An Incentive-Compatible Utilitarian Voting Procedure for Permanent Citizens' Assemblies *

Rajarshi Ghosh† and Marcus Pivato‡

August 25, 2025

Abstract

We consider a committee of voters randomly drawn from a larger population facing an infinite sequence of voting decisions, akin to a citizen jury. We propose a new voting mechanism for such juries where each voter has a privately known von Neumann-Morgenstern (vNM) utility function over social alternatives in each decision, and is asked to report a real-valued 'valuation' for each alternative of a decision. We further impose a probability of being removed from the committee for the next decision dependent on the report of a voter. If a voter is removed, then they are replaced by some non-committee member from the larger population. We show that when the voters' discount factor is large enough, imposing a probability equal to a scalar multiple of the Vickrey-Clarke-Groves tax leads to truthful revelation by the voters and consequently utilitarian efficient outcomes at a Bayesian Nash equilibrium.

1 Introduction

Citizen juries are committees of ordinary citizens selected by a government to advise them on a particular issue. They are already used around the world for a variety of issues such as waste management (Australia), constitutional reform (Ireland), political donations (Estonia) and mental health strategy (Canada). Citizen juries are usually selected for a particular issue and through a process known as *sortition*: a random selection from the larger population of citizens. Information regarding the issue is made available to them, and the jury members deliberate and deliver a verdict. Sometimes the verdict is given in the form of a decision e.g. to support or reject a policy proposal. This gives rise to the question of what voting rules could be desirable in such citizen juries. Voting in today's world normally entails declaring what a voter most prefers or ranking the alternatives in order of preference. A number of voting rules exist that allow voters to make such reports. The most commonly used voting rule is the *plurality rule* where each voter is allowed to declare their most preferred alternative and the alternative with the maximum votes is

*This research supported by Labex MME-DII (ANR11-LBX-0023-01) and CHOp (ANR-17-CE26-0003). We are grateful to Gorkem Celik, Antonin Macé as well as anonymous reviewers from COMSOC 2025 for many helpful suggestions.

†ESSEC Business School, CNRS, THEMA, F-95000 Cergy, France, rajarshi.ghosh@essec.edu

‡Centre d'Économie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne, CNRS, F-75013 Paris, France, marcuspivato@gmail.com

declared the winner. However, the plurality rule does not allow voters to express exactly "how much" they like an alternative and so the winning alternative might be one that is not liked too much by a narrow majority but hated by a large minority. Other voting rules, in which voters are only able to represent their ordinal preferences either fully or partially, have similar problems. These include the *Borda Rule*, in which voters rank the alternatives in their order of preference and *Approval Voting*, in which voters simply declare which alternatives they approve of and which ones they do not. The Borda rule, for example, faces the same problem as the plurality rule as well as other opportunities for strategic voting such as declaring the second-most preferred alternative as the least preferred alternative if the most likely winners are the first and second preferences.

The Gibbard-Satterthwaite Theorem says that no deterministic and non-dictatorial voting method is immune to tactical voting when there are at least three alternatives and all possible profiles of ordinal preferences are feasible. Thus, any incentive-compatible mechanism must impose some restriction on the domain of preference profiles, such as single-peakedness or quasilinearity. In our model, we assume that each voter has additively separable intertemporal expected-utility preferences over the outcomes of an infinite sequence of future decisions. We consider an infinite sequence of decisions because we assume that new decisions will continuously be made in the future. Incentive-compatibility is achieved by forcing each voter to trade off between her influence over the *present* decision and her potential influence over *future* decisions. Our mechanism allows each voter to report a *cardinal* utility function for each decision, and selects the alternative that maximizes the sum-total reported utility; assuming truthful reports, this is equivalent to maximizing utilitarian social welfare.

For the context of citizen juries facing an infinite sequence of decisions, we propose a class of "Sortition Mechanisms" that select a committee of voters from a larger population, and allow voters to express cardinal preferences. The Sortition Mechanisms we characterize can be applied to decisions with two or more discrete alternatives. To incentivize honest reports, the mechanisms impose a cost on the voters depending on what they report in the form of a tax on their probability of being allowed to remain in the committee for the next decision. So at the end of each decision, each voter is assigned a probability of being removed from the committee depending on what they reported, and then randomly removed and replaced with other members of the larger population. The Sortition Mechanism then proceeds in a manner similar to the previous decision for all subsequent decisions. Removal from the committee is costly for a voter because she will not be able to influence future decisions. Hence, she must make the trade-off between influencing the current decision and being able to participate in future ones. We consider some approximations and conditions on the way voters are selected into committee and show that a tax on the probability equal to a specific scalar multiple of the Vickrey-Clarke-Groves tax incentivizes truthful revelation in the form of Bayesian incentive-compatibility. Consequently, the class of mechanisms maximize utilitarian social welfare at a Bayesian Nash equilibrium. This consequently means that the case of a scalar multiple of the Pivotal Vickrey-Clarke-Groves tax also achieves these properties. In this case, a voter would be deterred from exaggerating their utility of an alternative because they would face a tax on their probability of participation in future decisions if they are pivotal in making that alternative win.

