

Designing Digital Voting Systems for Citizens

Achieving Fairness and Legitimacy in Participatory Budgeting

JOSHUA C. YANG*, Computational Social Science, ETH Zürich, Switzerland

CARINA I. HAUSLADEN, Computational Social Science, ETH Zürich, Switzerland

DOMINIK PETERS, CNRS, LAMSADE, Université Paris Dauphine - PSL, France

EVANGELOS POURNARAS, University of Leeds, UK

REGULA HÄNGGLI FRICKER, University of Fribourg, Switzerland

DIRK HELBING, Computational Social Science, ETH Zürich, Switzerland

Participatory Budgeting (PB) has evolved into a key democratic instrument for resource allocation in cities. Enabled by digital platforms, cities now have the opportunity to let citizens directly propose and vote on urban projects, using different voting input and aggregation rules. However, the choices cities make in terms of the rules of their PB have often not been informed by academic studies on voter behaviour and preferences. Therefore, this work presents the results of behavioural experiments where participants were asked to vote in a fictional PB setting. We identified approaches to designing PB voting that minimise cognitive load and enhance the perceived fairness and legitimacy of the digital process from the citizens' perspective. In our study, participants preferred voting input formats that are more expressive (like rankings and distributing points) over simpler formats (like approval voting). Participants also indicated a desire for the budget to be fairly distributed across city districts and project categories. Participants found the Method of Equal Shares voting rule to be fairer than the conventional Greedy voting rule. These findings offer actionable insights for digital governance, contributing to the development of fairer and more transparent digital systems and collective decision-making processes for citizens.

CCS Concepts: • **Applied computing** → **E-government; Voting / election technologies; Decision support systems; • Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: participatory budgeting, digital democracy, collective decision-making, explainable AI, trust, legitimacy

CONTENTS

1	Introduction	2
2	Related Work	4
3	Formal Model	7
4	Experimental Design	8
5	Results: Voting Input Formats	12
6	Results: Aggregation Methods	18
7	Conclusions and Future Work	21
A	Additional Tables	28
B	Experiment Interface	31

*Corresponding author

Authors' addresses: Joshua C. Yang, joyang@ethz.ch, Computational Social Science, ETH Zürich, Switzerland; Carina I. Hausladen, Computational Social Science, ETH Zürich, Switzerland, carinah@ethz.ch; Dominik Peters, CNRS, LAMSADE, Université Paris Dauphine - PSL, France, dominik.peters@lamsade.dauphine.fr; Evangelos Pournaras, University of Leeds, UK, e.pournaras@leeds.ac.uk; Regula Hänggli Fricker, University of Fribourg, Switzerland, regula.haenggli@unifr.ch; Dirk Helbing, Computational Social Science, ETH Zürich, Switzerland, dhelbing@ethz.ch.

1 INTRODUCTION

Participatory budgeting (PB) is a democratic process that allows citizens to influence government spending in their community. This democratic innovation has emerged as a great success: Hundreds of cities now regularly let their residents vote over how the city government will spend a designated part of its budget (De Vries, Nemec, and Špaček 2022; Wampler, McNulty, and Touchton 2021). Originating in Brazil during the 1980s as a process focused on fostering community deliberation, PB has since gained international momentum (Cabannes 2004; Ryan 2021). In its currently predominant form, PB proceeds by asking residents to submit project proposals, and then passes to a formal *voting stage*, where residents vote for the projects they would like to see implemented. The most popular projects are thereby selected for implementation, subject to budget constraints.

Digitalisation has transformed PB in the recent decades for multiple reasons. Digital platforms can support the scaling of participatory processes and broaden participation, without some of the limitations of purely offline, face-to-face deliberations (Davies et al. 2021). Online processes make it easy for cities to configure the platform for their respective use case and to reduce costs (Palacin et al. 2024). Digital PB, supported by digital means that allow for large-scale participation, also creates space for collective intelligence, which refers to conditions under which a diverse group of people outperform individuals, often even experts (Page 2007).

The transition to a digitalisation of democratic processes, however, is not without challenges. Notably, when relying solely on digital tools, the “digital divide” (Cullen 2001) poses a risk to equal access to democratic deliberation. For example, Lev-On (2022) finds significant differences in the budgeting outcomes reached in physical versus remote settings. Also, a large number of projects to be voted on in a PB can make it challenging for voters to take decisions (Laruelle 2021). Furthermore, algorithmic decisions can be perceived as less fair and trustworthy than human decisions (Lee 2018). Thus, to realise the full potential of collective intelligence, digital democracy systems should be thoughtfully designed from a human-centred perspective, emphasising trust and legitimacy (Helbing et al. 2023).

How do cities currently decide how to set up their digital PB systems? According to a list of PB votes in cities around the world collected by Peters (2022), there is much heterogeneity in the rules cities use, and in particular in their choice of the *voting input format*. For example, many cities use a form of *approval voting*, where voters can vote for several projects (but only once per project), typically up to a certain maximum (e.g., up to 3 projects in Rome 2019 and Toulouse 2022, up to 5 in Montreal 2021 and Cambridge, MA, 2016–2021, up to 6 in Grenoble 2019–2022, up to 10 in Lyon 2022, up to 15 in Warsaw 2020–2024). Other cities give voters a certain number of *points* that they can assign to various projects, and allow them to give several points to the same project (e.g., 5 points in Strasbourg 2019–2021 and Gdańsk 2016–2022, 7 points in Toulouse 2019). Yet other cities allow voters to select a certain number of projects and then *rank* them in order of their preference (e.g., 3 projects in Brest 2022 and Kraków 2019–2021, 5 projects in Gdynia 2016–2019).

However, it is not clear how city officials and legislators decide on these voting input formats and their parameters. There are few academic studies that they can draw on to see which formats are easiest to use, which formats best express voter preferences, and which formats voters prefer to use. The only such study we are aware of is a recent MTurk user study (Fairstein, Benadè, and Gal 2023), but it has two significant limitations: (1) Several of the input formats it studied are popular among voting theorists but not used by cities (e.g., ranking all 20 projects, or distributing 100 points across projects), and therefore the results aren’t directly informative for the choices that city officials make. (2) Fairstein, Benadè, and Gal (2023) used a between-group design, where each participant only voted using a single voting format. As a consequence, they could not ask participants for their preferences among input formats.

We asked 180 participants to vote in a fictional PB program. Participants were drawn from the joint subject pool of students of the two major universities in Zurich (ETH Zurich and the University of Zurich). The fictional PB program involved 24 projects, whose descriptions we adapted from proposals of Zurich’s 2022 PB (“*Stadtidée*”).

Via an online interface, we let participants vote over the same list of projects 6 times in a row, using a different input format each time. We chose to study the following formats: selecting any number of projects, selecting 5 projects, distributing 5 points across projects, distributing 10 points, ranking 5 projects, and distributing 10 points across 5 chosen projects. After the vote, we asked participants how easy they found each format to use (approval and ranking were rated easiest) and how well it captured their preferences (ranking and distributing 10 points were rated best). We also asked participants which input format they would recommend a city to use for their PB, by giving a preference ordering of the input formats that they had just used. Ranking 5 projects was a clear favourite.

The choice of input format is only the first decision that needs to be made to organise a PB vote. The second decision is which voting aggregation rule to use, i.e., the rule that, given the votes, decides which projects are winning (Rey and Maly 2023). At this point, most cities around the world use a simple aggregation method that follows a majority rule to select projects according to the number of votes they accumulated, which is known as the *Greedy rule* in the computer science literature. This rule designates the projects with the highest vote count as winners, until the budget runs out (Talmon and Faliszewski 2019). Recent literature has identified important downsides of this rule (Boehmer et al. 2023; Faliszewski, Flis, et al. 2023; Peters, Pierczyński, and Skowron 2021):

- The rule is *majoritarian*: a plurality of voters can monopolise the entire budget, while the interests of smaller groups may be ignored.
- The rule tends to favour projects in the city centre and to neglect less densely populated areas.
- The rule has low sensitivity to changes in project costs, leading to strategic bundling and cost inflation.

Researchers have developed an alternative voting rule known as the *Method of Equal Shares* (MES) (Peters, Pierczyński, and Skowron 2021) that supports *proportional representation* in PB and thereby creates fair opportunities for small groups and projects to win (Aziz, Lee, and Talmon 2018). MES works by assigning each voter an equal share of the available budget, and then only selects a project as a winner if the supporters of that project together have enough budget to cover the project's costs. This approach strives for every voter to have an equal influence on the voting outcome, whenever possible. In 2023, MES has been implemented for the first time for PBs in the cities of Aarau (Switzerland) and Wieliczka and Świecie (Poland).

While there are well-developed *theoretical* arguments for using MES (in particular, formal fairness guarantees, see, e.g., Brill, Forster, et al. 2023; Peters, Pierczyński, and Skowron 2021), and while the implementations in Switzerland and Poland appear to have been successes,¹ there are no studies on the questions of which voting rule is preferred by voters and whether voters feel that fairness in the sense of proportional representation is a desirable feature of a voting rule for PB.

To answer these questions, in a second part of our study, we presented participants with a fictional outcome of a PB program, involving the same 24 projects they had seen in the first part of the study. Two outcomes were displayed, one obtained by “Method A” and the other by “Method B”. (These outcomes correspond to the outcome of the Greedy rule and the Method of Equal Shares, respectively.) We asked participants how satisfied they were with each outcome and how fair they perceived them to be. We then divided the participants into three groups to see how their views would change in response to different types of *explanations* of the outcomes they saw, mirroring ways in which cities might choose to communicate the voting results.

- The first group was given a detailed explanation of how the two methods work, explaining how Method A selects the projects with the most votes, while Method B equally distributes the available budget.
- The second group was shown graphs with statistical information about the distribution of voter utility. For example, they saw that, for Method A, voters on average approved a higher fraction of the spending. They

¹See, for example, news media coverage (Rzeczpospolita, 2023-04-28, *Czy czeka nas rewolucja w budżetach obywatelskich?* (Are we facing a revolution in participatory budgeting?); ArgoviaToday, 2023-07-05, *Das passiert mit den 50'000 Franken in Aarau* (What happens with the CHF 50,000 in Aarau)), or the analysis of Boehmer et al. (2023).

also saw that, under Method B, a larger share of voters got at least one of the projects which they voted for, funded.

- The third group was shown graphs with statistical information about how the two methods distributed the budget across city districts and project categories (transportation, culture, or nature), and how the relevant percentages compared to the vote shares of the categories and districts.

After going through the explanation, we asked participants again what they thought of the two methods. Their answers indicated that they thought Method B (i.e., MES) provided fairer outcomes.

We hope that our study provides action-guiding information for governments when deciding how to design their PB digital voting platform, how to choose the voting rule, and how to communicate about the voting outcome. We also believe that our results are of interest to the subject of digital democracy more generally, as they underline the value of asking participants for their preferences about the way democratic institutions should be structured.

2 RELATED WORK

In the literature review below, we examine existing research on digital PB, particularly in relation to issues of trust, fairness, and legitimacy, as well as relevant research from Computational Social Choice, with a focus on prior behavioural experiments regarding PB.

2.1 Evolution and Challenges of Digital Participatory Budgeting

PB is a democratic process where community members directly decide how to spend a portion of public budget. It allows citizens to identify, discuss, and prioritise public projects, and gives them the power to make real decisions about how to spend money (Wampler 2000). A significant evolution in PB over the past decade has been the shift towards online participation. As Wampler, McNulty, and Touchton (2021) point out, efforts to include greater participation and overcome the high costs associated with large-scale face-to-face participation are at the core of digital PB programs. In cities like Helsinki, Paris, Madrid, and Barcelona, the adoption of digital platforms enables the collection of a much larger volume of votes and facilitates the use of more sophisticated voting inputs and/or aggregation methods.

However, the digitalisation of PB is not without challenges. A primary concern is the “Digital Divide”, as noted by Cullen (2001), which poses a risk to equal access and participation in PB for all citizens due to the difference in their digital skills. Concerns about trust in online voting, vote manipulation, and vote security are also paramount. These challenges underscore the importance of adopting inclusive strategies and robust security measures in digital PB to support participation and trust in the process.

Note that PB is a form of “multi-candidate election”. In electoral processes that involve voters selecting from more than two options (Merrill 1984), a common challenge is managing “information overload”, which can significantly impact democratic decision-making processes (Roetzel 2019). A connected idea is the “Choice Overload” dilemma, where an excess of options can overwhelm voters, as explored by Iyengar and Lepper (2000) and Laruelle (2021). An illustrative case is Tartu’s PB program, which demonstrated that limiting project choices can enhance voter participation. Eventually, this led to a city-imposed cap of 25 projects (Mærøe et al. 2021).

2.2 Trust, Fairness, and Legitimacy in the Digital Age

As the digital era advances, trust and legitimacy emerge as pivotal elements in the political landscape, especially in the context of decision-making processes and the digital transformation of cities. Helbing et al. (2023) emphasises the role of digital technologies in promoting participatory democracy, stressing the importance of accessibility and fairness to avoid reinforcing existing power imbalances.

Political trust refers to the confidence that stakeholders have in institutions, specifically with regard to acting competently and ethically (Uslaner 2002). *Fairness* as a term comprising many meanings. In the PB context, *distributive fairness* and *proportional fairness* can achieve that resource allocation proportionally reflects collective community preferences, where no subgroup is able to achieve a more favourable outcome by reallocating their budget share (Fain, Goel, and Munagala 2016).

