TU WIEN Informatics

Cost Bias in Participatory Budgeting

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieurin

im Rahmen des Studiums

Logic and Computation

eingereicht von

Felicia Schmidt, BSc
Matrikelnummer 12130505

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Univ.Prof. Dr.techn. Stefan Woltran
Mitwirkung: Dr.techn. Jan Maly, MSc
Oliviero Nardi, MSc

Wien, 30. September 2024

Felicia Schmidt   Stefan Woltran

Technische Universität Wien
A-1040 Wien ▪ Karlsplatz 13 ▪ Tel. +43-1-58801-0 ▪ www.tuwien.at

TU WIEN Informatics

Cost Bias in Participatory Budgeting

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieurin

in

Logic and Computation

by

Felicia Schmidt, BSc
Registration Number 12130505

to the Faculty of Informatics

at the TU Wien

Advisor: Univ.Prof. Dr.techn. Stefan Woltran
Assistance: Dr.techn. Jan Maly, MSc
Oliviero Nardi, MSc

Vienna, September 30, 2024

Felicia Schmidt Stefan Woltran

Erklärung zur Verfassung der Arbeit

Felicia Schmidt, BSc

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.
Ich erkläre weiters, dass ich mich generativer KI-Tools lediglich als Hilfsmittel bedient habe und in der vorliegenden Arbeit mein gestalterischer Einfluss überwiegt. Im Anhang „Übersicht verwendeter Hilfsmittel“ habe ich alle generativen KI-Tools gelistet, die verwendet wurden, und angegeben, wo und wie sie verwendet wurden. Für Textpassagen, die ohne substantielle Änderungen übernommen wurden, habe ich jeweils die von mir formulierten Eingaben (Prompts) und die verwendete IT- Anwendung mit ihrem Produktnamen und Versionsnummer/Datum angegeben.

Wien, 30. September 2024

Felicia Schmidt

Acknowledgements

My Logic and Computation journey started with a phone call to Stefan Woltran, where I was very nervous about being the right fit for the program, while he was being calm and assured that I will be doing just fine. It is only fitting that this degree now ends with him as my thesis supervisor. Thank you very much for your continuing support and trust in me. If I had met somebody else over the phone that day, I would have maybe never been brave enough to take the leap.

I have endless gratitude for both Jan Maly and Oliviero Nardi. Starting from the first project and now throughout this entire thesis, you have both been nothing but supportive, knowledgeable, reliable, personable, excitable about any successes, and absolutely invaluable in your feedback. It cannot be overstated how much the honest and encouraging work environment helped me during the whole run of this thesis and also when sparking my excitement for COMSOC. I have learned so much from the both of you and I know I will continue to do so for the next few years. Therefore, thank you especially to Jan Maly for opening the COMSOC door to me.

Thank you very much to my parents, my brother Vincent and the rest of my family for their never-ending both practical and emotional support, through academia and life in general. Thank you to Annika for still having no idea what I actually study while having all the knowledge about me as a person. Thank you to Alice, for your constant friendship, but also for your laptop. Without you, this work would still be unfinished, in so many ways. Last but not least, thank you to all the other lovely people I get to call my friends. You are too many to list, and for that I am truly grateful.

Kurzfassung

Das relativ junge Forschungsfeld Computational Social Choice (COMSOC) untersucht die informatischen Aspekte kollektiver Entscheidungsfindung, wie die Schwierigkeit einen Gewinner zu bestimmen und Abstimmungsmanipulation. Eine Wahlart, die von COMSOC ausführlich erforscht wurde und ihren Ursprung im Südamerika der 1980er hat, heißt Participatory Budgeting (PB). Die Gemeinde bestimmt ein Budget, es werden Projekte vorgeschlagen, die mit diesem Budget umgesetzt werden könnten, und die Wähler:innen können dann angeben, welche dieser Projekte sie unterstützen. Wie bei jeder Abstimmung entscheiden Abstimmungsregeln über das Ergebnis. Mehrere solcher Abstimmungsregeln existieren, generell werden greedy-style Abstimmungsregeln verwendet. Vor kurzem wurde eine neue und proportionale Abstimmungsregel entwickelt, die Method of Equal Shares. Abstimmungsregeln sind gut erforscht, zum Beispiel mit Fokus auf Komplexität oder Proportionalität. Allerdings wurde in der COMSOC-Literatur bisher Bias in PB-Wahlen noch keine Aufmerksamkeit geschenkt. Wir formulieren das Konzept des Cost Bias in Participatory Budgeting. Dies meint die Idee, dass die Inklusion eines Projektes in das Wahlergebnis unangemessen stark davon abhängt, ob ein Projekt verhältnismäßig günstig oder teuer ist. Des Weiteren führen wir das Konzept der Proportionalität eines Projektes ein. Dies bezieht sich auf die Differenz zwischen dem Prozentsatz des Budgets, welches ein Projekt kosten würde, und dem Prozentsatz aller möglichen Zusprüche, die es erhalten hat. Wir samplen PB-Wahlen und erweitern das Samplen um Projektkosten, wobei wir beeinflussen können, wie hoch die Wahrscheinlichkeit ist, dass ein Projekt proportional ist. Wir nutzen die kreierten Wahlinstanzen für umfassende statistische Experimente, welche die mögliche Korrelation zwischen der Proportionalität eines Projektes und der Aufnahme ins Wahlergebnis messen. Diese gemessene Korrelation vergleichen wir für die Utilitarian Greedy Rule und MES. Unsere Ergebnisse zeigen einen eindeutigen Bias von MES zu proportionell günstigen Projekten, wenn die meisten Projekte wahrscheinlich proportional teuer sind. Dieser gemessene Bias ist im Einklang mit weiteren Häufigkeitsanalysen, die wir mit realen Datensätzen durchgeführt haben. Des Weiteren beweisen wir zwei neue theoretische Resultate: Für zwei spezifische Arten von PB-Instanzen gilt, dass bei einer greedy-style Abstimmungsregel ein günstigeres Projekt mindestens mit gleicher Wahrscheinlichkeit in das Wahlergebnis aufgenommen wird wie ein höherpreisiges.

Abstract

The relatively young research field of Computational Social Choice (COMSOC) studies the computational aspects of communal decision making, such as the difficulty of election winner determinations and voting manipulation. One election type that has been studied extensively in COMSOC and originated in South America in the 1980s is called Participatory Budgeting (PB). The municipality decides on a budget, projects are proposed that could be realized with that budget, and the voters can then indicate which of these projects they approve of. As with any election, a voting rule computes the election outcome. Several voting rules exist, commonly used are greedy-style voting rules. Fairly recently, a new and proportional voting rule has been developed, called the Method of Equal Shares. Voting rules have been studied extensively, for example in regards to computational effort or proportionality. However, up until now, bias in PB elections has not received any attention in COMSOC literature. We formulate the concept of cost bias in participatory budgeting, which refers to the idea that a possible inclusion of a project is unduly dependent on if a project is relatively cheap or relatively expensive. Furthermore, we introduce the concept of the proportionality of a project, meaning the difference between the percentage of the budget that a project would cost and the percentage of possible approvals it received. We sample PB elections, extending the procedure to also include project costs, where we can influence the probability that a project is proportional. We use these created election instances to perform extensive statistical experiments measuring a possible correlation between project proportionality and inclusion of that project in the election outcome. We compare the measured correlations for both the utilitarian greedy voting rule and MES. Our findings clearly show a measurable bias of MES towards proportionally inexpensive projects if most projects are disproportionately expensive. This measured bias aligns with additional frequency analysis experiments done on real-world election data. Furthermore, we prove two new theoretical results: For two specific PB instance types it holds that under a greedy-style voting rule a smaller cost project is at least as likely to be included in the outcome as a higher cost one.

Contents

CHAPTER 1

Introduction

Social choice theory studies the way that individuals' preferences can be aggregated into a collective decision. In other words, social choice theory is interested in elections and how they are decided via voting rules. Since collective decision making is an integral part of any community, social choice theory has been studied by political scientists, mathematicians and miscellaneous thinkers for centuries [FS21]. The field of computational social choice (COMSOC) is a fairly new research area, having only been formed as such in the early 2000s [FCE+16]. COMSOC, in a theoretical computer science approach to collective decision making, studies among other things the computational difficulty of voting rules, the mathematical complexity of manipulating them and aims to bring computer science algorithm design and analysis to the problems of social choice theory.

The field of COMSOC, albeit it so young, has already gained traction. Social choice online alone has opened up a whole new research area [End17, Chapter 20]. But not only the internet is giving food for thought, further examples include judgment aggregation, voting rule verification and multiwinner voting [End17, Chapters 7,14,2].

One election type that falls under multiwinner voting is called participatory budgeting (PB). PB elections are a form of democratic practice that originated in Brazil in the 1980s, more specifically in the very south in Porto Alegre [Sha07]. It is nowadays being practised all over the world, sometimes adapted to individual needs or circumstances[SHR08], including in Amsterdam and Warsaw. The idea is fairly simple: The municipality defines a budget, projects can be proposed that could be realized with that budget. Projects come with an attached proposed cost, for feasibility reasons this cost is upper bounded by the budget. All eligible voters can then vote on which projects they would like to see realized [Wam00]. The relevant vote type for this work is approval voting, where voters can approve of as many of the proposed projects as they like [AS21]. Non-approval does not translate to disapproval.

In an effort to democratize public decision making, participatory budgeting can therefore be a big step towards more direct citizen involvement, regarding interest, perceived representation and responsibility [Wam00]. COMSOC is hereby interested in the computational complexities of such an election as well as how formal fairness criteria apply.

As with any election, the question of how to best determine the, in this case multiple, winners immediately begs itself. Widely used is the utilitarian greedy rule [RM23]: Add a new project to the winning outcome each round. In each round, include the project with the highest number of approvals in the outcome, such that the combined project costs of the outcome do not surpass the given budget. A second rule with which one can calculate the outcome of a participatory budgeting election is called the Method of Equal Shares (MES) [PPS21]: The budget is divided equally among all the voters. The voters then use their allotted share of the budget to fund the projects they approve of. The MES outcome is also decided in rounds. In each round, the project with the highest number of approvals is chosen for the outcome, as long as it is still affordable by its voters' shares. The cost of the project gets distributed among its voters as evenly as possible. After each round, the remaining voters' shares are updated accordingly [RM23].

The Method of Equal Shares has been proven to be proportional [PPS21], whereas the utilitarian greedy method is known not to be. Proportionality in this context refers to the principle that a fraction of the voters of a certain size decide on how that equal fraction of the budget is allocated. This proportionality definition stems from the idea that it is fair for an interest group of voters to decide on its proportional share of the outcome. In a PB election, the outcome is defined by how the budget is allocated. If proportionality is not a given, minority voter groups might get disenfranchised.

Well-researched and evidence-based claims for the Method of Equal Shares are an integral part of PB research. This thesis aims to contribute to this. We aim to study bias in PB elections in regards to project costs and how they influence the election outcome. In other words, do voting rules favor cheaper or more expensive projects, disregarding popularity? Our focus lies on a comparative analysis between the widely used utilitarian greedy voting rule and the Method of Equal Shares.

