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; Rey and Maly, 2023).

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). The common approach to choosing a voting rule is the *normative approach*, where we encode normatively desirable properties of voting rules, e.g., properties related to fairness, in the form of axioms and then analyse to what extent a given voting rule will satisfy those axioms. In this paper, we instead explore what is known as the *epistemic approach*—or *truth-tracking approach*—to the analysis and design of voting rules for PB. Both

approaches play a central role in the broader field of computational social choice (Brandt, Conitzer, Endriss, Lang and Procaccia, 2016). But in the specific context of PB, the epistemic approach may, at first sight, seem less natural a choice. So let us briefly outline why we nonetheless believe that it is important to explore its potential.

One reason is that, even in the familiar setting of PB with citizens voting for projects they personally would like to receive public funding, we might interpret these votes as evidence for objective quality. Indeed, whether a given project is a success will often become clear only some time *after* it has been realised: Will local citizens actually use the compost bins? Will the number of mosquitoes go down? Will the new speed camera reduce the number of accidents?[1] So we might 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.

Furthermore, PB as a means to select projects to receive public funding is but one example of a selection process of *costly alternatives*. For some other such processes—that are, at least mathematically, equivalent to PB—the existence of a ground truth may be more obvious. Consider, for instance, the case of the Eterna platform.[2] On this collaborative platform, 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. One can think of this as a PB process: the projects are the different protein foldings; their cost is the cost of synthesising them; the budget limit is the amount of money that is allocated to this process; and finally, the protein foldings submitted by a user constitute their approval ballot. Mathematically speaking, this is thus a well-defined PB process. Moreover, there is a clear ground truth here: a set of objectively most stable protein configurations. Interestingly, this is also the motivating example of the first epistemic analysis of multi-winner voting rules,[3] though in a setting without costs (Procaccia, Reddi and Shah, 2012).

A last example is that of a selection committee for research grant proposals. In such a committee, the members have to decide which of the grants should be funded. The grants typically have different costs, and there is a maximum amount of money that can be allocated. Then one might argue that the proposals have a "ground truth" probability of success that can only be observed *a posteriori*. At selection time, the committee members observe noisy signals regarding this ground truth through the reports of reviewers, and they need to make a decision based on this information.

If we have a clear idea how these noisy views on the ground truth are generated (votes

[1] All these examples are taken from projects that were brought to the vote in Toulouse 2019; see jeparticipe.metropole.toulouse.fr/processes/bp2019 for more details, and in particular the file "Catalogue des 30 idées soumises au vote". The data is also hosted on the website pabulib.org (Faliszewski, Flis, Peters, Pierczyński, Skowron, Stolicki, Szufa and Talmon, 2023), though the description of the projects is not available there.

[2] See eternagame.org and wikipedia.org/wiki/EteRNA for more information.

[3] In multi-winner voting with approval ballots (Lackner and Skowron, 2022), we ask each voter which of the candidates up for election they approve of and then, for some fixed k, need to choose a committee of k candidates. This scenario is mathematically equivalent to PB when all projects have the same cost.

that are cast in the context of PB, protein configurations that are proposed for the EterRNA platform, or approvals submitted for some grant proposals)—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 that rule will be less suitable in contexts where it is reasonable to assume that there exists some form of ground truth regarding the right decision to be taken.

Contribution. We first analyse the rule overwhelmingly used in practice around the world—the greedy cost approval rule—as well as 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). 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 multi-winner 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 on either 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 are able to identify noise models under which they are MLEs.

Still, the overall picture pained by our analysis is largely negative, demonstrating that it is difficult to design voting rules for PB with good epistemic performance, certainly if this requirement comes on top of other, particularly common normative, requirements. Identifying conditions under which more positive results can be achieved—and ideally conditions that can be satisfied for specific application scenarios where the epistemic perspective is most relevant—therefore presents itself as an important challenge for future work in the field.

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. Given that for every application domain 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 *normative* or *axiomatic approach* (Arrow, 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, assum-

ing that the votes received 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 multi-winner elections (Procaccia, Reddi and Shah, 2012; Caragiannis, Procaccia and Shah, 2016), 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 this is not the case for the greedy cost approval rule in the context of *indivisible* PB.

