Time-Constrained Participatory Budgeting Under Uncertain Project Costs

Dorothea Baumeister , **Linus Boes** and **Christian Laußmann**
Heinrich-Heine University Düsseldorf
{d.baumeister, linus.boes, christian.laussmann}@uni-duesseldorf.de

Abstract

In participatory budgeting the stakeholders collectively decide which projects from a set of proposed projects should be implemented. This decision underlies both time and monetary constraints. In reality it is often impossible to figure out the exact cost of each project in advance, it is only known after a project is finished. To reduce risk, one can implement projects one after the other to be able to react to higher costs of a previous project. However, this will increase execution time drastically. We generalize existing frameworks to capture this setting, study desirable properties of algorithms for this problem, and show that some desirable properties are incompatible. Then we present and analyze algorithms that trade-off desirable properties.

1 Introduction

As an introduction, we consider the following example, which will accompany us through our work.

Example 1. *The country "Participation Island" wants to address climate change. The government decides to provide* 200 *million dollars for new mobility projects that reduce emissions. Climate change will not wait forever, so these projects should be finished within the next five years. Companies send proposals for projects that each cost at most* 200 *million and can be realized within the given five years: bike sharing (BS); an express train route (ET); electric vehicle charging stations (EV); and development of fuel-cell vehicles (FV). The citizens of Participation Island are asked to vote on these projects by approving each project they like. The government then chooses which projects should be realized according to the voters' preferences and the constraints on time and money.*

This problem has been extensively studied in Knapsack and participatory budgeting literature. Yet, it is not realistic.

Example 1 (continued)**.** *While it is very predictable how much EV and BS will cost (because one knows the cost for each station/bike and the number of stations/bikes), there is some uncertainty about the other projects. It is known that the train route goes through undeveloped swamps, where it is unclear yet, how many foundations are needed. The company which proposed the project guarantees that the costs are between* 80 *and* 150 *million. Hydrogen technologies are still in the development and while the company is sure that they can develop the cars, they are unsure about the exact costs. They guarantee the costs to be between* 100 *and* 160 *million.*

The exact costs of a project are usually only known after finishing it. A common way to accommodate this uncertainty is to submit the project application with an *estimated* cost. However, if the estimation is far too expensive, the company will probably not get the bid. On the other hand, if it is far too low, the company is unable to realize the project and will have to ask for additional money.[1] So a more realistic approach is to provide a cost range, which moves the uncertainty risk to Participation Island. The question is now: *how should Participation Island select projects?* Of course, the limit of 200 million should not be exceeded, or at least this should be improbable. To reduce this risk, they could implement the projects sequentially, and wait for the exact costs before deciding which project to start next. But this is a slow process, and the time limit of five years should also not be exceeded.

Our Contribution. We develop a framework for time-constrained participatory budgeting under uncertain project costs based on the existing framework for approval-based participatory budgeting by Talmon and Faliszewski [2019]. We propose desirable properties for such budgeting methods and explore which properties can be combined and which are incompatible. In addition, we experimentally evaluate algorithms for uncertain project costs that sequentially start new projects using the information of the cost of already finished projects and thus optimize the utilization of the budget. Meanwhile, they do their *best effort* to satisfy as many desirable properties as possible. Furthermore, we analyze different forms of proportionality for uncertain project costs.

Related Work. First, our work is related to scheduling and project planning literature. Given a set of tasks (which may form a bigger project), this research area is about making a plan on how to complete *all* tasks while respecting their

[1] Depending on the contract, the company will have to pay from own resources. However, it is common that cost estimations may be exceeded by $10-20\%$ without further consultation with the client. And if the company has not enough own resources to pay for the extra cost, they may become bankrupt, and the client has to pay someone else for finishing the project.

dependencies and resource requirements. In the *Resource-Constrained Project Scheduling Problem* the goal is to minimize the completion time of the last project while projects have dependencies and resource requirements (e.g. workers), so parallelization of projects is not always possible. A similar problem is the *Time-Constrained Project Scheduling Problem* where the goal is to minimize the amount of extra resources needed to finish all projects in time. Both problems have been studied with uncertain project durations (see [Ma *et al.*, 2016] and [Moradi and Shadrokh, 2019]). However, the resource requirements are deterministic. Vaziri *et al.* [2007] analyze project planning where the time for tasks is uncertain and can be influenced by the resources allocated to that task.

Note that in comparison to the above problems, we do not aim at implementing all projects and we also want to take into account the voters' preferences. In fact, if there is enough budget available to implement all projects our problem is trivial. Instead, we want to implement a subset of projects that makes the stakeholders happy, while facing the budget and the time as resource constraints. This is related to project portfolio management where the question is *which* projects should be implemented, suspended, or canceled, to serve the overarching objective of an institution. For an introduction to project portfolio management see [Rad and Levin, 2006, Chapter 1]. Another difference is that projects in our model have no dependencies on other projects, and our resources money and time are bounded, unlike resources in scheduling.

Also related is the work of Pindyck [1993], who studies projects that require continuous investment, and the final cost is only known after completion. However, there is no time constraint, no preferences, and projects can be canceled.

