Epistemic Selection of Costly Alternatives: The Case of Participatory Budgeting

Simon Rey and Ulle Endriss

s.j.rey@uva.nl and u.endriss@uva.nl

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam

We initiate the study of voting rules for participatory budgeting using the so-called epistemic approach, where one interprets votes as noisy reflections of some ground truth regarding the objectively best set of projects to fund. Using this approach, we first show that both the most studied rules in the literature and the most widely used rule in practice cannot be justified on epistemic grounds: they cannot be interpreted as maximum likelihood estimators, whatever assumptions we make about the accuracy of voters. Focusing then on welfare-maximising rules, we obtain both positive and negative results regarding epistemic guarantees.

1 Introduction

The term *participatory budgeting* (PB) has been used to describe a range of mechanisms that directly involve citizens in public spending decisions (Cabannes, 2004). The basic idea is that people can vote on grassroots projects (e.g., building a playground or funding a cultural event), each of which has a certain cost. The most popular projects—that fit a given budget constraint—then get funded. In recent years, PB has flourished around the world, making it one of the most popular tools of participatory democracy. At the same time, it also has received a lot of attention in the literature on (computational) social choice (Aziz and Shah, 2020).

Given the votes of the citizens, it is not always obvious how to decide which projects to fund (in other words, there are many different voting rules one could consider using). In this paper, we explore what is known as the *epistemic approach*—or *truth-tracking approach*—to analysing and designing voting rules for PB. Whether a given project is a success will usually become clear only some time *after* it has been realised: Will the new park bench really be used? Will the new zebra crossing really reduce accidents? We think of the citizens casting their votes as agents with bounded rationality who enjoy a noisy view of this *ground truth*. They do not know what the best set of projects to fund is, but each of them is more likely to vote for a good rather than a bad set of projects.

PB is but one example of a selection process of costly alternatives. For some others, the existence of a ground truth may be more obvious. Consider, for instance, the case of the EterRNA platform.[1] On this collaborative platform—the motivating example of the first epistemic analysis of multiwinner voting (Procaccia, Reddi and Shah, 2012)—users can submit different ways of folding a given protein. A subset of the proposed configurations is then synthesised in a laboratory to determine which are most stable. Since not all configuration would induce the same cost to synthesise, this setting is mathematically equivalent to PB. Moreover, there is a clear ground truth here: an objectively most stable set of protein configurations.

If we have a clear idea how these noisy views on the ground truth are generated (votes that are cast in the context of PB or protein configurations that are proposed for the EterRNA platform)—that is, if we have a well-defined *noise model*—we, in principle, are able to design a voting rule that maximises the likelihood of returning the ground truth, i.e., the best set of projects that fit our budget. Of course, in practice we do not have access to this noise model. Still, if a natural voting rule turns out to be such a *maximum likelihood estimator* (MLE) for a natural choice of noise model, then we can interpret this as an argument for using that rule. Similarly, if we can prove that for a given rule there does not exist *any* noise model that would make that rule an MLE, then we should interpret this as an argument against using that rule.

Contribution. We first analyse the rules most studied in the recent literature—the Method of Equal Shares (Peters, Pierczynski and Skowron, 2021; Brill, Forster, Lackner, Maly and Peters, 2023) and the Sequential Phragmén Rule (Los, Christoff and Grossi, 2022)—as well as the rule overwhelmingly used in practice—the greedy cost approval rule. Using a necessary condition provided by Conitzer and Sandholm (2005), we prove that none of these rules can be interpreted as an MLE, even for instances where all projects have the same cost (corresponding to multiwinner voting instances).

We then turn to a family of rules that all satisfy the aforementioned necessary condition: additive argmax rules. These can be thought of as welfare-maximising rules. We focus on eight specific rules based either on utilitarian or Nash social welfare. In the case of utilitarian social welfare, we show that it is impossible to find a noise model for which the most natural rules would be MLEs in the general case. For Nash welfare, the picture is brighter, since for two rules we exhibit noise models under which they are MLEs.

Related work. The study of PB rules is part of social choice theory, which more generally deals with the design and analysis of voting rules for different kinds of scenarios. As for every scenario many different rules can be—and have been—devised, it can be hard for the decision maker to select the rule to be used. To assist in this choice, two main approaches have been developed. The first one, the *axiomatic approach* (Arrow,

[1] https://eternagame.org – https://wikipedia.org/wiki/EteRNA

1951; Thomson, 2001), tries to identify voting rules that satisfy certain normative requirements. The second one, the *epistemic approach* (Elkind and Slinko, 2016; Pivato, 2019), seeks out rules that can recover a ground truth, assuming that the votes are noisy estimates of that ground truth. It is the latter approach we follow here.

Formal work on PB to date instead has followed the axiomatic approach, with a special focus on fairness (Aziz, Lee and Talmon, 2018; Peters, Pierczynski and Skowron, 2021; Hershkowitz, Kahng, Peters and Procaccia, 2021; Lackner, Maly and Rey, 2021; Los, Christoff and Grossi, 2022; Brill, Forster, Lackner, Maly and Peters, 2023), incentive compatibility (Fain, Goel and Munagala, 2016; Freeman, Pennock, Peters and Vaughan, 2021; Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019; Rey, Endriss and de Haan, 2021), and monotonicity requirements (Talmon and Faliszewski, 2019; Baumeister, Boes and Seeger, 2020; Rey, Endriss and de Haan, 2020). We refer the reader to the survey by Rey and Maly (2023) for more details.

The epistemic approach has been first applied to the standard voting model (Young, 1995; Conitzer and Sandholm, 2005; Caragiannis, Procaccia and Shah, 2014). Later on, other social choice scenarios have been investigated through the epistemic lens, notably multiwinner elections (Procaccia, Reddi and Shah, 2012; Caragiannis, Procaccia and Shah, 2013), and judgment aggregation (Bovens and Rabinowicz, 2006; Bozbay, Dietrich and Peters, 2014; Terzopoulou and Endriss, 2019). To the best of our knowledge, the only epistemic study of PB is the one section dedicated to the topic by Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019), though in the context of divisible PB, *i.e.*, when projects can be partially funded. Interestingly, they show that *knapsack voting*—a PB rule that resembles the greedy cost approval rule in the divisible setting—can be interpreted as a maximum likelihood estimator, while we will see that the greedy cost approval rule cannot in the context of *indivisible* PB.

2 Preliminaries

In this section, we recall the standard model of PB, and introduce maximum likelihood estimators (MLEs).