2 Preliminaries

In this section, we recall the standard model of PB, used in much of the literature (see, e.g., Talmon and Faliszewski, 2019; Aziz and Shah, 2020; Rey and Maly, 2023), and then define what it means for a PB rule to be a maximum likelihood estimator (MLE).

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, i.e., of allocations $\pi \in \mathcal{A}(I)$ 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 multi-winner voting problems where one needs to elect a committee of some fixed 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

Under the truth-tracking perspective, we assume that, for every given instance I, there exists an objectively best feasible budget allocation, the so-called *ground truth* $\pi^\star$, and we would like every reasonable rule to select $\pi^\star$. The ground truth is not known, neither to the agents nor to 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 generated by the noise model, ballots are drawn identically and independently from $\mathcal{M}$.

Suppose the noise model $\mathcal{M}$ indicates how the voters form their preferences for any given ground truth.[4] Then a good rule should select the outcome that most likely would have generated the observed profile we observe 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:

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

[4] We stress that here we are making the (implicit) assumption that every voter's ballot is generated by the same noise model. This simplifying assumption, though common in the literature (Elkind and Slinko, 2016), does not account for the fact that different voters may have different expertise. Specifically, in the context of PB, a voter living in a specific district may be more likely to accurately judge the quality of a project based in the same district than a voter who lives elsewhere.

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.

As a first application of this lemma, let us consider what perhaps is the most widely used rule in real-world PB elections, namely the *greedy cost approval rule* F_{greed} (Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019; Aziz and Shah, 2020). Under this rule, we select projects in order of the number of approvals they receive, skipping over any projects that would lead us to exceed the budget.[5] Note that F_{greed} is exhaustive. Unfortunately, we can easily show that it cannot be interpreted as an MLE.

Proposition 2. *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 $b = 4$, and costs as shown in the table below. Moreover, consider two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$ in which the approval scores are as presented below.

	Cost	**Approval score in $\boldsymbol{A}$**	**Approval score in $\boldsymbol{A}'$**	**Approval score in $\boldsymbol{A} \oplus \boldsymbol{A}'$**
p_1	2	10	1	11
p_2	2	1	10	11
p_3	3	9	9	18

One can easily check that F_{greed} would return $\{\{p_1, p_2\}\}$ on both $\boldsymbol{A}$ and $\boldsymbol{A}'$. However, on the joint profile $\boldsymbol{A} \oplus \boldsymbol{A}'$, it would return $\{\{p_3\}\}$. Weak reinforcement is thus violated. The claim now follows from Lemma 1. □

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

[5] The name of the rule is linked to the fact that it can be thought of as approximating the maximum utilitarian social welfare under the cost satisfaction function (Rey and Maly, 2023). These concepts will be formally introduced in Section 4.2.

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 also 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 π using the following continuous process.*

Initially, π is the empty set. 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 together have at least $c(p)$ money available. As soon as, for a given t, the set $P_t^\star$ 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 $P_t^\star$).

Observe that, due to the fact that each voter's budget increases continuously, the approvers whose project p is selected at a given time t in fact will have accumulated *exactly* $c(p)$, not more, which justifies setting their budgets back all the way to 0. Also note that with the above stopping condition—required to guarantee a property known as priceability (Los, Christoff and Grossi, 2022)—Sequential Phragmén is not exhaustive. But it is exhaustive on unit-cost instances.

Unfortunately, whatever assumptions we are willing to make regarding the noise model generating the approval sets of voters, Sequential Phragmén cannot be used to track the ground truth, not even for unit-cost instances.

Proposition 3. *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 of $b = 3$. We consider two profile, $\boldsymbol{A}^1$ and $\boldsymbol{A}^2$, as presented below, where $\boldsymbol{A}^1$ is on the left and $\boldsymbol{A}^2$ is on the right.