Another closely related field is Knapsack, as we study maximizing the (additive) utility of a set of projects while facing a cost constraint. For a broad overview, we refer to the textbook by Kellerer *et al.* [2004], which includes a chapter on multidimensional Knapsack problems considering more than one resource (in our case cost and time). Setting aside the time dimension, Knapsack has been studied under uncertain weights similar to our model by Monaci and Pferschy [2013] and Monaci *et al.* [2013]. Following a concept by Bertsimas and Sim [2004], both aim to find *robust* solutions, that perform well even if the exact weights turn out to be unfavorable. The model by Goerigk *et al.* [2015] allows for querying the exact weight of a fixed number of items in order to find a good solution when weights are uncertain.

A collective variant of Knapsack, namely participatory budgeting, has gained some attention in computational social choice lately. For notation, we adopt the formal participatory budgeting framework for approval-based preferences, which was introduced by Talmon and Faliszewski [2019] and extended to irresolute budgeting rules by Baumeister *et al.* [2020]. We refer to the bookchapter by Aziz and Shah [2021] for a broad overview on participatory budgeting in the context of computational social choice. Gomez *et al.* [2016] present a broad model considering uncertainty for both, cost and utility, for every project. In contrast to ours, their model is purely stochastic (a set of projects is feasible if its expected cost is within the budget limit), and projects are implemented all at the same time. An important stream of research we follow is the concept of *proportionality* in collective decision making, where every voter should be represented equally by a given solution. Aziz *et al.* [2018] study a variety of axioms suitable for participatory budgeting, while Pierczyński *et al.* [2021] also provide a rule with desirable properties with respect to proportional representation. We generalize some of their results to work with uncertain costs.

2 Preliminaries

Throughout this paper, for $i, j \in \mathbb{N}$ we write $[i, j] = \{i, \dots, j\}$ and $[i] = [1, i]$. Let $A = \{a_1, \dots, a_m\}$ be the set of *projects* and each subset $B \subseteq A$ is a *bundle*.[2] In our model we assume uncertainty about the exact *cost* for every project. Therefore, each project is associated with a total of four cost functions $\widetilde{c} = (c_{\min}, c_{\max}, c, c_p)$, where $c_{\min}$ and $c_{\max}$ model lower and upper bounds on the project's costs, while c models the exact costs. Hence, all three functions map from A to $\mathbb{N}^+$ and for each project $a \in A$ it holds that $c(a) \in [c_{\min}(a), c_{\max}(a)]$. For simplicity, we abuse notation by denoting $c(B) = \sum_{a \in B} c(a)$ as the cost of a bundle B (analogously for $c_{\min}$ and $c_{\max}$). By $c_p(a, x)$ we denote the probability that project $a \in A$ costs at most $x \in \mathbb{N}^+$. Note that c_p is monotonic, $c_p(a, c_{\min}(a) - 1) = 0$, and $c_p(a, c_{\max}(a)) = 1$. Slightly abusing notation we write $c_p(B, x)$ to denote the probability that for a given bundle B the cost $c(B)$ is bounded by x. Finally, each project $a \in A$ takes time $\delta : A \to \mathbb{N}^+$ to finish, and we have an overall *time limit* $\tau \in \mathbb{N}^+$ at which all projects have to be finished, and a *budget limit* ℓ (also referred to as budget)[2], which is the available money to implement projects. We assume for no project $a \in A$ holds $\delta(a) > \tau$ or $c_{\max}(a) > \ell$.

The projects are evaluated by a set of *voters* $V = \{v_1, \dots, v_n\}$. Each voter v approves a subset of projects denoted by $\mathrm{app}_v \subseteq A$. By $s_v(B) = |B \cap \mathrm{app}_v|$ we denote the *satisfaction* of voter v with bundle B, i.e. the number of items in B approved by v.[3] We define $s(B) = \sum_{v \in V} s_v(B)$ as the *total satisfaction* of all voters. We assume $s(\{a\}) > 0$ for all $a \in A$, i.e., each project is approved by at least one voter.

Let $E = (A, V, \widetilde{c}, \ell, \delta, \tau) \in \mathcal{E}$ be a budgeting scenario with uncertain cost, where $\mathcal{E}$ is the set of all such scenarios. An *online budgeting method* $\mathcal{R}$ works in discrete time steps, and successively builds a *budgeting log* $L : A \to \mathbb{N} \cup \{\bot\}$ representing at which time step a project has been started, where $\bot$ denotes that the project will not be realized. The budgeting method has limited access to the cost function c. If $L(a) = t^*$, then $c(a)$ is available only after the project has been implemented, i.e. at time step $t^* + \delta(a)$. Obviously, the decision made at step t^* is fixed and may not be revised when more information is available. Formally, the output of a budgeting method $\mathcal{R}(E)$ is a budgeting log. Further, the set of *realized* projects for a budgeting log L is denoted by $R(L) = \{a \in A \mid L(a) \neq \bot\}$. On the other hand, an *offline* budgeting method may access the exact cost function

[2] We differ in terminology from Talmon and Faliszewski [2019], who refer to a project selection as *budget*.

[3] Also studied by Talmon and Faliszewski [2019]. Other satisfaction functions (e.g. cost based) do not fit varying costs that well.

c. Hence, an optimal solution can be precomputed and every project can be implemented simultaneously.