Towards that goal, we firstly put this work in the context of the current research status. Furthermore, we introduce the needed mathematical principles and other preliminaries. We then develop the first ever bias formulation for PB elections. We prove two theoretical results about bias in general instances of PB elections decided with differing greedy-style voting rules. We conduct several statistical experiments, including small scale experiments to gain insight into the relation of project cost and project popularity in real-world instances and large scale experiments on artificial instances to extensively collect data on bias in PB elections. We discuss our results as well as further possible research into this topic.

CHAPTER 2

Related Work

The Method of Equal Shares was first proposed by Peters and Skowron [PS20] in the context of approval-based committee voting. The authors pick up on a century-old feud between Scandinavian mathematicians Thiele and Phragmén [Jan16]. The mathematicians proposed different proposed for fair voting in committee elections that ensure proportional voter representation in said committee. Thiele argued for *Welfarism*, a concept that maximizes a chosen voter's utility function. The voting rule he developed is now known as Proportional Approval Voting (PAV). Phragmén was in favor of the proportionality concept already described above: A fraction of the voters decide on that equal fraction of the committee. Peters and Skowron showed that these two categories of proportionality are distinct. Furthermore, they designed the Method of Equal Shares to be a Phragmén-like rule and, importantly, proportional in that sense. Notably, however, the Method of Equal Shares acts as a compromise, since it is similar in some fairness properties to PAV. For example, a strong fairness axiom in PAV, Extended Justified Representation (EJR), meaning that a very large subgroup of voters that share preferences have several representatives in the outcome committee [ABC+15], is satisfied by the Method of Equal Shares. This is not the case with Phragmén-style rules. Furthermore, they showed that the Method of Equal Shares is polynomial-time computable.

Peters and Skowron together with Pierczyński [PPS21] built upon their work and introduced the Method of Equal Shares in a participatory budgeting context. They argue that concepts from committee elections can be generalized to the PB case, since a committee election can be seen as a PB election where all candidates have unit cost. Therefore, they introduce EJR in the participatory budgeting context: In the name of proportionality, no group of voters that share preferences is underserved. The authors formulate the Method of Equal Shares for approval-based PB elections, prove that it satisfies this new definition of EJR, and that the voting rule is still polynomial-time computable. Additionally, they show that PAV cannot be extended to the non-unit cost case and that this Thiele-method is therefore not usable for participatory budgeting elections.

Lackner and Skowron [LS23] wrote the book on multiwinner voting with approval preferences, giving in-depth characterizations of approval-based committee voting, its rules and their properties as well as putting a big focus on proportionality. Participatory budgeting is also included, even if not a focus.

Rey and Maly [RM23] performed a full survey into indivisible participatory budgeting research, indivisible meaning that projects are either fully funded or not at all. They concluded their survey as follows: Mostly relevant and mostly studied are approval ballots. Whether a non-approval constitutes a disapproval is unclear. Several fairness aspects can still be better researched or discussed, including the fact that large cohesive voter groups are not common in practice, and therefore not an ideal focus of fairness research. In addition to fairness criteria, other axioms from multiwinner voting could be generalized to further other comparisons of voting rules. Since PB elections are inherently dependent on voter participation, more and high-quality explainability in participatory budgeting can only ever be helpful. Furthermore, Rey and Maly took note of the fact that most PB research is done with Western standards of PB in mind, while non-Western participatory budgeting elections not only deserve research too but could also lead to new research directions.

Notably, bias in PB elections has not been studied so far. Bias in a PB elections needs a nuanced approach. Since no bias formulation exists in participatory budgeting as of yet, we have decided to focus on the measurable factor of project cost and how the outcome of two differing voting rules is influenced by it.

Bartocci et al. [BGME23] published an organized summary in 2023 how participatory budgeting research has gained traction over the last three decades. Their findings clearly show a trend towards more published papers regarding PB. This journey of PB research, as they call it, started in the Americas, as historically is to be expected. In their data set, 26% of the published work included a survey or statistical analysis on collected empirical data. Our work is aligned with a seemingly recent (starting 2014) trend of mixed approaches between theory and statistical work.

A widely-used PB resource, that also extensively features in this work, is the Pabulib[1] website. This open access data bank was created by Faliszewski et al. [FFP+23] and officially published in 2023. It offers an extensive back catalog of real-world participatory budgeting instances of European cities, mostly from the last decade. As such, it is an invaluable resource for even small-scale statistical experiments with PB data. Faliszewski et al. also formally introduced Pabutools[2]. The Python library is also publicly available and provides a standardized environment for working with PB election data. This includes calculating the election outcome with all common voting rules. Furthermore, the authors introduced Pabustats[3], which is a web application that helps users compare PB voting rules and how they apply to instances in the Pabulib data format *.pb*.

[1] `www.pabulib.org`
[2] `www.pypi.org/project/pabutools/`
[3] `www.pabulib.org/pabustats`

Boehmer et al. [BFJ+24] published a guide on numerical experiments on computational social choice elections and during that also performed some themselves. Specifically, the authors studied how numerical experiments in papers published in the three major relevant conferences (IJCAI, AAAI, AAMAS), for example including frequency analysis on whether the focus lay on approval or ordinal voting. Furthermore, and relevant to this work, they also focused on participatory budgeting elections, noting that the overwhelming number of election instances used from research were taken from the Pabulib website. The authors also performed a frequency analysis on those Pabulib-instances, specifically on the number of voters and the number of projects that occur.

Szufa et al. [SFJ+22] extensively researched different statistical cultures when sampling approval elections. Sampling models studied are the resampling, disjoint, noise and Euclidean models. These were compared in extensive statistical experiments, including to real-world PB elections taken from Pabulib. The authors propose using the resampling model to sample approval PB elections. Furthermore, they define parameters to create realistic approval instances. The paper was the basis for a Python package to sample approval elections with the specified models[4].

[4] www.pypi.org/project/prefsampling/

CHAPTER 3

Preliminaries

We formally introduce the election, voting rules and other mathematical concepts at the basis of this work. Unless otherwise disclosed, the nomenclature and definitions of participatory budgeting notions follow those of the survey by Rey and Maly [RM23].

3.1 Participatory Budgeting Elections

A participatory budgeting election is formally represented by an *instance*.

Definition 1 (Instance)**.** An *instance* of a participatory budgeting election is a triple $I =< P, c, b >$, where $P = \{p_1, ..., p_m\}$ is the set of *projects* up to vote, $c : P \rightarrow \mathbb{R}_{>0}$ is the *cost function* with the cost of a project given by $c(p) \in \mathbb{R}_{>0}$, and $b \in \mathbb{R}_{>0}$ the *budget limit*.

Essentially, the participatory budgeting instance is the information voters get presented with when going to cast their vote: What are the names of the projects being voted upon? How costly is each of them? What is the overall budget? For the sake of meaningful election results, only projects which are feasible inside the given budget are being put up to vote, and a budget has to be realistic. This first point is important since this work is interested in *indivisible* or *discrete* PB: A project is either fully funded as a result of the election or not at all. In other words, the inclusion in the winning outcome is a binary value.

Definition 2 (Profile)**.** A *profile* of a participatory budgeting election is a vector $A = (A_1, ..., A_n)$, with $N = \{1, ..., n\}$ as the set of *voters* taking part in the election, and A_i being the *ballot* of voter $i \in N$.

Their individual ballot is therefore how a voter expresses their private preferences over the set of projects. In our case these ballots are *approval ballots*, which are a widely used ballot style to gain insight into voter support of the respective projects. For each $i \in N$ $A_i : P \to 0, 1$, where 1 expresses a preference for the given project.

In summary, an election instance is what is being voted upon, and the insight into voter preferences gained through that election is the election profile. We will at a later point in this work sometimes refer to the combination of instance and profile as simply an election instance, this will be clear in context, however.

An election *outcome* is defined by a budget allocation $\pi \subseteq P$ such that $c(\pi) \leq b$. The condition ensures necessary feasibility. Election outcomes are determined by *voting rules*:

Definition 3 (Voting Rule)**.** A *voting rule* for a participatory budgeting election is defined as $R(I, A) \subseteq FEAS(I)$, where $FEAS(I) = \{\pi \subseteq P \mid c(\pi) \leq b\}$ is the set of all feasible budget allocations i.e. outcomes for the given instance I.

It is generally assumed that a PB voting rule is *resolute*, meaning that the rule always returns a single feasible budget allocation as not only an but *the* outcome. Considering this, the choice of voting rule is evidently crucial.

3.1.1 Voting Rules

There are several voting rules for participatory budgeting elections relevant to this work. As a whole, however, they can be divided into two categories: greedy-style voting rules and the Method of Equal Shares, which is the sole relevant Phragmén-style voting rule.

Greedy-style Voting Rules

What we refer to as greedy-style voting rules are formally voting rules maximizing voters' *welfare*. This idea of welfare is based upon the concept of voter *satisfaction*. As true preferences are inherently private, we can only assume about voters' individual and overall satisfaction. In general, however, it is reasonable that voters are more satisfied with an election outcome the more of the projects they have indicated to support are included. A *utilitarian* voting rule given an instance I and a feasible outcome π seeks high utilitarian social welfare, dependent on a utility function $\mu_i : 2^P \to \mathbb{R}_{\geq 0}$, where the utility function has to be guessed at based on the voters' ballots:

$$\text{Util-SW}(I, (\mu_i)_{i \in N}, \pi) = \sum_{i \in N} \mu_i(\pi)$$

This voting rule is what we have conducted our statistical experiments with. The utility function in practice gets switched out for a *satisfaction function*.

Definition 4 (Approval-Based Satisfaction Functions)**.** When given an instance I and a profile A, an *approval-based satisfaction function* is a mapping $sat : 2^P \to \mathbb{R}_{\geq 0}$ that satisfies the following conditions:

- inclusion monotonic: $sat(P) \geq sat(P')$ for all P, P' such that $P \supseteq P'$
- 0 only for the empty set: $sat(P) = 0$ if and only if $P = \emptyset$

Furthermore, the *satisfaction of a voter* $i \in N$ for a possible outcome π is

$$sat_i(\pi) = sat(\{p \in \pi | A_i(p) = 1\})$$

There are several different satisfaction functions, the relevant one for most of this work is the *cost satisfaction function*:

Definition 5 (Approval Score)**.** The *cost satisfaction function* defines voters' satisfaction as the cost of the selected and approved projects: $sat^{cost}(P) = c(P)$.

Towards further greedy-style voting rules for an election using approval ballots, we can utilize a project's overall *approval score*:

Definition 6 (Approval Score)**.** The *approval score* of a project $p \in P$ is defined as $app(p, A) = |\{i \in N | A_i(p) = 1\}|$.