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1^1	✓	✗	✗	✗
A_2^1	✓	✗	✓	✓
A_3^1	✗	✓	✓	✓
A_4^1	✗	✓	✓	✓
A_5^1	✗	✓	✓	✓
$b = 3$				

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1^2	✗	✓	✗	✗
A_2^2	✓	✗	✓	✗
A_3^2	✓	✗	✓	✗
A_4^2	✓	✗	✗	✓
A_5^2	✓	✗	✓	✓
$b = 3$				

We claim that on both profile $\boldsymbol{A}^1$ and profile $\boldsymbol{A}^2$, Sequential Phragmén outputs a single budget allocation, namely $\pi = \{p_1, p_3, p_4\}$. Remember that, in the unit-cost setting, Sequential Phragmén is exhaustive. The budget allocation π indeed is exhaustive.

For $\boldsymbol{A}^1$, after $1/2$ units of money have been distributed, both p_3 and p_4 will have been bought at a price of $1/4$ each. Once an additional $1/4$ units of money have been injected, project p_1 is bought as well, making π the unique budget allocation returned by Sequential Phragmén.

For $\boldsymbol{A}^2$, first p_1 is bought at a price of $1/4$, then p_3 at a price of $1/3$. Finally, after $1/3$ additional units of money have been distributed, project p_4 is bought (at that time, the supporters of project p_2 have collected $1/4 + 2/3 < 1$ money). The final outcome is thus indeed $\{\pi\}$.

Now, consider the joint profile $\boldsymbol{A}^3 = \boldsymbol{A}^1 \oplus \boldsymbol{A}^2$, and let us detail the run of Sequential Phragmén on I and $\boldsymbol{A}^3$. The first project to be bought is p_3 at price $1/7$. Then, once an extra $5/42$ units of money have been distributed, project p_1 is bought as well. At that time, the supporters of p_4 who do not approve of p_2 have no money since they approve of p_1. On the other hand, the only supporter of p_2 who does not approve of p_4 has a strictly positive amount of money. Project p_2 will then be the last project selected (after another $2/21$ units of money have been injected). Overall, the outcome will be $\{\{p_1, p_2, p_3\}\} \neq \{\pi\}$.

We have thus proven that the Sequential Phragmén rule fails weak reinforcement. Lemma 1 then concludes the proof. Note that all budget allocations we considered are exhaustive; the result thus also applies if we only focus on exhaustive ground truths. □

The same bad news apply to 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 $\mu(p) \in \mathbb{R}_{>0}$.[6] 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:*

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

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 ties between the projects in $P^\star$).

[6] 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.

Note that MES_μ 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_μ be interpreted as an MLE.

Proposition 4. *For any given satisfaction function μ, there is no noise model $\mathcal{M}$ such that MES_μ 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.

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 p_1 collectively have $3/2$ units of money, and the same is true for p_2. 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 need to pay $1/3$. The approvers of p are thus now left with $1/2 + 2 \cdot (1/2 - 1/3) = 1/2 + 1/3 < 1$, which is 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 shown that MES_μ fails weak reinforcement. Lemma 1 then concludes the proof. □

Note that for the proofs of both the results in this section, the outcomes on the initial profiles are always disjoint from the outcomes on the joint profiles. This implies that even resolute versions of the rules (obtained by introducing some form of tie-breaking), or any refinement, would also fail weak reinforcement.

4 Monotonic Argmax Rules

As we have seen, our first obstacle to finding rules that are MLEs is that none of the proportional rules we considered satisfy weak reinforcement. Aiming for more positive results, we will now follow a different approach: instead of checking whether known PB rules do satisfy weak reinforcement, we will start from rules we know satisfy it, and investigate their epistemic abilities. To this end, we will thus focus on *monotonic argmax rules*, a large class of rules, all of which satisfy weak reinforcement. We start by defining both what an argmax rule is and what makes such a rule monotonic.