Example 1 (continued)**.** *We have $A = \{BS, ET, EV, FV\}$ with following costs (in million), durations, and satisfactions.*

a	$c(a)$	$c_{\min}(a)$	$c_{\max}(a)$	$\delta(a)$	$s(\{a\})$
BS	40	37	42	1	5,000
ET	120	80	150	4	9,000
EV	59	59	61	1	6,000
FV	100	100	160	5	11,000

The budget limit is $\ell = 200$ and we have time $\tau = 5$. If we knew the exact costs, we would immediately start the projects BS, EV, and FV since with cost of 199 they fit in our budget and have the maximum number of 22,000 approvals. However, this bares a high risk if the cost is unknown. The projects could also cost up to 263 which is far beyond our budget. An online budgeting method could for instance do the following. Start ET in the first year and wait four years until it finishes. If we are lucky, it turns out that it costs at most 97 so that we can safely implement both BS and EV in the remaining year. If it costs more, we can implement at least BS or in some cases EV. Also, the following is possible. If it is very improbable that FV costs more than 139, we could also relatively safely begin FV and EV simultaneously in the first year. Note that a sequential implementation of FV and EV fails the time limit.

We assume each started project in a budgeting log will be implemented regardless of the final cost and completion time. It is often assumed that projects can be stopped if their cost becomes more than their value, i.e. one cannot gain money with the project (e.g. in [Pindyck, 1993]). We decided against canceling for two reasons. First, in our model the "value" of a project is not measured in the same currency as the cost; second, the projects we have in mind cannot simply be canceled (imagine an eternal construction site in the city center).

3 Properties of Online Budgeting Methods

A budgeting log, and thus also an online budgeting method, has rather weak requirements. For instance, it is allowed in a budgeting log to start arbitrarily many projects simultaneously, even if they will certainly exceed the budget limit; or to start projects so late that they cannot be completed in time. These issues are undesirable and should be avoided. In this section, we define some desirable properties a budgeting log (and the algorithm that generates it) should satisfy.

Definition 1. *Let $E \in \mathcal{E}$ be a budgeting scenario and L be a budgeting log with respect to E. L satisfies the following axioms if respective conditions are met.*

Punctuality (PU): *Every realized project finishes within the given time limit. Formally, for all $a \in A$ it holds that either $L(a) = \bot$ or $L(a) + \delta(a) \leq \tau$.*

α-Risk-assessment (α-RA): *A (set of) project(s) may only be started if the probability for exceeding the budget limit is at most α. Formally, for every $t \in [\tau]$, let $U_t = \{a \in A \mid L(a) \leq t < L(a) + \delta(a)\}$ be the running yet unfinished projects, and $F_t = \{a \in A \mid L(a) + \delta(a) \leq t\}$ the finished projects. For given $\alpha \in [0, 1)$, it holds that a set of projects S may only be started at time t if $c_p(U_t \cup S, \ell - c(F_t)) \geq 1 - \alpha$.*

κ-Limitation (κ-LI): *The budget limit may not be exceeded by a factor greater than κ. Formally, $c(R(L)) \leq \kappa\ell$.*

Exhaustiveness (EX): *There should be no project, which could have been implemented even with maximum cost without breaking feasibility. Formally, for $B = R(L)$ and every $a \in A \setminus B$, it holds that $c(B) + c_{\max}(a) > \ell$.*

Note that 0-risk-assessment and 1-limitation coincide. A budgeting method $\mathcal{R}$ satisfies some axiom χ if $\mathcal{R}(E)$ satisfies χ for every $E \in \mathcal{E}$ (assuming parallel universe tie-breaking).

We study punctuality as a property since it is very interesting to see what kind of restriction it is, and what algorithms are possible if we relax it. Risk-assessment and limitation can be interpreted as follows. The client (e.g. Participation Island) has $(\kappa - 1)\ell$ extra money as a security — for instance as a loan option — which should be used only if absolutely necessary. With a good risk-assessment (i.e. small α) it is improbable that the security is ever touched. Exhaustiveness has two interpretations. First, voters naturally expect that approved projects are realized if there is money left to do so safely. Second, it is common that the budget of a department may be reduced in the next period if it is not completely spent.

Independent of the above properties we want to maximize the satisfaction of the voters with the outcome. One key metric for the analysis of online optimization algorithms is the worst-case ratio between a solution found by an online algorithm and an optimal (satisfaction maximizing) solution with complete knowledge. This factor is known as *competitive ratio* (CR) (see Fiat and Woeginger [1998]).

Definition 2. *An online budgeting method $\mathcal{R}$ is σ-competitive (σ-CR) if there is a constant $\Delta \in \mathbb{R}$, such that for every $E \in \mathcal{E}$ and $\mathcal{B}_\ell = \{B \subseteq A \mid c(B) \leq \ell\}$ it holds that $s(R(\mathcal{R}(E))) + \Delta \geq \frac{1}{\sigma} \max_{B \in \mathcal{B}_\ell} s(B)$.*

So which combinations of properties are possible, and is there a perfect online budgeting method? Unfortunately, the answer is no, i.e. no method can satisfy all properties simultaneously for all combinations of parameters.