Simply put, the approval score can be interpreted the sum of voters who have indicated approval of project p. We can further use this approval score to define two voting rules for exact welfare maximization (some fixed tie-breaking method ensures resoluteness):

$$\text{MaxCard}(I, A) = \underset{\pi \in FEAS(I)}{\arg\max} \sum_{p \in \pi} app(p, A)$$

This rule chooses the feasible outcome with the highest approval score, disregarding any other factors. Factoring in the project cost is done using the following variation:

$$\text{MaxCost}(I, A) = \underset{\pi \in FEAS(I)}{\arg\max} \sum_{p \in \pi} app(p, A) \cdot c(p)$$

As such exact methods are computationally expensive, we make use of an approximation scheme to effectively maximize the outcome welfare during the execution of the voting rule.

Definition 7 (Greedy Scheme)**.** Assume an instance $I =< P, c, b >$ and a strict ordering $\triangleright$ over P. The *greedy scheme* $\text{Greed}(I, \triangleright)$ iteratively adds projects to an initially empty outcome π. Following the order $\triangleright$, a project p is included in the outcome if and only if $c(\pi \cup \{p\}) \leq b$. If there is no more project according to $\triangleright$, π is the outcome of $\text{Greed}(I, \triangleright)$.

Towards a greedy approximation of the MaxCard-rule, we assume the ordering $\triangleright$ to be given by the approval scores of the projects divided by their costs. In the order, $p \triangleright p'$ if and only if $\frac{app(p,A)}{c(p)} \geq \frac{app(p',A)}{c(p')}$:

$$\text{GreedCard}(I, A) = \{\text{Greed}(I, \triangleright) | \triangleright \text{ is compatible with } \frac{app}{c}\}$$

The approximation formulation for the MaxCost-rule asks for the order $\rhd$ to be compatible with the approval score *app* only, i.e. $p \rhd p'$ if and only if $app(p, A) \geq app(p', A)$. The scheme formulation follows:

$$\text{GreedCost}(I, A) = \{\text{Greed}(I, \rhd) | \rhd \text{ is compatible with } app\}$$

The Method of Equal Shares

The Method of Equal Shares (MES), like the greedy approximation schemes, computes its outcome over the course of rounds. Notably, every voter gets an equal share of the budget with which they can pay for projects they approve of to be included in the winning set of projects. Aiming for proportionality, the Method of Equal Shares prefers to include projects where the cost of that project are divided between its voters as equally as possible. Formally, MES is defined as thus:

Definition 8 (Method of Equal Shares)**.** When given an instance I and a profile A (consisting of approval ballots), the *Method of Equal Shares* (MES) iteratively constructs a budget allocation π from an initially empty one for a satisfaction function.

A voter's *load* $l_i : 2^P \to \mathbb{R}_{\geq 0}$ defines for each voter $i \in N$ how much money they have already spent. Initially, this is defined to be 0.

A voter's *contribution* γ_i for each voter $i \in N$ defines how much they would spend on a given project p for a given allocation π and a given scalar $\alpha \geq 0$:

$$\gamma_i(\pi, \alpha, p) = A_i(p) \cdot \min(\frac{b}{n} - l_i(\pi), \alpha \cdot sat(\{p\}))$$

Therefore, initially each voter receives $\frac{b}{n}$ of the money and only ever pays for projects they approve of.

For a given allocation π, a project p that is not currently included is defined as α-*affordable*, if it holds that

$$\sum_{i \in N} \gamma_i(\pi, \alpha, p) \geq c(p)$$

In other words, for a project to be α-affordable, the remaining shares of all voters approving of the project need to cover its cost. At any round of the rule, two cases are possible:

Case 1: No project is α-affordable for any α. The rule terminates and outputs the current budget allocation π.

Case 2: At least one project $p \in P \setminus \pi$ is α-affordable. Then α^*-affordable refers to a project that is α-affordable with the smallest α. An α^*-affordable project gets included in the outcome, $\pi \cup \{p\}$. To ensure correctness, all voters' loads are updated as well: $l_i(\pi \cup \{p\}) = l_i(\pi) + \gamma_i(\pi, \alpha, p)$.

Therefore, voters indeed pay their contribution if a project they approve of is included in the outcome. Choosing the α^*-affordable project works towards as equal of a pay distribution as possible. It is possible that no perfectly equal distribution is possible, because e.g. some supporters of the project have already spent a large chunk of their allotted currency. Optimizing for the lowest α guarantees that the ratio of what any contributor pays toward their satisfaction is minimal. Similarly to above, some tie-breaking method applies for resoluteness.

Since the Method of Equal Shares is known not to be *exhaustive*, i.e. utilizing the whole of the budget with its calculated outcome, *completion methods* exist. These methods aim to ensure that while budget is still free to use and projects are still feasible, the budget allocation does not yet halt. Completion methods therefore start when the original budget allocation is finished. The technique relevant to this work increases the budget by a given amount each round. The MES outcome is then computed again. This continues until an outcome overshoots the initial budget. The completion method then backtracks one round and uses the last feasible outcome inside the initial budget as an election result. Another alternative completion approach is utilizing a greedy-style rule for the leftover budget. The first method can be computationally expensive if the budget increments are small while the greedy-style rule is stronger in its neglect of wanted proportionality criteria. It is evident that the choice of completion method can hold some weight. Given that this work focuses on proportionality and the differences between greedy-style voting rules and the MES, the choice fell on the incremental method.

3.2 Statistics

As some statistical experiments are an important part of this work, some minor statistical preliminaries are necessary [Bos12].

3.2.1 Statistical Research Design

Essentially, the statistical research we performed followed the subsequent scheme:

Definition 9 (Descriptive Analysis)**.** *Descriptive Analysis* includes summaries like the mean and median values of what is being studied, as well as standard deviations and value ranges.

Such data, often displayed in histograms, is useful for the visualization of the distribution of project costs or how project costs relate to project approvals in real-world instances. Following this investigation, the next step is *regression analysis*. We are mostly interested in how the outcome of an election could and does relate to other factors, which is why our regression analysis is making use of a fitting model:

Definition 10 (Simple Logistic Regression)**.** *Simple Logistic Regression* studies a possible relationship between a single predictor variable and a binary outcome.

The goal of any such analysis is to form a hypothesis about statistical correlations. Such a hypothesis then needs to be extensively tested. Testing includes comparing the results of different instance groups to decide on the model that best fits what we want to measure. The *0-hypothesis* states, that a reduced, in other words not highly complicated model fits the data better than a highly complicated one. The goal is therefore to start with a simple model and only make specifications as truly required. If a model has been decided upon, it needs to be thoroughly checked:

Definition 11 (Robustness Checks)**.** *Robustness checks* of a statistical model ensure that the model does not rely on outliers to produce statistically meaningful results.

Essentially, robustness checks come down to orderly repetition. It needs to be ensured that the circumstances of these checks are not only reproducible but the same for each check of the model. Furthermore, to combat outliers in the measured statistical values that could distort the interpretation of results, averaging the results over several if not all runs of the tests is a common and sensible tool.

3.2.2 Point-biserial Coefficient

The point-biserial coefficient (PBC) measures the correlation between a numerical and a binary value. It therefore lends itself perfectly to measuring the correlation between our proportionality notion (difference between percentage of budget a project would use and the percentage of the possible approvals it received) and whether or not the project was included in the outcome (1 if yes, 0 otherwise). The statistical values of the PBC lie between 1 and -1. The PBC statistical value is 0 if there is no correlation. Positive values indicate a positive correlation (the higher the difference the more likely a project is to be in the outcome), and vice versa. Every PBC statistical value stands in relation to a p-value calculated at the same time. The p-value gives information about the statistical relevance of the determined value. Generally, a p-value below 0.06 is seen as statistically relevant. For any value higher than this, the measured correlation cannot be seen as necessarily accurate. Any reasonings based upon a PBC statistical value that is in relation to a p-value above 0.05 can at most be helpful in the context of very general trends, if even.

CHAPTER 4

Bias in Participatory Budgeting

Bias commonly refers to the idea that some factor distorts the result of a procedure in some measurable way. In our case, the result this refers to is the outcome of a participatory budgeting election. Essentially, we are asking whether there is a factor that holds more sway in the election decision than it should have. We want the result of a voting rule to be decided by project approvals, apart from obvious constraints like budget feasibility. Towards proving - mathematically or empirically - whether this is the case or not, we have decided to focus on project cost as the potentially distorting factor.

For completion's sake without any actual guarantee of exhaustiveness, other bias models are also conceivable: Length of project description, socioeconomic background of voters directly impacted by the project on a daily average, number of times a project has been up for vote before. However, all of these specific ideas veer towards less measurable and more into the area of political science.

Intuitively, bias in PB elections regarding the project costs means the following: While we want the election to be decided by project approvals alone, do the project costs hold measurable influence? Since an outcome is decided by approvals, and you can assume that approvals are not independent of cost, this approval-cost-relationship is going to be our starting point. We will see that empirically, the project approvals are closely linked to its cost. While no bias formulation can be expected to be easy, this evident connection between what we want the outcome to be dependent on and what we want it to be independent of creates an obvious complexity.

As mentioned, we will mostly think about bias in PB elections as whether or not the project cost influences the likelihood that the project is part of the election outcome. An alternative bias formulation would be whether the likelihood of getting 1 unit of funding is influenced by the project cost. Given the scope of this work and the fact that inclusion in the outcome is easily measurable, we have decided to focus on the former.

Such a likelihood purely between project cost and outcome inclusion can easily be calculated with an existing instance. Given the relationship between approvals and cost, however, only focusing on that likelihood does not necessarily produce any meaningful results. We therefore focus on a project's *proportional cost*, comparing the percentage of the budget a project would use and the percentage of approvals it got. An expensive project under that definition would therefore use an disproportionately high amount of the budget, or in other words, be more costly than popular. The question of if the project costs influence the likelihood of outcome inclusion can then be approached via an investigation of a possible correlation between proportional cost and outcome inclusion. This lends itself well to an empirical approach to our bias question. Since we aim to compare possible bias in different participatory budgeting voting rules, having a clear correlation available for juxtaposition is practical as well.

Towards meaningful theoretical results, formulations based on likelihoods and statistical correlations have shown themselves to be beyond the scope of this work. We therefore take advantage of specialized instances of PB elections, where the correlation between project approvals and project cost is the same for all projects. We can then argue about how a specific voting rule handles two distinct projects that have differing costs. This illustrates how the voting rule is influenced by the project costs specifically. On the flip side, this model does not lend itself well to statistical experiments, given that a real-world instance seldomly resembles such an artificial one.

In summation, no clear-cut, general bias formulation for project costs in participatory budgeting elections is possible, at least along our lines of argumentation. This does not mean, however, that no sensible and furthering formalisms can and have been created. Specifically general instances hold even more potential for theoretical reasoning about project cost bias. In terms of statistical measurements, given the nature of how project approvals are distributed, our formulation seems sensible. A precise mathematical expression remains a goal, of course, but the presented ones still hold clear value when arguing about the given bias objective, from both the theoretical and the empirical standpoints. In addition to that, both concepts are independent of voting rules and notably lend itself to comparisons between such rules.

CHAPTER 5

Theoretical Results

We prove new theoretical results regarding bias for two types of general instances in participatory budgeting elections: project approvals being fully dependent or fully independent on project cost.