Definition 5 (Monotonic 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:*

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

An argmax rule defined via the function f is called monotonic *if for every two profiles $\boldsymbol{A}$ and $\boldsymbol{A}'$ and every two budget allocations π and π', the following two conditions hold:*

1. $$\left.\begin{array}{r} f(I,\boldsymbol{A},\pi) < f(I,\boldsymbol{A},\pi') \\ f(I,\boldsymbol{A}',\pi) < f(I,\boldsymbol{A}',\pi') \end{array}\right\} \Longrightarrow f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi) < f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi'),$$

2. $$\left.\begin{array}{r} f(I,\boldsymbol{A},\pi) = f(I,\boldsymbol{A},\pi') \\ f(I,\boldsymbol{A}',\pi) = f(I,\boldsymbol{A}',\pi') \end{array}\right\} \Longrightarrow f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi) = f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi').$$

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

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

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

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

Consider a budget allocation $\pi \in F(I,\boldsymbol{A})$. Since F is an argmax rule, for all $\pi' \in \mathcal{A}(I)\setminus F(I,\boldsymbol{A})$, we know that $f(I,\boldsymbol{A},\pi) > f(I,\boldsymbol{A},\pi')$. This also holds for $\boldsymbol{A}'$. By the definition of a monotonic argmax rule, we immediately get that $f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi) > f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi')$.

Moreover, for any two budget allocations $\pi,\pi' \in F(I,\boldsymbol{A})$, we have $f(I,\boldsymbol{A},\pi) = f(I,\boldsymbol{A},\pi')$. The same also holds for $\boldsymbol{A}'$. Since F is monotonic, we thus immediately get that $f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi) = f(I,\boldsymbol{A}\oplus\boldsymbol{A}',\pi')$.

Overall, we proved that (i) no budget allocation that was winning under $\boldsymbol{A}$ or $\boldsymbol{A}'$ is dominated under $\boldsymbol{A}\oplus\boldsymbol{A}'$, and (ii) that all budget allocations that were winning under $\boldsymbol{A}$ and $\boldsymbol{A}'$ all have the same score. It is then immediate that $F(I,\boldsymbol{A}\oplus\boldsymbol{A}') = F(I,\boldsymbol{A})$. □

In what follows we will introduce and study several concrete examples of monotonic argmax rules, based either on the Nash or on the utilitarian social welfare.

4.1 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 the fair allocation 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).

4.1.1 Cardinality and Cost Satisfaction

We first consider two common measures of satisfaction: based on the cardinality and on the cost of approved and selected projects (Talmon and Faliszewski, 2019; Rey and

Maly, 2023). This gives rise to two argmax rules, F^N_{app} and F^N_{cost}, defined via the following functions:

$$f^N_{app}(I, \boldsymbol{A}, \pi) = \begin{cases} \sum_{A \in \boldsymbol{A}} \log(|A \cap \pi|) & \textit{if for all } A' \in \boldsymbol{A}, |A' \cap \pi| \neq 0, \\ 0 & \textit{otherwise.} \end{cases}$$

$$f^N_{cost}(I, \boldsymbol{A}, \pi) = \begin{cases} \sum_{A \in \boldsymbol{A}} \log(c(A \cap \pi)) & \textit{if for all } A' \in \boldsymbol{A}, c(A' \cap \pi) \neq 0, \\ 0 & \textit{otherwise.} \end{cases}.$$

It can be checked that these two argmax rules are indeed monotonic. From Proposition 5, we thus know that they satisfy weak reinforcement. Note that one would need to take the exponential of these function to make the Nash social welfare appear.

Interestingly, under the common and natural 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:

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

where, $Z^{Ncost}_{\pi^\star}$ 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, of course, is open to debate. On the one hand, this avoids having to assume extremely high probabilities for correctly identifying particularly expensive projects (in the case where the relationship would not be linear). On the other hand, the probability of generating a ballot that is completely wrong (a ballot 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. However, for this intuitive connection to hold, it should be the case that the normalisation factor $Z^{Ncost}_{\pi^\star}$ is independent of $\pi^\star$. So let us look at it in more detail.

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

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

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

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{Ncost}}(A \mid \pi^\star, I) = 1 \quad \Longleftrightarrow \quad Z^{Ncost}_{\pi^\star} = \sum_{A \subseteq \mathcal{P}} c(|A \cap \pi^\star|).$$

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:

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

This result tells us that the normalisation factor of the noise model $\mathcal{M}_{Ncost}$ depends on the ground truth, meaning 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 since not all feasible budget allocations have the same cost.