2.1 Participatory Budgeting

A PB problem is described by an instance $I = \langle \mathcal{P}, c, b \rangle$ where $\mathcal{P}$ is the set of available *projects*, $c : \mathcal{P} \to \mathbb{N}$ is the *cost function*—mapping any given project $p \in \mathcal{P}$ to its cost $c(p) \in \mathbb{N}$—and $b \in \mathbb{N}$ is the *budget limit*. We write $c(P)$ instead of $\sum_{p \in P} c(p)$ for sets of projects $P \subseteq \mathcal{P}$. For a given PB instance, we ask several *agents* to each submit an *approval ballot* $A \subseteq \mathcal{P}$, resulting in a vector $\boldsymbol{A}$ of ballots, one for each agent. Such a vector of approval ballots is called a *profile*. Given two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$, we use $\boldsymbol{A} \oplus \boldsymbol{A}'$ to denote the profile obtained by concatenating them.

Given an instance $I = \langle \mathcal{P}, c, b \rangle$, we need to select a subset of projects $\pi \subseteq \mathcal{P}$ to implement. Such a *budget allocation* π has to be *feasible*, i.e., we require $c(\pi) \leq b$. Let $\mathcal{A}(I) = \{\pi \subseteq \mathcal{P} \mid c(\pi) \leq b\}$ be the set of feasible budget allocations for I. Moreover, let

$\mathcal{A}_{EX}(I)$ be the set of all *exhaustive* budget allocations for I, that is, every $\pi \in \mathcal{A}(I)$ for which there is no project $p \in \mathcal{P} \setminus \pi$ such that $c(\pi \cup \{p\}) \leq b$.

Computing budget allocations is done by means of *PB rules.* An *irresolute* PB rule F is a function that takes as input a PB instance I and a profile $\boldsymbol{A}$ over I, and that returns a nonempty set of feasible budget allocations $F(I, \boldsymbol{A}) \subseteq \mathcal{A}(I)$. A rule is *exhaustive* if $F(I, \boldsymbol{A}) \subseteq \mathcal{A}_{EX}(I)$ for all I and $\boldsymbol{A}$.

Some of our results will apply only to unit-cost instances. An instance $I = \langle \mathcal{P}, c, b \rangle$ is a *unit-cost* instance if there exists an $\ell \in \mathbb{N}$ such that (i) $c(p) = \ell$ for all projects $p \in \mathcal{P}$ and (ii) $b \equiv 0 \bmod \ell$. This restriction is particularly interesting since unit-cost instances are equivalent to multiwinner voting problems where one needs to elect a committee of size k (Faliszewski, Skowron, Slinko and Talmon, 2017). Candidates can be thought of as projects of cost ℓ, so under a budget limit of $b = k \cdot \ell$ exhaustive budget allocations correspond to such committees.

2.2 The Truth-Tracking Perspective

According to the truth-tracking perspective, there exists an objectively best feasible budget allocation for every instance that is the outcome that every reasonable rule should select. Such a budget allocation is called the *ground truth* and is denoted by $\pi^\star$. The ground truth is not known, neither by the agents, nor by the decision maker. We will thus assess the quality of PB rules based on their ability to retrieve the ground truth given noisy votes.

Formally, a *noise model* $\mathcal{M}$ is a generative model that produces random approval ballots for a given instance and ground truth. We represent it as a probability distribution over all approval ballots. For a given instance $I = \langle \mathcal{P}, c, b \rangle$, ground truth $\pi^\star \in \mathcal{A}(I)$, and approval ballot $A \subseteq \mathcal{P}$, we denote by $\mathbb{P}_{\mathcal{M}}(A \mid \pi^\star, I)$ the probability for the noise model $\mathcal{M}$ to generate ballot A given I and $\pi^\star$. For profiles, ballots are drawn identically and independently from $\mathcal{M}$.

Suppose the noise model $\mathcal{M}$ indicates how the voters form their preferences. Then, a good rule should select the outcome that most likely would have generated the observed profile if it were the ground truth plugged into $\mathcal{M}$. This is the *maximum likelihood estimator* (MLE) for $\mathcal{M}$.

Definition 1 (Maximum likelihood estimators)**.** *For a noise model $\mathcal{M}$, the likelihood of a profile $\boldsymbol{A}$ over the instance I and a budget allocation $\pi \in \mathcal{A}(I)$ is defined as:*

$$L_{\mathcal{M}}(\boldsymbol{A}, \pi, I) = \prod_{A \in \boldsymbol{A}} \mathbb{P}_{\mathcal{M}}(A \mid \pi, I).$$

A PB rule F is said to be the MLE for $\mathcal{M}$, if for every instance I and every profile $\boldsymbol{A}$ we have:

$$F(I, \boldsymbol{A}) = \underset{\pi^\star \in \mathcal{A}(I)}{\operatorname{argmax}}\, L_{\mathcal{M}}(\boldsymbol{A}, \pi^\star, I).$$

In the context of the standard model of voting theory, Conitzer and Sandholm (2005) identified a necessary condition for a voting rule to be interpretable as an MLE: it should satisfy what we are going to call *weak reinforcement*. This result straightforwardly carries over to the PB setting.

Definition 2 (Weak reinforcement)**.** *A PB rule F is said to be satisfying weak reinforcement if and only if, for every instance I and every two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$, we have:*

$$F(I, \boldsymbol{A}) = F(I, \boldsymbol{A}') \implies F(I, \boldsymbol{A} \oplus \boldsymbol{A}') = F(I, \boldsymbol{A}).$$

Lemma 1 (Conitzer and Sandholm, 2005)**.** *If a PB rule F does not satisfy weak reinforcement, then there exists no noise model $\mathcal{M}$ for which F is the MLE.*

Note that this result applies for any set of possible ground truths, so also if we assume the ground truth to be exhaustive.

3 Proportional PB Rules

A large part of recent research on PB has been devoted to the study of *proportional* rules, *i.e.*, rules that treat groups of agents fairly. In this section we focus on the most prominent ones—Sequential Phragmén (Los, Christoff and Grossi, 2022; Brill, Forster, Lackner, Maly and Peters, 2023) and approval-based variants of the Method of Equal Shares (MES) (Peters, Pierczynski and Skowron, 2021)—and show that they cannot be interpreted as MLEs.

Definition 3 (Sequential Phragmén)**.** *Given an instance I and a profile $\boldsymbol{A}$, the Sequential Phragmén rule constructs budget allocations using the following continuous process.*

Voters receive money in a virtual currency. They all start with a budget of 0 and that budget continuously increases as time passes. At time t a voter will have received an amount t of money. For any time t, let $P_t^\star$ be the set of projects $p \in \mathcal{P}$ for which the approvers altogether have more than $c(p)$ money available. As soon as, for a given t, $P_t^\star$ is non-empty, if there exists a $p \in P_t^\star$ such that $c(\pi \cup \{p\}) > b$, the process stops; otherwise one project from $P_t^\star$ is selected, the budget of its approvers is set to 0, and the process resumes.