Theorem 3. *For any fixed $\alpha < 1$, no online budgeting method simultaneously satisfies α-risk-assessment, punctuality, and exhaustiveness.*

Proof. Consider an odd budget limit $\ell \geq 3$, and $A = \{a_1, a_2, a_3, \ldots\}$. Each project $a_i \in A$ has minimum cost $c_{\min}(a_i) = (\ell-1)/2$, maximum cost $c_{\max}(a_i) = (\ell+1)/2$, and takes time $\delta(a_i) = \tau$. Note that a set of two projects exceeds the budget limit if and only if both projects have maximum cost. Let each project have maximum cost with probability greater than $\sqrt{\alpha}$. Due to α-risk-assessment a budgeting method can only start one project at the first time step, say a_1. By punctuality it is impossible to start another project. Since $c(a_1) = \ell/2$ is possible, there are instances for which $c(a_1) + c_{\max}(a_2) \leq \ell$, thus exhaustiveness is violated. $\square$

Similar holds for κ-limitation as long as $\kappa < m = |A|$.

Theorem 4. *For any fixed $\kappa < m$, there exists no online budgeting method simultaneously satisfying κ-limitation, punctuality, and exhaustiveness.*

Proof. Consider $E \in \mathcal{E}$ with $A = \{a_1, \dots, a_m\}$, $m \geq 3$, and $\ell > m$. Each project $a_i \in A$ takes time $\delta(a_i) = \tau$ to realize, has cost $c(a_i) = 1$, and maximum cost $c_{\max}(a_i) = \ell - m$. By exhaustiveness, all projects must be realized, since for each $a_i \in A$ holds $c(A \setminus \{a_i\}) + c_{\max}(a_i) = \ell$ and by punctuality, all projects must be started simultaneously. However, this decision has to be made without knowing the exact cost. Let $E' \in \mathcal{E}$ be equivalent to E, except for having maximum cost as exact cost for each project. An online budgeting method that implements all projects to satisfy exhaustiveness and punctuality might end up spending $c_{\max}(A) = m \cdot (\ell - m)$. Thus, to start all projects, it cannot be better than $\frac{m\cdot(\ell-m)}{\ell} = \left(m - \frac{m^2}{\ell}\right)$-limited. By choosing ℓ large, we can approach m to any fixed value $\kappa < m$. $\square$

Interestingly, for $\kappa \geq m$, above properties are compatible and can be satisfied by a 1-competitive algorithm. Since $c_{\max}(a) \leq \ell$ holds for every $a \in A$, we can implement every project at the first time step, only violating risk-assessment.

Observation 5. *There exists a* 1*-competitive online budgeting method satisfying* m*-limitation, punctuality, and exhaustiveness.*

However, trading off some desirable properties, it is possible to achieve 0-risk-assessment (or equivalently 1-limitation) together with either punctuality or exhaustiveness.

Theorem 6. *There is an* m*-competitive online budgeting method satisfying* 0*-risk-assessment (and thus* 1*-limitation) and either punctuality or exhaustiveness.*

Proof. First, we start the most valuable project, say a_1, which has by definition maximum cost of at most ℓ. This way we achieve m-competitiveness already, since $s(A) \leq m \cdot s(\{a_1\})$. If we want to achieve punctuality, we stop now. For exhaustiveness, we sequentially add those projects, that can be safely added without exceeding the budget. $\square$

A competitive ratio of m is bad. Yet, for fixed α this factor cannot be improved if the projects' cost intervals are large.

Theorem 7. *For any online budgeting method that satisfies* α*-risk-assessment for a fixed* $\alpha < 1$*, the competitive ratio is in* $\Omega(m)$*. If* $c_{\max}(a) = c_{\min}(a) + 1$ *holds for all* $a \in A$*, the competitive ratio is in* $\Omega(2)$*.*

Proof. Consider $E \in \mathcal{E}$ with $A = \{a_1, a_2, a_3, \dots, a_m\}$ and a set of voters, such that each project $a_i \in A$ yields the same (additive) satisfaction $s(a_i) = \lambda \in \mathbb{N}^+$. Let $\ell = m$, $c_{\min}(a_i) = 1$ and $c_{\max}(a_i) = m$ for every $a_i \in A$. Further, for each a_i we set the probability that a_i costs exactly m to α (thus, $c_p(a_i, \ell - 1) = 1 - \alpha$). An online algorithm with α-risk-assessment cannot start more than one project at the same time because for every pair of projects $a_i \neq a_j$ it holds $c_p(\{a_i, a_j\}, \ell) \leq c_p(a_i, \ell - 1) \cdot c_p(a_j, \ell - 1) = (1 - \alpha)^2 < 1 - \alpha$. Thus it starts at most one project, for example a_1. Revealing $c(a_1) = m$ and $c(a_i) = 1$ for $i \in [2, m]$, an offline algorithm may select the optimal solution $B = A \setminus \{a_1\}$ with $s(B) = \lambda \cdot (m - 1)$, while the online algorithm yields a satisfaction of $s(\{a_1\}) = \lambda$. Overall we deduce a competitive ratio of $(m - 1) \in \Omega(m)$.

Properties	**CR**	**Ref**
PU, EX, m-LI	1-CR	Obs. 5
0-RA, 1-LI, PU	m-CR (up to 2-CR)	Thm. 6, 8
0-RA, 1-LI, EX	m-CR (up to 2-CR)	Thm. 6, 8

Table 1: Summary of our possibility results regarding the combination of axioms for online budgeting methods.