Are there specific cases for which the normalisation factor is independent of the ground truth? Yes, namely 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 of the noise model $\mathcal{M}_{Ncost}$ for unit-cost instances.*

Proof. Let I be a unit-cost instance I. Consider any two exhaustive budget allocations $\pi, \pi' \in \mathcal{A}_{EX}(I)$. Since we have $|\pi| = |\pi'| = c(\pi) = c(\pi')$, Lemma 6 entails that $Z_{\pi}^{Ncost} = Z_{\pi'}^{Ncost}$. For any profile $\boldsymbol{A}$, we have then:

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

The last line follows form the fact that F_{cost}^N is exhaustive.

Given that F_{app}^N and F_{cost}^N coincide on unit-cost instances, this also applies to F_{app}^N. $\square$

The fact that F_{app}^N and F_{cost}^N are MLEs for $\mathcal{M}_{Ncost}$ only under some restricted hypothesis is the first hint of a general impossibility result. Indeed, we can 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 of $\mathcal{M}$, not even on unit-cost instances.*

Proof. Consider an instance I with two projects p_1 and p_2 of cost 1, and a budget limit of $b = 2$. 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 two equalities should be satisfied:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_A^{p_1} = 1 \quad \Leftrightarrow \quad \mathbb{P}_{\emptyset}^{p_1} + \mathbb{P}_{p_1}^{p_1} + \mathbb{P}_{p_2}^{p_1} + \mathbb{P}_{p_1,p_2}^{p_1} = 1, \tag{1}$$

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_A^{p_1,p_2} = 1 \quad \Leftrightarrow \quad \mathbb{P}_{\emptyset}^{p_1,p_2} + \mathbb{P}_{p_1}^{p_1,p_2} + \mathbb{P}_{p_2}^{p_1,p_2} + \mathbb{P}_{p_1,p_2}^{p_1,p_2} = 1. \tag{2}$$

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 (2) from (1), we get:

$$(\mathbb{P}_{p_2}^{p_1} - \mathbb{P}_{p_2}^{p_1,p_2}) + (\mathbb{P}_{p_1,p_2}^{p_1} - \mathbb{P}_{p_1,p_2}^{p_1,p_2}) = 0. \tag{3}$$

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 (3). 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. □

This impossibility result concludes this part of our analysis. We will now consider "normalised" satisfaction functions.

4.1.2 Normalised Satisfaction

In the hope of overcoming Theorem 8, we also consider normalised variants of the cardinality and cost satisfaction functions. In these variants, the satisfaction of an agent is expressed in terms of the proportion of the outcome that satisfies them. The rules induced by these normalised satisfaction functions are denoted by $\tilde{F}_{app}^N$ and $\tilde{F}_{cost}^N$, and they are the argmax rules defined in terms of the following functions:

$$\tilde{f}_{app}^N(I, \boldsymbol{A}, \pi) = \begin{cases} \sum_{A \in \boldsymbol{A}} \log(|A \cap \pi|) - \log(|\pi|) & \textit{if for all } A' \in \boldsymbol{A}, |A' \cap \pi| \neq 0, \\ 0 & \textit{otherwise.} \end{cases}$$

$$\tilde{f}_{cost}^N(I, \boldsymbol{A}, \pi) = \begin{cases} \sum_{A \in \boldsymbol{A}} \log(c(A \cap \pi)) - \log(c(\pi)) & \textit{if for all } A' \in \boldsymbol{A}, c(A' \cap \pi) \neq 0, \\ 0 & \textit{otherwise.} \end{cases}.$$

These rules are inspired by the concept of *relative satisfaction* introduced by Lackner, Maly and Rey (2021). However, while in the cited work the authors normalised the satisfaction with respect to the ballot, we define it here with respect to the budget allocation.