In this study, we pay particular attention to *perceived fairness*. This means a procedure or outcome that is “perceived by individuals or collectives as fair according to previous norms or standards” (Peiró, Martínez-Tur, and Moliner 2014). In this way, our findings can be realistically applied to real-world social settings and connected to the idea of political trust and legitimacy.

Legitimacy is the perceived acceptability of an institution’s actions, often gauged by their adherence to democratic principles, policy outcomes, trust, and procedural fairness (Beetham 1991; Scharpf 1999). *Output legitimacy* reflects citizens’ acceptance of the results of their government’s policies and laws, which faces various challenges in the digital age. For example, the survey by Hartley (2021) reveals a dichotomy: while citizens appear to appreciate the benefits of smart cities, they tend to distrust their privacy and security measures, overall feeling dissatisfied with their participation in policy-making. This underscores the challenge of balancing democratic procedures against policy outcomes (Strebel, Kübler, and Marcinkowski 2019), i.e., of balancing input and output legitimacy (Scharpf 2000). Altogether, this calls for the concept of “throughput legitimacy”, which evaluates the efficiency and quality of the processes that transform inputs into outputs (Schmidt 2013).

In the digital governance context, Hänggli, Pournaras, and Helbing (2021) advocate a human-centred approach, which marries inclusive voting mechanisms with the integration of diverse opinions and the synergy between digital tools and real-life interactions. The human-centric approach does not only aim to bolster legitimacy and trust in governance, but also emphasises the critical need for further research into democratic innovations that enhance citizen participation in the digital age.

Another relevant concept is *Explainable AI* (XAI). As algorithmic decision-making becomes commonplace in people’s everyday lives, XAI has emerged as a crucial approach to address the issue of trust in algorithmic “black boxes”. In discussions on XAI, a fundamental distinction exists between “explanation proper” and “explanation as justification” (Kenny et al. 2021; Sørmo, Cassens, and Aamodt 2005) The “explanation proper” (also known as *pre-hoc explanation* or *model transparency*) focuses on explaining how an AI model works. This type of explanation aims to make the internal mechanisms and decision-making processes of the model inherently understandable. On the other hand, “explanation as justification” (or *post-hoc explanation*) seeks to justify or explain after the fact why the system produced a given answer. This approach does not require the model itself to be transparent, but instead uses various techniques to interpret and explain the model’s outputs.

The distinction between pre-hoc and post-hoc explanations reflects a trade-off between accuracy and comprehensibility in XAI. Pre-hoc explanations prioritise inherent understandability. However, requiring that such explanations are available might limit the complexity and performance of algorithms. Post-hoc explanations focus on making outcomes easily understandable. However, such explanations might be too simple to fully capture the model’s decision-making process. We will apply these XAI approaches to our study of voting in PB, and compare 3 explanation types of the voting process and outcome.

2.3 Computational Social Choice Literature on PB

Researchers in Computational Social Choice study voting from a computational perspective (Brandt et al. 2016). Participatory budgeting is interesting from this point of view, as the selection of a budget-feasible set of winning projects is a combinatorial collective decision problem. In particular, the approach belongs to the class of multi-winner voting problems (Faliszewski, Skowron, et al. 2017), which have received much attention in recent years. For an overview of recent work on participatory budgeting, see the survey by Rey and Maly (2023).

The seminal paper of Benadè, Shah, et al. (2021) compared different voting input formats in the framework of *implicit utilitarian voting*, where the goal was to find a winning set that maximises utilitarian social welfare despite eliciting only partial information about voter preferences. Benadè, Shah, et al. (2021) proved that a certain variant of approval voting (randomised threshold approval voting) allowed for good outcomes on their worst-case measure. They also studied *knapsack voting* (where voters were asked to select a highest-value set of projects that fits in the budget constraint (Goel et al. 2019)), which performed much worse, as well as some input formats based on complete rankings of projects. However, the voting rules implicitly proposed in their work would be difficult to explain and justify to voters. In addition, the input formats studied are different from the ones that are common in practical use.

Most other work in the Computational Social Choice literature focuses on the design and evaluation of different voting rules, i.e., methods of aggregating votes and deciding on sets of winning projects. Researchers have studied the properties of the standard Greedy method and its variants (Baumeister, Boes, and Seeger 2020; Los, Christoff, and Grossi 2022; Talmon and Faliszewski 2019) as well as related rules that maximise social welfare objectives (Fluschnik et al. 2019; Hershkowitz et al. 2021; Laruelle 2021).

An issue with the Greedy method and related approaches is that they are majoritarian and may neglect minority interests, which motivates the search for a PB voting rule that provides *proportional representation* (Aziz, Lee, and Talmon 2018). Peters, Pierczyński, and Skowron (2021) proposed the *Method of Equal Shares* (MES), a simple voting rule based on the idea of dividing the available budget equally between the voters and then electing popular projects when their supporters have sufficient remaining budget to fund the projects. They proved that MES provides a proportional representation guarantee called EJR (“extended justified representation”, Aziz, Brill, et al. 2017). Several other papers established additional formal proportionality properties of MES (Brill, Forster, et al. 2023; Brill and Peters 2023; Los, Christoff, and Grossi 2022). Faliszewski, Flis, et al. (2023) simulated MES on datasets from real-world PB and found that it produces fairer outcomes than the Greedy method. MES has since been implemented in PB voting in Aarau, Wieliczka, and Świecie (see equalshares.net/elections).

2.4 Prior Experimental Work

The work on voting in PB we mentioned above has only used theoretical and simulation approaches. Overall, there has been little work testing these ideas with human subjects. Below, we discuss three relevant studies.

In an unpublished study with 1,200 subjects from Amazon Mechanical Turk, Benadè, Itzhak, et al. (2018) evaluated the efficiency and usability of different input formats in PB. They found that k -approval voting imposes a low cognitive burden on voters, although the participants themselves did not always perceive this to be the case. The study was framed as the participants being stranded on a desert island – what items they would wish to have in such a situation. Therefore, the applicability to PB in cities is unclear. In addition, some of the input formats studied were far removed from those used in practice.

Fairstein, Benadè, and Gal (2023) recruited more than 1,800 subjects from Amazon Mechanical Turk. In a similar setup to our study, they had each participant enter a vote in a fictional PB. Some participants saw a PB with 10 projects, others saw 20 projects. Each participant voted only once, using one of six input formats: Select up to 5 projects, select projects up to the budget limit (“knapsack vote”), rank all projects, rank all projects in order of “value for money”, distribute 100 points between projects, and select all projects you would give more than 10 points if you were to divide 100 points (“threshold approval”). The voting interface displayed a category symbol for each project, as well as a fictional city map showing project locations. The authors measured the time taken by participants, and asked participants to answer the following questions on a 1–5 scale: “How easy did you find the voting task?” and “How well did the input format capture your preferences?”. They found that selecting up to 5 projects was the input format that took the least time. It also scored highest on both rating questions. Participants further rated the points input format highly, even though it took more time to use. Our

study differs from this by letting the same participants vote several times, using different input formats, so that participants can compare them. We also used different input formats that are closer to the ones in practical use.

The two studies discussed above focused on the choice of the input format. Regarding the choice of the voting rule (i.e., how to aggregate the votes into a set of winning projects), in an unpublished paper, Rosenfeld and Talmon (2021) present a study with 215 subjects aimed at identifying what non-experts deem as “fair” or “desirable” outcomes in PB programs. In particular, they showed participants a small vote matrix of how 5 voters assigned utilities to 5 projects, and then asked which set of projects they would choose to fund, if placed in the position of social planner. They also asked participants to express their opinions on statements such as “Every voter should get at least one of her approved projects funded (if possible)” or “The most approved projects should be funded”. However, this study did not feature the two voting rules that we focus on (Greedy and MES).

Finally, note that a domain related to participatory budgeting is that of *threshold public goods*, where a cost threshold must be met for community-wide benefits to be realised. These have been a focus of studies in behavioural economics (Corazzini, Cotton, and Valbonesi 2015; Croson and Marks 2000).

3 FORMAL MODEL

To understand the setup of our study, we need to formally define the two voting rules we are interested in – Greedy and MES – and how they interact with different voting input formats. We are working in a setting where we have discrete projects that can either be fully funded or not funded at all, and we are subject to a budget bound. This is the model of *bounded discrete PB* (Aziz and Shah 2021). Let $P = \{p_1, \dots, p_m\}$ be the set of available projects, and let B be the budget limit. Each project $p \in P$ has a cost $c_p > 0$. A set $W \subseteq P$ of projects is *feasible* if $\sum_{p \in W} c_p \leq B$. Our goal is to select a feasible set of winning (i.e., funded) projects based on *votes*.

Let $N = \{1, \dots, n\}$ be the set of voters. These voters submit their preferences using a *voting input format*, and we will translate their submissions into an *additive utility function*, so that each voter $i \in N$ has their preferences represented by a function $u_i : P \rightarrow \mathbb{R}_{\geq 0}$. Here, $u_i(p)$ reflects the (imputed) utility that voter i obtains from project p being selected. Then $u_i(W) := \sum_{p \in W} u_i(p)$ is the (imputed) overall utility that voter i enjoys when W is the set of winning projects.

There are two main schemes for translating voting inputs into utility functions that have been discussed in the literature, known as *cost utilities* and *cardinality utilities* (Brill, Forster, et al. 2023; Rey and Maly 2023). Intuitively, when using approval voting, under cost utilities a voter’s satisfaction is measured as the *total amount of funding* that went to approved projects previously voted for, while under cardinality utilities it is the *number* of approved projects that were selected for funding. Formally, the cost utility translation scheme defines utilities as follows, for the contexts of approval voting and of point voting, respectively:

$$u_i(p) = \begin{cases} c_p & \text{if } i \text{ approves } p \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad u_i(p) = \begin{cases} c_p \cdot s_i(p) & \text{if } i \text{ assigns } s_i(p) \text{ points to } p \\ 0 & \text{if } i \text{ assigns no points to } p \end{cases}$$

The cardinality utility translation scheme, in contrast, defines utilities as follows:

$$u_i(p) = \begin{cases} 1 & \text{if } i \text{ approves } p \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad u_i(p) = \begin{cases} s_i(p) & \text{if } i \text{ assigns } s_i(p) \text{ points to } p \\ 0 & \text{if } i \text{ assigns no points to } p \end{cases}$$

Note that we treat ranking input as a type of point voting, by converting ranks into points. Specifically, in one of our input formats (“S5R”), we ask participants to rank 5 projects, telling them before this means that 5 points will be assigned to their top choice, 4 points to their second choice, and so on.

We can now formally define the two voting rules that we study in this paper.

Greedy Rule. The Greedy rule sorts the projects in order of decreasing ratio of total utility to cost, i.e., $\sum_{i \in N} u_i(p)/c_p$. It starts with an empty set $W = \emptyset$ of winners, and goes through the sorted list project by project, starting with the highest-valued project. For each project p , it checks whether it can be added to W without violating the budget constraints, i.e., if $\sum_{p' \in W} c_{p'} + c_p \leq B$. If so, it adds p to W .

Method of Equal Shares (MES). MES assigns each voter $i \in N$ an equal budget $b_i \geq 0$ (the value of which will be determined later). For $q \geq 0$, we say that a project is q -affordable if its cost can be funded by voters paying at most q monetary units per unit of utility, i.e., if

$$c_p = \sum_{i \in N} \min(q \cdot u_i(p), b_i).$$

The rule starts with an empty set $W = \emptyset$ of winners. It then checks if any remaining project is q -affordable for any q . If no project is affordable, the rule finishes and returns W . Otherwise, it selects a project p that is q -affordable with minimum q and adds it to W . It also reduces voters' budgets to account for the cost of p by setting $b_i \leftarrow b_i - \min(q \cdot u_i(p), b_i)$ for each $i \in N$. The value of the starting budget is determined by starting with a starting budget of B/n , running the rule, and repeatedly increasing the starting budget by 1 unit if this is possible, i.e., if the cost of the selected set W does not overshoot the budget limit. This procedure of budget increase is known as "Add1" in the literature (Faliszewski, Flis, et al. 2023).

In our formal setup, the rule used by most cities running PB – select the projects that received the highest total number of points – is equivalent to the Greedy rule with cost utilities. (Note that cost utilities multiply the number of points by the project cost, but then the definition of the Greedy rule divides that value by the project cost.) Recently, some cities have started using MES with cost utilities. Because the use of cost utilities is standard in implementations, when we refer to Greedy and MES, we implicitly mean their cost-utility versions. We refer to the cardinality-utility versions as *Economical Greedy* and *Economical MES* because these rules proceed in order of vote count divided by project cost.

While studying voting input formats, we will explore the effect of adopting a format on the output according to each of the four aggregation methods (Greedy and MES, with cost utilities and with cardinality utilities). While studying participants' rating of outcomes and how it changes when presenting different explanations, we will focus on the outcomes obtained by the cost-utility variants of Greedy and MES (since these are the standard implementations).

4 EXPERIMENTAL DESIGN

We performed a study in which participants could vote in several ways in a fictional PB program and express their views and preferences. The experiment was pre-registered on AEA ACT Registry,² approved by the ETH Zurich Ethics Commission (approval number: EK 2022-N-143) and conducted in March 2023 via the online platform Qualtrics in English language. All participants provided informed consent.