5.1 Project Approvals Proportional to Project Cost

We define a *random costs-dependent* instance of participatory budgeting to consist of randomly sampled project approvals for a given set of projects. Furthermore, the project costs are dependent on these approvals in the following way: $\forall p \in P$ let $|N(p)| = x_p \cdot c(p)$, where x_p is a project-specific parameter sampled randomly with some distribution function x. Assume two projects p_1, $p_2 \in P$ s.t. $c(p_1) < c(p_2)$.

Proposition 1. *When determining the election outcome of a random costs-dependent instance with the greedy cardinality welfare rule maximizing the ratio of approvals to cost, the likelihood of p_1 being chosen for the outcome cannot be smaller than that of p_2.*

Proof. Let $\mathcal{F}$ be the set of feasible projects that can still be added to the outcome. Since the set of feasible projects is subject to change over the rounds of the voting rule, let $\mathcal{F}_r$ be the set of feasible projects in round r of the voting rule. Let b_r be the remaining budget at round r. If in the first round r_0 the project p_1 is not feasible, i.e. $p_1 \notin \mathcal{F}_{r_0}$, because $c(p_1) > b_{r_0}$, it holds that $p_2 \notin \mathcal{F}_{r_0}$ as well. In general, the proposition holds if $p_2 \notin \mathcal{F}_{r_0}$. We therefore assume for further argument that both projects are feasible in the beginning.

The given greedy-style voting rule optimizes the following criterion:

$$\max \frac{N(p)}{c(p)}$$

For the likelihood $\mathcal{L}_{p_1,r_k}$ that project p_1 is chosen for the outcome in exactly round r_k the following holds:

Case 1: $p_1 \notin \mathcal{F}_{r_k}$, $\mathcal{L}_{p_1,r_k} = 0$.

Case 2: $p_1 \in \mathcal{F}_{r_k}$, the likelihood depends on how well p_1 optimizes the given criterion in comparison to the other feasible projects. Specifically, p_1 is chosen if $\forall p \in \mathcal{F}_{r_k} \setminus \{p_1\}$ it holds that $\frac{N(p)}{c(p)} < \frac{N(p_1)}{c(p_1)}$. By construction, $\frac{N(p_1)}{c(p_1)} = x_{p_1}$. Therefore, $\mathcal{L}_{p_1,r_k}$ is fully and only dependent on the x_p values of the feasible projects.

Following the equal construction of p_1 and p_2, $\mathcal{L}_{p_2,r_k}$ is therefore fully and only dependent on the x_p values of the feasible projects. As all of those values are sampled in the same way, the likelihood that p_1 or p_2 are most optimal and chosen in some round r_k is also the same, if they are both in $\mathcal{F}_{r_k}$. Therefore, for any round r_k where a project p_j is feasible we can calculate

$$\mathcal{L}^{\geq}_{p_j,r_k} = \mathcal{L}_{p_j,r_k} \cdot (1 - \mathcal{L}^{\geq}_{p_j,r_i-1}) + \mathcal{L}^{\geq}_{p_j,r_i-1}$$

where $\mathcal{L}_{p_j,r_i}$ depends on the maximal set P_O s.t. $P_O \subseteq \mathcal{F}_{r_i} \setminus \{p_j\}$ and $\forall p \in P_O$ $\frac{N(p)}{c(p)} > \frac{N(p_j)}{c(p_j)}$. The base case for r_0 is defined analogously. If both projects are feasible in round r_{k+1}, $\mathcal{L}_{p_1,r_{k+1}} = \mathcal{L}_{p_2,r_{k+1}}$ since at any round r_k the likelihoods are equal, and it still holds in r_{k+1} that the likelihoods are both solely dependent on x_{p_1} and x_{p_2}, which were sampled equally.

It remains to be shown that it is not possible for p_2 to be in $\mathcal{F}_{r_{k+1}}$ while p_1 is not. The set of feasible projects changes from round to round. After each round r_k, exactly two things are possible:

Case 1: The set of feasible projects is empty, i.e. all the budget has been allocated. The outcome calculation is completed.

Case 2: Some budget has been allocated and $\mathcal{F}_{r_{k+1}} \subset \mathcal{F}_{r_k}$ gets created. Exactly two types of projects are being excluded in $\mathcal{F}_{r_{k+1}}$:

- The project $p_O \in \mathcal{F}_{rk}$ that was chosen to be in the outcome in round r_k.
- The set of projects $P_e \subset \mathcal{F}_{rk} \setminus \{p_O\}$ where $\forall p_e \in P_e$ it holds that $c(p_e) > b_{r_{k+1}}$, which are now newly infeasible at round r_{k+1}.

Our argument stops the moment one of the two projects is chosen. We can therefore assume that p_O was neither p_1 nor p_2. It is true that if $p_1 \in P_e$ since $c(p_1) < c(p_2)$ it has to hold that $p_2 \in P_e$. The likelihoods were equal up to this point and are equally 0 now, the argument holds. If $p_2 \in P_e$ it does not automatically imply that $p_1 \in P_e$. Therefore, the likelihood that p_1 is chosen for the outcome is equal to that of p_2 for all rounds where they are both in the set of feasible projects, and larger in the case where p_2 is not feasible anymore while p_1 still is.

□

5.2 Project Approvals Uniformly At Random

We define a *random costs-independent* instance of participatory budgeting to consist of randomly sampled project approvals for a given set of projects with some distribution function x. Similarly to above, we therefore assume x_p to be a project-specific parameter. Furthermore, the project costs are sampled fully independently from the project approvals. Assume two projects p_1, $p_2 \in P$ s.t. $c(p_1) < c(p_2)$.

Proposition 2. *When determining the election outcome of a random costs-independent instance with the greedy cost welfare rule maximizing the number of approvals, the likelihood of p_1 being chosen for the outcome cannot be smaller than that of p_2.*

Proof. Let $\mathcal{F}$ be the set of feasible projects that can still be added to the outcome. Since the set of feasible projects is subject to change over the rounds of the voting rule, let $\mathcal{F}_r$ be the set of feasible projects in round r of the voting rule. Let b_r be the remaining budget at round r. If in the first round r_0 the project p_1 is not feasible, i.e. $p_1 \notin \mathcal{F}_{r_0}$, because $c(p_1) > b_{r_0}$, it holds that $p_2 \notin \mathcal{F}_{r_0}$ as well. In general, the proposition holds if $p_2 \notin \mathcal{F}_{r_0}$. We therefore assume for further argument that both projects are feasible in the beginning.

The given greedy-style voting rule chooses the project with the highest number of approvals for the outcome. Since the election instance was created randomly, the project approvals and therefore costs are fully independent from the project names, and we can assume lexicographic tie-breaking without loss of generality.

For the likelihood $\mathcal{L}_{p_1,r_k}$ that project p_1 is chosen for the outcome in exactly round r_k the following holds:

Case 1: $p_1 \notin \mathcal{F}_{r_k}$, $\mathcal{L}_{p_1,r_k} = 0$.

Case 2: $p_1 \in \mathcal{F}_{r_k}$, the likelihood depends on how many approval votes p_1 has received. Specifically, p_1 is chosen if $\forall p \in \mathcal{F}_{r_k} \setminus \{p_1\}$ it holds that $N(p) < N(p_1)$. By construction, $N(p_1)$ is defined by x_{p_1}. Therefore, $\mathcal{L}_{p_1,r_k}$ is fully and only dependent on the x_p values of the feasible projects.

Following the equal construction of p_1 and p_2, $\mathcal{L}_{p_2,r_k}$ is therefore fully and only dependent on the x_p values of the feasible projects. As all of those values are sampled in the same way, the likelihood that p_1 or p_2 are most optimal and chosen in some round r_k is also the same, if they are both in $\mathcal{F}_{r_k}$.

Therefore, for any round r_k where a project p_j is feasible we can calculate

$$\mathcal{L}^{\geq}_{p_j,r_k} = \mathcal{L}_{p_j,r_k} \cdot (1 - \mathcal{L}^{\geq}_{p_j,r_i-1}) + \mathcal{L}^{\geq}_{p_j,r_i-1}$$

where $\mathcal{L}_{p_j,r_i}$ depends on the maximal set P_O s.t. $P_O \subseteq \mathcal{F}_{r_i} \setminus \{p_j\}$ and $\forall p \in P_O\ N(p) > N(p_j)$. The base case for r_0 is defined analogously.

If both projects are feasible in round r_{k+1}, $\mathcal{L}_{p_1,r_{k+1}} = \mathcal{L}_{p_2,r_{k+1}}$ since at any round r_k the likelihoods are equal, and it still holds in r_{k+1} that the likelihoods are both solely dependent on x_{p_1} and x_{p_2}, which were sampled equally.

It remains to be shown that it is not possible for p_2 to be in $\mathcal{F}_{r_{k+1}}$ while p_1 is not. The set of feasible projects changes from round to round. After each round r_k, exactly two things are possible:

Case 1: The set of feasible projects is empty, i.e. all the budget has been allocated. The outcome calculation is completed.

Case 2: Some budget has been allocated and $\mathcal{F}_{r_{k+1}} \subset \mathcal{F}_{r_k}$ gets created. Exactly two types of projects are being excluded in $\mathcal{F}_{r_{k+1}}$:

- The project $p_O \in \mathcal{F}_{rk}$ that was chosen to be in the outcome in round r_k.
- The set of projects $P_e \subset \mathcal{F}_{rk} \setminus \{p_O\}$ where $\forall p_e \in P_e$ it holds that $c(p_e) > b_{r_{k+1}}$, which are now newly infeasible at round r_{k+1}.

Our argument stops the moment one of the two projects is chosen. We can therefore assume that p_O was neither p_1 nor p_2. It is true that if $p_1 \in P_e$ since $c(p_1) < c(p_2)$ it has to hold that $p_2 \in P_e$. The likelihoods were equal up to this point and are equally 0 now, the argument holds. If $p_2 \in P_e$ it does not automatically imply that $p_1 \in P_e$. Therefore, the likelihood that p_1 is chosen for the outcome is equal to that of p_2 for all rounds where they are both in the set of feasible projects, and larger in the case where p_2 is not feasible anymore while p_1 still is.

$\square$

CHAPTER 6

Statistical Experiments

In line with what we have proven for bias regarding project costs in PB elections, we conducted experiments on real-world and artificial instances. Specifically, we were interested in how project costs, project approvals and election outcomes relate. We compared the results for the greedy utilitarian welfare rule using the cost satisfaction function and the Method of Equal Shares with the cost satisfaction function under the use of the described voter budget increment completion method (budget increase 1).

6.1 Approval-Cost Proportionality

The *proportionality* of a project is referring to the difference between the percentage of the budget a project would need, and the percentage of the possible approvals it got in the vote. A project is said to be perfectly *proportional* if that difference is 0.

$$prop_p = \frac{c(p)}{b} - \frac{N(p)}{n}$$

By definition, a positive proportionality would constitute a project being proportionally expensive and vice versa. As argued before, we are using this idea of proportionality to make our idea of bias more meaningful. Using fractions ensures comparability between different voting instances. Furthermore, it seems reasonable that a more expensive project would impact more people and therefore be of interest to more voters. Simply investigating the percentage of a budget a project would use and how that effects inclusion in the election outcome is not nonsensical, of course, but for a complete picture acknowledging the relationship between cost and approvals is necessary.