It is worth noting that the rules $\tilde{F}_{app}^N$ and $\tilde{F}_{cost}^N$ can lead to extreme behaviours. For example, consider an instance with budget limit b, that we assume to be even, and a set of projects $\mathcal{P} = \{p^\star\} \cup \{p_1, \ldots, p_b\}$ of arbitrary cost lower than b. Consider the two-agent profile $\boldsymbol{A}$ such that:

$$A_1 = \{p^\star\} \cup \{p_1, p_3, \ldots, p_{b-1}\} \qquad A_2 = \{p^\star\} \cup \{p_2, p_4, \ldots, p_b\}.$$

Then, according to both 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 public money than use it on projects they do not approve. This would correspond to associating a strong rejection, rather than indifference, with the action of not approving a project.

Note that this example also implies that the rules are not exhaustive, even on unit-cost instances.

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. Plugging it into the definition of $\mathcal{M}_{Ncost}$, we obtain the following expression for any instance I, approval ballot A, and ground truth $\pi^\star$:

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

From this, it should be clear that $\tilde{F}^N_{cost}$ is the MLE of $\mathcal{M}_{Ncost}$.

Theorem 9. *The rule $\tilde{F}^N_{cost}$ is the MLE of the noise model $\mathcal{M}_{Ncost}$.*

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

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

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

$$\operatorname*{argmax}_{\pi \in \mathcal{A}(I)} L_{\mathcal{M}_{Ncost}}(\boldsymbol{A}, \pi, I) = \operatorname*{argmax}_{\pi \in \mathcal{A}(I)} \prod_{A \in \boldsymbol{A}} \frac{c(A \cap \pi)}{c(\pi)} = \tilde{F}^N_{cost}(I, \boldsymbol{A}).$$

Thus, maximising the likelihood under $\mathcal{M}_{Ncost}$ is the same as maximising the social welfare in the sense of $\tilde{F}^N_{cost}$. □

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 denoted by $\mathcal{M}_{Napp}$. It is such that for any instance I, approval ballot A, and ground truth $\pi^\star$, 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 suitable normalisation factor.

Using a similar proof technique as above, we show that $\tilde{F}^N_{app}$ is the MLE of $\mathcal{M}_{Napp}$.

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

Proof. Let us first compute the exact value of the normalisation factor $Z^{Napp}_{\pi^\star}$. For the noise model $\mathcal{M}_{Napp}$ to be a well-defined probability distribution, the following must hold:

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

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

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

This immediately implies that $\tilde{F}^N_{app}$ is the MLE of $\mathcal{M}_{Napp}$. □

This concludes our study of the Nash social welfare. We now turn to the more standard measure of the social welfare, the utilitarian social welfare.

4.2 Utilitarian Social Welfare

Let us now turn to the analysis of monotonic argmax rules defined in terms of utilitarian social welfare. As we did before, we will consider different satisfaction functions.

4.2.1 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:

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

It should be clear that these two rules are monotonic argmax rule (since the score they use is additive) and thus both satisfy weak reinforcement. Note also that these two rules are exhaustive.

Following an idea developed by Conitzer and Sandholm (2005) for scoring rules (in the standard voting framework), we introduce the noise model $\mathcal{M}_{app}$. It is defined such that for any instance $I = \langle\mathcal{P}, c, b\rangle$, approval ballot $A \subseteq \mathcal{P}$, and ground truth $\pi^\star \in \mathcal{A}(I)$:

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

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

$\mathcal{M}_{app}$ 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 exponentially with the size of the intersection between A and the ground truth.

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 the case that the approval maximising rule is the MLE of $\mathcal{M}_{app}$. 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 must 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:

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

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 consist 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:

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

where the last two lines are derived from the binomial expansion. □

The normalisation factor of $\mathcal{M}_{app}$ thus depends on the ground truth, since not all feasible budget allocations have the same cardinality. We thus cannot conclude that the approval maximising rule is the MLE of this noise model.

Interestingly, this is not the case on unit-cost instances when considering exhaustive budget allocations.

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

Proof. Let I be a unit-cost instance. For any 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:

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

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

F_{app} thus coincides 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}. □

But this result is only half satisfactory. 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 of $\mathcal{M}$, 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, this also applies to F_{cost}. □

It is interesting to note that the same result applies if one maximises the total relative satisfaction as introduced by Lackner, Maly and Rey (2021), i.e., the sum over all agents of the number (or cost) of approved and selected projects divided by the number (or cost) of approved projects. Indeed, one can easily check that for all relevant profiles this maximisation rule coincides with F_{app}^N. This trivially also applies when using the Nash social welfare with relative satisfaction.