A total of 180 participants were sourced from the joint subject pool of the two major universities in Zurich: ETH Zurich and the University of Zurich. To be eligible for participation, individuals had to be registered students at one of these universities. We could therefore assume that participants were overall familiar with the city of Zurich. (In response to the question "to what extent do you feel connected with the city of Zürich?", 108 participants said "a lot" or "a great deal", 52 said "a moderate amount", 17 said "a little", 3 said "not at all"). Our general impression was that the participants were very engaged during the study. Many participants volunteered detailed thoughts in free text fields, where we invited them to explain their responses.

The study started with several screens giving a high-level explanation of what participatory budgeting is about and how it works. When asked if participants knew about PB before today, 16 answered "yes" and 164 answered

²<https://www.socialscienceregistry.org/trials/11021>

“no”. We then familiarised participants with the idea of voting in PB by going through a brief tutorial. In this test run, the participants were situated in a hypothetical city with 3 different districts and exposed to the concept of different voting inputs and voting aggregation methods.

The main task of the participants was to collectively decide on allocating 60,000 CHF to various projects within the city of Zurich, as part of a fictional PB program. Participants were made aware that their payoff in the experiment would be linked to their gains in the simulated PB, thus providing a tangible incentive for their decision-making within the study. The description was: “*On top of the show-up fee, you can win an extra 0-20 CHF in this democratic game. The higher the amount of budget you can win for your selection of projects in the collective outcome determined by all 180-200 participants, the more you will be paid in this experiment.*” Note that this induces cost utilities, though we did not observe systematically higher vote counts for the more expensive projects. With the variations in payoffs (CHF 20), we chose this approach to encourage the participants to take the task seriously.

To mirror a realistic PB setting, where voters have real preferences for urban projects, participants were asked to identify the city district in Zurich that they “feel connected to the most” (Nord [North], Ost [East], Süd [South], West), and to indicate how much they care about different categories of urban projects (Transportation and traffic organisation; Culture; Nature and urban green areas), by dividing 100 points between the three categories according to their intensity of preference. (For simplicity of analysis, we identified participants with the category they assigned the highest rating to.) This initial part was designed to situate the voters in the city and prompt them to think about realistic urban concerns in order to simulate a real PB setting.

Participants were then shown a list of 24 projects and asked to think about their preferences. The 24 projects are shown in [Figure 11](#). The majority of project descriptions was adapted from the PB “Stadtidee” that Zurich ran in 2022, with some additional projects created to have an equal number of projects in all 4 districts, all 3 categories, and all 2 different cost levels. The projects were presented in a table, showing a project’s ID, its name, its district, its category, and its cost. For simplicity, each project was assigned a cost of either CHF 5,000 or CHF 10,000. Projects were displayed sorted by district and then by category. For each district-category combination, there were exactly 2 projects (a cheap and an expensive one), so that overall there were 6 projects per districts and 8 projects per category. Projects were shown in the same order across all screens to reduce participant workload.

The experiment was designed to answer two research questions, regarding the choice of voting input format and the choice of aggregation rule, respectively.

4.1 Research Question 1: What are voters’ preferred voting input formats, and how does the choice of voting input formats impact the collective outcome?

Once participants had reviewed and understood the 24 project descriptions on a preview page, participants proceeded to cast their votes, using six voting input formats for PB. We chose these formats for being similar to those used by cities in practice. [Figure 10](#) in the Appendix shows screenshots of the different interfaces.

- SN: *Select any number of projects*. This is the standard approval voting (Brams and Fishburn 1978), without any constraints imposed, except that we required the selection of at least one project. The prompt was: “You can select all the projects you approve of. Your vote will be perfectly valid no matter if you vote for many or for few projects.”
- S5: *Select 5 projects*. This is 5-approval voting, mirroring a commonly used input format. We decided to require participants to select exactly 5 projects, to make responses better comparable.
- D5: *Distribute 5 points*. This screen displayed a slider next to each project, allowing participants to select a value between 0 and 5 for each project (with sliders starting at 0). The system enforced the chosen values to sum up to 5. The prompt was: “In this vote, you distribute 5 points among the projects you like. You can concentrate your points on one project, or spread them across a few projects.”

- D10: *Distribute 10 points*. Using the same setup as for D5, but with sliders allowing values between 0 and 10, enforcing a total of 10 points assigned.
- S5R: *Select 5 and rank*. This screen had a drag-and-drop interface allowing participants to rank up to 5 projects. The prompt was: “In this vote, you first drag the top 5 projects that you like to the box, then rank them in order. The top project will get 5 points; the second project, 4 points; the third, 3 points; and so on.”
- S5D10: *Distribute 10 points across 5 projects*. This screen showed the participants only the 5 projects they had ranked in S5R and asked them to distribute 10 points across them using the same interface as for D10. We re-used the selection of 5 projects from the prior step to reduce participant workload.

Note that all participants used the input formats in the same order, so that their experience followed a logical sequence to avoid confusion. However, when interpreting results, we need to keep in mind that participants may become more used to expressing their preferences for later input formats.

After submitting their votes, participants evaluated each of the input formats with regard to several criteria.

- *Easiness*: “How difficult do you find using each of the voting input methods?” (5-point scale from “extremely easy” to “extremely difficult”).
- *Expressiveness*: “How well does each voting input method capture your actual preferences for these projects?” (5-point scale from “does not describe my preference” to “clearly describes my preference”).
- *Recommendation*: “What voting input method would you recommend a city to use in a Participatory Budgeting program?” (by ranking the six input formats in the order of preference).

We then asked participants to explain the reasons for their evaluations, in particular the reasons why they chose to rank an input format in the first or last place with regard to the Recommendation question.

Finally, we asked participants to report how important various project characteristics were to them when deciding which projects to vote for. The characteristics were “the district/location”, “the categories/topics”, “the cost of the project”, and “how likely they can win”. For each characteristic, we provided a five-point scale from “not important at all” to “extremely important”.

4.2 Research Question 2: How do voter perceptions of fairness and trustworthiness vary across aggregation methods such as Greedy and MES, and what effect do different explanations have on these perceptions?

In order to investigate voter perception of different voting outcomes, participants were shown simulated outcomes which were calculated using the two aggregation methods Greedy and MES. These outcomes concerned the same PB instance with 24 projects that participants were familiar with after their votes.

We generated two simulated voting instances using a random procedure, each based on the votes of 200 fictitious voting agents using the S5 and D5 voting input format, respectively.

- Instance 1 was obtained by having agents select a set of 5 projects uniformly at random from all 5-subsets of the 24 available projects.
- Instance 2 was designed to be “polarised”, with agents preferring certain projects. For each agent, we assigned them randomly to one of the 4 districts and randomly to one of the 3 categories (independently, in each case uniformly at random). We then identified the 2 projects matching the agent’s assigned district-category preference, and gave them a weight of 6, with the remaining 22 projects getting a weight of 1. We then randomly sampled a project 5 times (i.i.d., with probabilities proportional to the weights, with replacement), giving a point to a project when sampled. Note that the weights translated to an expected 1.76 out of 5 points assigned to focus projects.

For each instance, we then converted the resulting votes in D5 format to utilities, using the cost utility translation scheme, and we used the Greedy rule and MES to compute the sets of winning outcomes. As expected, the two

rules produced relatively similar outcomes in Instance 1, but quite distinct outcomes in Instance 2, due to its polarised nature (see [Figures 11](#) and [12](#)). We divided participants into two groups, each seeing one of the two instances in this part of the experiment. Participants were shown a result table that displayed the outcomes, labelled A and B (corresponding to Greedy and MES, respectively, but participants were not told these names). This table also highlighted the districts and categories of the projects, allowing participants a comparative view. These outcome tables are shown in [Figures 11](#) and [12](#) in the Appendix.

We instructed participants that these results were simulated (and therefore in particular not based on their earlier votes). We then asked participants to rate the outcomes:

- “Considering how you voted, how satisfied are you with these two different outcomes?” (5-point scale from “extremely dissatisfied” to “extremely satisfied”),
- “How fair do you perceive the two outcomes?” (5-point scale from “extremely unfair” to “extremely fair”),
- “Say outcome A and B were produced using aggregation method A and B, respectively. Based on your intuition, how trustworthy do you perceive these two aggregation methods?” (5-point scale from “strongly disagree” to “strongly agree”).

We also asked participants to explain the reasons for their answers.

One might be hesitant to adopt voting rules such as MES because they are harder to explain to voters than the straightforward Greedy rule. Thus, we explored how participants would react to different ways of explaining the voting outcomes. We divided participants into three groups, and each group was shown a different type of explanation. Inspired by different approaches in XAI, we designed three explanations, which consisted either of a description of the methods used to compute the winning sets of projects, or of statistics illustrating the resulting voting outcomes.

- *Mechanism explanation*: Connected to the XAI idea of model transparency, participants in this group received a walk-through of how the two voting rules work, i.e., of the mechanism by which the Greedy and MES outcomes were computed. For each method, we displayed an about 100-word description defining the method, together with a simple infographic (for Method A, showing a list of projects sorted by vote count; for Method B, showing a diagram explaining how the budget flows from city to voters and from voters to projects, similar to a graphic used on the website equalshares.net).
- *Individual utility distribution explanation*: In this group, we showed participants a graph of the distribution of cost utilities produced by the Greedy and MES outcomes for the 200 simulated voters. Adapting the notion of post-hoc explanation in XAI, the graph showed monetary amounts on the x -axis, with each of the two outcomes inducing a line that, for each monetary amount, indicated the number of simulated voters who approved winning projects, whose cost summed up to that amount. Inspecting these graphs, it was clear that Outcome A left a larger number of voters with utility 0, while also giving some voters very high utility; in contrast, Outcome B reduced the number of voters at both extremes and ensured that most voters had a moderate amount of utility. This reflects the behaviour of the two voting rules when simulated on real PB data (Faliszewski, Flis, et al. 2023).
- *Category and district explanation*: Participants in this group were shown how the two outcomes divide the budget between projects from different districts and different categories. For districts, they were shown three pie charts in a row, the first one showing the division of the population of 200 simulated voters between the 4 districts; the second one showing the proportions of budget spent by Outcome A in the districts, and the third one showing this for Outcome B. For categories, they were shown a similar row of three pie charts. Consistent with findings from simulations on real PB data (Faliszewski, Flis, et al. 2023), Outcome B reflected the proportions of the underlying population distribution better than Outcome A.

After these explanation treatments (referred to as “Mechanism”, “Individual”, and “Group” treatment, respectively, in [Section 6.2](#)), participants were asked how well they understood Method A and Method B. We then

Table 1. Number of projects selected from 24 projects by voters across different voting input formats (for point distribution formats, the number of projects receiving 1 or more points). The table displays the mode (along with its frequency), the mean, median, and standard deviation for each voting input format based on 180 responses.

Acronym	Input Format	Mode (Freq)	Mean	Median	Std. Dev.
SN	Select any number of projects	8 (20)	10.60	10	4.73
S5	Select 5 projects	5 (180)	5.00	5	0.00
D5	Distribute 5 points	4 (73)	3.51	4	1.02
D10	Distribute 10 points	5 (49)	5.24	5	1.85
S5R	Select 5 and rank	5 (179)	4.98	5	0.30
S5D10	Distribute 10 points for 5 projects	5 (146)	4.67	5	0.81

Table 2. Example of a typical voting pattern in our lab experiment (participant ID: 1PW9NBML). In columns SN and S5, the ticks represent approval. For columns D5, D10, and S5D10, the numbers indicate the points allocated to the respective projects. In the S5R column, the numbers signify the order of ranks, with #1 being the highest rank. The participant’s choices align with the modal number of selections observed in each of the six voting input formats.

Project	SN	S5	D5	D10	S5R	S5D10
17 Bike Lanes on Seefeldstrasse	✓	✓	2 pts	4 pts	#1	4 pts
16 Multicultural Festival at Sechseläutenplatz	✓	✓	1 pts	3 pts	#2	3 pts
6 More Night Buses to Oerlikon	✓	✓	1 pts	1 pts	#3	1 pts
7 Free Open Badi Space in Wollishofen	✓	✓	1 pts	1 pts	#4	1 pts
14 More Trees in Bellevue Sechseläutenplatz	✓	✓		1 pts	#5	1 pts
22 Sustainable Cooking Workshop with Kids	✓					
12 Car Sharing System for Young People	✓					
5 Safe Bike Paths around Irchel Park	✓					

prompted them once more to evaluate the methods in terms of fairness and trustworthiness. See [Figure 13](#), [Figure 14](#), and [Figure 15](#) in the Appendix for screenshots of these displays.

5 RESULTS: VOTING INPUT FORMATS

5.1 Individual Voting Patterns

We begin by comparing how voters voted in the six voting input formats. [Table 1](#) shows how many projects they voted for, in the sense of selecting a project (SN, S5, S5R) or giving 1 or more points to a project (D5, D10, S5D10). The unconstrained SN format results in an average selection of 10.6 projects with a considerable standard deviation of 4.73. For S5, the interface enforced a selection of exactly 5 projects. For S5R, it was possible to select fewer than 5 projects, but all but 1 participant followed the written instruction to select (and rank) exactly 5 projects. For D5, voters on average distributed their points over 3.51 distinct projects, with a mode and median of 4 projects (i.e., a plurality of voters gave 2 points to one project and 1 point each to three projects). The D10 format similarly showed a preference for spreading points, but less evenly, with an average distribution over 5.24 projects with 5 being the mode. For S5D10, a quarter of voters gave 0 points to at least one of the selected 5 projects.