From our point of view, a project being funded that uses, along the line of the Method of Equal Shares, its voters' allotted part of the budget is neutral in terms of bias. The chosen aspect of project proportionality therefore lends itself quite nicely to our idea.

6.1.1 Proportionality in Real-World Instances

To get an idea of this proportionality in practice and therefore justify our argumentation, we performed a frequency analysis of proportionality in several real-world instances taken from Pabulib. The chosen instances are all instances of participatory budgeting elections with approval voting in the span of the last decade. Beside that no special criteria applied.

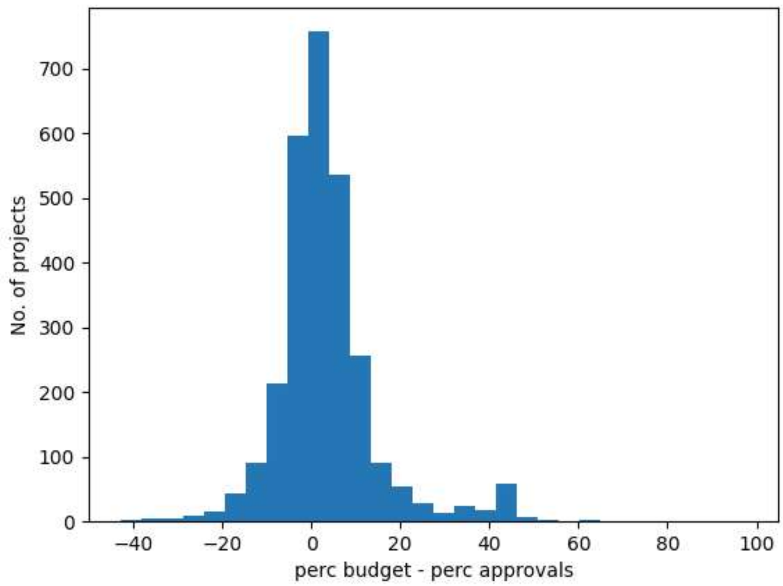

Figure 6.1: Frequency analysis of cost-approval-proportionality in real-world instances

We can see in Figure 6.1 that from the thousands of projects in our query group, the large majority of those projects is roughly, if not exactly proportional. In other words, the normal distribution around the difference 0 dominates. There does exist a group of outliers that does not meet our proportionally requirements but is instead much more costly than popular.

Intuitively, this makes sense. As argued before, higher cost projects will likely have a bigger influence and therefore impact more people directly. It is only logical that more voters might approve of such a project. The same argumentation holds for relatively inexpensive projects, which likely affect way less of the populace. The outliers can have several reasons: Either, the project is very expensive for the theoretical impact it would have, or that impact is ideologically unwanted by the majority of people, or the project is deemed to be not well-thought out and as such unpopular. Just by definition, any project that is first and foremost very unpopular despite of cost will have a large positive project proportionality.

6.1.2 Sampling Proportional Instances

To be able to come to conclusions dependent on our chosen proportionality definition, we needed instances where that proportionality is not only known but also controllable. We therefore utilized the preference profile sampling tool created by Szufa et al. [SFJ+22] to sample the instances we used for further experimentation.

Our goal is to create instances that share the same approval profile, while differing in the project cost. In this way the impact of the cost explicitly can be measured in a controlled way. Following our line of argumentation, the instances specifically differ in how likely it is that the project cost of an individual project are directly proportional to its approvals. Instances are created as follows:

- **Sampling approval profile** Using the aforementioned sampling tool, we sample one approval profile. We are using the disjoint resampling method with parameters the authors have shown to lead to similar approval votes as real-world instances from Pabulib ($p = 0.75, \Phi = 0.125$) [SFJ+22]. The parameters for number of projects (20) and number of voters (1000) are taken from the most frequent values in Pabulib instances worked out by Boehmer et al. [BFJ+24]. Using these parameters, we use the budget associated to similar real-world instances on Pabulib. The specific budget does not matter, however, since we are using the same budget value for all instances, and furthermore only ever talk about project costs in terms of budget percentage.

- **Proportionality probability model** In line with our frequency analysis on real-world instances, our main decider is that a project can either be proportional or more expensive than it is popular. The proportionality probability gives the likelihood that the former is the case. If a project is chosen to be disproportionate, the project cost is sampled from a normal distribution that is limited by and centered around the would-be proportional cost and the budget limit (standard deviation is 20000 for a budget of 600000). This is a simplification of the distribution we could observe in real-world instances. While a consistent distribution model is necessary, a specific type of distribution is not strictly required, since we are only ever interested in whether a project is proportionally expensive in a binary and not a quantified sense. As can be observed in Figure 6.2b, however, this specific distribution can model project costs close to what we could see in Pabulib instances.

- **Deciding on proportionality probability step** We decide on three parameters: How many instances we create (40), the initial likelihood that a project is proportional (generally chosen as 1, fully proportional), and what the difference in that likelihood is to be between the created instances. This proportionality probability step assumes that the likelihood for a project to be proportional decreases by that step for the next created instance. In simple terms, our instance projects overall become less and less proportional with every created instance.

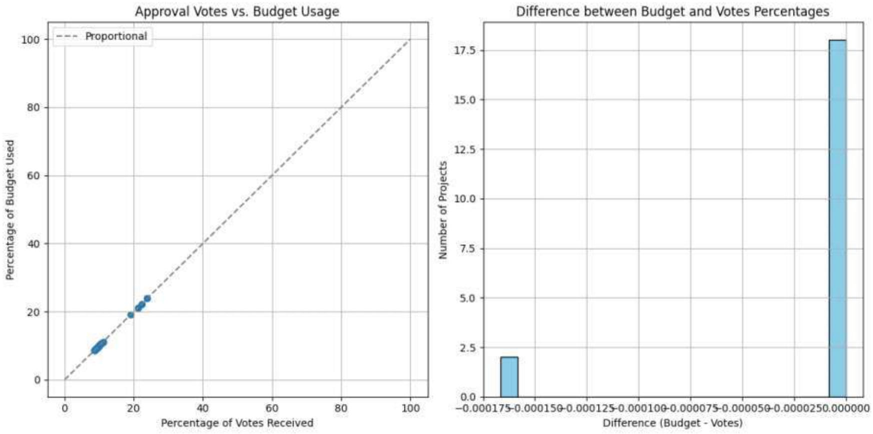

(a) Proportionality probability 0.9

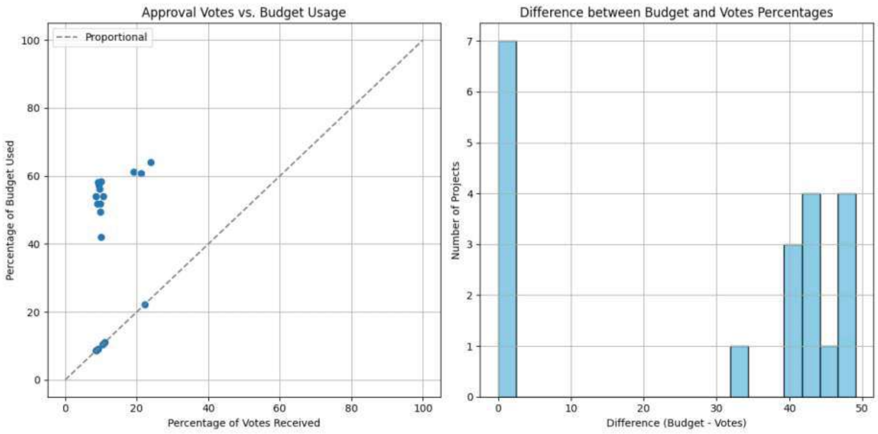

(b) Proportionality probability 0.32

Figure 6.2: Frequency analysis of cost-approval-proportionality for instances differing only in the given probability.

- **Creating cost profiles** We use these parameters and the described model to create cost profiles that together with the sampled approval profile build our instances.

- **Creating proportionality histograms** To further a good understanding of our created instances, we create a histogram of frequency analysis for project proportionality for each instance as well as a scatter plot. Two of those created plots can be seen in Figures 6.2a and 6.2b. Instances with a high proportionality probability have such minor differences in percentages that the plotting only makes sense in context with low proportionality probability instances. This can be observed in Figure 6.2a, where the differences can be interpreted as 0 for all projects. For such instances with a proportionality probability below 0.5, the disproportional projects dominate, as is logical. We can see that our chosen normal distribution for determining the expensive costs does ensure close-to-reality values for the specific proportionality of the outliers. Real-world instances can be assumed to have a proportionality probability closer to 1 and certainly above 0.5.

- **Saving the instances** In the name of traceability and repeatability, we use the Pabutools library to correctly save the created instances as pb-files.[1]

6.2 Correlation between Project Proportionality and Inclusion in Outcome

Towards meaningful statistical results of our bias formulation, we aim to measure if there exists a correlation between our idea of project proportionality and whether or not the project is included in the outcome.

6.2.1 Experiments with Differing Proportionality Probability

Firstly, it was important whether any correlation between our chosen factors exists. Furthermore, we wanted to determine the environment in which we could measure such a correlation best, if it existed.

To achieve this target, we used the described sampling setup to create groups of instances that differ in the initial proportionality probability and the proportionality probability steps. For each group of instances, we calculated the outcomes with both the utilitarian greedy method with the cost satisfaction function and the Method of Equal Shares using the described completion method. These outcomes were then used to determine the PBC statistical value and the p-value for correlation between first simply the percentage of the budget a project would use and outcome inclusion. The point-biserial coefficient is by definition the logical correlation measure to use, given what measures we are investigating.

When looking at instances where all project costs are definitely disproportionately large (proportionality probability 0) as in the 40 instances from Figure 6.3, no significant correlation for the Method of Equal Shares can be shown. For the greedy rule at least some instances suggest a positive correlation between the percentage of the budget and outcome. Since there is no clear trend for a low p-value, however, no statistical significance can surely be assumed.

For the group of instances where the proportionality probability for each project decreases by 0.025 from initial probability 1 for every instance as in Figure 6.4, it can be observed that the Method of Equal Shares has a more negative correlation, i.e. it holds that for instances including proportionally expensive projects, the Method of Equal Shares is more likely to include cheap projects in the outcome, while the correlation using the greedy method is less severe. It is also relevant that the p-value for MES is very low for

[1]All main source code for this work is publicly available at `https://github.com/duzilia/CostBiasInPB`.