The outcome of Sequential Phragmén is the set of all budget allocations constructed by the procedure above (for all possible ways of breaking ties between the projects in $P_t^\star$).

Note that with this stopping condition (required to guarantee a property known as priceability), Sequential Phragmén is not exhaustive. But it is exhaustive on unit-cost instances.

Proposition 2. *There exists no noise model $\mathcal{M}$ such that Sequential Phragmén is the MLE for $\mathcal{M}$, not even on unit-cost instances with the additional assumption that the ground truth is exhaustive.*

Proof. Consider an instance I with four projects denoted by p_1, p_2, p_3 and p_4, all of cost 1, and budget limit $b = 3$.

Consider profiles $\boldsymbol{A}^1 = (\{p_1\}, \{p_1, p_3, p_4\}, \{p_2, p_3, p_4\}, \{p_2, p_3, p_4\}, \{p_2, p_3, p_4\})$ and $\boldsymbol{A}^2 = (\{p_2\}, \{p_1, p_3\}, \{p_1, p_3\}, \{p_1, p_4\}, \{p_1, p_3, p_4\})$. We claim that on both $\boldsymbol{A}^1$ and

$\boldsymbol{A}^2$ the Sequential Phragmén rule outputs $\pi = \{\{p_1, p_3, p_4\}\}$. For $\boldsymbol{A}^1$, after $1/2$ units of money have been distributed both p_3 and p_4 would have been bought at price $1/4$. Once an additional $1/4$ of money has been injected, project p_1 would be bought. For $\boldsymbol{A}^2$, first p_1 is bought at price $1/4$, then p_3 at price $1/3$. Finally, when $1/3$ additional units of money will have been distributed, project p_4 will be bought (at that time project p_2 has collected $1/4 + 2/3 < 1$ money). Now, let $\boldsymbol{A}^3 = \boldsymbol{A}^1 \oplus \boldsymbol{A}^2$. On $\boldsymbol{A}^3$, the first project to be bought is p_3 at price $1/7$. Then, once an extra $5/42$ amount of money has been distributed, project p_1 would be bought. At that time, no agents with a non-zero budget approve of p_4 but not p_2. Project p_2 will then be the last project selected (after another $2/21$ money has been injected). Overall the outcome would be $\pi' = \{\{p_1, p_2, p_3\}\} \neq \pi$. Note that both π and π' are exhaustive.

We have thus proved that Sequential Phragmén fails weak reinforcement. Lemma 1 then concludes the proof. □

The same holds for the MES-based rules. These rules are parametrised by a measure of the satisfaction of the voters. We call *satisfaction function* (on singletons), any mapping from projects p to satisfaction levels $\mu(p) \in \mathbb{R}_{>0}$.[2] We can now define MES with respect to satisfaction function μ.

Definition 4 (MES$_\mu$)**.** *Given an instance I, a profile $\boldsymbol{A} = (A_1, \ldots, A_n)$ with n agents, and a satisfaction function μ, MES$_\mu$ constructs budget allocations π, initially empty, iteratively as follows. Every agent i is initially assigned a budget $b_i = b/n$ of virtual money. Given a budget allocation π, a project $p \in \mathcal{P} \setminus \pi$ is said to be α-affordable, for $\alpha \in \mathbb{R}_{\geq 0}$ if:*

$$\sum_{i \mid p \in A_i} \min(b_i, \alpha \cdot \mu(p)) \quad \geq \quad c(p).$$

At a given round with current budget allocation π, if no project is α-affordable for any α, MES$_\mu$ terminates. Otherwise, let $P^\star$ be the set of projects that are α^*-affordable for a minimum α^*. The rule selects one project $p \in P^\star$ (π is updated to $\pi \cup \{p\}$), and every approver i of p sees their budget reduced by $\min(b_i, \alpha \cdot \mu(p))$.

The outcome of MES$_\mu$ is the set of all budget allocations constructed by the procedure above (for all possible ways of breaking tie between the projects in P).

Note that MES$_\mu$ fails to be an exhaustive rule, for any satisfaction function μ and even on unit-cost instances.

We show that for no μ can MES$_\mu$ be interpreted as an MLE.

Proposition 3. *For any given satisfaction function μ, there exists no noise model $\mathcal{M}$ such that MES$_\mu$ is the MLE for $\mathcal{M}$, not even on unit-cost instances.*

Proof. Consider an instance I with two projects denoted by p_1 and p_2, both of cost 1, and a budget limit $b = 2$. Let μ be an arbitrary satisfaction function.

[2] Note that Brill, Forster, Lackner, Maly and Peters (2023) give a more complete definition of satisfaction functions. Since we only need to discuss satisfaction of single projects here, we simplified the definition.

Consider the two profiles $\boldsymbol{A}^1 = (\{p_1\}, \{p_2\})$ and $\boldsymbol{A}^2 = (\{p_1, p_2\}, \{p_1, p_2\})$. We claim that on both of these profiles, MES_μ would return $\pi = \{\{p_1, p_2\}\}$. Indeed, on $\boldsymbol{A}^1$ both agents receive 1 unit of money and can both afford the project they approve of. On $\boldsymbol{A}^2$ both agents approve of all the projects and can afford them. Now, for $\boldsymbol{A}^3 = (\{p_1\}, \{p_2\}, \{p_1, p_2\}, \{p_1, p_2\}) = \boldsymbol{A}^1 \oplus \boldsymbol{A}^2$, we claim that MES_μ would return either $\{\{p_1\}\}$, $\{\{p_2\}\}$, or $\{\{p_1\}, \{p_2\}\}$. Here, the initial budget is $1/2$ for each agent. Thus, the approvers of either p_1 and p_2 collectively have $3/2$ units of money. Since $\mu(p) > 0$ for both p_1 and p_2 (by definition of satisfaction functions) and since we can choose α arbitrarily large, this implies that MES_μ would select either p_1 or p_2 in the first round. Let $p^\star$ be the selected project and p the other project. To buy $p^\star$, all its approvers paid $1/3$. The approvers of p are thus now left with $1/2 + 2 \cdot (1/2 - 1/3) = 1/2 + 1/3 < 1$, not enough to afford p. There is thus no way for MES_μ to return $\{\{p_1, p_2\}\}$ on $\boldsymbol{A}^3$.

We have thus proved that MES_μ fails weak reinforcement. Lemma 1 then concludes the proof. □

Note that for both proofs, we have $\pi \cap \pi' = \emptyset$, meaning that even resolute versions of the rules (obtained by introducing some form of tie-breaking) would fail weak reinforcement.