[Table 2](#) offers an example of what a typical vote looks like, by showing the answers of a particular participant on all six formats. This participant selected projects exactly in line with the mode number of selections from

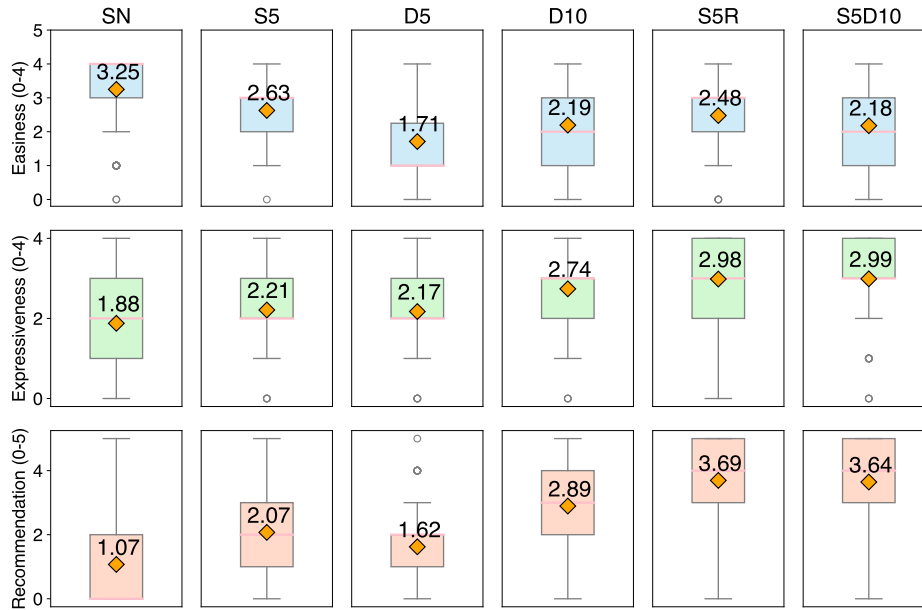


Fig. 1. Overview of voter perceptions of different voting formats. The top row shows how voters ranked the formats (Recommendation), quantified via the rank assigned to each format (Borda’s Method) with higher values indicating greater recommendation. The middle and bottom rows display the Expressiveness and Easiness, respectively, derived from 5-point Likert scale responses. Box plots indicate the inter-quartile ranges, with medians highlighted in pink and means represented by orange diamonds.

Table 1 for each input format. For instance, under D10, the participant allocated points primarily to presumably more prioritised projects: “Bike Lanes on Seefeldstrasse” received 4 points. In the ranking format S5R, we observe a clear hierarchy in their preferences. Notably, some projects like “Sustainable Cooking Workshop with Kids” and “Car Sharing System for Young People” were only selected when there were fewer constraints on the selection process.

Participants generally voted consistently across input formats. For S5, 148 participants (82%) chose either a superset or a subset of projects they selected for SN. 145 participants (81%) assigned each project at least as many points under D10 as under D5. 73 participants (41%) selected the same 5 projects in S5 as in S5R/S5D10, and for 139 participants (77%) the choice set differed in at most 1 project. For 119 participants (66%), whenever they ranked a project above another in S5R, they gave it at least as many points in S5D10.

5.2 Perception of Voting Input Formats

After the voting process for all six input formats, we asked participants to assess each input format on 5-point Likert scales in terms of Expressiveness and Easiness. The results of a Mann-Whitney U test are presented in Table 4 and Table 5 of the Appendix. For Expressiveness, S5D10, S5R, and D10 received high ratings, suggesting that participants felt formats involving a ranking or a distribution of points were representing their preferences better. For Easiness, SN was viewed as easiest, with a mean rating of 3.25, while D5 was considered the most challenging, with a mean rating of 1.71.

Table 3. Participants used an open-answer text field to explain their rationale for ranking an input format first or last. We categorised these responses into six themes and counted the number of responses for each theme.

	SN	S5	D5	D10	S5R	S5D10		SN	S5	D5	D10	S5R	S5D10
Clarity	4	4	1	11	33	33	Lack of Clarity	105	3	9	1	0	1
Granularity	1	0	0	14	14	34	Not Granular	3	6	33	1	0	2
Narrow Choices	0	6	1	1	19	27	Overwhelming	38	3	3	0	1	4
Flexibility	4	1	0	10	4	7	Lack of Specificity	19	2	1	2	2	1
Ease	3	6	0	3	23	8	Complexity	2	4	2	3	2	1
Visual Repr.	1	1	0	0	6	3	Randomness	9	0	0	0	0	0

(a) Reasons for ranking an input format *first*.

(b) Reasons for ranking an input format *last*.

We then asked participants to rank the formats based on their Recommendations for adoption in a city’s PB program. For analysis, we translated the rankings into points using Borda’s Method. (1st rank: 5 points, 2nd rank: 4 points, and so forth) (Brams and Fishburn 2002). As shown in Figure 1, S5D10 and S5R were ranked highest, obtaining a similar number of points (see also Table 6 in the Appendix). Both formats significantly outperformed D10, the third-ranked format (Wilcoxon signed-rank test $W = 4613$, $p = 2.8 \times 10^{-7}$). SN was ranked as least recommended ($W = 5974$, $p = 0.0017$, see Table 7 in the Appendix).³

We performed an OLS regression to see how Expressiveness and Easiness ratings predict the Recommendation rank. (Table 8 in the Appendix). It turns out that Expressiveness is a strong positive predictor of Recommendation with a coefficient of 0.89, while Easiness has a statistically insignificant effect with a coefficient of 0.006. This analysis highlights that Expressiveness, not Easiness of use, appears to be the key factor influencing whether participants recommend an input format.

Reasons for Recommendations of Voting Input Formats. To gain a deeper understanding of the reasons for choosing the ranking of input formats, we analysed the open-ended responses from participants, who were asked to write up their reasons for ranking a format the highest or the lowest.

We did a thematic evaluation by coding frequent responses and, after this, grouping similar codes into themes. From all the responses obtained, we identified six prevalent themes that explained the reasons for ranking a voting input either first or last. Table 3 breaks down how frequently each type of reason was given for each input format. For the most preferred input format, the main reasons given were clarity (the ability to clearly express preferences), granularity (differentiating degrees of preference), a narrowed choice set (thereby preventing overload), flexibility, ease of use, and clear preference visualisation. The main reasons given for the least preferred input format were a lack of clarity (inability to prioritise), a lack of granularity (such as too few points for preferences), feeling overwhelmed, lack of specificity, complexity of use, as well as randomness (concern of non-deliberate voting). The detailed explanations of these themes are listed in Table 9 of the Appendix.

S5R and S5D10 were ranked highest most often, with participants justifying this due to these formats’ clarity and granularity. SN was ranked lowest most often, with participants criticising a perceived lack of clarity. The results suggest that participants desire the ability to express the granularity of their preferences, but they also seek clear instructions on how to convey these nuances. Too much freedom in expressing preferences can be perceived as overwhelming.

³To reduce participant workload and ensure a logical progression, all participants used the input formats in the same order. Thus, one might expect that input formats encountered later would be rated higher due to greater familiarity with the voting process. However, participants actually did not rate later formats as easier to use than earlier formats.




Fig. 2. Box plots illustrating the percent of votes that participants assigned to their self-identified preferred district and project category. The medians are highlighted in red.

5.3 Drivers of Voting Behaviour Across Different Voting Input Formats

Before the voting, we asked voters to identify the district and project category they most identified with. To what extent do these preferences drive their votes? Figure 2 shows what fraction of points participants placed on projects from their identified district and from their preferred category, across various input formats. To compute these fractions, we assume SN and S5 to assign 1 point to each selected project, and S5R to assign points according to Borda’s Method. The detailed mean, median and standard deviation statistics are shown in Table 10 of the Appendix. For each input format, we note a similar or equal median of the fraction given to district and category projects. Since there were 4 districts and 3 categories, i.e., more projects for each category than for each district, this suggests a stronger emphasis on locations than on topic preferences among voters. Looking at the input formats individually, approval voting formats such as SN and S5 appear to induce less focus on projects with favoured characteristics. In contrast, voting input formats based on point distribution display a higher median voting percentage on their own district or category, all 50% and above. This concentration of votes is particularly pronounced for D5, with 60% of the votes allocated to their own districts and category. This may be due to the limitation of having only 5 points in D5, where 33 out of 180 voters also reported to have “insufficient points to capture preferences accurately” from the previous thematic analysis (Table 3b). In D10, where voters had more points to distribute, the concentration was reduced.

In summary, approval voting tends to encourage voters to disperse their votes beyond their top self-reported interests. Conversely, point-distribution and ranking systems apparently encourage voters to allocate more weight to their top self-reported interests.

We also asked participants to self-report which project characteristics influenced their project selections, using a five-point scale from “not at all important” (assigned a value of 1) to “extremely important” (value 5). Project category and location were assigned high ratings on average (4.13 and 3.29), while cost and a project’s likelihood of winning were seen as much less important (with mean ratings of 2.04 and 1.99).

5.4 Differences in Voting Outcomes Using Different Voting Input Formats

Does the choice of the voting input format affect the outcome of the vote? This is an important question to answer, to make an informed decision about input formats. We study this question from two perspectives: first,




Fig. 3. Voting distribution across projects based on six distinct input formats (SN, S5, D5, D10, S5R, S5D10). The projects are ordered by their mean popularity across the 6 inputs.

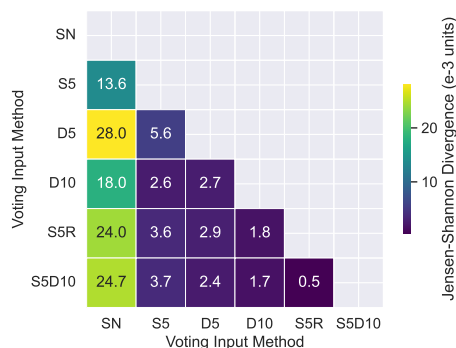


Fig. 4. Heat map of Jensen-Shannon divergence values comparing the distribution of votes between different voting input formats. The higher the value, the larger the difference between the distributions.

we consider the summed number of votes or points assigned to the projects and how they differ between formats; second, we compute the outcomes of the Greedy and MES aggregation rules and compare them.

Distribution of Votes and Points. For each of the six voting input formats, we compute the total number of votes or points assigned to each project, by summing across all 180 votes. For the approval formats (SN, S5), projects receive 1 point for each approval. For the point distribution formats (D5, D10, S5D10), projects receive the number of assigned points. For ranking (S5R), we use Borda’s Method as described in Section 4.1. Then, we normalise the total points of the projects by dividing by the sum, to obtain a point *distribution* indicating the percentage of votes that each project received.

Figure 3 arranges projects in descending order based on the average percentage of votes they received across input formats, from left to right. We then plot the distribution induced by each of the input formats. Notably, five of the six input formats give a very similar distribution, with SN being the outlier and giving a much flatter distribution (and in particular not as readily identifying the strongest projects in terms of the overall vote share). Visually, the plot suggests that project rankings are consistent across these input formats, implying that the collective preference remains largely unaffected by the format used.

To quantify this difference, as a measure of their (dis)similarity we computed the *Jensen-Shannon (JS) divergence* between distributions (a symmetric version of KL divergence). Figure 4 shows the resulting values. SN has a high JS divergence of on average 18.04×10^{-3} to the other input formats, confirming that its distribution of votes is quite distinct from the other formats. On the other hand, the formats “Distribute 10 points” (D10) and “Select 5 projects” (S5) have low average JS divergences of 4.46×10^{-3} and 4.85×10^{-3} respectively, suggesting they are more similar to other voting formats in terms of the voting distribution. Since S5 is particularly easy to use and to explain to voters, the observation that it induces a similar vote distribution as the more involved formats makes it an attractive choice.

Outcomes of Voting Rules. We computed the outcomes of the Greedy and MES aggregation methods under various budget settings, as shown in Figure 5. For instance, in the Expensive Projects section, the first column indicates the funding situation of project 14. It is funded by Greedy with a total budget of only 10,000 CHF. In contrast, it




Fig. 5. Visualisation of funded projects based on varying voting input formats, aggregation methods, and total budgets. Each row displays which of the 24 projects were selected by a specific combination of input format and aggregation method. The different shades of green or blue represent the projects that would be funded if the total budget were 10,000, 20,000, 60,000, or 90,000 CHF, from dark to light respectively. Projects with the lightest shade were not funded in all 4 budget scenarios. Projects with a cost of 5,000 CHF are on the left, and those with 10,000 CHF on the right. The heatmap display follows Fairstein, Benadè, and Gal (2023, Figure 5).




Fig. 6. Swarm plot of Hamming distances showing how much the voting outcome according to an aggregation method changes when we change the voting input format. Each point represents the Hamming distance between a pair of outcomes (as illustrated in Figure 5) obtained from two different voting input formats, averaged over nine values of the budget limit (10,000 to 50,000 CHF in increments of 5,000 CHF). Horizontal lines mark the average Hamming distance for each method, taken over all pairs, thus quantifying the variability under changing input formats for a given aggregation method.

is funded under the MES method only when the budget in the PB program reaches 20,000 CHF. Moreover, if the economical variants of Greedy or MES are applied, project 14 is funded only when the budget reaches 60,000 CHF. We also quantify the difference between the outcome bundles and show it in Figure 6, with each point indicating the Hamming distance between the outcomes using two different voting inputs.