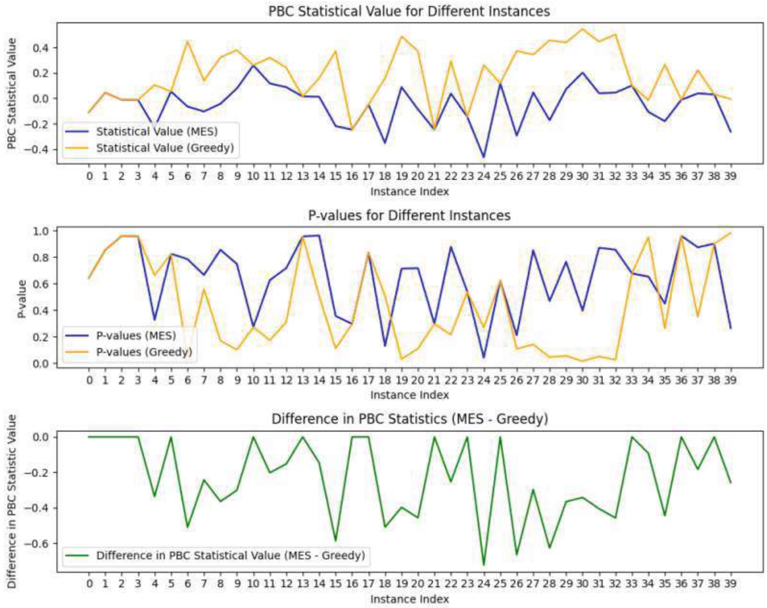

Figure 6.3: Correlation coefficient for percentage of budget and outcome, all projects of all instances disproportionately expensive, proportionality probability 0

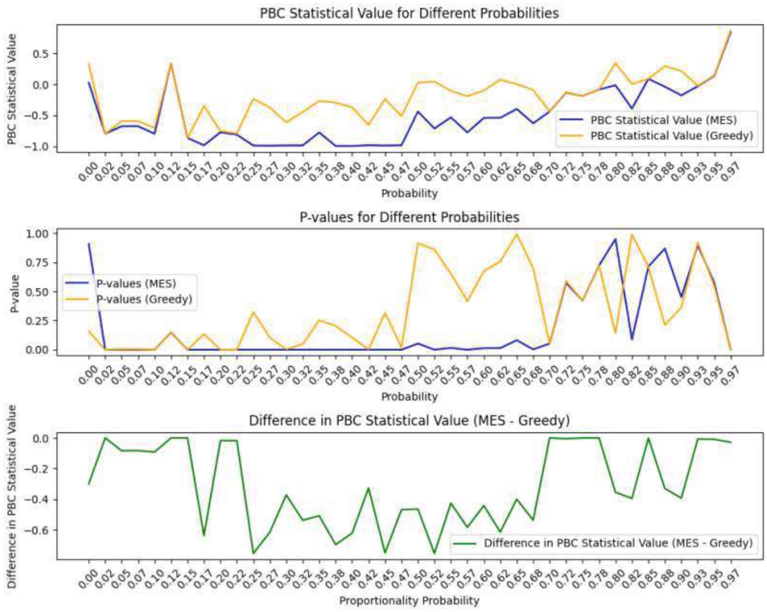

Figure 6.4: Correlation coefficient for percentage of budget and outcome, step in proportionality probability -0.025, initially 1

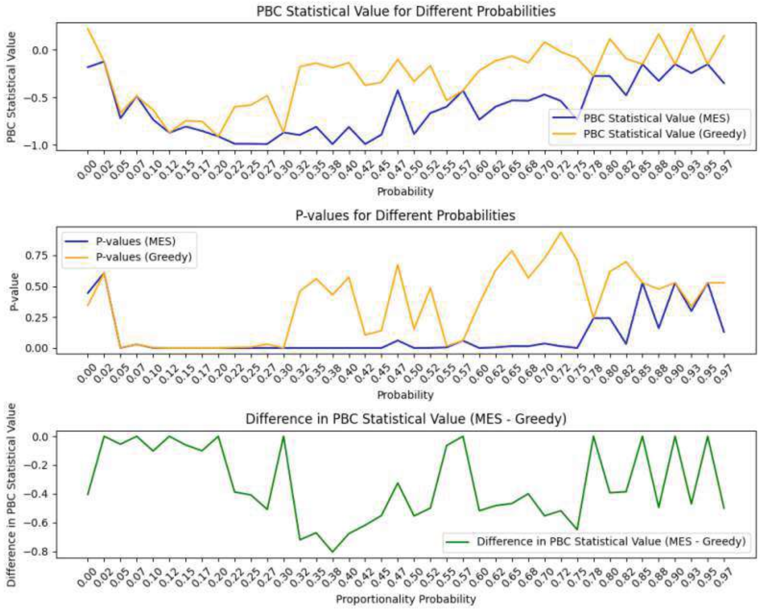

Figure 6.5: Correlation coefficient for difference between percentage of budget and percentage of approvals and outcome, step in proportionality probability -0.025, initially 1

most instances, which makes the correlation statistically significant. The p-value for the greedy-style rule often suggests that the PBC statistical value could be due to chance.

The case where the project cost is perfectly proportional to its number of approvals results in a very high statistical significance for positive correlation for both rules. It is notable that the correlation is very slightly stronger for the greedy rule.

All of these results firstly show that statistically relevancy is not an immediate given. Some statistically relevant results can be achieved, however. Notably, the Method of Equal Shares correlation shown in Figure 6.4 intuitively makes sense. The difference between what would be the proportional cost a voter has to pay for a cheap project and what is now the disproportional amount is quite small for projects that are still low cost.

Furthermore, we measured the correlation between the percentage difference of budget and approvals and whether or not a project is in the outcome, what we have described as our main bias idea. Similar to the result above, it is statistically significant that the Method of Equal Shares has a strong negative correlation, see Figure 6.5. Intuitively, this is correct: The more proportional a project cost to its approvals is the less a voter has to pay more than their proportional share to include that project in the outcome. Additionally, some evidently significant values for the less proportional instances can also be measured for the greedy-style rule. The PBC statistical value is less obviously negative, while still firmly below 0. By nature, a greedy rule will prefer popular projects. Over the rounds of the outcome calculation, the remaining budget will get smaller and smaller. A bias toward disproportionately inexpensive projects is only logical, since equally popular higher cost projects have a higher likelihood to be infeasible in later rounds. The fact that this bias only clearly shows up in the latter half of instances also makes sense, since a statistically significant result requires a certain amount of precedence. This is simply not realistic with mostly proportional projects in an instance.

It is also relevant to note, however, that the difference between these correlation values and the ones measured between percentage of budget and inclusion in outcome is not significant. Some difference exist, but the trends of both the PBC statistical value and the p-value can only be described as very similar. This lends itself quite nicely to our experiments, since we can easily lay our main focus on one measure, in this case our difference-based notion of bias, while at the same time being able to reason with the general findings for the correlation calculated solely on percentage of budget used.

6.2.2 Large Scale Experiments

We used our knowledge that some significant results are possible and built upon that. Towards meaningful results about our chosen bias formulation, we repeated the described experiments with 100 approval profiles (groups of instances) all created with the same parameters. Specifically, the number of votes is 1000, the number of projects 20, the budget 500000. The number of instances created for each approval profile is 40. The probability step is -0.025 starting at the initial proportionality probability 1.

Similarly as above, we calculated the election outcomes and used these to determine the PBC statistical and p-values. To reason about generality, we then used those values to calculate their average.

When measuring a possible correlation between the percentage of the budget a project uses and whether or not it is in the outcome, we achieve some statistically relevant results for the Method of Equal Shares, see Figure 6.6. The plots show a clear significance for MES results with proportionality probability between 0.018 and 0.5 and a strong negative correlation for those instances. In other words, for most instances where at least half of the projects are likely disproportionately expensive, it holds that the higher the cost the lower the likelihood for the project to be in the Method of Equal Shares outcome of the election. Similarly as above, this makes sense. The PBC statistical value for greedy is much more evened out around 0 when compared with MES. Also like above, and therefore expected, is that the results for the greedy-style voting rule are not defined to be statistically relevant. The difference in p-values is for most instances very large in between the voting rules. This clear difference in statistical relevance can be interpreted to mean that one of the voting rules, here the greedy-style one, is measurably less biased. MES on the other hand shows a clear bias towards proportionally inexpensive projects.

The closest any of the two correlation values of the chosen voting rules ever are, is when the project costs are almost perfectly proportional to their approvals. In this case, the correlation is also positive. This is likely due to the fact that expensive but proportional projects are affordable in both voting rules. In a greedy sense especially

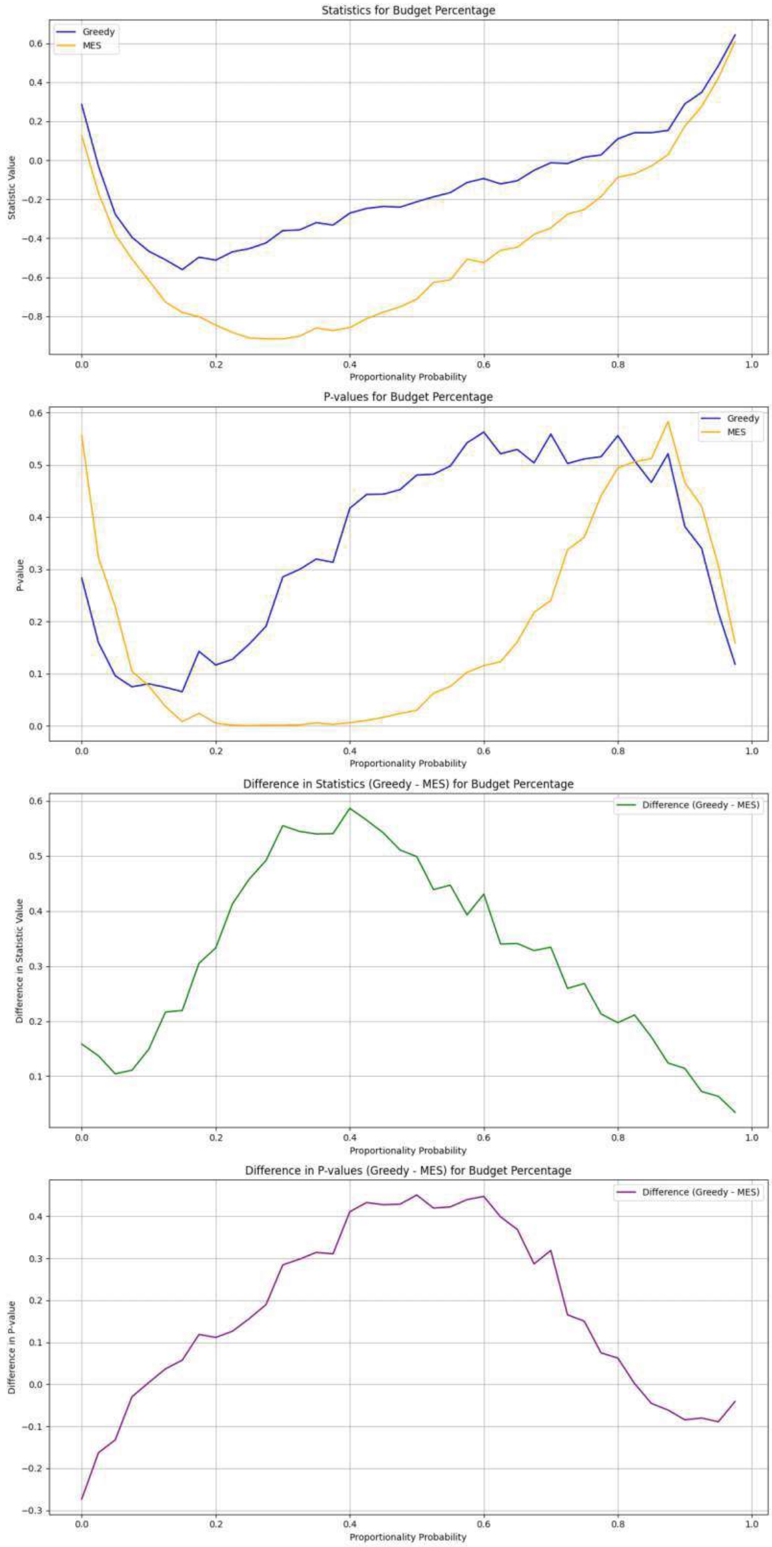

Figure 6.6: Average of correlation coefficient for difference between percentage of budget and outcome, step in proportionality probability −0.025, initially 1