The rest of this paper is organized as follows. Section 2 reviews the related literature and highlights out contribution to it. Section 3 sets up the model and defines the mechanism while Section 4 contains our main results. Section 5 concludes and the Appendix contains the proofs of our results.

2 Review of the Literature

The problem of incentivizing agents to reveal their true preferences has been studied by economists for several decades. Under the assumption of quasi-linear preferences, the Vickrey-Clarke-Groves family of mechanisms achieves dominant strategy incentive compatibility and hence utilitarian optimality (Clarke, 1971; Groves, 1973). They do this by allowing each voter to report her utility for each alternative as a *willingness-to-pay* in terms of money or some other numéraire with an exogeneous value. The winning alternative is defined to be the alternative with the maximum sum-total reported utility. By transferring to each voter the sum of the reported utilities of the *other* voters for the winning alternative, the mechanism causes each voter to internalize the externalities of her participation and incentivizes her to report her true utility function. The mechanism maximizes utilitarian social welfare as a result. However, budget balance cannot be achieved so there may be a net monetary transfer between the voters and the government. The Expected Externality mechanism, on the other hand, achieves budget balance at the cost of a weaker Bayesian incentive-compatibility (d'Aspremont and Gérard-Varet, 1979) by defining the transfers to be the *expected value* of those in the Vickrey-Clarke-Groves(VCG) family of mechanisms. The Quadratic Voting mechanism, which requires voters to pay the *square* of the valuations they report for each alternative, is asymptotically utilitarian-efficient for a large voting population, but is only applicable to binary decisions (Weyl, 2012, 2017; Goeree and Zhang, 2017; Lalley and Weyl, 2019; Tideman and Plassmann, 2017). Voting at quadratic cost has been extended to apply to decisions of two or more discrete alternatives with the Quadratic Transfer Mechanism which asymptotically achieves utilitarian optimality for a large electorate (Eguia et al., 2023).

Despite the merits of mechanisms such as the VCG mechanisms and Quadratic Voting, they require that voters have preferences that are *quasi-linear* in money, and they require the voters to make monetary transfers. Thus, a larger influence on the outcome of the election can be achieved by the capability to make larger monetary transfers. The obvious problem with this is that rich people will have a disproportionately high influence on the outcome. To address these concerns, a number of mechanisms using an "artificial currency" have been developed. These include the Hylland-Zeckhauser mechanism (Hylland and Zeckhauser, 1980; Benjamin et al., 2013), Storable Votes (Casella, 2005), the linking mechanism (Jackson and Sonnenschein, 2007; Hortala-Vallve, 2010), some versions of Quadratic Voting (Casella and Macé, 2021; Casella and Sanchez, 2022) and the Quadratically Normalized Utilitarian Voting mechanism (Ghosh and Pivato, 2024). For example, in Casella's version of Quadratic Voting, each voter votes on multiple simultaneous binary decisions with a fixed budget of "voting points" which can be allocated across these decisions. This creates an opportunity cost of allocating these points to one decision instead of another thus disincentivizing the exaggeration of preferences. The Quadratically Normalized Utilitarian Voting mechanism uses this intuition and a weighted quadratic cost function to achieve an asymptotic dominant strategy incentive-compatibility and utilitarian efficiency for a large electorate and a large number of decisions. Another example is the Linking Mechanism, where there is a sequence of decisions and each voter is allowed to report a particular preference intensity only as often as the expected frequency of it being drawn from the type distribution. The mechanism is asymptotically Bayesian incentive-compatible and ex-ante efficient for a large sequence of decisions. Upper bounds in the welfare loss for a finite sequence of decisions have also been established in the literature (Ball and Kattwinkel, 2024).

There is also a stream of literature on ex-ante utilitarian optimal voting without the expression of cardinal preferences. A number of works exist for the case of binary decisions. Schmitz and Tröger (2012) show that if utilities are stochastically independent across agents and the distribution is symmetric, majority voting is ex-ante optimal among all anonymous and incentive-compatible rules. Azrieli and Kim (2014) shows that only a specific class of qualified majority rules defined by them is anonymous, incentive compatible and ex-ante efficient. Grüner and Tröger (2019) introduces voter participation costs and defines a linear voting rule to be a rule that weights alternatives. So if a voter votes in favour of an alternative, the score of that alternative will increase by a constant. They show that any ex-ante efficient mechanism is a linear rule. Finally, Gershkov et al. (2016) considers the case of two or more discrete alternatives and single-crossing preferences and proposes a sequential qualified majority voting rule. Each alternative is either approved or disapproved using qualified majority voting. If it is disapproved then the next alternative is voted on. They specify the majority thresholds and other conditions under which the mechanism is dominant strategy incentive compatible and ex-ante efficient.