Examining Figures 5 and 6, we can make the following observations:

- *Greedy's Popularity Representation:* As Greedy simply funds the projects with the highest number of votes, its outcome table visually reflects the relative popularity of projects. The roughly equal funding scenario observed between cheaper and more expensive projects suggests that voters were not particularly cost conscious. For example, when the total budget was 90,000 CHF, Greedy with S5 or D5 input would fund 7 expensive projects and only 4 cheap projects in Figure 5. This indicates that voters did not show a stronger preference for cheaper projects when casting their votes and were not overly cost-conscious or frugal in their selections.
- *Impact of the Input Format:* Although outcomes that use the SN input format exhibit distinct patterns compared to other input formats, there is no significant evidence suggesting that the choice of input affects

the robustness or stability of the outcome in the experiment. Excluding the SN input, variations in input format do not substantially impact which projects are winning.

- *Impact of the Aggregation Methods:* As [Figure 5](#) visually demonstrates, Greedy tends to fund more expensive projects at the expense of cheaper ones, whereas the two economical variants are inclined to fund more inexpensive projects, as expected. MES strikes a balance between the two. Regarding the robustness of aggregation methods to varying the voting input format, the average Hamming distance shown in [Figure 6](#) indicates that outcomes using different inputs tend to diverge more when Greedy (1.69) and Economical Greedy (1.79) are employed, compared to MES (1.13) and Economical MES (1.24) which show noticeably lower averages. It is also clear that the choice of the aggregation method influences the outcome more than the choice of the voting input.

6 RESULTS: AGGREGATION METHODS

In this section, we discuss how participants rated the outcomes produced by the two different aggregation methods: the Greedy rule (“Method A”) and MES (“Method B”). We will analyse how participants rated their fairness and trustworthiness, and how well they understood the methods, as well as how these ratings were influenced by different types of explanations of the outcome.

Translating Likert scales into numerical values from 1 (extremely dissatisfied/unfair) to 5 (extremely satisfied/fair), participants rated their satisfaction with the Greedy outcome to be 2.83 on average, compared to 3.44 for MES, with 79% of participants rating MES higher than Greedy. For fairness (before explanation), Greedy received a rating of 2.99, and MES received 3.61, with 87% of participants considering MES fairer than Greedy. After explanation, fairness ratings were 2.97 for Greedy and 3.93 for MES. Therefore, across both dimensions – satisfaction and fairness – participants consistently rated MES higher than Greedy, with all differences being statistically significant (Wilcoxon test, $p < 0.001$, see [Table 11](#) in the Appendix).

6.1 Correlation Between Utility, Satisfaction, and Perceived Fairness

Participants might choose their ratings of voting outcomes according to two conceptually distinct types of reasoning: (1) based on how well they personally liked the projects chosen (e.g., because a voting outcome selected more of the projects they voted for in the first part of the experiment); (2) based on global properties of a voting outcome such as whether it chooses projects with high vote counts or whether it splits the available budget reasonably across districts and categories.

To distinguish the importance of these two approaches, our analysis in [Figure 7](#) shows how their satisfaction and fairness ratings depended on the participants’ utility resulting from voting outcomes (measured as the number of projects they selected in the S5 format among those selected for funding). This relationship is crucial to understand the legitimacy of participatory budgeting outcomes: Satisfaction ratings indicate whether the outcomes align with the participants’ expectations (outcome legitimacy), while fairness ratings assess the level of equality (throughput legitimacy).

For Satisfaction, [Figure 7](#) shows that the Greedy rule outcomes display a stronger correlation between utility and satisfaction compared to MES. For fairness ratings, the outcomes from the Greedy rule suggest that voters tend to perceive an outcome as fairer if it results in higher utility. This significant correlation indicates that, within the Greedy method, fairness perceptions are closely linked to success in terms of utility. In contrast, MES does not show a statistically significant correlation like this, suggesting that perceived fairness associated with MES outcomes is not tied to individual success in terms of utility in the PB program. In the Appendix, we present further results in terms of cost utility, using also formats other than S5, see [Tables 12 to 15](#).




Fig. 7. Correlation between Utility and Voter Satisfaction & Fairness, using S5 voting input and cardinality utility. Mean utility and satisfaction are listed on the bottom right of each plot.

6.2 Explanations of Voting Aggregation Methods and Outcomes

To assess the explainability of the two voting aggregation methods employed, we divided participants into three distinct treatment groups, each receiving different explanations as detailed in Section 4.2. Afterwards, participants rated how well they understood Method A and Method B, and were asked to rate the two aggregation methods again in terms of fairness and trustworthiness.

Among the “Mechanism” participants, it was evident that the underlying algorithm of the Method of Equal Shares was more difficult to grasp, with only 49% claiming to understand “a lot” or “a great deal”, compared to 87% for the Greedy method.

Interestingly, despite the relative complexity and lower levels of comprehension of MES within the Mechanism group, its perceived trustworthiness and fairness ratings overall surpass those of the Greedy method. Despite the complexity and lower comprehension levels of MES within the Mechanism group, it garners higher trustworthiness and fairness ratings overall than the Greedy method.

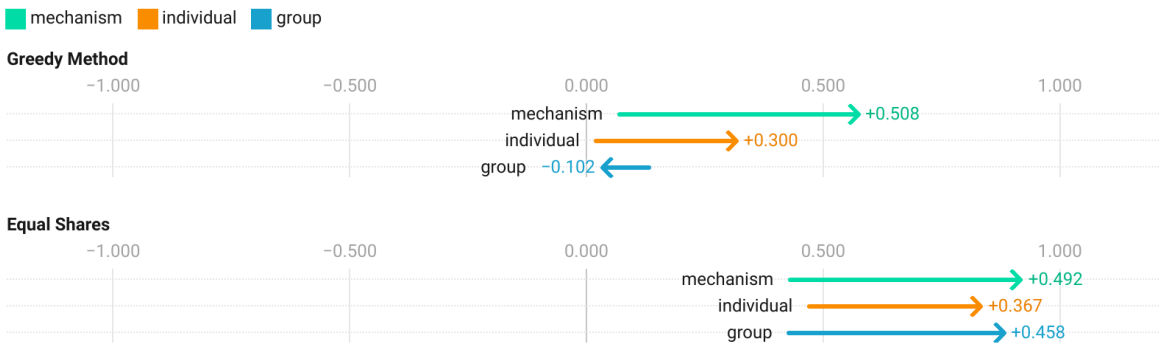


Fig. 8. Change in perceived trustworthiness of aggregation methods in different explanation treatment groups.




Fig. 9. Change in the perceived fairness of aggregation methods in different explanation treatment groups.

The results represented in Figures 8 and 9 show the difference in the ratings of fairness and trustworthiness of Greedy and MES before and after the explanation treatment.⁴ In the figure, we reoriented the 5-point Likert scales to range from -2 to 2, with 0 being a neutral sentiment and a higher rating indicating a more positive evaluation. To have a better understanding of the overall shift of sentiment before and after the explanation, we determined the means of the 180 rating for each question.

- In the Mechanism group, to which the detailed workings of the aggregation methods were explained, there was a slight increase in the perceived fairness of the Greedy method post-explanation, from approximately 0.11 to 0.16. However, as shown in Figure 9, the perceived fairness of MES improved markedly, from 0.54 to 1.03. Trustworthiness ratings for both methods also show substantial increases, with MES experiencing an increase from 0.43 to 0.92 (Figure 8). This is the most effective explanation in terms of boosting the trustworthiness of both algorithms.
- In the Individual treatment, where participants were shown the utility distributions among the (simulated) voters, the Greedy method's initially negative fairness rating shifted to a slightly positive rating after the explanation. In contrast, MES saw a small increase from 0.63 to 0.80. However, trustworthiness for both methods showed a substantial improvement.
- In the Group treatment, where the budget distribution was shown over districts and categories, the results were different. The Greedy method's fairness rating fell from a near-neutral rating to -0.31 post-explanation, while the fairness rating of MES showed an improvement from 0.64 to 0.97 (Figure 9). Trustworthiness for the Greedy method declined slightly, while the trustworthiness of MES improved from 0.42 to 0.88 (Figure 8).

In summary, while the details provided in the explanations of the different voting methods varied across the three treatment groups, one can observe a consistent trend in Figures 8 and 9: the post-explanation trustworthiness and fairness ratings of MES are higher across all treatment groups.

A further insight from the data is that higher perceived trustworthiness does not necessarily equate to increased perceived fairness. Despite the overall increase in trustworthiness ratings post-explanation, fairness ratings varied. This is especially the case in the Mechanism group. While understanding the Greedy method significantly increases the trustworthiness of the method, the fairness rating remains around the same level. This suggests that,

⁴Before being shown the explanation, participants rated trustworthiness on a scale ranging from "strongly disagree" to "strongly agree". After the explanation, the scale ranged from "not at all trustworthy" to "very trustworthy". Changes in response should be seen in the light of this mismatch. For fairness, the labels were identical before and after the explanation.

increasing “model transparency” and a better understanding of a system can contribute to the trustworthiness of an algorithm, but it does not always lead to increased perceptions of fairness.

7 CONCLUSIONS AND FUTURE WORK

7.1 Discussion

Our study provides critical insights for the design of digital PB systems, with a particular focus on how different aspects of the voting process affect both the outcomes and the participants’ perceptions of these.

Regarding voting input formats, we found that, in the multi-winner election setting, there are differences from the outcome perspective between, on the one hand, an unstructured voting input format (SN, selecting any number of projects) and, on the other hand, semi-structured voting input formats (S5, D5, D10, S5R, S5D10). However, the choice between different semi-structured voting input formats (ranked or cumulative formats) have a less pronounced effect on the outcome. This supports the findings of Fairstein, Benadè, and Gal (2023) in their comparison between k -approval voting and other voting inputs. However, our findings reveal that voters have strong preferences between voting input formats. Our experimental participants told us that permitting voters to select any number of projects came with a lack of clarity, and that the format of distributing 5 points may lack granularity. Participants showed a preference for voting input formats using rankings or a higher number of points (e.g., 10), suggesting that these formats offer a stronger sense of engagement in the participatory process. The most favoured formats involved ranking selected projects (S5D) or allocating points among selected projects (S5D10). These formats strike a balance between structured decision-making and choice flexibility, making them suitable as recommendations for digital participatory budgeting interfaces that prioritise voter perception in collective decision-making.

Regarding voting aggregation rules, voters showed a strong intuition for outcomes that provide proportional representation and prioritise fairness. These features were significantly correlated with the overall satisfaction with an outcome. Accordingly, there is now experimental support for recent efforts in academia to design, study, and advocate collective decision-making mechanisms that enhance fairness and proportionality, such as MES. When we asked participants to compare the outcomes of the Greedy rule and of MES, they more readily understood the Greedy outcome. However, this difference was reduced after showing participants an explanation or key statistics about the outcomes. Participants tended to assign higher satisfaction ratings to outcomes which funded more of the projects that they personally supported. Furthermore, they tended to assign high fairness ratings to the MES outcome independent of their personal utility, thereby arguably promoting a more community-centric perspective. This observation also suggests that trustworthiness, fairness, and satisfaction are interconnected, hinting at the broader notion that legitimacy is a latent construct that emerges from the combination and interaction of various factors.

In terms of explaining aggregation algorithms, compared to Greedy, MES was perceived as more complex but more trustworthy and fair. Our results also suggest that, in the context of PB, pre-hoc explanations are more effective in improving trust in algorithms than post-hoc explanations. Additionally, increasing “model transparency” can contribute to an algorithm’s trustworthiness, but does not always lead to enhanced perceptions of fairness, especially when a less fair algorithm becomes more transparent to voters through explanation.

Overall, the voting input and voting aggregation can be seen as the “front-end” and “back-end” of a voting process, respectively. While the “front-end” voting input format may not change the overall result as much, it influences voter perception of their contribution to the outcome as the main point of interaction. In contrast, the “back end” voting aggregation method, though less visible to voters, is critical for a fair outcome. An effective voting system consists of a voting input that allows voters to freely express their choices, an aggregation method that supports proportionality, and an explanation that clearly explains both the mechanism and the outcome of the voting process.

7.2 Practical Relevance

In 2023, the city of Aarau, Switzerland, implemented its first PB program (stadtidee.aarau.ch), for distributing a 50,000 CHF budget among 33 projects. Preliminary results from our laboratory experiment directly informed the setup of the PB process. In particular, the city decided to use a 10-point voting system similar to D10 and to use the Method of Equal Shares to aggregate votes. About 1,800 of approximately 22,000 residents voted. Before the vote, the city website featured an explanation of how the Method of Equal Shares works, inspired by the explanation in our *Mechanism explanation* treatment. After the vote, the [city website](#) showed statistics of how the budget was divided between districts, inspired by the explanation in our *Group explanation* treatment. The key takeaways included: (1) the Method of Equal Shares funded 17 projects, in contrast to just 7 (on average more expensive) projects which would have been chosen by the Greedy method; (2) unlike the Greedy method, the Method of Equal Shares allowed all districts to receive some portion of the budget.