4 The Greedy Cost Approval Rule

Perhaps the most widely used rule in real-world PB processes is the *greedy cost approval rule* (Aziz and Shah, 2020).[3] Before defining it, we define the *approval score* of a project p for a profile $\boldsymbol{A}$ as $n_p^{\boldsymbol{A}} = |\{A \in \boldsymbol{A} \mid p \in A\}|$.

Definition 5 (Greedy cost approval rule)**.** *Given a strict ranking $\rhd$ on the set of all projects, GREED($\rhd$) is the budget allocation we obtain when we examine all projects in order of $\rhd$ and select a project whenever the total cost of the selected projects does not exceed the budget limit.*

For a given instance I and profile $\boldsymbol{A}$, the greedy cost approval rule F_{greed} returns the set of all GREED($\rhd$) for any $\rhd$ such that, for all $p, p' \in \mathcal{P}$, we have $p \rhd p'$ whenever $n_p^{\boldsymbol{A}} > n_{p'}^{\boldsymbol{A}}$.

Note that greedy cost approval rule is exhaustive.

Being extensively used in practice, this rule deserves a special focus in our formal analysis. However, we can easily show that it cannot be interpreted as an MLE.

Proposition 4. *There exists no noise model $\mathcal{M}$ such that the greedy cost approval rule F_{greed} is the MLE for $\mathcal{M}$.*

Proof. Consider an instance I with three projects denoted by p_1, p_2 and p_3, a budget limit of $b = 2$, and costs as shown in the table below. Moreover, consider two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$ with the following approval scores:

[3]The name is linked to the fact that this rule approximates the utilitarian social welfare used with cost satisfaction function (Rey and Maly, 2023).

	Cost	App. score in $\boldsymbol{A}$	App. score in $\boldsymbol{A}'$	App. score in $\boldsymbol{A} \oplus \boldsymbol{A}'$
p_1	1	4	1	5
p_2	1	1	4	5
p_3	2	3	3	6

One can check that greedy cost approval returns $\{\{p_1, p_2\}\}$ on both $\boldsymbol{A}$ and $\boldsymbol{A}'$. However, on $\boldsymbol{A} \oplus \boldsymbol{A}'$, the rule returns $\{\{p_3\}\}$. Greedy cost approval thus violates weak reinforcement and the claim immediately follows from Lemma 1. □

Observe that in the counterexample used in the proof there is a single winner. Hence, no refinement of F_{greed}—such as, say, the *leximax rule* (Rey, Endriss and de Haan, 2020)—will satisfy weak reinforcement either. This also applies for resolute variants. We discuss the case of unit-cost instances in Section 7.1.

5 Additive Argmax Rules

It does not seem easy to find rules that satisfy weak reinforcement. For the remainder of the paper we shall focus on what we shall call *additive argmax rules*, as we can show that they all satisfy weak reinforcement.

Definition 6 (Argmax rules)**.** *A PB rule F is called an argmax rule if there exists a function f, taking as input an instance I, a profile $\boldsymbol{A}$, and a budget allocation π and returning a number $f(I, \boldsymbol{A}, \pi) \in \mathbb{R}$, such that for all instances I and all profiles $\boldsymbol{A}$, we have:*

$$F(I, \boldsymbol{A}) = \underset{\pi \in \mathcal{A}(I)}{\operatorname{argmax}} f(I, \boldsymbol{A}, \pi).$$

Definition 7 (Additive argmax rules)**.** *An argmax rule defined via the function f is called additive if for every two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$ and every budget allocation π, we have:*

$$f(I, \boldsymbol{A} \oplus \boldsymbol{A}', \pi) = f(I, \boldsymbol{A}, \pi) + f(I, \boldsymbol{A}', \pi).$$

Note that *every* rule is an argmax rule—f can be the indicator function on the outcome of the rule for a given instance and profile—but not all are additive.

Are the additive argmax rules good candidates for being MLEs? Yes, they are, as we can show that additive argmax rules all satisfy weak reinforcement.

Proposition 5. *Every additive argmax rule satisfies weak reinforcement.*

Proof. Consider the additive argmax rule F defined with respect to the function f. Let I be an instance and $\boldsymbol{A}$ and $\boldsymbol{A}'$ two profiles over I such that $F(I, \boldsymbol{A}) = F(I, \boldsymbol{A}')$. Let us show that we also have $F(I, \boldsymbol{A} \oplus \boldsymbol{A}') = F(I, \boldsymbol{A})$.

Remember that F is additive. This implies the following:

$$\begin{aligned}
&\max_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A} \oplus \boldsymbol{A}', \pi) \\
&\qquad = \max_{\pi \in \mathcal{A}(I)} \big(f(I, \boldsymbol{A}, \pi) + f(I, \boldsymbol{A}', \pi) \big) \\
&\qquad \leq \max_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A}, \pi) + \max_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A}', \pi). \qquad (1)
\end{aligned}$$

Moreover, given that $F(I, \boldsymbol{A}) = F(I, \boldsymbol{A}')$, we have:

$$\operatorname*{argmax}_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A}, \pi) = \operatorname*{argmax}_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A}', \pi).$$

Hence, the same budget allocations achieve the maximum value of the two terms appearing on the righthand side of inequality (1). So the inequality in fact is an equality and these same budget allocations also maximise the score for the concatenated profile. We thus have: $F(I, \boldsymbol{A}) = \operatorname{argmax}_{\pi \in \mathcal{A}(I)} f(I, \boldsymbol{A} \oplus \boldsymbol{A}', \pi) = F(I, \boldsymbol{A} \oplus \boldsymbol{A}')$. □

Next, we introduce and study several concrete examples of additive argmax rules. We consider two types of rules, based either on Nash or utilitarian social welfare (Moulin, 1988).

6 Nash Social Welfare

We first study rules based on the Nash social welfare. It tries to reach balanced outcomes by measuring the score of a budget allocation as the product of the agents' levels of satisfaction. The concept of Nash social welfare provides strong guarantees in the context of fair division of indivisible items (Caragiannis, Kurokawa, Moulin, Procaccia, Shah and Wang, 2019). It has also been identified as an appealing rule for PB (Rosenfeld and Talmon, 2021).

6.1 Cardinality and Cost Satisfaction

We first consider two usual measures of satisfaction: based on the cardinality and on the cost of approved and selected projects (Talmon and Faliszewski, 2019; Peters, Pierczynski and Skowron, 2021; Brill, Forster, Lackner, Maly and Peters, 2023). This gives rise to two additive argmax rules, F^N_{app} and F^N_{cost}, defined via the following functions:

$$f^N_{app}(I, \boldsymbol{A}, \pi) = \prod_{A \in \boldsymbol{A}} |A \cap \pi|,$$

$$f^N_{cost}(I, \boldsymbol{A}, \pi) = \prod_{A \in \boldsymbol{A}} c(A \cap \pi).$$

Stated in this form, it might not be immediately obvious that these rules are additive. It becomes clear once we express the above products as sums of logarithms.