The literature on sortition covers a variety of topics. In a setting where voters have a preference order over alternatives and the outcome is determined in a deterministic manner, Saran and Tumennasan (2013) show that social choice functions that are implementable at a Nash equilibrium can also be implemented after taking a random sample of voters from the population. They show that this is possible under their "p-monotonicity" condition on the voters' preferences and some conditions on the way voters are sampled. In another paper, Saran and Tumennasan (2019) obtain a similar result for a Bayesian Nash equilibrium under a "Bayesian monotonicity" condition and conditions on how the voters are sampled. They also show implementability in other environments without any monotonicity conditions. Other work on sortition includes the establishment of bounds on the deviation from optimal welfare for the entire population when a randomly selected committee makes the decision. This has been done for the majority rule (Meir et al., 2021) as well as voting mechanisms with pair-wise comparisons of alternatives (Anagnostides et al., 2022). Finally, another problem is that of obtaining committees that are representative of the larger population. In a setting where there is bias in the characteristics of volunteers (e.g. gender, education etc) willing to participate in the committee, Flanigan et al. (2020) et al proposes an algorithm that selects committees satisfying certain quotas on these characteristics while ensuring that each agent in the population has an almost equal probability of selection. Flanigan et al. (2021b) also proposes algorithms that achieve such committees. However, their algorithms instead ensure that probabilities of selections approximately maximize either a Maximin, Leximin or Nash Welfare objective function respectively. Ebadian et al. (2022) measure the representation of an agent by taking the distance between the q^{th} closest committee member and the agent where each agent's location on a metric space indicates their preferences. They go on to show trade-offs between representation and fairness (i.e. equal probability of selection). Their 'RandomReplace' algorithm performs better than uniform selection (i.e. selecting each possible committee with equal probability) in managing this trade-off.

We propose a class of mechanisms in the context of citizen juries that are Bayesian incentive-compatible and utilitarian efficient under conditions on the way the committees are selected and on the discount factor of the voters. We do so without requiring money as a numéraire or restrictions on preferences such quasi-linearity. Our results apply to decisions with more than two discrete alternatives.

3 Model and Mechanism

We consider a population of agents denoted by the set $\mathcal{M} = \{1, 2, ...M\}$ from which a committee with a set of voters $\mathcal{I} = \{1, 2, ..I\}$ is selected. The committee is to vote in a set of decisions indexed by $\mathbb{N}$. There are an infinite number of decisions and they are sequential i.e. one decision is voted on and the outcome is decided before the next decision is voted on. For each decision $n \in \mathbb{N}$, the set of alternatives is $\mathcal{A} = \{1, 2, ...A\}$. We assume an equal number of alternatives in each decision for the sake of analytical tractability as it helps us obtain a recursive Bellman equation type expression for a voter's expected utility function. An example of a decision could be an issue such as "Should there be a new metro line introduced from location X to location Y?" where the alternatives would refer to being either for or against it. For any voter i in any decision n, we denote her type by $\theta_n^i \in \Theta$ where Θ is some measurable type space. Her vNM utility function for a decision n is given by $u : \Theta \times \mathcal{A} \to [0, 1]$. We also assume that each voter has the same discount factor $\delta \in [0, 1]$. For each voter, their type is drawn from some distribution after the result of the previous decision and before the voter votes on decision n. Let f be the distribution from which the type of each agent in the entire population is drawn from *independently and identically across decisions and agents.* Moreover, we assume this distribution, the discount factor and the members of the committee are known to everyone. In each decision $n \in \mathbb{N}$, each voter $i \in \mathcal{I}$ is allowed to report a type $\hat{\theta}_n^i \in \Theta$ i.e. a type that may not be their true type.

We next define the class of *Sortition Mechanisms* as voting mechanisms where committee members are randomly selected from a larger population, decisions are made using a voting rule and then committee members are randomly replaced with agents from the larger population before the next decision. More formally, we define a Sortition Mechanism for a set of M agents over an infinite sequence of decisions indexed by $\mathbb{N}$ to be a mechanism $S = (\phi^s, d, \phi^r)$ where $\phi^s : \mathcal{M} \to \Delta(\mathcal{M})$ determines the probability of each agent in the population $\mathcal{M}$ of being selected in the committee before each decision. Moreover, $d : \Theta^I \to \Delta(\mathcal{A})$ is a decision rule that determines the outcome of a decision depending on the committee members' reports, and $\phi^r : \Theta^I \to \Delta(\mathcal{I})$ is a rule for removing current committee members after a decision depending on their reports for the decision.