7.3 Limitations and Future Work

One potential limitation of our study is the non-representativeness of our participant sample, as all participants were university students in Zurich, which may somewhat limit the generalisability of our findings. Another potential limitation is the lack of full randomisation in the design of the experiment. Given the high complexity of the experiment, in order to reduce the cognitive load and avoid confusion, we chose not to randomise the order of voting input methods and the displayed order of projects. This decision might inherently introduce some bias into participants' perceptions. Furthermore, the simulated nature of our fictional voting may not fully capture the complexities of real-world PB processes. As a result, some of our findings could, in principle, be context-sensitive, affected by specific interface designs, the setting of random outcomes, or the number of projects.

Moving forward, a multifaceted approach to research is recommended:

- Continue connecting Social Choice theories to experimental case studies and real-world applications. In particular, we look forward to the publication of the in-depth evaluation of case study in Aarau, Switzerland.
- Extend the research to larger and more diverse populations, and to various socio-economic and cultural contexts, thereby increasing the generalisability of our findings.
- Investigate the impact of different interface designs and presentation formats on voter engagement and decision-making. Understanding these aspects could inform the development of more intuitive and accessible PB platforms.
- Conduct longitudinal studies to grasp the evolution of voter behaviour and satisfaction over time, offering insights into the long-term implications and efficacy of different PB designs.
- Investigate human priorities, perceptions of project costs, and utility functions in real-world PB settings.
- Explore the criteria used by voters to filter projects before voting and the impact of voting rules on these decisions.

In conclusion, our study enhances the understanding of digital participatory budgeting systems, especially in relation to the impact of voting input formats and aggregation methods on fairness and voter perception. It points towards the necessity of further empirical research and of applying insights to real-world PB settings. Future work will focus on refining these systems for greater efficacy and inclusivity, while aiming to move from theoretical explorations to practical applications in digital governance.

7.4 Acknowledgements

This study was financed by the Swiss National Science Foundation (SNSF) as part of the National Research Programme NRP 77 Digital Transformation, project no. 187249 [Applicants: Regula Hänggli Fricker, Evangelos Pournaras, Dirk Helbing]. Evangelos Pournaras is also supported by the UKRI Future Leaders Fellowship (MR/W009560/1): '*Digitally Assisted Collective Governance of Smart City Commons-ARTIO*'. We would further like to

thank the Decision Lab of ETH Zürich for their great support in carrying out the lab experiment. Special thanks also go to the city of Aarau, especially the Aarau Stadtidee team, and Thomas Wellings from the team from the University of Leeds for their insightful suggestions and their commitment to scientific collaboration. Last but not least, we would like to thank three anonymous reviewers for their useful feedback that helped us improve the presentation of our paper.

REFERENCES

- Aziz, Haris, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh (2017). “Justified representation in approval-based committee voting”. In: *Social Choice and Welfare* 48.2, pp. 461–485. DOI: [10.1007/s00355-016-1019-3](https://doi.org/10.1007/s00355-016-1019-3) (→ p. 6).
- Aziz, Haris, Barton E. Lee, and Nimrod Talmon (2018). “Proportionally representative participatory budgeting: Axioms and algorithms”. In: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 23–31. URL: <https://www.ifaamas.org/Proceedings/aamas2018/pdfs/p23.pdf> (→ pp. 3, 6).
- Aziz, Haris and Nisarg Shah (2021). “Participatory budgeting: Models and approaches”. In: *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer, pp. 215–236. DOI: [10.1007/978-3-030-54936-7_10](https://doi.org/10.1007/978-3-030-54936-7_10) (→ p. 7).
- Baumeister, Dorothea, Linus Boes, and Tessa Seeger (2020). “Irresolute approval-based budgeting”. In: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1774–1776. URL: <https://ifaamas.org/Proceedings/aamas2020/pdfs/p1774.pdf> (→ p. 6).
- Beetham, David (1991). *The Legitimation of Power*. Macmillan. DOI: [10.1007/978-1-349-21599-7](https://doi.org/10.1007/978-1-349-21599-7) (→ p. 5).
- Benadè, Gerdus, Nevo Itzhak, Nisarg Shah, Ariel D. Procaccia, and Kobi Gal (2018). *Efficiency and usability of participatory budgeting methods*. Unpublished. URL: http://www.cs.toronto.edu/~nisarg/papers/pb_usability.pdf (→ p. 6).
- Benadè, Gerdus, Nisarg Shah, Ariel D. Procaccia, and Kobi Gal (2021). “Preference elicitation for participatory budgeting”. In: *Management Science* 67.5, pp. 2813–2827. DOI: [10.1287/mnsc.2020.3666](https://doi.org/10.1287/mnsc.2020.3666) (→ p. 6).
- Boehmer, Niclas, Piotr Faliszewski, Łukasz Janeczko, Dominik Peters, Grzegorz Pierczyński, Šimon Schierreich, Piotr Skowron, and Stanisław Szufa (2023). *Evaluation of project performance in participatory budgeting*. arXiv: [2312.14723](https://arxiv.org/abs/2312.14723) [cs.GT] (→ p. 3).
- Brams, Steven J. and Peter C. Fishburn (1978). “Approval voting”. In: *American Political Science Review* 72.3, pp. 831–847. DOI: [10.2307/1955105](https://doi.org/10.2307/1955105) (→ p. 9).
- Brams, Steven J. and Peter C. Fishburn (2002). “The Borda count and its real-world applications”. In: *Handbook of Social Choice and Welfare*. Vol. 1, pp. 123–148. DOI: [10.1016/S1574-0110\(02\)80009-1](https://doi.org/10.1016/S1574-0110(02)80009-1) (→ p. 14).
- Brandt, Felix, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (2016). *Handbook of Computational Social Choice*. Cambridge University Press (→ p. 5).
- Brill, Markus, Stefan Forster, Martin Lackner, Jan Maly, and Jannik Peters (2023). “Proportionality in approval-based participatory budgeting”. In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5524–5531. DOI: [10.1609/aaai.v37i5.25686](https://doi.org/10.1609/aaai.v37i5.25686) (→ pp. 3, 6, 7).
- Brill, Markus and Jannik Peters (2023). “Robust and verifiable proportionality axioms for multiwinner voting”. In: *Proceedings of the 24th ACM Conference on Economics and Computation (ACM EC)*, p. 301. DOI: [10.1145/3580507.3597785](https://doi.org/10.1145/3580507.3597785) (→ p. 6).
- Cabannes, Yves (2004). “Participatory budgeting: a significant contribution to participatory democracy”. In: *Environment and Urbanization* 16.1, pp. 27–46. DOI: [10.1177/095624780401600104](https://doi.org/10.1177/095624780401600104) (→ p. 2).
- Corazzini, Luca, Christopher Cotton, and Paola Valbonesi (2015). “Donor coordination in project funding: Evidence from a threshold public goods experiment”. In: *Journal of Public Economics* 128, pp. 16–29. DOI: [10.1016/j.jpubeco.2015.05.005](https://doi.org/10.1016/j.jpubeco.2015.05.005) (→ p. 7).
- Crosan, Rachel T. A. and Melanie Beth Marks (2000). “Step returns in threshold public goods: A meta- and experimental analysis”. In: *Experimental Economics* 2.3, pp. 239–259. DOI: [10.1023/A:1009918829192](https://doi.org/10.1023/A:1009918829192) (→ p. 7).
- Cullen, Rowena (2001). “Addressing the digital divide”. In: *Online Information Review* 25.5, pp. 311–320. DOI: [10.1108/14684520110410517](https://doi.org/10.1108/14684520110410517) (→ pp. 2, 4).

- Davies, Jonathan, Miguel Arana-Catania, Rob Procter, Felix Anselm Van Lier, and Yulan He (2021). “A mixed-methods ethnographic approach to participatory budgeting in Scotland”. In: *Proceedings of the 2021 Conference on Information Technology for Social Good (GoodIT)*, pp. 317–320. DOI: [10.1145/3462203.3475891](https://doi.org/10.1145/3462203.3475891) (→ p. 2).
- De Vries, Michiel S, Juraj Nemec, and David Špaček (2022). *International Trends in Participatory Budgeting: Between Trivial Pursuits and Best Practices*. Palgrave Macmillan. DOI: [10.1007/978-3-030-79930-4](https://doi.org/10.1007/978-3-030-79930-4) (→ p. 2).
- Fain, Brandon, Ashish Goel, and Kamesh Munagala (2016). “The core of the participatory budgeting problem”. In: *Web and Internet Economics (WINE)*. Vol. 10123. Lecture Notes in Computer Science. Springer, pp. 384–399. DOI: [10.1007/978-3-662-54110-4_27](https://doi.org/10.1007/978-3-662-54110-4_27) (→ p. 5).
- Fairstein, Roy, Gerdus Benadè, and Kobi Gal (2023). “Participatory budgeting designs for the real world”. In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5633–5640. DOI: [10.1609/aaai.v37i5.25699](https://doi.org/10.1609/aaai.v37i5.25699) (→ pp. 2, 6, 17, 21).
- Faliszewski, Piotr, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon (2023). “Participatory budgeting: Data, tools and analysis”. In: *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2667–2674. DOI: [10.24963/ijcai.2023/297](https://doi.org/10.24963/ijcai.2023/297) (→ pp. 3, 6, 8, 11).
- Faliszewski, Piotr, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon (2017). “Multiwinner voting: A new challenge for social choice theory”. In: *Trends in Computational Social Choice*. AI Access, pp. 27–47. URL: <https://archive.illc.uva.nl/COST-IC1205/BookDocs/Chapters/TrendsCOMSOC-02.pdf> (→ p. 5).
- Fluschnik, Till, Piotr Skowron, Mervin Triphaus, and Kai Wilker (2019). “Fair knapsack”. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1941–1948. DOI: [10.1609/aaai.v33i01.33011941](https://doi.org/10.1609/aaai.v33i01.33011941) (→ p. 6).
- Goel, Ashish, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto (2019). “Knapsack voting for participatory budgeting”. In: *ACM Transactions on Economics and Computation* 7.2, pp. 1–27. DOI: [10.1145/3340230](https://doi.org/10.1145/3340230) (→ p. 6).
- Hänggli, Regula, Evangelos Pournaras, and Dirk Helbing (2021). “Human-centered democratic innovations with digital and participatory elements”. In: *Proceedings of the 22nd Annual International Conference on Digital Government Research (DG.O)*, pp. 227–233. DOI: [10.1145/3463677.3463708](https://doi.org/10.1145/3463677.3463708) (→ p. 5).
- Hartley, Kris (2021). “Public trust and political legitimacy in the smart city: A reckoning for technocracy”. In: *Science, Technology, & Human Values* 46.6, pp. 1286–1315. DOI: [10.1177/0162243921992864](https://doi.org/10.1177/0162243921992864) (→ p. 5).
- Helbing, Dirk, Sachit Mahajan, Regula Hänggli Fricker, Andrea Musso, Carina I. Hausladen, Cesare Carissimo, Dino Carpentras, Elisabeth Stockinger, Javier Argota Sanchez-Vaquerizo, Joshua C. Yang, Mark C. Ballandies, Marcin Korecki, Rohit K. Dubey, and Evangelos Pournaras (2023). “Democracy by design: Perspectives for digitally assisted, participatory upgrades of society”. In: *Journal of Computational Science* 71, p. 102061. DOI: [10.1016/j.jocs.2023.102061](https://doi.org/10.1016/j.jocs.2023.102061) (→ pp. 2, 4).
- Hershkowitz, D. Ellis, Anson Kahng, Dominik Peters, and Ariel D. Procaccia (2021). “District-fair participatory budgeting”. In: *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5464–5471. DOI: [10.1609/aaai.v35i6.16688](https://doi.org/10.1609/aaai.v35i6.16688) (→ p. 6).
- Iyengar, Sheena S. and Mark R. Lepper (2000). “When choice is demotivating: Can one desire too much of a good thing?” In: *Journal of Personality and Social Psychology* 79.6, pp. 995–1006. DOI: [10.1037/0022-3514.79.6.995](https://doi.org/10.1037/0022-3514.79.6.995) (→ p. 4).
- Kenny, Eoin M., Eoin D. Delaney, Derek Greene, and Mark T. Keane (2021). “Post-hoc explanation options for XAI in deep learning: The insight centre for data analytics perspective”. In: *Pattern Recognition: ICPR International Workshops and Challenges (Part III)*. Springer, pp. 20–34. DOI: [10.1007/978-3-030-68796-0_2](https://doi.org/10.1007/978-3-030-68796-0_2) (→ p. 5).
- Laruelle, Annick (2021). “Voting to select projects in participatory budgeting”. In: *European Journal of Operational Research* 288.2, pp. 598–604. DOI: [10.1016/j.ejor.2020.05.063](https://doi.org/10.1016/j.ejor.2020.05.063) (→ pp. 2, 4, 6).