Property	**Incompatible**	**Ref**
PU	$\{\alpha$-RA, EX$\}$, $\{m'$-LI, EX$\}$	Thm. 3, 4
α-RA	$\{$PU, EX$\}$	Thm. 3
m'-LI	$\{$PU, EX$\}$	Thm. 4
EX	$\{\alpha$-RA, PU$\}$, $\{m'$-LI, PU$\}$	Thm. 3, 4

Table 2: Summary of our impossibility results regarding the combination of axioms for online budgeting methods ($m' < m$).

For bounded uncertainty by $c_{\max}(a) = c_{\min}(a) + 1$ for all $a \in A$, we can use a similar argument. Let $\ell \geq 2$ be an even number, $A = \{a_1, a_2, a_3, \dots, a_m\}$ with $c_{\min}(a_i) = \ell/2$ and $c_{\max}(a_i) = \ell/2 + 1$. Again, we assume equal utility of λ for every project. We set $c_p(a_i, \ell/2) = 1 - \alpha$, so the probability that two projects can be implemented within ℓ is $(1 - \alpha)^2 < 1 - \alpha$. Let an online budgeting method implement a_1 first (due to α-RA it cannot implement two projects). Revealing $c(a_1) = \ell/2 + 1$ and $c(a_2) = c(a_3) = \ell/2$, the optimal solution is $\{a_2, a_3\}$, yielding a competitive ratio of $\frac{2\lambda}{\lambda}$. $\square$

Theorems 6 and 8 show that these bounds are tight.

Theorem 8. *If the uncertainty on the cost is bounded by a small factor* c^**, that is,* $c_{\max}(a) - c_{\min}(a) < \frac{c_{\max}(a')}{m} = c^*$ *for all* $a, a' \in A$*, there is a* 2*-competitive method satisfying* 0*-risk-assessment and either punctuality or exhaustiveness.*

Proof. We use the optimal offline method to retrieve an optimal bundle B, assuming the lower bound cost for each project, i.e. $\sum_{b \in B} c_{\min}(b) \leq \ell$.
Case 1: If $\sum_{b \in B} c_{\max}(b) \leq \ell$, we are done.
Case 2: Otherwise, since implementing B may exceed the budget limit, we remove the least valuable project $a \in B$ and implement $B' = B \setminus \{a\}$ at the first time step. It holds that $\ell \geq \sum_{b \in B} c_{\min}(b) \geq \sum_{b \in B} c_{\max}(b) - |B|c^* \geq \sum_{b \in B} c_{\max}(b) - c_{\max}(a) = \sum_{b \in B'} c_{\max}(b)$, due to $|B|c^* \leq |B| \cdot \frac{c_{\max}(a)}{m} \leq c_{\max}(a)$. On the other hand, since a is the least valuable project in B, the satisfaction with B' is at least $s(B') \geq s(B) \cdot \frac{|B'|}{|B|} = s(B) \cdot \frac{|B|-1}{|B|}$. The worst competitive ratio of two is achieved if $|B| = 2$, since $|B| = 1$ is already covered by case 1. We can now return B' which satisfies punctuality or we add projects to B' until we achieve exhaustiveness. Note that in both cases there is no risk for exceeding the budget limit. $\square$

Table 1 summarizes the lower bound competitive ratios any online algorithm can achieve while satisfying given axiomatic properties, while Table 2 summarizes incompatible axioms.

4 Best Effort Online Budgeting

The possibility and impossibility results show that there is no perfect online budgeting method, i.e., no method satisfies all properties and has a good competitiveness. So the only way is to design "best effort" algorithms that trade-off properties and have a good competitiveness in most cases.

We propose the online budgeting method **Best effort exhaustiveness (BEE)** that trades exhaustiveness against punctuality, risk-assessment, and limitation. That is, the method guarantees punctuality, α-risk-assessment, and κ-limitation for given τ, α, κ, but not exhaustiveness. However, it tries to be as exhaustive as possible. With small modifications, we get the **Best effort punctuality (BEP)** method which trades punctuality against exhaustiveness, risk-assessment, and limitation. Both algorithms generalize a common greedy algorithm for Knapsack (see [Kellerer *et al.*, 2004]) to our setting.

We test both algorithms using real data from the Participatory Budgeting Library (see [Stolicki *et al.*, 2020]), modified to fit our uncertainty scenario. BEE performs better in terms of both exhaustiveness and competitiveness the more we increase τ. If projects have durations in $[1, 10]$, the performance is already remarkably good at time limits between 20 and 30. However, our experiments with the BEP method imply that guaranteed exhaustiveness results in massive unpunctuality. Also, κ-limitation plays a role for the exhaustiveness, however, the role of α-risk-assessment is negligible unless α is almost 0.

5 Proportionality

Apart from maximizing the overall utility, another way to satisfy voters is a proportional distribution of the realized projects among them. There is a variety of proportionality axioms in the literature. We focus on *justified representation*, considering *BPJR-L* by Aziz *et al.* [2018] and *Extended Justified Representation* by Pierczyński *et al.* [2021]. Assuming uncertainty over the exact projects' costs, we provide two relaxations for proportionality axioms: *ex ante* and *ex post*. For the relaxations, we assume the upper cost bound is given for all projects (ex ante) or not implemented projects (ex post).