We are interested in finding Sortition Mechanisms that yield efficient decisions in a Bayesian Nash equilibrium. We follow a utilitarian notion of efficiency as defined below:

Definition 1 *A Sortition Mechanism S is utilitarian efficient if:*

$$d(\boldsymbol{\theta}_n) = \arg\max_{a \in \mathcal{A}} \sum_{i \in \mathcal{I}} u(\theta_n^i, a) \quad \forall\, \boldsymbol{\theta}_n \in \Theta^I,\, n \in \mathbb{N}. \tag{1}$$

In the definition above, $\boldsymbol{\theta}_n = (\theta_n^i)_{i \in \mathcal{I}}$ is the vector of true types in decision n. Therefore, our notion of efficiency is that the winning alternative in each decision should maximize the total welfare of the committee. Utilitarian efficiency is usually obtained by first incentivizing voters to truthfully reveal their type in equilibrium so that the efficient outcome is computed accurately. Therefore, we aim to find Bayesian incentive-compatible Sortition Mechanisms and to define this, let $s_n^i : \Theta \to \Theta$ be a voter i's strategy as a function of her type in decision n and let it belong to the set of strategies available to the voter denoted by Σ. Moreover, let $\boldsymbol{s}_n^i = (s_k^i)_{k=n}^{\infty}$ be the infinite sequence of strategies for i starting from decision n. Let us further denote the case of a truthful sequence of strategies by $\bar{\boldsymbol{s}}_n^i = (\theta_k^i)_{k=n}^{\infty}$. Finally, let us denote her total expected utility from decision n as well as all future decisions by $\bar{U}_n : \Theta^I \times \Sigma^{\{n, n+1, \dots\}} \to \mathbb{R}^+$ where $\Sigma^{\{n, n+1, \dots\}}$ denotes the set of possible strategies in all future decisions combined. This expected utility is dependent on the true type of the voter in that decision as well as her vector of strategies. Moreover, since we are interested in truthful equilibria, this expected utility is computed by fixing the other voters' type reports to truthful reports from decision n onwards. As the types of the future decisions are not known in advance, computation of the expected utility will require the beliefs on the type space set up in the model, thus meaning that we have a Bayesian game. We therefore define Bayesian incentive-compatible Sortition Mechanisms as follows:

Definition 2 *A Sortition Mechanism $S = (\phi^s, d, \phi^r)$ is Bayesian incentive-compatible if:*

$$\bar{U}_n(\theta_n^i, \boldsymbol{\theta}_n^{-i}, \bar{\boldsymbol{s}}_n^i) \geqslant \bar{U}_n(\theta_n^i, \boldsymbol{\theta}_n^{-i}, \boldsymbol{s}_n^i) \quad \forall\, \boldsymbol{s}_n^i \in \Sigma^{\{n, n+1, \dots\}},\, (\theta_n^i, \boldsymbol{\theta}_n^{-i}) \in \Theta^I, n \in \mathbb{N} \text{ and } i \in \mathcal{I}.$$

In the definition above, $\boldsymbol{\theta}_n^{-i} = (\theta_n^j)_{j \neq i,\, j \in \mathcal{I}}$ are the types of the voters other than voter i in decision n. Therefore, a Sortition Mechanism is Bayesian incentive-compatible if every voter maximizes their expected utility by revealing their true type when the other voters are truthful and this is so for each decision.

We will show that a set of mechanisms satisfying Bayesian incentive-compatibility and utilitarian efficiency fall into a sub-class of sortition mechanisms which we will call the class of *Vickrey-Clarke-Groves(VCG) Sortition Mechanisms*. In this class of mechanisms, the decision rule is such that the winning alternative in each decision n is the alternative that maximizes the total reported utility and is given by:

$$a^*(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i}) \quad := \quad \arg\max_{a \in \mathcal{A}} \sum_{i \in \mathcal{I}} u(\hat{\theta}_n^i, a), \tag{2}$$

where $\hat{\boldsymbol{\theta}}_n^{-i} = (\hat{\theta}_n^j)_{j \neq i,\, j \in \mathcal{I}}$ are the types reported by the voters other than voter i in decision n. When the outcome is decided, each voter i is assigned a probability of being removed from the committee and is randomly replaced by someone else from the larger population for the next decision. In the VCG Sortition Mechanisms, this probability of being removed or remaining in the committee is determined by the type reports of the voters in the committee. In case the voter randomly remains in the committee for the next decision, they will have a probability of remaining in the committee two decisions in the future as well, depending on what voters report in the next decision. The mechanism proceeds in a similar manner for all members of the population and all decisions.

Now let $G : \Theta^{I-1} \to \mathbb{R}$ be an arbitrary function (determined by the mechanism designer), and define the "tax" function $t : \Theta^I \to \mathbb{R}$ by:

$$t(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i}) \quad := \quad G(\hat{\boldsymbol{\theta}}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\hat{\theta}_n^j, a^*(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i})). \tag{3}$$

Thus, $t(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i})$ is a function $G(\hat{\boldsymbol{\theta}}_n^{-i})$ that is independent of the reported type of a voter i, minus the total reported utility of the voters *other* than voter i (which is dependent on voter i's reported type through the decision outcome $a^*(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i})$. This is the tax from the Vickrey-Clarke-Groves family of mechanisms so an example would be where $G(\hat{\boldsymbol{\theta}}^{-i}) = \sum_{j \neq i,\, j \in \mathcal{I}} u(\hat{\theta}^j, a^*(\hat{\boldsymbol{\theta}}^{-i}))\ \forall\, \hat{\boldsymbol{\theta}} \in \Theta^I$ where $a^*(\hat{\boldsymbol{\theta}}^{-i})$ is the winner when voter i's report is excluded. This would lead to the tax being the amount by which a voter is pivotal.