- Lee, Min Kyung (2018). “Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management”. In: *Big Data & Society* 5.1, p. 2053951718756684. DOI: [10.1177/2053951718756684](https://doi.org/10.1177/2053951718756684). eprint: <https://doi.org/10.1177/2053951718756684>. URL: <https://doi.org/10.1177/2053951718756684> (→ p. 2).
- Lev-On, Azi (2022). “Watching participatory budgeting events or attending them produce different distributive outcomes”. In: *Online Information Review* 46.2, pp. 244–255. DOI: [10.1108/OIR-01-2020-0033](https://doi.org/10.1108/OIR-01-2020-0033) (→ p. 2).
- Los, Maaïke, Zoé Christoff, and Davide Grossi (2022). “Proportional budget allocations: Towards a systematization”. In: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 398–404. DOI: [10.24963/ijcai.2022/57](https://doi.org/10.24963/ijcai.2022/57) (→ p. 6).
- Mærøe, Anders Røsten, Alexander Norta, Valentyna Tsap, and Ingrid Pappel (2021). “Increasing citizen participation in e-participatory budgeting processes”. In: *Journal of Information Technology and Politics* 18.2, pp. 125–147. DOI: [10.1080/19331681.2020.1821421](https://doi.org/10.1080/19331681.2020.1821421) (→ p. 4).
- Merrill, Samuel (1984). “A comparison of efficiency of multicandidate electoral systems”. In: *American Journal of Political Science* 28.1, pp. 23–48. DOI: [10.2307/2110786](https://doi.org/10.2307/2110786) (→ p. 4).
- Page, Scott E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies (New Edition)*. Princeton University Press. URL: <http://www.jstor.org/stable/j.ctt7sp9c> (visited on 03/20/2024) (→ p. 2).
- Palacin, Victoria, Samantha McDonald, Pablo Aragón, and Matti Nelimarkka (2024). “Configurations of Digital Participatory Budgeting”. In: *ACM Transactions on Computer-Human Interaction* 31.2. DOI: [10.1145/3635144](https://doi.org/10.1145/3635144) (→ p. 2).
- Peiró, Jose María, Vicente Martínez-Tur, and Carolina Moliner (2014). “Perceived fairness”. In: *Encyclopedia of Quality of Life and Well-Being Research*. Springer, pp. 4693–4696. DOI: [10.1007/978-94-007-0753-5_2125](https://doi.org/10.1007/978-94-007-0753-5_2125) (→ p. 5).
- Peters, Dominik (2022). *List of participatory budgeting votes — Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/w/index.php?title=List_of_participatory_budgeting_votes&oldid=1126104283 (→ p. 2).
- Peters, Dominik, Grzegorz Pierczyński, and Piotr Skowron (2021). “Proportional participatory budgeting with additive utilities”. In: *Advances in Neural Information Processing Systems*. Vol. 34, pp. 12726–12737. URL: <https://proceedings.neurips.cc/paper/2021/file/69f8ea31de0c00502b2ae571fbab1f95-Paper.pdf> (→ pp. 3, 6).
- Rey, Simon and Jan Maly (2023). *The (computational) social choice take on indivisible participatory budgeting*. arXiv: [2303.00621](https://arxiv.org/abs/2303.00621) [cs.GT] (→ pp. 3, 5, 7).
- Roetzel, Peter Gordon (2019). “Information overload in the information age: a review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development”. In: *Business Research* 12.2, pp. 479–522. DOI: [10.1007/s40685-018-0069-z](https://doi.org/10.1007/s40685-018-0069-z) (→ p. 4).
- Rosenfeld, Ariel and Nimrod Talmon (2021). *What should we optimize in participatory budgeting? An experimental study*. arXiv: [2111.07308](https://arxiv.org/abs/2111.07308) [cs.MA] (→ p. 7).
- Ryan, Matt (2021). *Why Citizen Participation Succeeds or Fails*. Bristol University Press. DOI: [10.2307/j.ctv1rnpjh7](https://doi.org/10.2307/j.ctv1rnpjh7) (→ p. 2).
- Scharpf, Fritz W. (1999). *Governing in Europe: Effective and Democratic?* Oxford University Press. DOI: [10.1093/acprof:oso/9780198295457.001.0001](https://doi.org/10.1093/acprof:oso/9780198295457.001.0001) (→ p. 5).
- Scharpf, Fritz W. (2000). “Interdependence and democratic legitimation”. In: *Disaffected Democracies*. Princeton University Press. Chap. 5, pp. 101–120. DOI: [10.1515/9780691186849-009](https://doi.org/10.1515/9780691186849-009) (→ p. 5).
- Schmidt, Vivien A. (2013). “Democracy and legitimacy in the European Union revisited: Input, output and ‘throughput’”. In: *Political Studies* 61.1, pp. 2–22. DOI: [10.1111/j.1467-9248.2012.00962.x](https://doi.org/10.1111/j.1467-9248.2012.00962.x) (→ p. 5).
- Sørmo, Frode, Jörg Cassens, and Agnar Aamodt (2005). “Explanation in case-based reasoning—Perspectives and goals”. In: *Artificial Intelligence Review* 24.2, pp. 109–143. DOI: [10.1007/s10462-005-4607-7](https://doi.org/10.1007/s10462-005-4607-7) (→ p. 5).

- Strebel, Michael Andrea, Daniel Kübler, and Frank Marcinkowski (2019). “The importance of input and output legitimacy in democratic governance: Evidence from a population-based survey experiment in four West European countries”. In: *European Journal of Political Research* 58.2, pp. 488–513. DOI: [10.1111/1475-6765.12293](https://doi.org/10.1111/1475-6765.12293) (→ p. 5).
- Talmon, Nimrod and Piotr Faliszewski (2019). “A framework for approval-based budgeting methods”. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2181–2188. DOI: [10.1609/aaai.v33i01.33012181](https://doi.org/10.1609/aaai.v33i01.33012181) (→ pp. 3, 6).
- Uslaner, Eric M. (2002). *The Moral Foundations of Trust*. Cambridge University Press. DOI: [10.1017/CBO9780511614934](https://doi.org/10.1017/CBO9780511614934) (→ p. 5).
- Wampler, Brian (2000). *A guide to participatory budgeting*. International Budget Partnership. URL: <https://internationalbudget.org/wp-content/uploads/A-Guide-to-Participatory-Budgeting.pdf> (→ p. 4).
- Wampler, Brian, Stephanie McNulty, and Michael Touchton (2021). *Participatory budgeting in global perspective*. Oxford University Press. DOI: [10.1093/oso/9780192897756.001.0001](https://doi.org/10.1093/oso/9780192897756.001.0001) (→ pp. 2, 4).

A ADDITIONAL TABLES

A.1 Voting Input Formats

Table 4. Mann-Whitney U test results for easiness rating on input formats.

Comparison	<i>U</i> -statistic	<i>p</i> -value
SN vs S5	22696.0	3.71×10^{-12}
SN vs D5	26650.0	5.54×10^{-28}
SN vs D10	24857.5	8.25×10^{-20}
SN vs S5R	22880.5	1.25×10^{-12}
SN vs S5D10	24518.5	1.86×10^{-18}
S5 vs D5	23512.5	1.37×10^{-14}
S5 vs D10	20198.0	2.09×10^{-05}
S5 vs S5R	17424.0	0.19200
S5 vs S5D10	20069.5	4.29×10^{-05}
D5 vs D10	12056.0	1.29×10^{-05}
D5 vs S5R	10204.5	3.61×10^{-10}
D5 vs S5D10	12376.0	6.01×10^{-05}
D10 vs S5R	13717.0	0.00907
D10 vs S5D10	16341.5	0.88216
S5R vs S5D10	18670.0	0.00974

Table 5. Mann-Whitney U test results for the rating of Expressiveness as a function of input formats.

Comparison	<i>U</i> -statistic	<i>p</i> -value
SN vs S5	13260.0	0.00213
SN vs D5	13515.5	0.00501
SN vs D10	9998.5	8.99×10^{-11}
SN vs S5R	8757.0	9.68×10^{-15}
SN vs S5D10	8548.0	1.58×10^{-15}
S5 vs D5	16823.0	0.50425
S5 vs D10	10849.0	7.99×10^{-09}
S5 vs S5R	8903.0	1.37×10^{-14}
S5 vs S5D10	8761.0	2.41×10^{-15}
D5 vs D10	10444.0	7.98×10^{-10}
D5 vs S5R	8779.5	6.25×10^{-15}
D5 vs S5D10	8533.0	5.12×10^{-16}
D10 vs S5R	13172.5	0.00111
D10 vs S5D10	13563.0	0.00406
S5R vs S5D10	16737.5	0.56367

Table 6. Average ranks and standard deviations of the six voting input formats studied. Smaller average rank values indicate higher popularity. Collected from 180 participants.

Input Format	Average Rank	Std. Dev.
Select 5 and rank	2.31	1.32
Select 5 and distribute 10 points	2.36	1.34
Distribute 10 points	3.11	1.39
Select 5 projects	3.93	1.36
Distribute 5 points	4.38	1.20
Select any number of projects	4.93	1.71

Table 7. Input formats ranked by participants in the order of recommendation for use of PB by a city. The formats were ordered by average rank, and adjacent formats were compared using the Wilcoxon signed-rank test for significant rank differences.

Comparison	<i>W</i> -statistic	<i>p</i> -value
S5R > S5D10	7965.5	0.7920
S5D10 > D10	4613.0	2.78×10^{-7}
D10 > S5	5067.0	8.95×10^{-6}
S5 > D5	6170.5	0.0040
D5 > SN	5973.5	0.0017

Table 8. Regression results for Recommendation of voting input formats.

Variable	Coefficient	<i>p</i> -value
Constant	0.266	0.043*
Expressiveness	0.890	<0.001***
Easiness	0.006	0.879
Model Statistics		
R-squared	0.327	
F-statistic (Prob.)	261.9 (<0.001***)	

Table 9. Reasons for ranking a voting method as either first or last. Participants provided their rationale using an open-answer text field. These responses were subsequently categorised by the authors into specific themes, and the percentage of responses calculated for each theme.

First	%
Preference Clarity	48
Granularity	35
Limited choices	30
Ease of process	24
Flexibility	14
Visual Representation	6
Last	%
Lack of Preference Clarity	66
Overwhelming Choices	27
Limited Allocation	25
Lack of Specificity	15
Complexity	8
Randomness	5

Table 10. Statistical summary of percent of vote going to projects of the participant's selected district or selected preferred category, by input format.

Input	District			Category		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
SN	37.42	33.33	19.61	44.76	42.86	15.04
S5	44.78	40.00	23.58	48.33	40.00	20.91
D5	53.22	60.00	33.50	53.33	60.00	28.36
D10	50.39	50.00	30.02	50.61	50.00	23.09
S5R	51.33	53.33	27.15	51.22	53.33	22.68
S5D10	50.94	50.00	29.06	51.72	50.00	24.40

A.2 Aggregation Methods

Table 11. Statistical summary of Satisfaction and Fairness ratings and Wilcoxon test results for MES being rated higher than Greedy.

Metric	Mean	Median	Std. Dev.	<i>p</i> -value
Satisfaction (Greedy)	2.83	3.00	1.15	
Satisfaction (MES)	3.44	4.00	1.04	3.18×10^{-7}
Fairness Before (Greedy)	2.99	3.00	1.03	
Fairness Before (MES)	3.61	4.00	0.87	6.58×10^{-7}
Fairness After (Greedy)	2.97	3.00	1.07	
Fairness After (MES)	3.93	4.00	0.77	2.66×10^{-14}

Table 12. Correlation (and *p*-value) between satisfaction and cardinality utility for different voting methods.

Aggregation Method	SN	S5	D5	D10	S5R	S5D10
MES	-0.07 (0.3454)	0.16 (0.0342)	0.32 (0.0000)	0.34 (0.0000)	0.28 (0.0002)	0.29 (0.0001)
Greedy	0.23 (0.0019)	0.42 (0.0000)	0.42 (0.0000)	0.47 (0.0000)	0.47 (0.0000)	0.42 (0.0000)

Table 13. Correlation (and *p*-value) between satisfaction and cost utility for different voting methods.

Aggregation Method	SN	S5	D5	D10	S5R	S5D10
MES	-0.04 (0.5604)	0.20 (0.0063)	0.31 (0.0000)	0.33 (0.0000)	0.27 (0.0002)	0.29 (0.0001)
Greedy	0.26 (0.0003)	0.43 (0.0000)	0.43 (0.0000)	0.49 (0.0000)	0.48 (0.0000)	0.43 (0.0000)

Table 14. Correlation (and *p*-value) between fairness and cardinality utility for different voting methods.

Aggregation Method	SN	S5	D5	D10	S5R	S5D10
MES	0.05 (0.4948)	0.02 (0.7795)	0.07 (0.3764)	0.10 (0.1973)	0.07 (0.3175)	0.07 (0.3707)
Greedy	0.17 (0.0195)	0.21 (0.0039)	0.23 (0.0020)	0.25 (0.0009)	0.22 (0.0030)	0.18 (0.0130)

Table 15. Correlation (and *p*-value) between fairness and cost utility for different voting methods.

Aggregation Method	SN	S5	D5	D10	S5R	S5D10
MES	0.05 (0.5021)	0.02 (0.8209)	0.04 (0.6189)	0.07 (0.3546)	0.05 (0.5084)	0.06 (0.4496)
Greedy	0.17 (0.0206)	0.19 (0.0104)	0.23 (0.0022)	0.23 (0.0022)	0.22 (0.0030)	0.18 (0.0133)

B EXPERIMENT INTERFACE

B.1 Voting Interface

A. Select any number of projects

Here, in this vote, you can select all the projects you approve of. Your vote will be perfectly valid no matter if you vote for many or for few projects.