As an aside, observe that the greedy cost approval rule F_{greed} discussed in Section 2.2 coincides with F_{app} on unit-cost instances. Thus, both Proposition 12 and Theorem 13 apply to F_{greed} as well, and we can refine Proposition 2 as follows.

Corollary 14. *Under the assumption that the ground truth is exhaustive, the greedy cost approval rule F_{greed} is the MLE for $\mathcal{M}_{app}$ for unit-cost instances.*

Moreover, for unconstrained ground truths, there is no noise model $\mathcal{M}$ such that F_{greed} is the MLE for $\mathcal{M}$, not even on unit-cost instances.

These insights apply also for any refinements of these rules, such as the *leximax rule* introduced by Rey, Endriss and de Haan (2020).

4.2.2 Normalised Satisfaction

Let us conclude our formal analysis by briefly mentioning the utilitarian social welfare with the normalised satisfaction functions. For the same reasons as for $\tilde{F}^N_{app}$ and $\tilde{F}^N_{cost}$, the corresponding utilitarian rules are not exhaustive. Analysing the epistemic status of these rules however turns out to be rather intricate, even on unit-cost instances. Indeed, it is less clear what a suitable noise model might look like, especially due to the complications related to the potential normalisation factor. Exploring these rules remains an interesting open problem.

5 Conclusion

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

There is still some work to be done regarding the study of MLEs in the context of PB. The most obvious task would be to answer the open questions corresponding to the two missing cells in Table 1. In the case of our positive results, one could try to come up with more natural noise models for which the rules are MLEs, e.g., noise models for which not all the supersets of the ground truth have the same probability. Our work also shows a certain tension between efficiency requirements (exhaustiveness) and truth-tracking ability: the two rules that we proved to be MLEs (for the general case) both fail exhaustiveness. This interaction deserves further study. Similarly, we found that the most studied rules that enforce proportional representation fail to satisfy weak reinforcement. It would be interesting to investigate whether or not this constitutes a formal incompatibility.

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 might be the case 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 as well. For instance, it could be interesting to study PB rules with respect to their sample complexity (Caragiannis, Procaccia and Shah, 2016) or their robustness against noise (Caragiannis, Kaklamanis, Karanikolas and Krimpas, 2022)—a criterion that is somewhat similar to the MLE requirement but easier to satisfy.

	Sequential Phragmén	**MES$_\mu$** (for all μ)	**Greedy Cost Approval**	**Approval-max** Standard $\sum$	**Approval-max** Standard $\prod$	**Approval-max** Normalised $\sum$	**Approval-max** Normalised $\prod$	**Cost-approval-max** Standard $\sum$	**Cost-approval-max** Standard $\prod$	**Cost-approval-max** Normalised $\sum$	**Cost-approval-max** Normalised $\prod$
Unit-cost$_{EX}$	✗	–	✓	✓	✓	–	–	✓	✓	–	–
Unit-cost	✗	✗	✗	✗	✗	?	✓	✗	✗	?	✓
General case	✗	✗	✗	✗	✗	?	✓	✗	✗	?	✓

Table 1: Summary of results: The sum Σ and product Π 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; this is indicated by a bar –. Remember that Sequential Phragmén is exhaustive on unit-cost instances. Question marks ? indicate open problems.

All of these constitute interesting directions for future work on a topic that is still very much under-studied and deserving of further attention.

Acknowledgments. We are grateful to multiple anonymous reviewers for the valuable feedback we received on this work.

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.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge University Press.

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. 2022. Evaluating approval-based multiwinner voting in terms of robustness to noise. *Journal of Autonomous Agents and Multiagent Systems* 36 (2022), 1.

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. 2014. Modal Ranking: A Uniquely Robust Voting Rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 616–622.

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.

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, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory budgeting: Data, Tools, and Analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2667–2674.

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).*

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

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).*

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.