Let us also define:

$$\bar{u}_n^i \quad := \quad \mathbb{E}_{(\tilde{\theta}_n^r, \tilde{\boldsymbol{\theta}}_n)}[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^r, \tilde{\boldsymbol{\theta}}_n)], \tag{4}$$

where $\tilde{\theta}_n^i$ is the random variable that generates voter i's true type in decision n, while $\tilde{\boldsymbol{\theta}}_n = (\tilde{\theta}_n^i)_{i \in \mathcal{I}}$ is the random variable for all the committee members in decision n. Thus formula (4) is what voter i's expected utility in decision n would be if voter r were in the committee instead of i, and all the voters in the committee revealed their true types. This expected utility will be the same for any $i \in \mathcal{I}$ and $n \in \mathbb{N}$ as the voters' types are i.i.d across voters and decisions. Therefore, we can drop indices i and n leaving the notation as $\bar{u}$. Moreover, we can use this and the previous definition to define a *"Tax on Probability of Participation (TPP)"* i.e. we define a 'tax' in terms of the probability that a voter will be removed from the committee for the next decision. Each voter's probability of participation is set to 1 before each decision is made and the TPP is subtracted from this probability of participation after the decision is made. The voter then remains in the committee with the new probability of participation. If they are then randomly chosen to remain, then their probability of participation is reset to 1 for the next decision, and the mechanism continues in a fashion similar to the last decision.

We define the Tax on the Probability of Participation $T : \Theta^I \to [0, 1]$ as a specific *scalar multiple of a VCG tax* as follows:

$$T(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i}) \quad := \quad \frac{(1-\delta) t(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i})}{\delta \left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n} \left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}) \right] - \bar{u} \right]}, \quad \text{for all } i \in \mathcal{I}. \tag{5}$$

Therefore, the probability that voter i will remain in the committee after decision n will be $1 - T(\hat{\theta}_n^i, \hat{\boldsymbol{\theta}}_n^{-i})$. As the tax on the probability of participation is a scalar multiple of a VCG tax, a voter is more likely to be removed from the committee if they report more 'intense' preferences i.e. a utility function that assigns a very large value to one alternative relative to the others. Note that this scalar with which the VCG tax is multiplied in (5) will be the same for all voters and decisions because of the i.i.d types and identical discount factors.

Finally, we define a VCG Sortition Mechanism S_{VCG} to be a Sortition Mechanism with $d(\hat{\boldsymbol{\theta}}_n) = a^*(\hat{\boldsymbol{\theta}}_n)$ according to (2) for all $\hat{\boldsymbol{\theta}}_n \in \Theta^I$ and $n \in \mathbb{N}$, and $\phi^r = (T(\hat{\theta}^i_n, \hat{\boldsymbol{\theta}}^{-i}_n))_{i \in \mathcal{I}}$ according to (5) for all $n \in \mathbb{N}$. Therefore, our model considers the class of VCG Sortition Mechanisms which is a subset of the class of Sortition Mechanisms. The class of VCG Sortition Mechanisms would proceed in a manner described by Figure 1.

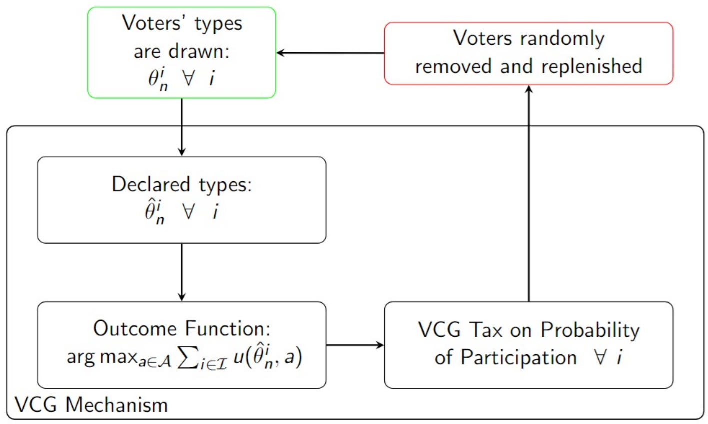

Figure 1: Flowchart of a VCG Sortition Mechanism.

We aim to show that a sub-class of VCG Sortition Mechanisms are Bayesian incentive-compatible (in a manner similar to d'Aspremont and Gérard-Varet (1979)) and efficient under some conditions and an approximation on the way agents are selected into the committee. We consider the case where $M \gg I$ i.e. when the population size is much larger than the committee. In this case, the probability that a voter removed from the committee after a decision will be re-selected into the committee again in the future is negligible. This is akin to small citizen juries selected from a very large population to advise the government. The size of the committee is so small relative to the size of the larger population that the probability of being selected into a committee is negligible. Therefore, the approximation we make is that we neglect the probability of re-selection.