Definition 9. *Consider a budgeting scenario $E = (A, V, \widetilde{c}, \ell, \delta, \tau) \in \mathcal{E}$ and a bundle $B \subseteq A$. We define the following axioms in different variants, distinguishable by respective cost functions. Let $c' : A \to \mathbb{N}^+$ be a cost function. Typically, for known exact cost, we consider $c' = c$ for both of the following axioms. We study two relaxations for our setting with uncertain cost. That is, for the* ex ante *variant, we study $c' = c_{\max}$ and for the* ex post *variant, for $a \in A$ we study $c'(a) = c(a)$ if $a \in B$ and $c'(a) = c_{\max}(a)$, otherwise.*

A bundle B satisfies following axioms if the respective condition holds, while a budgeting method $\mathcal{R}$ satisfies an axiom if it holds for every $R(\mathcal{R}(E))$ with $E \in \mathcal{E}$.

BPJR-L: *For all $k \in [\ell]$ there exists no set of voters $V' \subseteq V$ with $\frac{|V'|}{n} \geq \frac{k}{\ell}$ such that $c'\left(\bigcap_{v \in V'} \mathrm{app}_v\right) \geq k$, but there is a set $T \subseteq \bigcap_{v \in V'} \mathrm{app}_v$ with $c'\left(\left(\bigcup_{v \in V'} \mathrm{app}_v\right) \cap B\right) < c'(T) \leq \frac{\ell \cdot |V'|}{n}$.*

Extended Justified Representation (EJR): *For all $V' \subseteq V$ and $T \subseteq \bigcap_{v \in V'} \mathrm{app}_v$ it holds that, if $c'(T) \leq \frac{\ell \cdot |V'|}{n}$, then there exists some $v \in V'$ with $|\mathrm{app}_v \cap B| \geq |T|$.*

Informally, a group of voters should not be worse off, if they could spend their proportional share of the budget on projects they collectively approve of. If every voter had exactly ℓ/n to spend, BPJR-L states that if a group of voters can afford a set of unanimously approved projects T, the funds spent on projects no one of the group approves should be lower than $\ell - c'(T)$. EJR on the other hand states that not all voters of a group should get less projects implemented than they could afford with joint funds.

Other axioms can be defined analogously in an ex ante and ex post version. In many cases stronger variants imply weaker variants. In case of EJR, consider some bundle that satisfies EJR. If the cost function is altered in a way that items only become more expensive, we observe that EJR is still satisfied (even though the price increase may now exceed the budget). Surprisingly, similar implications do not hold for BPJR-L.

Observation 10. *EJR implies ex post EJR and ex post EJR implies ex ante EJR.*

Theorem 11. *BPJR-L does not imply ex post (or ex ante) BPJR-L.*

Proof. Consider $E \in \mathcal{E}$ with three projects $A = \{a_1, a_2, a_3\}$, two voters $V = \{v_1, v_2\}$, and $\ell = 3$. The first two projects are known to be unit cost in advance, i.e., $c(a_i) = c_{\max}(a_i) = 1$ for $i \in [2]$. For a_3 it holds that $c(a_3) = 2$ and $c_{\max}(a_3) = 3$. For $i \in [2]$, voter v_i approves both a_i and a_3.

First, we will show that the bundle $B = \{a_1, a_2\}$ satisfies BPJR-L. There are three nonempty subsets of voters $V_1 = \{v_1\}$, $V_2 = \{v_2\}$ and $V_3 = \{v_1, v_2\}$. Let $T_i = \bigcap_{v \in V_i} \mathrm{app}_v$ and $U_i = \bigcup_{v \in V_i} \mathrm{app}_v$. For $i \in [2]$ it holds that $c(T_i) = c(U_i \cap B) = 1$. For V_3 it holds that $c(T_3) = c(U_3 \cap B) = 2$. Yet, ex post (and ex ante) BPJR-L is violated. For $k = 3$ it holds that $|V_3| = \frac{kn}{\ell}$, $c_{\max}(T_3) = 3 = k$ but $c(U_3 \cap B) = c_{\max}(U_3 \cap B) = 2 < k$. ☐

Theorem 12. *There is no online budgeting method satisfying α-risk-assessment and ex post BPJR-L.*

Proof. Consider two projects $A = \{a_1, a_2\}$ with $c_{\min}(a_i) = 1$ and $c_{\max}(a_i) = 2$ for $i \in [2]$. Further, let $c_p(a_i, 1) < 1 - \alpha$ and $\ell = 2$. If some budgeting method $\mathcal{R}$ selects both projects simultaneously, α-risk-assessment is violated. If $\mathcal{R}$ selects one project $a \in A$ first, revealing $c(a) = 1$ yields two options. $B = \{a\}$ violates BPJR-L and implementing the remaining project violates α-risk-assessment. ☐

To compute a feasible outcome satisfying an ex ante proportionality axiom, we can assume maximum cost for each project and compute a bundle that satisfies the strong variant of the axiom. For the actual cost the result may not be exhaustive, but we may implement it instantly. Aziz *et al.* [2018] showed that a feasible outcome satisfying BPJR-L is guaranteed to exist (although hard to compute). Regarding EJR, Peters and Skowron [2020] recently introduced an aggregation method for committee elections, called *Rule X* and Pierczyński *et al.* [2021] showed that a generalization of Rule X for participatory budgeting satisfies EJR (although they consider cardinal utilities instead of approval based preferences).