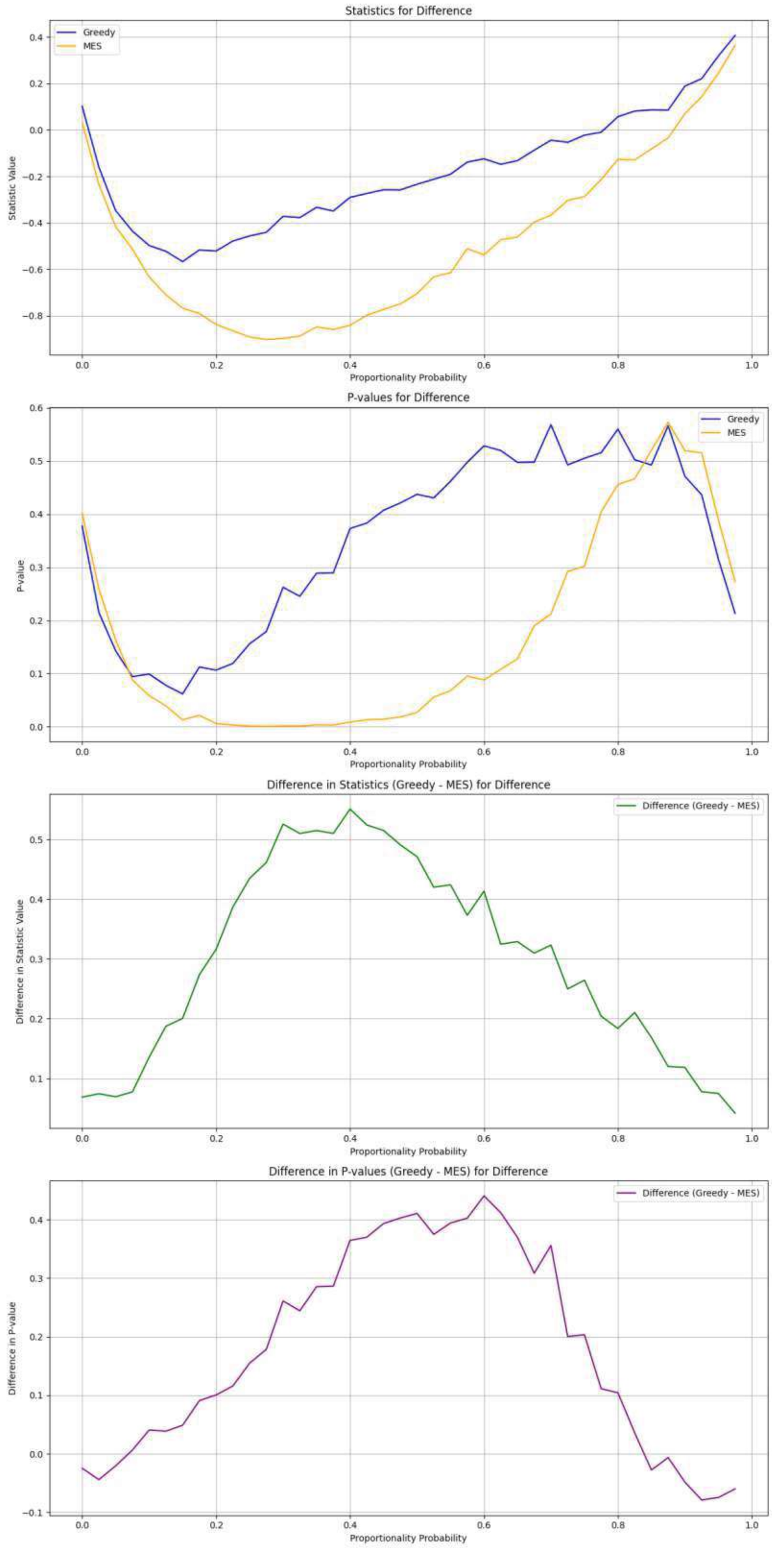

Figure 6.7: Average of correlation coefficient for difference between percentage of budget and percentage of approvals and outcome, step in proportionality probability -0.025, initially 1

it is only reasonable that very popular projects are chosen, as long as the cost is still within the remaining budget. In general, of course, this correlation cannot be said to be statistically relevant, but the trend goes towards relevancy for both the greedy-style voting rule and the Method of Equal Shares.

We also investigated the correlations for our bias formulation using the difference of relative cost and popularity. There are almost no discernible differences in any values when comparing the above correlation metrics to the ones calculated for the difference between the percentage of the budget and the percentage of the votes used and whether or not a project is in the outcome. We have seen in the smaller scale experiments already that any difference is small. When averaging over 100 instances, the reduction of any small difference towards basically 0 is to be expected.

6.3 Further Statistics

We performed some more minor statistics to get a fuller picture. They do not make statements about our chosen bias formulations, but are interesting nonetheless. The used instances are the same for all experiments. We used 40 instances from Pabulib, all having between 500 and 1500 voters and 10 to 50 projects. Note that the artificial instances described above are also within that spectrum. The chosen instances are of high quality, meaning that the meta data is well-kept and easily usable.

6.3.1 Cost Distribution in Real-World Instances

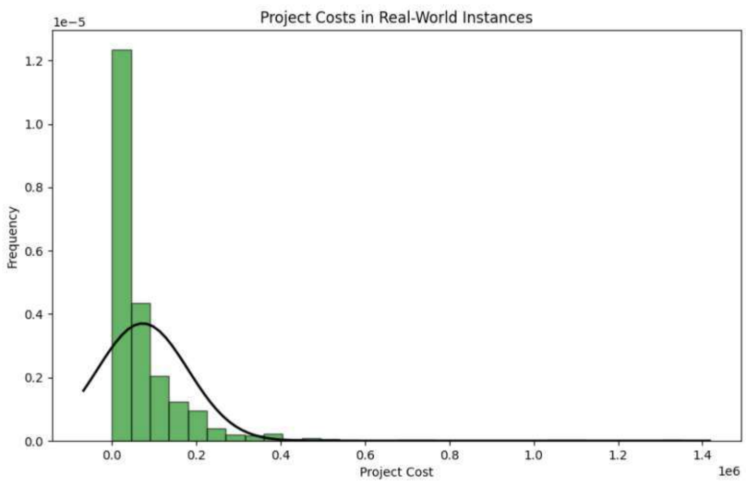

Figure 6.8: Frequency analysis of project costs in 40 real-world instances

As is visible in Figure 6.8, project costs are most often quite low. The average budget limit of all used instances is roughly 792, 000, therefore more costly projects could theoretically

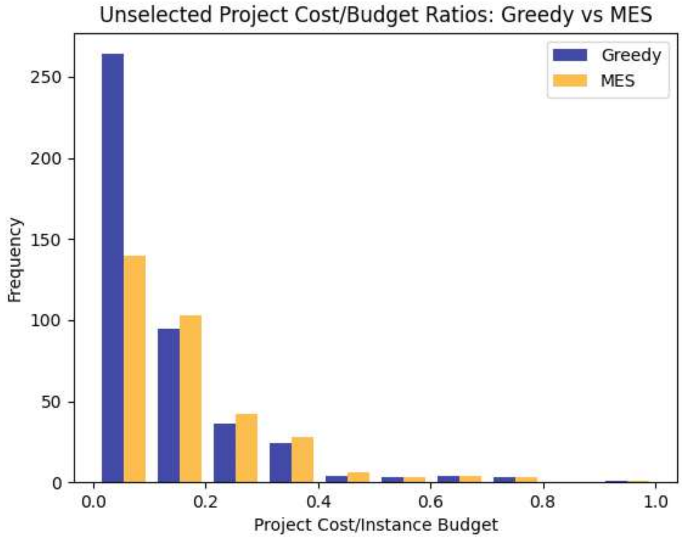

Figure 6.9: Frequency analysis of project cost/budget in unselected projects

be feasible. There is no direct reason as to why projects that use only a little of the budget are so prevalent. A cost limit for projects is not always a rule. It is possible that the reasons lie in political science more so than computer science. It is e.g. conceivable that experience has shown that more small projects are able to garner the support they need to be included in the outcome, making them more attractive to propose.

6.3.2 Cost/Budget Ratios of Unsuccessful Projects

We calculated the utilitarian greedy and MES outcomes (cost satisfaction function, budget increment completion method with increment 1) for all chosen 40 real-world instances. We then compared the ratio of project cost over the instance budget for the projects which were not included in the respective outcomes.

Figure 6.9 shows that the greedy voting rule has much more unselected relatively inexpensive projects than MES. This aligns with our findings in the main experiments, that MES is slightly more biased towards inexpensive projects. For truly expensive projects the greedy rule and MES show no significant difference. MES leaves some slightly more expensive projects out of the outcome, which probably follows from having included more cheaper ones.

When taking a look at the corresponding plot of that same ratio for the successfully chosen projects, all the values align as well, see Figure 6.10. MES includes more inexpensive projects while the greedy voting rule has some more costly projects in its outcome. It follows from the fact that a more costly project reserves more of the outcome that the difference in frequency for more costly projects is less severe than that of inexpensive ones.

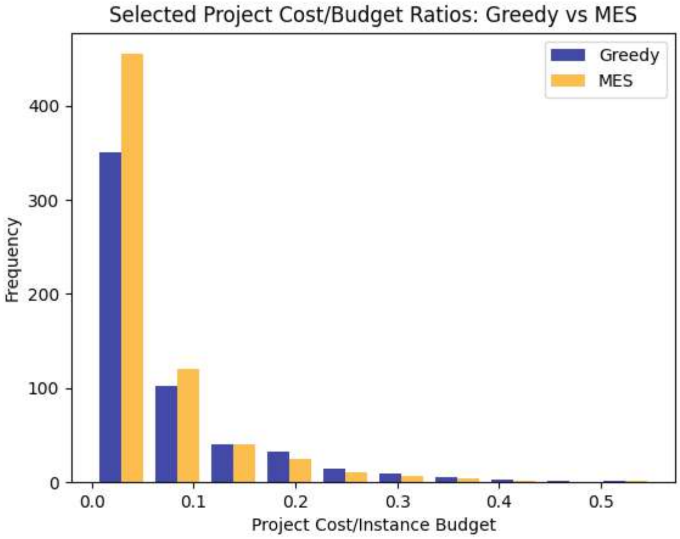

Figure 6.10: Frequency analysis of project cost/budget in selected projects

6.3.3 Popularity of Unsuccessful Projects

We used the same projects and same voting rules to compare how popular unchosen projects were depending on the voting rule. To make popularity comparable, we counted by the ratio of approvals a project got over how many approvals would have been possible (maximal number of approval votes any project could have achieved).