4 Main results

The VCG Sortition Mechanisms assign a probability of that an agent is selected into the committee, but this selection process is exogenously set by the mechanism designer. However, the probability of removal from the committee is set according to (5), so the Tax on the Probability of Participation defined by these equations needs to be between 0 and 1. As we will be examining equilibria with truthful revelation, we first establish conditions under which this is true when voters reveal their types truthfully. These are described in the following lemma, where we establish conditions under which truthful equilibria will have TPP between 0 and 1 for each voter.

Lemma 1 *For any $n \in \mathbb{N}$ and $i \in \mathcal{I}$, $0 \leqslant T(\theta^i_n, \boldsymbol{\theta}^{-i}_n) \leqslant 1$ for all $\boldsymbol{\theta}_n \in \Theta^I$ if and only if the following conditions are satisfied simultaneously:*

a)

$$G(\boldsymbol{\theta}_n^{-i}) \quad \geqslant \quad \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})), \quad \textit{for all } \boldsymbol{\theta}_n \in \Theta^I. \tag{6}$$

b)

$$\bar{u} + \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n^{-i}}[G(\tilde{\boldsymbol{\theta}}_n^{-i})] \quad < \quad \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}))\right]. \tag{7}$$

c) $\delta \geqslant \underline{\delta}$*, for all* $\boldsymbol{\theta}_n \in \Theta^I$*, where:*

$$\underline{\delta} \quad := \quad \frac{t(\theta_n^i, \boldsymbol{\theta}_n^{-i})}{\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] + t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) - \bar{u}}. \tag{8}$$

Therefore, when voters vote truthfully, the TPP function defined by (5) is guaranteed to be between 0 and 1 when G is set appropriately by the mechanism designer according to Lemma 1. Note that if Lemma 1 is satisfied for some $n \in \mathbb{N}$ and $i \in \mathcal{I}$, then it is satisfied for all $n \in \mathbb{N}$ and $i \in \mathcal{I}$ due to the i.i.d types of the voters. Part (a) says G must always be greater than or equal to the total welfare of the other voters in the committee, while part (b) says the expected welfare of the committee must be at least as large as the sum of the expectation of G and a voter's expected utility without participation. Finally, part (c) says the discount factor must be greater than or equal to $\underline{\delta}$, which is itself determined by the choice of G. The proof of Lemma 1 has been left in the appendix, where we also show that $\underline{\delta} < 1$ if G satisfies (6) and (7). Therefore, there will always be some $0 \leqslant \delta \leqslant 1$ such that the TPP is between 0 and 1 when the voters are truthful and G satisfies (6) and (7).

4.1 Efficiency of VCG Sortition Mechanisms

We now proceed to characterize the incentive compatibility and efficiency of a sub-class of VCG Sortition Mechanisms which satisfy the conditions in Lemma 1 under our approximation on the probability of future re-selection into the committee after being removed. For the sake of describing this approximation, let $\phi^s(m)$ denote the probability that an agent $m \in \mathcal{M}$ will be selected into the committee for a decision. To obtain our results, we consider the case that $\lim_{M \to \infty} \phi^s(m) = 0$ for all $m \in \mathcal{M} \backslash \mathcal{I}$ i.e. the committee selection process is such that the probability that an agent that was not in the committee in the previous decision will be selected into the committee for the forthcoming decision, tends to zero for a large population size. Therefore, we consider the case of a large population relative to the size of the committee i.e. $M \gg I$ and we neglect this probability of selection altogether in our computations. By doing so, we can compute a voter i's expected utility when the other voters are truthful in today's as well as all future decisions. Let us consider a voter i in the committee voting on some decision with an index n. This can be thought of as the decision to be made today. Then after neglecting the probability of re-selection into the committee if she is removed at some point, the voter i's expected utility for participation in the committee in round n is:

$$\bar{U}_n(\theta_n^i, \boldsymbol{\theta}_n^{-i}, \boldsymbol{s}_n^i) = u(\theta_n^i, a^*(s_n^i(\theta_n^i), \boldsymbol{\theta}_n^{-i})) + (1 - T(s_n^i(\theta_n^i), \boldsymbol{\theta}_n^{-i}))\delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}_{n+1}}[\bar{U}_{n+1}(\tilde{\theta}_{n+1}^i, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}, \boldsymbol{s}_{n+1}^i)]$$
$$+ T(s_n^i(\theta_n^i), \boldsymbol{\theta}_n^{-i}) \frac{\delta}{1-\delta} \mathbb{E}_{(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1})}[u(\tilde{\theta}_{n+1}^i, a^*(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}))]. \tag{9}$$