Observation 13. *There is an online budgeting method, satisfying* 0*-risk-assessment, punctuality and ex ante BPJR-L (respectively ex ante EJR).*

We slightly adjust Rule X to also work with uncertain cost and show subsequently that our variant satisfies ex post EJR.

Definition 14. *Rule X for uncertain cost ($\mathcal{R}_X$) works as follows. Every voter $v \in V$ is given a (real valued) individual budget of $b_v = \ell/n$, which they can use to implement projects sequentially. We start with an empty bundle $B = \emptyset$ and fund exactly one project in each iteration. The cost will be deducted from supporting voters' funds. For $\rho_{\max} > 0$ a project $a \in A \setminus B$ is $\rho_{\max}$-affordable if the following equation holds.*

$$\sum_{v \in V | a \in \text{app}_v} \min(b_v, \rho_{\max}) = c_{\max}(a) \tag{1}$$

$\mathcal{R}_X$ implements the project a^* with the lowest $\rho_{\max}$-affordability, waits for it to finish to obtain the exact cost, and finally withdraws the required funds from approving voters. That is, we replace the upper cost bound $c_{\max}(a^*)$ in Equation (1) with $c(a^*)$ and calculate the ρ-affordability for a^*. Then for every voter $v \in V$ with $a^* \in \text{app}_v$, b_v is set to $\max(0, b_v - \rho)$. If there is no $\rho_{\max}$-affordable project left, Rule X returns the corresponding budgeting log.

By definition, $\mathcal{R}_X$ always satisfies 0-risk-assessment (and 1-limitation) but fails punctuality. Following Pierczyński *et al.* [2021], Rule X (and thus $\mathcal{R}_X$) fails exhaustiveness.

Theorem 15. *$\mathcal{R}_X$ satisfies ex post EJR (and ex ante EJR).*

Proof. Let $E = (A, V, \widetilde{c}, \ell, \delta, \tau) \in \mathcal{E}$ and assume that $B = R(\mathcal{R}_X(E))$ violates ex post EJR. Consider the cost function c' with $c'(a) = c(a)$ if $a \in B$ and $c'(a) = c_{\max}(a)$ if $a \in A \backslash B$. Then by assumption there is a set of voters $V' \subseteq V$ and a set of projects $T \subseteq \bigcap_{v \in V'} \text{app}_v$ with $c'(T) \leq \frac{\ell \cdot |V'|}{n}$ and for all $v \in V'$ it holds that $|\text{app}_v \cap B| < |T|$. We investigate what led to the violation of EJR, by simulating $\mathcal{R}_X$ with (ex post) knowledge of the exact cost for projects in B. We index the elements in $T \backslash B = \{t_1, \dots, t_k\}$ in a way that $c'(t_i) \leq c'(t_j)$ for all $i < j$. Let every voter $v \in V'$ split her initial individual budget $b_v = \ell/n$ into $k + 1$ piles b_v^j, such that for $j \in [k]$ the pile $b_v^j = \frac{c'(t_j)}{|V'|}$ is v's equal share for funding t_j (w.r.t. V'). The leftover funds $b_v^0 = \frac{\ell}{n} - \frac{c'(T)}{|V'|}$ will be reserved to fund projects in $B \cap T$. We execute $\mathcal{R}_X$ again, but this time we try to keep track of which project is financed with which pile of funds. We will show, that after execution either some voter $v \in V'$ has no funds left, but at least $|T|$ preferred projects in the outcome, or every voter has some pile left, which could have been used to implement a project in $T \setminus B$.

By assumption, each voter $v \in V'$ helps funding $T \cap B$ and at most $k - 1$ additional projects. We let each voter $v \in V'$ pay for projects in $T \cap B$ with the leftover funds b_v^0. In the following numbering we skip projects in T (each voter in V' has reserved funds for $T \cap B$ and $T \setminus B$ wont be budgeted). For each $v \in V'$, let a_v^j be the j-th project, v helps funding (in addition to $T \cap B$). We let v pay her share for a_v^j with b_v^j. Note that $c_{\max}(a_v^j) \leq b_v^j$. Otherwise, either t_j has a lower $\rho_{\max}$-affordability and would have been funded instead, or some voter $v' \in V'$ has not enough budget left to pay her (full) share. In the latter case, consider the first iteration that leads to some $v' \in V'$ being bankrupt. Then v' has paid all the other projects with respective dedicated piles, resulting in at least k projects being funded (in addition to the projects $T \cap B$). This is a contradiction to EJR being violated, since $|\text{app}_{v'} \cap B| \geq k + |T \cap B| = |T|$. Following $c_{\max}(a_v^j) \leq b_v^j$, our assumption can only hold if every voter $v \in V'$ has at least one pile untouched. This is a contradiction, as $\mathcal{R}_X$ could implement project t_1 with a total of $|V'|$ piles. □

Punctuality cannot be added without violating ex post EJR.