Note that, under the usual assumption that all projects are approved by at least one agent, these two rules are exhaustive.

We start our investigation by introducing a noise model, denoted by $\mathcal{M}_{Ncost}$, for which for all I, A and $\pi^\star$, we have:

$$\mathbb{P}_{\mathcal{M}_{Ncost}}(A \mid \pi^\star, I) = \frac{1}{Z_{\pi^\star}^{Ncost}} c(A \cap \pi^\star),$$

where, $Z_{\pi^\star}^{Ncost}$ is a suitable normalisation factor ensuring that $\mathcal{M}_{Ncost}$ is a well-defined probability distribution.

Under this noise model, the probability of generating a given ballot A increases with the cost of the ground-truth projects in A. The intuition here is that voters may reflect more carefully on expensive projects and thus are more likely to make correct choices for them. Moreover, the probability of generating A increases linearly in the "quality" of A. How realistic this is, is open to debate. On the one hand, this avoids having to assume extremely high probabilities for correctly identifying particularly expensive projects. On the other hand, the probability of generating a ballot that is completely wrong (in the sense of not including even a single ground-truth project) is zero.

Under $\mathcal{M}_{Ncost}$, maximising the likelihood would be similar to maximising the cost-approval Nash social welfare of a budget allocation. For this to yield an MLE result, we need $Z_{\pi^\star}^{Ncost}$ to be independent of $\pi^\star$. Let's look at it then.

Lemma 6. *For the noise model $\mathcal{M}_{Ncost}$ to be a well-defined probability distribution, it should be the case that*

$$Z_{\pi^\star}^{Ncost} = 2^{|\mathcal{P}|-1} c(\pi^\star).$$

Proof. Consider any instance $I = \langle \mathcal{P}, c, b \rangle$, $A \subseteq \mathcal{P}$ and $\pi^\star \in \mathcal{A}(I)$. For $\mathcal{M}_{Ncost}$ to be a probability distribution, it should be the case that $Z_{\pi^\star}^{Ncost} = \sum_{A \subseteq \mathcal{P}} c(|A \cap \pi^\star|)$. Remember that there are $2^{|\mathcal{P}|}$ subsets of projects and that any project $p \in \mathcal{P}$ appears in exactly half of them. Each time a project $p \in \pi^\star$ appears in a subset $A \subseteq \mathcal{P}$, its contribution to the value of $Z_{\pi^\star}^{Ncost}$ is exactly $c(p)$. We thus have:

$$Z_{\pi^\star}^{Ncost} = \sum_{p \in \pi^\star} 2^{|\mathcal{P}|-1} c(p) = 2^{|\mathcal{P}|-1} c(\pi^\star). \qquad \square$$

This result tells us that the normalisation factor of the noise model $\mathcal{M}_{app}$ depends on the ground truth, the consequence being that the value of the likelihood is impacted by the ground truth one is considering when computing the MLE. In particular, we cannot conclude that the Nash cost-approval maximising rule is the MLE for this noise model.

Are there specific cases for which the normalisation factor is independent of the ground truth? Yes, for unit-cost instances, as then all exhaustive allocations have the same cost.

Proposition 7. *Under the assumption that the ground truth is exhaustive, both F_{app}^N and F_{cost}^N are the MLE for the noise model $\mathcal{M}_{Ncost}$ for unit-cost instances.*

Proof. For a unit-cost instance I, by Lemma 6, for every two exhaustive budget allocations $\pi, \pi' \in \mathcal{A}_{EX}(I)$, we have $Z_{\pi}^{Ncost} = Z_{\pi'}^{Ncost}$. Finally, for any profile $\boldsymbol{A}$, we have then:

$$\begin{aligned}\operatorname*{argmax}_{\pi \in \mathcal{A}_{EX}(I)} L_{\mathcal{M}_{Ncost}}(\boldsymbol{A}, \pi, I) &= \operatorname*{argmax}_{\pi \in \mathcal{A}_{EX}(I)} \prod_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{Z_{\pi}^{Ncost}} \\ &= \operatorname*{argmax}_{\pi \in \mathcal{A}_{EX}(I)} \prod_{A \in \boldsymbol{A}} c(A \cap \pi).\end{aligned}$$

The last line follows from the fact that F_{app}^{N} is exhaustive.

Given that on unit-cost instances F_{app}^{N} and F_{cost}^{N} coincide, the result also applies to F_{app}^{N}. □

The fact that F_{app}^{N} and F_{cost}^{N} are MLEs for $\mathcal{M}_{app}$ only on some restricted instances is the first hint of a general impossibility result. Indeed, we can actually show that there are no noise models for which these rule are MLEs.

Theorem 8. *There is no noise model $\mathcal{M}$ such that either F_{app}^{N} or F_{cost}^{N} is the MLE for $\mathcal{M}$, not even on unit-cost instances.*

Proof. Consider an instance I with two projects p_1 and p_2 of cost 1, and with $b = 2$. Let $\mathcal{M}$ be a generic noise model, and denote by $\mathbb{P}_{A}^{\pi}$ the value of $\mathbb{P}_{\mathcal{M}}(A \mid \pi, I)$ for any A and π. To simplify notation, we omit braces around sets.

For the noise model $\mathcal{M}$ to be a well-defined probability distribution, the following should be satisfied:

$$\mathbb{P}_{\emptyset}^{p_1} + \mathbb{P}_{p_1}^{p_1} + \mathbb{P}_{p_2}^{p_1} + \mathbb{P}_{p_1,p_2}^{p_1} = 1, \tag{2}$$

$$\mathbb{P}_{\emptyset}^{p_1,p_2} + \mathbb{P}_{p_1}^{p_1,p_2} + \mathbb{P}_{p_2}^{p_1,p_2} + \mathbb{P}_{p_1,p_2}^{p_1,p_2} = 1. \tag{3}$$