The first term on the right-hand side of (9) is voter i's utility for today's decision, which depends on the winning alternative of the decision. This in turn depends on the strategy of voter i. The second term is the probability of remaining in the committee for decision $n+1$ (given by $(1 - T(s_n^i(\theta_n^i), \boldsymbol{\theta}_n^{-i}))$) times the discounted expected utility from participating in the committee in round $n+1$ ($\delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}_{n+1}}[\bar{U}_{n+1}(\tilde{\theta}_{n+1}^i, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}, \boldsymbol{s}_{n+1}^i)]$). Finally, the last term is the discounted expected utility from all future decisions if the voter is *removed* from the committee for decision $n+1$. It is $T(s_n^i(\theta_n^i), \boldsymbol{\theta}_n^{-i})$ (i.e. the probability of being removed), times the discounted future expected utility from all future decisions. If the voter is removed from the committee, she will have a negligible probability of being re-selected, so her future discounted utility, in this case, will be her future discounted expected without her participation in any future decision. However, her expected utility in each decision without her participation is the same because of the i.i.d types. Therefore, the future discounted expected utility without participation boils down to $\delta \mathbb{E}_{(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1})}[u(\tilde{\theta}_{n+1}^i, a^*(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}))] + \delta^2 \mathbb{E}_{(\tilde{\theta}_{n+2}^r, \tilde{\boldsymbol{\theta}}_{n+2})}[u(\tilde{\theta}_{n+2}^i, a^*(\tilde{\theta}_{n+2}^r, \tilde{\boldsymbol{\theta}}_{n+2}^{-i}))] + \ldots = \frac{\delta}{1-\delta} \mathbb{E}_{(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1})}[u(\tilde{\theta}_{n+1}^i, a^*(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}))]$. From the Bellman equation in (9), we see that any voter i will have the same expected utility function in each decision except for a change in indices. Therefore for simplicity, we can drop the index to obtain and replace $\mathbb{E}_{(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1})}[u(\tilde{\theta}_{n+1}^i, a^*(\tilde{\theta}_{n+1}^r, \tilde{\boldsymbol{\theta}}_{n+1}^{-i}))]$ with $\bar{u}$ to obtain an expected utility function of the form:

$$\bar{U}(\theta^i, \boldsymbol{\theta}^{-i}, \boldsymbol{s}^i) = u(\theta^i, a^*(s^i(\theta^i), \boldsymbol{\theta}^{-i})) + (1 - T(s^i(\theta^i), \boldsymbol{\theta}^{-i}))\delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i}, \boldsymbol{s}^i)]$$
$$+ \frac{\delta T(s^i(\theta^i), \boldsymbol{\theta}^{-i})\bar{u}}{1-\delta}. \tag{10}$$

A voter i voting in any decision maximizes their expected utility given in (10) with respect to their strategy $\boldsymbol{s}^i \in \Sigma^{\{1,2,\ldots\}}$. In the following main result, we show that for each voter, revealing their true type in each decision maximizes their expected utility. As a consequence, the sub-class of VCG Sortition Mechanisms is efficient. Here is our main result.

Theorem 1 *Define G according to (6) and (7), and $\underline{\delta}$ according to (8). If $\delta \geqslant \underline{\delta}$, then VCG Sortition Mechanisms are Bayesian incentive-compatible and utilitarian-optimal.*

Therefore, VCG Sortition Mechanisms with the appropriately defined G are incentive-compatible and efficient when the discount factor is large enough.

4.2 Efficiency of the Pivotal Sortition Mechanism

Theorem 1 characterizes a class of VCG Sortition Mechanisms that are efficient. It can be observed that setting $G(\boldsymbol{\theta}^{-i}) = \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta^j, a^*(\boldsymbol{\theta}^{-i}))\ \forall\, \boldsymbol{\theta} \in \Theta^I$ where $a^*(\boldsymbol{\theta}^{-i})$ is the winner when voter i's report is excluded, satisfies (6) and (7). We call this the *Pivotal Sortition Mechanism* because according to the definition of the TPP in (5), a voter will only have a strictly positive tax on their probability of participating in the next decision if they are pivotal in changing the outcome. We state a corollary result on this following Theorem 1.

Corollary 1 *Define $\underline{\delta}$ according to (8). If $\delta \geqslant \underline{\delta}$, then the Pivotal Sortition Mechanism is Bayesian incentive-compatible and utilitarian-optimal.*

Proof: Note that $\sum_{j \neq i,\, j \in \mathcal{I}} u(\theta^j, a^*(\boldsymbol{\theta}^{-i})) \geqslant \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta^j, a^*(\theta^i, \boldsymbol{\theta}^{-i}))$ for all $\boldsymbol{\theta} \in \Theta^I$, because the total utility of the committee members other than voter i when voter i's report is excluded must be at least as large as their total utility with her participation according to the definition of $a^*(\theta^i, \boldsymbol{\theta}^{-i})$. Therefore, condition (6) is satisfied. Moreover, we have:

