}$ might for each project $p \in \pi^\star$ include p in A with probability 0.6 and for each project $p \in \mathcal{P} \setminus \pi^\star$ include p in A with probability 0.3.

If a voting rule F is an MLE for a noise model $\mathcal{M}$ for which we have reason to believe that it is an accurate representation of how voters produce ballots, then this would constitute a strong argument for using F. But this is a very high bar: modelling the complex dynamics of voters choosing their ballots based on their imperfect perception of the world is challenging. Instead, we will be asking a more fundamental question: For a given voting rule F, does there exist *any noise model at all* for which F is an MLE?

4 IMPOSSIBILITY RESULTS

Our main results are impossibility results showing—quite surprisingly—that many voting rules cannot be interpreted as MLEs.

A first step in proving these impossibilities is the insight that an important result by Conitzer and Sandholm [6] for the classical setting of voting with ranked preferences can be adapted to the setting of PB. It states that a necessary condition for a rule F to be an MLE is for F to satisfy the axiom of *weak reinforcement*, which says that, if F returns the same allocation π for profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$ submitted by two disjoint electorates, then F must return π also when applied to the union of $\boldsymbol{A}$ and $\boldsymbol{A}'$. This allows us to prove:

Proposition 1. *Neither the Greedy Rule nor the Method of Equal Shares is an MLE (for any noise model).*

Let us illustrate the violation of weak reinforcement for the Greedy Rule. Consider an instance with projects $\mathcal{P} = \{p_1, p_2, p_3\}$, costs $c(p_1) = c(p_2) = 2$ and $c(p_3) = 3$, and budget limit $b = 4$. Suppose p_1 gets 10 approvals in $\boldsymbol{A}$ and just 1 in $\boldsymbol{A}'$, for p_2 it is the other way round, and p_3 gets 9 approvals in both profiles. Then, for both $\boldsymbol{A}$ and $\boldsymbol{A}'$, the Greedy Rule first picks the project with 10 approvals and then can still afford the second cheap project, so selects p_1 and p_2 in both cases. But for the joint profile it will select only p_3, which now has 18 approvals. The proof for the Method of Equal Shares is similar (we refer to the full paper [11] for all proofs).

Our welfare-maximising rules *do* satisfy weak reinforcement. Nevertheless, also here we obtain impossibility results:

Theorem 2. *For the cost-based definition of individual satisfaction, neither the rule maximising utilitarian social welfare nor the one maximising Nash social welfare is an MLE (for any noise model).*

The proof proceeds by exploring constraints on any noise model that could potentially be used to rationalise the voting rules under consideration, in the context of a carefully chosen PB instance.

5 POSSIBILITY RESULTS

Our impossibility results notwithstanding, we managed to obtain possibility results that apply under certain restrictive assumptions:

Proposition 3. *For unit-cost instances and under the assumption that the ground truth is an exhaustive budget allocation, the Greedy Rule as well as the rules maximising utilitarian and Nash social welfare for cost-based satisfaction are MLEs (for some noise model).*

For instance, for the Greedy Rule, the noise model rationalising it generates ballot A given ground truth $\pi^\star$ with a probability proportional to $2^{|A \cap \pi^\star|}$. As the Method of Equal Shares cannot guarantee outcomes to be exhaustive, assuming the ground truth to be exhaustive would not be meaningful in this context.

We also obtained a possibility result that applies without any restrictive assumptions, albeit only for a voting rule specifically designed for this purpose. This is the rule maximising Nash social welfare with individual satisfaction for allocation π by a voter with ballot A defined as $c(A \cap \pi)/c(\pi)$, i.e., the cost of the selected projects she approves of relative to the total cost of all selected projects.

Theorem 4. *For relative-cost-based individual satisfaction, the rule maximising Nash social welfare is an MLE (for some noise model).*

For ground truth $\pi^\star$, the noise model rationalising the rule generates ballot A with a probability proportional to $c(A \cap \pi^\star)$. This might be realistic when voters assess expensive projects more carefully and thus are more likely to classify them correctly. Despite Theorem 4, we stress that we advocate for this rule only in very specific circumstances, as it favours not exhausting the budget over funding projects of limited popularity—an unusual choice in PB.

6 OPEN PROBLEMS

We started to explore the potential of the epistemic approach in PB, but many questions remain open: Which other natural voting rules can be rationalised as MLEs (for some noise model)? What are natural noise models and what are the corresponding voting rules? None of the *proportional* rules we analysed is an MLE [11]—is this a coincidence or the sign of a broader impossibility? Finally, it would be interesting to study voting rules for PB with respect to their sample complexity [5] or their robustness against noise [4].

[2] In the full paper we also consider a second way of defining individual satisfaction, namely as $|\pi \cap A|$, i.e., as the *number* of selected projects our voter approves of [11].

REFERENCES

[1] Haris Aziz and Nisarg Shah. 2020. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer-Verlag.

[2] Markus Brill, Stefan Forster, Martin Lackner, Jan Maly, and Jannik Peters. 2023. Proportionality in Approval-Based Participatory Budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*.

[3] Yves Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.

[4] Ioannis Caragiannis, Christos Kaklamanis, Nikos Karanikolas, and George A. Krimpas. 2022. Evaluating Approval-Based Multiwinner Voting in Terms of Robustness to Noise. *Journal of Autonomous Agents and Multiagent Systems* 36, 1 (2022).

[5] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2016. When do Noisy Votes Reveal the Truth? *ACM Transactions on Economics and Computation* 4, 3 (2016), 15:1–15:30.

[6] Vincent Conitzer and Tuomas Sandholm. 2005. Common Voting Rules as Maximum Likelihood Estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 145–152.

[7] Edith Elkind and Arkadii Slinko. 2016. Rationalizations of voting rules. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 8.

[8] Martin Lackner and Piotr Skowron. 2022. *Approval-Based Committee Voting*. Springer-Verlag.

[9] Dominik Peters, Grzegorz Pierczynski, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*.

[10] Ariel D. Procaccia, Sashank J. Reddi, and Nisarg Shah. 2012. A Maximum Likelihood Approach for Selecting Sets of Alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 695–704.

[11] Simon Rey and Ulle Endriss. 2025. Epistemic Selection of Costly Alternatives: The Case of Participatory Budgeting. *Journal of Autonomous Agents and Multiagent Systems* 39, 1 (2025).

[12] Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv preprint arXiv:2303.00621* (2023).

[13] Zoi Terzopoulou and Ulle Endriss. 2019. Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*. 298–311.

[14] Wikipedia. 2025. Entry “EteRNA”. https://en.wikipedia.org/wiki/EteRNA. Last visited on 2025-01-03.

[15] H. Peyton Young. 1995. Optimal Voting Rules. *Journal of Economic Perspectives* 9, 1 (1995), 51–64.