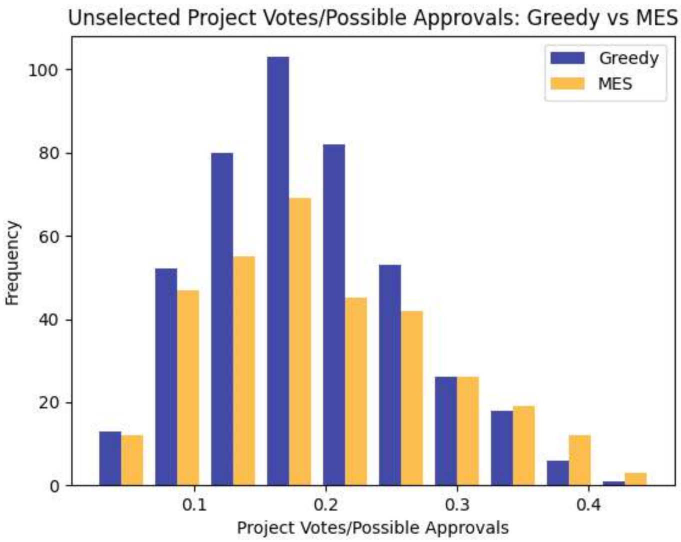

Figure 6.11: Frequency analysis of project votes/possible approvals in unselected projects

The plot of Figure 6.11 shows a clear tendency for the greedy rule to leave projects that can be deemed unpopular out of the outcome more frequently than the Method of Equal Shares does. This seems reasonable, given the definition of the greedy rule. MES focus does not lie on popularity, therefore it is slightly more likely for MES to exclude a fairly popular project in its outcome.

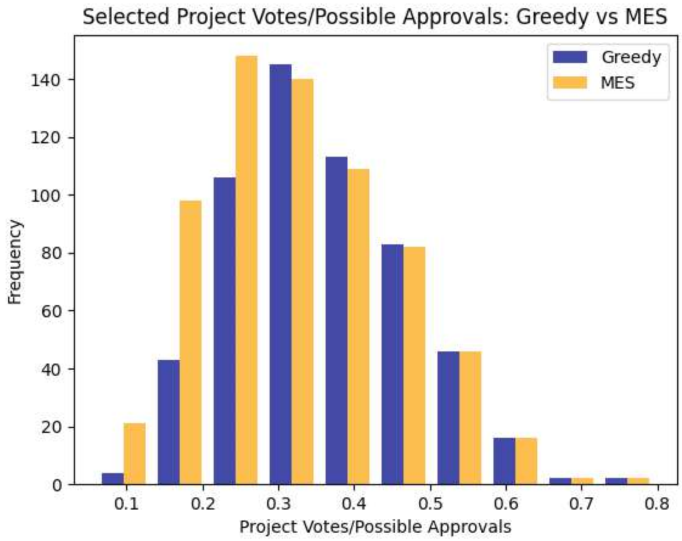

Figure 6.12: Frequency analysis of project votes/possible approvals in selected projects

When looking at the opposite statistic of the amount of selected seemingly unpopular projects in Figure 6.12, the trend that MES chooses more unpopular ones than greedy fits with that. It also fits with the trend observed earlier that MES chooses more lower cost projects. With the idea of proportionality in mind, it is reasonable that small projects get chosen by MES when they have proportionally enough, albeit a small amount of approvals. The greedy voting rule includes some more popular projects. Interestingly, the very popular projects seem to be included equally by both voting rules. This might stem from the phenomena that we have seen where most projects are of proportional cost and where high costs projects are infrequent, leading to a lower likelihood for enough of those projects to make a frequent difference in outcome inclusion.

CHAPTER 7

Discussion

As the first to study bias in PB elections, we were aiming to define a notion of cost bias, meaning whether the relative expense of a project unduly influenced its likelihood of being in the outcome. Furthermore, we wanted to study that bias more closely, in regards to both theoretical and empirical results. We have proven two new theoretical results for different general instances when using greedy-style voting rules. For both our models, the likelihood that the less costly project is being chosen for the election outcome cannot be smaller than that of a more expensive project. These results show that the likelihood of whether or not a project is included in the outcome is not independent from its cost for greedy-style voting rules. Cost bias in participatory budgeting is therefore proven to exist in certain settings and deserving of further study. Our cost bias formulation for experimentation relies on the idea of a project's proportionality, which denotes the difference between the percentage of the budget it requires and the percentage of possible approvals it received.

We have made use of the approval sampling for elections from recent literature [SFJ$^+$22], which among other election types sampled approval profiles for PB, and built upon that to also sample project costs. Our adjustable sampling approach takes the proportionality of real-world project costs into account, ensuring that all manner of cost proportional instances can be produced, realistic or improbable.

Our extensive statistical experiments show that for the Method of Equal Shares, a negative cost bias by our definition is measurable in a statistically relevant way if it holds that the proportionality probability is between 0.1 and 0.5. In other words, if at least half of the projects but not all are disproportionately expensive, the Method of Equal Shares favours projects that are comparatively cheap given their number of approvals. The utilitarian greedy voting rule does not show such a clear bias towards proportionally inexpensive projects, although a small such trend can be observed.

The smaller statistical experiments we performed on the data of real-world elections are in line with these findings. Firstly, most projects are inexpensive and proportional, meaning very large projects with a very large backing are almost non-existent. When comparing the possible expense of unselected projects, the greedy-style voting rule leaves out cheaper projects much more often than the Method of Equal Shares. The other way around is also true: MES is more likely to include low-cost projects, which is exactly the statistical result of our experiments. Furthermore, the greedy-style voting rule more often leaves out projects with little approval when compared to MES. By definition of the greedy-style rule, this is only sensible. MES, which relies heavily on proportionality, chooses these projects with small approval more often. This fits with the usual proportionality of real-world projects and our measured high likelihood of MES electing an inexpensive project.

In summary: Some greedy-style voting rules are shown to be biased against higher cost projects for specific types of instances. Greedy voting rules in real-world instances by definition first choose the projects with most approvals. Such projects are unlikely to use very large parts of the budget, but leave several more inexpensive projects unchosen nevertheless. MES in general is biased towards inexpensive, and consequently usually small projects, measurably so if most projects are disproportionately expensive. Such a bias is in itself neutral. The choice of voting rule for the election therefore depends on whether or not small projects are wanted in the election outcome.

One could argue that because a bias towards smaller projects increases the number of projects in the outcome, voter satisfaction is higher than with fewer larger projects. This could be true in several ways, making the argumentation not unreasonable: With more projects in the outcome, more of a voter's approved projects are likely to be, increasing the voter's satisfaction. Furthermore, if approval groupings are sparse, more projects in the outcome increase the likelihood that a voter sees at least one project realized that they approve of. Since the execution of a project can be time- and work-consuming, however, real-world elections might favour less projects to reduce planning and organizational overhead. Once again, the voting rule decision can only ever be handled in a case-by-case basis. Our findings, however, now enable a more informed decision-making process.

For further research building upon our findings, several questions remain or have opened up: Can more theoretical results for cost bias be shown, not only for greedy-style voting rules but for the Method of Equal Shares? Can we combine the approval and cost sampling techniques even more efficiently to sample PB elections as close to realistic as possible? Can we extend or change the statistical model to also show some statistically relevant cost bias in the utilitarian greedy voting rule? Can we prove or measure in a relevant way that voter satisfaction increases with the inclusion of more relatively small projects in the election outcome?

Overview of Generative AI Tools Used

No generative AI Tools were used in this work.

List of Figures

Bibliography

[ABC+15] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence, January 25-30 (AAAI 2015), Austin, Texas, USA*, pages 784–790. AAAI Press, 2015.

[AS21] Haris Aziz and Nisarg Shah. *Participatory Budgeting: Models and Approaches*, pages 215–236. Springer International Publishing, Cham, 2021.

[BFJ+24] Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Andrzej Kaczmarczyk, Grzegorz Lisowski, Grzegorz Pierczynski, Simon Rey, Dariusz Stolicki, Stanislaw Szufa, and Tomasz Was. Guide to Numerical Experiments on Elections in Computational Social Choice. *CoRR*, abs/2402.11765, 2024.

[BGME23] Luca Bartocci, Giuseppe Grossi, Sara Giovanna Mauro, and Carol Ebdon. The journey of participatory budgeting: A systematic literature review and future research directions. *International Review of Administrative Sciences*, 89(3):757–774, 2023.

[Bos12] Sarah Boslaugh. *Statistics in a Nutshell*. O'Reilly Media, Sebastopol, CA, 2012.

[End17] Ulle Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.

[FCE+16] Brandt Felix, Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, New York, NY, USA, 2016.

[FFP+23] Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. Participatory budgeting: Data, tools and analysis. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI 2023), 19th-25th August 2023, Macao, SAR, China*, pages 2667–2674. ijcai.org, 2023.

[FS21] Marc Fleurbaey and Maurice Salles. *A Brief History of Social Choice and Welfare Theory*, pages 1–16. Springer International Publishing, 2021.

[Jan16] Svante Janson. Phragmén's and Thiele's election methods. *CoRR*, abs/1611.08826, 2016.

[LS23] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer Briefs in Intelligent Systems. Springer, 2023.

[PPS21] Dominik Peters, Grzegorz Pierczynski, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021 (NeurIPS 2021), December 6-14, 2021, virtual*, pages 12726–12737, 2021.

[PS20] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *The 21st ACM Conference on Economics and Computation, Virtual Event (EC '20), Hungary, July 13-17, 2020*, pages 793–794. ACM, 2020.

[RM23] Simon Rey and Jan Maly. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *CoRR*, abs/2303.00621, 2023.

[SFJ+22] Stanislaw Szufa, Piotr Faliszewski, Łukasz Janeczko, Martin Lackner, Arkadii Slinko, Krzysztof Sornat, and Nimrod Talmon. How to sample approval elections? In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI 2022), Vienna, Austria, 23-29 July 2022*, pages 496–502. ijcai.org, 2022.

[Sha07] Anwar Shah, editor. *Participatory Budgeting*. The International Bank for Reconstruction and Development / The World Bank, Washington DC, 2007.

[SHR08] Yves Sintomer, Crsten Herzberg, and Anja Röcke. Participatory Budgeting in Europe: Potentials and Challenges. *International Journal of Urban and Regional Research*, 32(1):164–178, 2008.

[Wam00] B. Wampler. *A Guide to Participatory Budgeting*. International Budget Partnership, 2000.