Now, on the single-agent profile $\boldsymbol{A} = (\emptyset)$, F_{app}^{N} returns $\mathcal{A}(I)$. So for F_{app}^{N} to be the MLE of $\mathcal{M}$, we must have $\mathbb{P}_{\emptyset}^{p_1} = \mathbb{P}_{\emptyset}^{p_1,p_2}$. Moreover, on $\boldsymbol{A} = (\{p_1\})$, we have $F_{app}^{N}(I, \boldsymbol{A}) = \{\{p_1\}, \{p_1, p_2\}\}$, so $\mathbb{P}_{p_1}^{p_1} = \mathbb{P}_{p_1}^{p_1,p_2}$. Using these two equalities and by subtracting (3) from (2), we get:

$$(\mathbb{P}_{p_2}^{p_1} - \mathbb{P}_{p_2}^{p_1,p_2}) + (\mathbb{P}_{p_1,p_2}^{p_1} - \mathbb{P}_{p_1,p_2}^{p_1,p_2}) = 0. \tag{4}$$

Now, since $F_{app}^{N}(I, (\{p_2\})) = \{\{p_2\}, \{p_1, p_2\}\}$, we must have $\mathbb{P}_{p_2}^{p_1,p_2} > \mathbb{P}_{p_2}^{p_1}$. For $\boldsymbol{A} = (\{p_1, p_2\})$, we have $F_{app}^{N}(I, \boldsymbol{A}) = \{\{p_1, p_2\}\}$. We can then derive $\mathbb{P}_{p_1,p_2}^{p_1,p_2} > \mathbb{P}_{p_1,p_2}^{p_1}$. These two last inequalities contradict (4). It is then impossible for F_{app}^{N} to be the MLE of $\mathcal{M}$ on I. From the unit-cost assumption, it is clear that this also applies to F_{cost}^{N}. □

6.2 Relative Satisfaction

We also consider "relative" variants of our two rules, where the satisfaction of an agent is expressed in terms of the proportion of the outcome that satisfies her. We denote these rules by $\tilde{F}^N_{app}$ and $\tilde{F}^N_{cost}$, and they are the argmax rules defined in terms of the following functions:

$$\tilde{f}^N_{app}(I, \boldsymbol{A}, \pi) = \prod_{A \in \boldsymbol{A}} \frac{|A \cap \pi|}{|\pi|},$$
$$\tilde{f}^N_{cost}(I, \boldsymbol{A}, \pi) = \prod_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{c(\pi)}.$$

These rules are inspired by the concept of *relative satisfaction* introduced by Lackner, Maly and Rey (2021). However, while the denominator was defined with respect to the ballot in their work, we define it here w.r.t. the budget allocation.

Note that these rules can lead to extreme behaviours. For example, consider an instance with budget limit b that is even and a set of projects $\mathcal{P} = \{p^\star\} \cup \{p_1, \ldots, p_b\}$. Consider the two agent profile $\boldsymbol{A}$ such that $A_1 = \{p^\star\} \cup \{p_1, p_3, \ldots, p_{b-1}\}$ and $A_2 = \{p^\star\} \cup \{p_2, p_4, \ldots, p_b\}$. According to the relative rules, selecting just $p^\star$ is better than anything else. Even if this can seem extreme, these rules can still be justified when considering voters who would rather save the public money than use it on projects they do not approve (this corresponds to associating a strong rejection, rather than indifference, with the action of not approving a project). Note that this implies that the rules are not exhaustive.

Let us first investigate the rule $\tilde{F}^N_{cost}$. We will continue using the noise model $\mathcal{M}_{Ncost}$ introduced earlier. Recall the expression we found for the normalisation factor $Z^{Ncost}_{\pi^\star}$ in Lemma 6. We now plug it into the definition of $\mathcal{M}_{Ncost}$ so that for all I, A and $\pi^\star$, we have:

$$\mathbb{P}_{\mathcal{M}_{Ncost}}(A \mid \pi^\star, I) = \frac{1}{2^{|\mathcal{P}|-1}} \frac{c(A \cap \pi^\star)}{c(\pi^\star)}.$$

Using this expression, we can show that the relative Nash cost-approval maximising rule is the MLE for $\mathcal{M}_{Ncost}$.

Theorem 9. *The relative Nash cost-approval maximising rule $\tilde{F}^N_{cost}$ is the MLE for the noise model $\mathcal{M}_{Ncost}$.*

Proof. Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance. The likelihood of $\boldsymbol{A}$ and $\pi \in \mathcal{A}(I)$ under the noise model $\mathcal{M}_{Ncost}$ is:

$$L_{\mathcal{M}_{Ncost}}(\boldsymbol{A}, \pi, I) = \left(\frac{1}{2^{|\mathcal{P}|-1}}\right)^{|\boldsymbol{A}|} \prod_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{c(\pi)}.$$

Since the first multiplicative factor in the above expression is constant over all budget allocations, we have:

$$\operatorname*{argmax}_{\pi \in \mathcal{A}(I)} L_{\mathcal{M}_{Ncost}}(\boldsymbol{A}, \pi, I) = \operatorname*{argmax}_{\pi \in \mathcal{A}(I)} \prod_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{c(\pi)}. \qquad \square$$

We have finally been able to find a PB rule that can be interpreted as an MLE. In the following we will show a similar result for $\tilde{F}^N_{app}$. For this rule we introduce a new noise model: $\mathcal{M}_{Napp}$. It is such that for any I, A and $\pi^\star$, we have:

$$\mathbb{P}_{\mathcal{M}_{Napp}}(A \mid \pi^\star, I) = \frac{1}{Z^{Napp}_{\pi^\star}} |A \cap \pi^\star|,$$

where $Z^{Napp}_{\pi^\star}$ is a normalisation factor.

The proof techniques we used above also work for $\mathcal{M}_{Napp}$.

Theorem 10. *The relative Nash approval maximising rule* $\tilde{F}^N_{app}$ *is the MLE for the noise model* $\mathcal{M}_{Napp}$.

Proof. One can easily check that for the noise model $\mathcal{M}_{Napp}$ to be a well-defined probability distribution, we must have:

$$Z^{Napp}_{\pi^\star} = 2^{|\mathcal{P}|-1} |\pi^\star|.$$

We thus have: Hence, given a profile $\boldsymbol{A}$, we have:

$$\begin{aligned}\operatorname*{argmax}_{\pi \in \mathcal{A}(I)} L_{\mathcal{M}_{Napp}}(\boldsymbol{A}, \pi) &= \operatorname*{argmax}_{\pi \in \mathcal{A}(I)} \prod_{A \in \boldsymbol{A}} \frac{1}{2^{|\mathcal{P}|-1}} \frac{|A \cap \pi^\star|}{|\pi^\star|} \\ &= \operatorname*{argmax}_{\pi \in \mathcal{A}(I)} \prod_{A \in \boldsymbol{A}} \frac{|A \cap \pi|}{|\pi|}.\end{aligned}$$

□

7 Utilitarian Social Welfare

Let us now turn to the analysis of additive argmax rules defined in terms of utilitarian social welfare.