001 | 5,000 CHF | Nord | Nature

002 | 10,000 CHF | Nord | Nature

003 | 5,000 CHF | Nord | Culture

004 | 10,000 CHF | Nord | Culture

(a) Screenshot of the SN format interface.

C. Spread 5 points

Now, vote again in the same program. In this vote, you distribute 5 points among the projects you like. You can concentrate your points on one project, or spread them across a few projects.

0 1 2 3 4 5

001 | 5,000 CHF | Nord | Nature | 1

002 | 10,000 CHF | Nord | Nature | 3

003 | 5,000 CHF | Nord | Culture | 1

004 | 10,000 CHF | Nord | Culture | 0

(c) Screenshot of the D5 format interface.

E. Select 5 and rank

Now, vote again in the same program. In this vote, you first drag the top 5 projects that you like to the box, then rank them in order. The top project will get 5 points; the second project, 4 points; the third, 3 points; and so on.

Items

001 | 5,000 CHF | Nord | Nature

013 | 5,000 CHF | Ost | Nature

004 | 10,000 CHF | Nord | Culture

My top 5 projects

1 Bird Houses for Zurich-Höngg
019 | 5,000 CHF | West | Nature

2 Car-free Sunday Festivals on Mutschellenstrasse
010 | 10,000 CHF | Süd | Culture

(e) Screenshot of the S5R format interface.

B. Select 5 projects

Now, vote again in the same program. This time, select exactly 5 projects.

001 | 5,000 CHF | Nord | Nature

002 | 10,000 CHF | Nord | Nature

003 | 5,000 CHF | Nord | Culture

004 | 10,000 CHF | Nord | Culture

(b) Screenshot of the S5 format interface.

D. Spread 10 points

Now, vote again in the same program. In this vote, you have to distribute 10 points among the projects you like.

0 1 2 3 4 5 6 7 8 9 10

001 | 5,000 CHF | Nord | Nature | 2

002 | 10,000 CHF | Nord | Nature | 4

003 | 5,000 CHF | Nord | Culture | 4

004 | 10,000 CHF | Nord | Culture | 0

(d) Screenshot of the D10 format interface.

F. Select 5 and distribute 10 points

Now these are the 5 projects you previously selected. Distribute 10 points across the 5 projects.

My top 5 projects

0 1 2 3 4 5 6 7 8 9 10

001 | 5,000 CHF | Nord | Nature | 1

002 | 10,000 CHF | Nord | Nature | 3

003 | 5,000 CHF | Nord | Culture | 4

004 | 10,000 CHF | Nord | Culture | 1

005 | 5,000 CHF | Nord | Transportation | 1

(f) Screenshot of the S5D10 format interface.

Fig. 10. Screenshots of voting input format screens.

B.2 Simulated Voting Outcomes

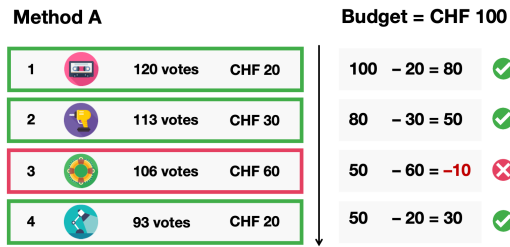
A	B	Project	Id	District	Category	Cost	Votes
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Music Studio at Kulturbahnhof Affoltern	004	Nord	Culture	10000	30
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Public Bicycle Moving Trailer to be Borrowed	023	West	Transportation	5000	27
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Planting Workshops at Oerlikon	001	Nord	Nature	5000	25
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Free Bike Fixing Workshops	011	Süd	Transportation	5000	25
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Car Sharing System for Young People	012	Süd	Transportation	10000	25
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Dingtheke: Community Things Exchange Library in Wipkingen	021	West	Culture	5000	25
<input checked="" type="checkbox"/>	<input type="checkbox"/>	More Night Buses to Oerlikon	006	Nord	Transportation	10000	24
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Free Open Badi Space in Wollishofen	007	Süd	Nature	5000	24
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Environmental Film Session for Kids	009	Süd	Culture	5000	24
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Bike Lanes on Seefeldstrasse	017	Ost	Transportation	5000	24
<input type="checkbox"/>	<input type="checkbox"/>	Car-free Langstrasse	024	West	Transportation	10000	24
<input type="checkbox"/>	<input type="checkbox"/>	Sunday Festivals on Mutschellenstrasse	010	Süd	Culture	10000	23
<input type="checkbox"/>	<input type="checkbox"/>	Transforming City Spaces under Trees into Gardens	013	Ost	Nature	5000	23
<input type="checkbox"/>	<input type="checkbox"/>	More Trees in Bellevue & Sechseläutenplatz	014	Ost	Nature	10000	23
<input type="checkbox"/>	<input type="checkbox"/>	Sustainable Cooking Workshop with Kids	022	West	Culture	10000	23
<input type="checkbox"/>	<input type="checkbox"/>	Footpath Gardens around Train Stations	002	Nord	Nature	10000	22
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Stuff-swap Closet: New life for old stuff	015	Ost	Culture	5000	22
<input type="checkbox"/>	<input type="checkbox"/>	Speed bumps in the City and the Lake Front	018	Ost	Transportation	10000	22
<input type="checkbox"/>	<input type="checkbox"/>	Wet Biotope as Learning Garden in Frankental	020	West	Nature	10000	22
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Bird Houses for Zurich-Höngg	019	West	Nature	5000	21
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Safe Bike Paths around Irchel Park	005	Nord	Transportation	5000	20
<input type="checkbox"/>	<input type="checkbox"/>	A Neighborhood Garden for Wiedikon	008	Süd	Nature	10000	19
<input type="checkbox"/>	<input type="checkbox"/>	Multicultural Festival at Sechseläutenplatz	016	Ost	Culture	10000	18
<input type="checkbox"/>	<input type="checkbox"/>	Kid Festival at Leutschenpark	003	Nord	Culture	5000	16

Fig. 11. Simulated outcome table for the fictional Zürich PB with voters having randomly sampled preferences (uniform distribution, “Instance 1”). Columns A and B show the projects selected using the Greedy rule and the Method of Equal Shares (MES), respectively.

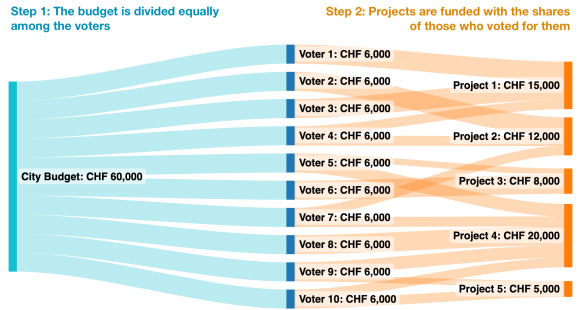
A	B	Project	Id	District	Category	Cost	Votes
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Car-free Langstrasse	024	West	Transportation	10000	28
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Music Studio at Kulturbahnhof Affoltern	004	Nord	Culture	10000	27
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Multicultural Festival at Sechseläutenplatz	016	Ost	Culture	10000	27
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Public Bicycle Moving Trailer to be Borrowed	023	West	Transportation	5000	26
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Safe Bike Paths around Irchel Park	005	Nord	Transportation	5000	25
<input checked="" type="checkbox"/>	<input type="checkbox"/>	More Night Buses to Oerlikon	006	Nord	Transportation	10000	25
<input type="checkbox"/>	<input type="checkbox"/>	Sunday Festivals on Mutschellenstrasse	010	Süd	Culture	10000	25
<input type="checkbox"/>	<input type="checkbox"/>	Sustainable Cooking Workshop with Kids	022	West	Culture	10000	25
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Environmental Film Session for Kids	009	Süd	Culture	5000	24
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Bird Houses for Zurich-Höngg	019	West	Nature	5000	24
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Wet Biotope as Learning Garden in Frankental	020	West	Nature	10000	24
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Bike Lanes on Seefeldstrasse	017	Ost	Transportation	5000	23
<input type="checkbox"/>	<input type="checkbox"/>	Speed bumps in the City and the Lake Front	018	Ost	Transportation	10000	23
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Free Open Badi Space in Wollishofen	007	Süd	Nature	5000	22
<input type="checkbox"/>	<input type="checkbox"/>	More Trees in Bellevue & Sechseläutenplatz	014	Ost	Nature	10000	22
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Planting Workshops at Oerlikon	001	Nord	Nature	5000	21
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Kid Festival at Leutschenpark	003	Nord	Culture	5000	21
<input type="checkbox"/>	<input type="checkbox"/>	Car Sharing System for Young People	012	Süd	Transportation	10000	21
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Stuff-swap Closet: New life for old stuff	015	Ost	Culture	5000	21
<input type="checkbox"/>	<input type="checkbox"/>	Dingtheke: Community Things Exchange Library in Wipkingen	021	West	Culture	5000	21
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Free Bike Fixing Workshops	011	Süd	Transportation	5000	20
<input type="checkbox"/>	<input type="checkbox"/>	Transforming City Spaces under Trees into Gardens	013	Ost	Nature	5000	20
<input type="checkbox"/>	<input type="checkbox"/>	Footpath Gardens around Train Stations	002	Nord	Nature	10000	19
<input type="checkbox"/>	<input type="checkbox"/>	A Neighborhood Garden for Wiedikon	008	Süd	Nature	10000	18

Fig. 12. Simulated outcome table for the fictional Zürich PB with voters having randomly sampled preferences (with strong preferences based on district and categories, “Instance 2”). Columns A and B show the projects selected using the Greedy rule and the Method of Equal Shares (MES), respectively.

B.3 Explanations




(a) Illustration used for the Greedy method, together with the text: “In Method A, the projects are selected simply based on the numbers of votes and total available budget. In each step, the method picks the most popular project that the remaining budget can fund. If the remaining available budget is not enough to fund the project, then the project is skipped. For example, the third most popular project in the following diagram is skipped as the remaining 50 CHF is not enough to fund the 60 CHF project. The process ends when no projects can be funded with the remaining budget.”



(b) Illustration used for MES, together with the text: “In the vote calculation of Method B, the total budget is divided so that each voter is assigned a hypothetical budget share to fund projects. The budget share assigned to a voter can only be used to fund projects that the voter has voted for. Method B goes through all project proposals, beginning with the projects with the highest number of votes. It selects a project if it can be funded using the budget shares of those who voted for the project. The method divides the cost of a project equally among its supporters.”


Fig. 13. Illustrations used for the *Mechanism Explanation Group* participants.

The following diagram shows the distribution of budget won by voters in the simulation.



	Outcome A	Outcome B
The average amount of budget won	8,508 CHF	8,204 CHF
% voters winning at least 1 project	79.6%	92.8%
% voters winning > 10,000 CHF	27.6%	16.0%

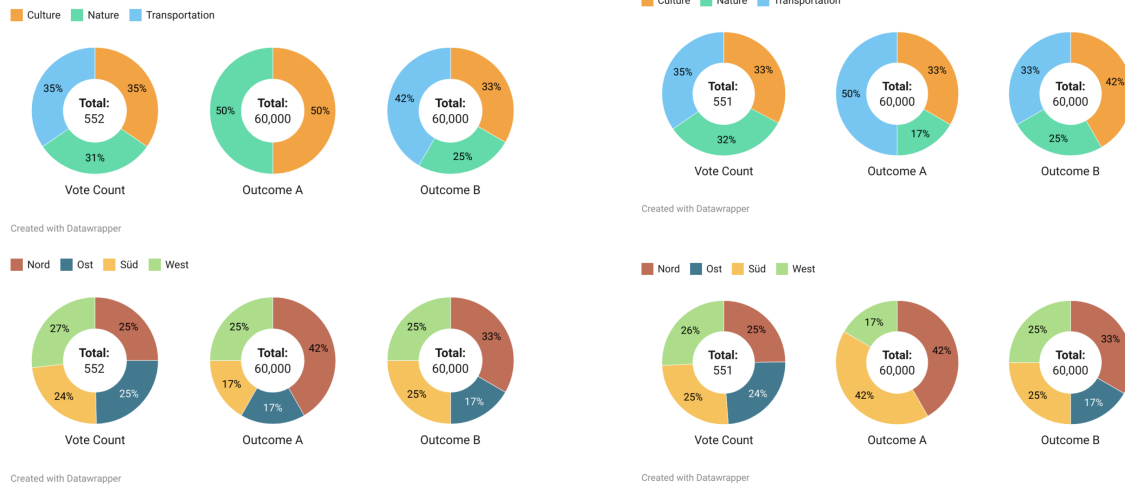
(a) Version for Instance 1.



	Outcome A	Outcome B
The average amount of budget won	8,702 CHF	8,011 CHF
% voters winning at least 1 project	64.1%	89.0%
% voters winning > 10,000 CHF	29.2%	18.2%

(b) Version for Instance 2.

Fig. 14. Illustration used for the *Individual Explanation Group* (depending on the instance they were shown). The distribution diagram shows the number of voters who “won” a particular amount of budget, in the cost utility sense, together with a table offering related statistics.



(a) Version for Instance 1.

(b) Version for Instance 2.

Fig. 15. Illustration used for the *Group Explanation Group* (depending on the instance they were shown). The diagram shows the budget distributions across different districts and categories, offering insights into the allocated funds for each interest group in the simulated outcomes.