Theorem 16. *There is no online budgeting method satisfying α-risk-assessment, punctuality and ex post EJR.*

Proof. Consider the following example with $A = \{a_1, a_2\}$ with $c_{\min}(a_i) = 1$ and $c_{\max}(a_i) = 2$ for $i \in [2]$. The budget limit is set to $\ell = 3$ and a single voter approves both projects. We choose c_p, such that $c_p(A, 3) < 1 - \alpha$. Finally, projects need to be implemented at the first time step, i.e., $\delta(a_i) = \tau$ for $i \in [2]$. Let $B \subseteq A$ be the realized projects by an online budgeting method $\mathcal{R}$. We study three cases assuming punctuality is satisfied. If $|B| = 0$, B clearly fails ex post EJR. If $|B| = 1$, B might fail ex post EJR in case $c(B) = 1$. $\mathcal{R}$ may not select $B = A$ while respecting α-risk-assessment, as the probability of $c(A) = 4$ is greater than α. □

6 Conclusions And Outlook

In a participatory budgeting campaign with uncertain costs one has to trade-off between the desirable properties punctuality, exhaustiveness, risk-assessment, limitation, but also voter satisfaction and proportionality. In a way, this confirms the old project management principle "fast, cheap, good, pick two." Being punctual and within the cost limit makes proportionality or a good competitive ratio usually impossible.

In some applications, projects compete for resources like machines or workers. Thus not all projects can run simultaneously (see Hans *et al.* [2007]). As a next step, we propose to extend our model in that direction. Additionally, more realistic satisfaction functions could model that a voter's satisfaction with a project may be affected by discovering the exact cost. Dependencies between parameters could be especially interesting for time and money, allowing for the possibility to speed up a project by spending more money. Finally, our model only considers uncertainty in one of two resources. Swapping respective functions to study uncertain finishing times instead, we are interested in which implications still hold (although uncertainty concepts for proportionality are more reasonable for the cost dimension). A more general model may explore uncertainty of both dimensions.

Acknowledgements

The authors would like to thank Jérôme Lang and the anonymous reviewers for their helpful comments. This work was supported in part by the project "Online Participation," funded by the NRW Ministry for Innovation, Science, and Research and by the DFG grants RO-1202/21-1 and BA-6270/1-1.

References

[Aziz and Shah, 2021] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In T. Rudas and P. Gábor, editors, *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*, pages 215–236. Springer, 2021.

[Aziz *et al.*, 2018] Haris Aziz, Barton E. Lee, and Nimrod Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 23–31. IFAAMAS, 2018.

[Baumeister *et al.*, 2020] Dorothea Baumeister, Linus Boes, and Tessa Seeger. Irresolute approval-based budgeting. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1774–1776. IFAAMAS, 2020.

[Bertsimas and Sim, 2004] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.

[Fiat and Woeginger, 1998] Amos Fiat and Gerhard J. Woeginger. *Online Algorithms: The State of the Art*, volume 1442. Springer, 1998.

[Goerigk *et al.*, 2015] Marc Goerigk, Manoj Gupta, Jonas Ide, Anita Schöbel, and Sandeep Sen. The robust knapsack problem with queries. *Computers & Operations Research*, 55:12–22, 2015.

[Gomez *et al.*, 2016] Javier Gomez, D. Rios Insua, and César Alfaro. A participatory budget model under uncertainty. *European Journal of Operational Research*, 249(1):351–358, 2016.

[Hans *et al.*, 2007] Erwin W. Hans, Willy Herroelen, Roel Leus, and Gerhard Wullink. A hierarchical approach to multi-project planning under uncertainty. *Omega*, 35(5):563–577, 2007.

[Kellerer *et al.*, 2004] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack Problems*. Springer, 2004.

[Ma *et al.*, 2016] Weimin Ma, Yangyang Che, Hu Huang, and Hua Ke. Resource-constrained project scheduling problem with uncertain durations and renewable resources. *International Journal of Machine Learning and Cybernetics*, 7(4):613–621, 2016.

[Monaci and Pferschy, 2013] Michele Monaci and Ulrich Pferschy. On the robust knapsack problem. *SIAM Journal on Optimization*, 23(4):1956–1982, 2013.

[Monaci *et al.*, 2013] Michele Monaci, Ulrich Pferschy, and Paolo Serafini. Exact solution of the robust knapsack problem. *Computers & Operations Research*, 40(11):2625–2631, 2013.

[Moradi and Shadrokh, 2019] Hadi Moradi and Shahram Shadrokh. A robust scheduling for the multi-mode project scheduling problem with a given deadline under uncertainty of activity duration. *International Journal of Production Research*, 57(10):3138–3167, 2019.

[Peters and Skowron, 2020] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 793–794, 2020. Extended version arXiv:1911.11747.

[Pierczyński *et al.*, 2021] Grzegorz Pierczyński, Piotr Skowron, and Dominik Peters. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34, 2021.

[Pindyck, 1993] Robert S. Pindyck. Investments of uncertain cost. *Journal of Financial Economics*, 34(1):53–76, 1993.

[Rad and Levin, 2006] Parviz F. Rad and Ginger Levin. *Project Portfolio Management: Tools and Techniques*. International Institute for Learning, 2006.

[Stolicki *et al.*, 2020] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. Pabulib: A participatory budgeting library. *arXiv preprint arXiv:2012.06539*, 2020.

[Talmon and Faliszewski, 2019] Nimrod Talmon and Piotr Faliszewski. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, volume 33, pages 2181–2188. AAAI Press, 2019.

[Vaziri *et al.*, 2007] Kabeh Vaziri, Paul G. Carr, and Linda K. Nozick. Project planning for construction under uncertainty with limited resources. *Journal of Construction Engineering and Management*, 133(4):268–276, 2007.