7.1 Cardinality and Cost Satisfaction

Following Talmon and Faliszewski (2019), we define the approval maximising rule F_{app} and the cost-approval maximising rule F_{cost} as the argmax rules determined by f_{app} and f_{cost}, respectively:

$$\begin{aligned}f_{app}(I, \boldsymbol{A}, \pi) &= \sum_{A \in \boldsymbol{A}} |A \cap \pi|, \\ f_{cost}(I, \boldsymbol{A}, \pi) &= \sum_{A \in \boldsymbol{A}} c(A \cap \pi).\end{aligned}$$

Note that these two rules are exhaustive.

Following an idea of Conitzer and Sandholm (2005), we define the noise model $\mathcal{M}_{app}$ such that for any $I = \langle \mathcal{P}, c, b \rangle$, $\pi^\star \in \mathcal{A}(I)$, and approval ballot $A \subseteq \mathcal{P}$:

$$\mathbb{P}_{\mathcal{M}_{app}}(A \mid \pi^\star, I) = \frac{1}{Z^{app}_{\pi^\star}} \prod_{p \in \mathcal{P}} 2^{\mathbb{1}_{p \in A \cap \pi^\star}} = \frac{1}{Z^{app}_{\pi^\star}} 2^{|A \cap \pi^\star|},$$

where $Z^{app}_{\pi^\star}$ is a suitable normalisation factor ensuring that $\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{app}}(A \mid \pi^\star, I) = 1$. $\mathcal{M}_{app}$ is a particularly simple manifestation of what we would expect to see in a noise model: any possible ballot might be generated in principle, but the probability of generating ballot A increases (significantly) with the number of ground-truth projects in A.

With this noise model, maximising the likelihood may appear to have the same effect as maximising the approval score of a budget allocation. It could then be that the approval maximising rule is the MLE for $\mathcal{M}_{app}$. However, for this to hold, one has to have a closer look at the normalisation factor.

Lemma 11. *For the noise model $\mathcal{M}_{app}$ to be a well-defined probability distribution, it should be the case that:*

$$Z^{app}_{\pi^\star} = 2^{|\mathcal{P}|} \left(\frac{3}{2}\right)^{|\pi^\star|}.$$

Proof. Consider any instance $I = \langle \mathcal{P}, c, b \rangle$. Let $A \subseteq \mathcal{P}$ be an approval ballot and $\pi^\star \in \mathcal{A}(I)$ a ground truth. For $\mathcal{M}_{app}$ to be a probability distribution, it should be the case that:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{app}}(A \mid \pi^\star, I) = 1 \iff Z^{app}_{\pi^\star} = \sum_{A \subseteq \mathcal{P}} 2^{|A \cap \pi^\star|}.$$

Let's do some combinatorics. For $k \in \{0, \dots, |\pi^\star|\}$, how many subsets of $\mathcal{P}$ will intersect with $\pi^\star$ on exactly k projects? A suitable subset will consists of k projects from $\pi^\star$ that make up the intersection and any number $j \in \{0, \dots, |\mathcal{P}| - |\pi^\star|\}$ of projects from $\mathcal{P} \setminus \pi^\star$ that do not have any impact on the intersection. Each such subset of projects contributes 2^k to the value of $Z^{app}_{\pi^\star}$. We thus have:

$$\begin{aligned} Z^{app}_{\pi^\star} &= \sum_{k=0}^{|\pi^\star|} 2^k \sum_{j=0}^{|\mathcal{P}|-|\pi^\star|} \binom{|\pi^\star|}{k} \binom{|\mathcal{P}| - |\pi^\star|}{j} \\ &= \sum_{k=0}^{|\pi^\star|} \binom{|\pi^\star|}{k} 2^k \sum_{j=0}^{|\mathcal{P}|-|\pi^\star|} \binom{|\mathcal{P}| - |\pi^\star|}{j} \\ &= 2^{|\mathcal{P}|-|\pi^\star|} \sum_{k=0}^{|\pi^\star|} \binom{|\pi^\star|}{k} 2^k = 2^{|\mathcal{P}|} \left(\frac{3}{2}\right)^{|\pi^\star|}. \end{aligned}$$

□

The normalisation factor of $\mathcal{M}_{app}$ thus depends on the ground truth. We cannot thus conclude that the approval maximising rule is the MLE for this noise model. This is not the case on unit-cost instances.

Proposition 12. *Under the assumption that the ground truth is exhaustive, both F_{app} and F_{cost} are the MLE for the noise model $\mathcal{M}_{app}$ for unit-cost instances.*

Proof. For a unit-cost instance I, every two exhaustive budget allocations π and $\pi' \in$

$\mathcal{A}_{EX}(I)$, by virtue of Lemma 11, we have $Z_{\pi}^{app} = Z_{\pi'}^{app}$. So, for any profile $\boldsymbol{A}$, we have:

$$\begin{aligned}\underset{\pi \in \mathcal{A}_{EX}(I)}{\text{argmax}}\ L_{\mathcal{M}_{app}}(\boldsymbol{A}, \pi, I) &= \underset{\pi \in \mathcal{A}_{EX}(I)}{\text{argmax}} \prod_{A \in \boldsymbol{A}} \frac{1}{Z_{\pi}^{app}} 2^{|A \cap \pi|} \\ &= \underset{\pi \in \mathcal{A}_{EX}(I)}{\text{argmax}}\ 2^{\sum_{A \in \boldsymbol{A}} |A \cap \pi|} \\ &= F_{app}(I, \boldsymbol{A}).\end{aligned}$$

The last line follows from the fact that F_{app} is exhaustive.

F_{app} coincides thus with the MLE on I for the noise model $\mathcal{M}_{app}$. Moreover, since F_{app} and F_{cost} coincide on unit-cost instances, the result also applies to F_{cost}. □

Can we find an impossibility result similar to the one we had for F_{app}^N and F_{cost}^N? It is actually easy to see that the proof we gave for Theorem 8 also works for both F_{app} and F_{cost}.

Theorem 13. *There is no noise model $\mathcal{M}$ such that either F_{app} or F_{cost} is the MLE for $\mathcal{M}$, not even on unit-cost instances.*

Proof. Consider the instance I used in the proof of Theorem 8. We claim that for all profiles that are relevant for the proof, F_{app} and F_{app}^N coincide. We list them below:

$$\begin{aligned} F_{app}(I, (\{p_1\})) &= \{\{p_1\}, \{p_1, p_2\}\} = F_{app}^N(I, (\{p_1\})), \\ F_{app}(I, (\{p_2\})) &= \{\{p_2\}, \{p_1, p_2\}\} = F_{app}^N(I, (\{p_2\})), \\ F_{app}(I, (\{p_1, p_2\})) &= \{\{p_1, p_2\}\} = F_{app}^N(I, (\{p_1, p_2\})), \\ F_{app}(I, (\emptyset)) &= \mathcal{A}(I) = F_{app}^N(I, (\emptyset)). \end{aligned}$$

Given that on unit-cost instances F_{app} and F_{cost} coincide, the result also applies to F_{cost}. □

We conclude with the observation that the greedy cost approval rule and F_{app} coincide on unit-cost instances. Thus, both Proposition 12 and Theorem 13 apply to the former as well.

7.2 Relative Satisfaction

Let us conclude by discussing the relative variants of the utilitarian rules. These two argmax rules, denoted by $\tilde{F}_{app}$ and $\tilde{F}_{cost}$, are defined as expected via $\tilde{f}_{app}(I, \boldsymbol{A}, \pi) = \sum_{A \in \boldsymbol{A}} \frac{|A \cap \pi|}{|\pi|}$ and $\tilde{f}_{cost}(I, \boldsymbol{A}, \pi) = \sum_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{c(\pi)}$.

For the same reasons as for $\tilde{F}_{app}^N$ and $\tilde{F}_{cost}^N$, these two rules are not exhaustive. Analysing the epistemic status of these rules however turns out to be rather intricate, even on unit-cost instances. Indeed, for the relative variants of the utilitarian rules it is less clear what a suitable noise model might look like, especially due to the complications related to the potential normalisation factor. We leave the analysis of these rule as interesting open problems.

	Sequential Phragmén	MES_μ For all μ	Greedy Cost Approval	Approval-max Standard $\sum$	Approval-max Standard $\prod$	Approval-max Relative $\sum$	Approval-max Relative $\prod$	Cost-approval-max Standard $\sum$	Cost-approval-max Standard $\prod$	Cost-approval-max Relative $\sum$	Cost-approval-max Relative $\prod$
Unit-cost$_{EX}$	✗	–	✓	✓	✓	–	–	✓	✓	–	–
Unit-cost	✗	✗	✗	✗	✗	?	✓	✗	✗	?	✓
General case	✗	✗	✗	✗	✗	?	✓	✗	✗	?	✓

Table 1: Summary of the results. The sum $\sum$ and product $\prod$ symbols represent the utilitarian and the Nash variant of a welfare-based rule. A check-mark ✓ indicates that there exists a noise model for which the rule is an MLE and a cross-mark ✗ the fact that it is impossible to find such a noise model. The $_{EX}$ subscript signifies that we make the additional assumption that the ground truth is exhaustive. This assumption would not be meaningful for non-exhaustive rules. Remember that Sequential Phragmén is exhaustive on unit-cost instances.

8 Conclusion

We have initiated the study of PB through the truth-tracking lens. For a total of eleven rules, we investigated whether they could be interpreted as MLEs. Whenever they could not, we tried to identify specific conditions under which they would serve as MLEs. All our results are summarised in Table 1.

There is still quite some work to be done regading the study of MLEs in the context of PB. Filling out the missing cells in Table 1 is one thing. Our work also shows some kind of tension between efficiency requirements (exhaustiveness) and truth-tracking ability: the two rules that we proved to be MLEs both fail exhaustiveness. This interaction deserves further study. Then, on top of the MLE concept, the epistemic approach also offers several other ways of studying voting rules. As we have seen, finding rules that are MLEs is not an easy task. It could be that this is simply too demanding a requirement. Instead, other criteria that have been studied in the literature on epistemic social choice could be applied to the PB setting. For instance, it could be interesting to study PB rules with respect to their sample complexity (Caragiannis, Procaccia and Shah, 2013) or their robustness against noise (Caragiannis, Kaklamanis, Karanikolas and Krimpas, 2020)—a criterion that is somewhat similar to the MLE requirement but easier to satisfy. All of these constitute interesting directions for future work on a topic that is still very much under-studied and deserving of further attention.

References

Kenneth J. Arrow. 1951. *Social Choice and Individual Values*. Cowles Foundation.

Haris Aziz, Barton E. Lee, and Nimrod Talmon. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 23–31.

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.

Dorothea Baumeister, Linus Boes, and Tessa Seeger. 2020. Irresolute Approval-Based Budgeting (Extended Abstract). In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1774–1776.

Luc Bovens and Wlodek Rabinowicz. 2006. Democratic Answers to Complex Questions—An Epistemic Perspective. *Synthese* 150, 1 (2006), 131–153.

Irem Bozbay, Franz Dietrich, and Hans Peters. 2014. Judgment Aggregation in Search for the Truth. *Games and Economic Behavior* 87 (2014), 571–590.

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)*.

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

Ioannis Caragiannis, Christos Kaklamanis, Nikos Karanikolas, and George A. Krimpas. 2020. Evaluating Approval-Based Multiwinner Voting in Terms of Robustness to Noise. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 74–80.

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 1–32.

Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2013. When do Noisy Votes Reveal the Truth?. In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*. 143–160.

Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2014. Modal Ranking: A Uniquely Robust Voting Rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 616–622.

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.

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.

Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*. 384–399.

Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 2, 27–47.

Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2021. Truthful Aggregation of Budget Proposals. *Journal of Economic Theory* 193 (2021), 105234.

Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting: Voting Mechanisms for Participatory Budgeting. *ACM Transaction on Economics and Computation* 7, 2 (2019), 8:1–8:27.

D. Ellis Hershkowitz, Anson Kahng, Dominik Peters, and Ariel D. Procaccia. 2021. District-Fair Participatory Budgeting. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*.

Martin Lackner, Jan Maly, and Simon Rey. 2021. Fairness in Long-Term Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*.

Maaike Los, Zoé Christoff, and Davide Grossi. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*.

Hervé Moulin. 1988. *Axioms of Cooperative Decision Making.* Cambridge University Press.

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.*

Marcus Pivato. 2019. Realizing Epistemic Democracy. In *The Future of Economic Design*, Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker (Eds.). Springer-Verlag, 103–112.

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.

Simon Rey, Ulle Endriss, and Ronald de Haan. 2020. Designing Participatory Budgeting Mechanisms Grounded in Judgment Aggregation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

Simon Rey, Ulle Endriss, and Ronald de Haan. 2021. Shortlisting Rules and Incentives in an End-to-End Model for Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*.

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

Ariel Rosenfeld and Nimrod Talmon. 2021. What Should We Optimize in Participatory Budgeting? An Experimental Study. *arXiv preprint arXiv:2111.07308* (2021).

Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2181–2188.

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.

William Thomson. 2001. On the Axiomatic Method and its Recent Applications to Game Theory and Resource Allocation. *Social Choice and Welfare* 18, 2 (2001), 327–386.

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