$$
\begin{aligned}
G(\boldsymbol{\theta}^{-i}) &= \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta^j, a^*(\boldsymbol{\theta}^{-i})) \text{ for all } \boldsymbol{\theta} \in \Theta^I \\
\Rightarrow \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[G(\tilde{\boldsymbol{\theta}}^{-i})] &= \mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[\sum_{j \neq i,\, j \in \mathcal{I}} u(\tilde{\theta}^j, a^*(\tilde{\boldsymbol{\theta}}^{-i}))\right] \\
\Rightarrow \bar{u} + \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[G(\tilde{\boldsymbol{\theta}}^{-i})] &\underset{*}{\leqslant} \mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}^i, a^*(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i}))\right].
\end{aligned}
$$

where we get (*) because the LHS is the expected total utility of the committee if one voter's report is excluded and RHS is the expected total utility when everyone's reports are included. The expected total utility when everyone's reports are included must be greater so our choice of G satisfies (7) as well. Since conditions (6) and (7) are satisfied, Corollary 1 follows from Theorem 1. □

Besides being incentive-compatible and efficient, the Pivotal Sortition Mechanism has the property that a voter will face a zero tax on the probability of participation unless they are pivotal in changing the outcome. This may be preferred to an alternative definition of G where a voter would face a TPP even when they have not influenced the outcome.

5 Conclusion

We consider a case akin to a citizen jury, where a committee faces an infinite sequence of future decisions. The committee is randomly selected from a larger population - a process known as sortition - and asked to deliver a verdict on an issue. We propose a sub-class of VCG Sortition Mechanisms to allow voters to express their cardinal preferences and incentivize truthful voting behaviour at a Bayesian Nash equilibrium when the discount factor is large enough. Consequently, the winning alternative at a Bayesian Nash equilibrium will be utilitarian efficient. The Pivotal Sortition Mechanism is an example of a mechanism with these properties as well as the additional property that a voter will only face a TPP when they are pivotal in changing the outcome. We obtain our results under an approximation on the way voters are selected into the committee. We assume that the size of the committee is much smaller than that of the larger population and that once a voter is removed from the committee, their probability of being re-selected into the committee is negligible. Therefore, an interesting avenue of improvement over this work could be the characterization of VCG Sortition Mechanisms without assuming a very large population and neglecting the probability of re-selection. VCG Sortition mechanisms that are efficient in such a general scenario may be suitable for a wider variety of applications as well such as committees in organizations tasked with hiring new members or committees generally making decisions on the well-being of their organization.

A Appendix

Proof of Lemma 1 From (5), we see that for a particular choice of the function G, we will have $T(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \geqslant 0$ for all $\boldsymbol{\theta}_n \in \Theta^I$ if one of the following two cases holds:

Case I. $t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \geqslant 0$ for all $\boldsymbol{\theta}_n \in \Theta^I$, and $\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} > 0$.

Case II. $t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) < 0$ for all $\boldsymbol{\theta}_n \in \Theta^I$, and $\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} < 0$.

We shall deal with these cases separately.

Case I: $t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \geqslant 0$ for all $\boldsymbol{\theta}_n \in \Theta^I$, and $\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} > 0$.

In this case, we have:

$$t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \geqslant 0 \text{ for all } \boldsymbol{\theta}_n \in \Theta^I \underset{(3)}{\Longleftrightarrow} G(\boldsymbol{\theta}_n^{-i}) \geqslant \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})) \text{ for all } \boldsymbol{\theta}_n \in \Theta^I, \tag{A1}$$

which is inequality (6) in part (a) of the lemma. Moreover, we have:

$$\begin{aligned}
&\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} > 0 \\
\underset{(3)}{\Longleftrightarrow} &\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - G(\tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} > 0 \\
\Longleftrightarrow &\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}[G(\tilde{\boldsymbol{\theta}}_n^{-i})] + \bar{u} < \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}))\right],
\end{aligned} \tag{A2}$$

which is inequality (7) in part (b) of the lemma.

Case II: $t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) < 0 \,\forall\, \boldsymbol{\theta}_n \in \Theta^I$, and $\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} < 0$.

In this case, we have:

$$\begin{aligned} t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) < 0 \,\forall\, \boldsymbol{\theta}_n \in \Theta^I \quad &\underset{(3)}{\iff} \quad G(\boldsymbol{\theta}_n^{-i}) < \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})) \text{ for all } \boldsymbol{\theta}_n \in \Theta^I \\ &\Rightarrow \quad \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}[G(\tilde{\boldsymbol{\theta}}_n^{-i})] < \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{j \neq i,\, j \in \mathcal{I}} u(\tilde{\theta}_n^j, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}))\right] \\ &\Rightarrow \quad \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}[G(\tilde{\boldsymbol{\theta}}_n^{-i})] + \bar{u} < \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}))\right] \end{aligned} \tag{A3}$$

where the last inequality is due to the fact that a voter i's expected utility with their participation in the decision must be greater than their expected utility without their participation.

The second condition of Case II implies that:

$$\begin{aligned} &\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} < 0 \\ \underset{(3)}{\iff} &\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - G(\tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u} < 0 \\ \iff &\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}[G(\tilde{\boldsymbol{\theta}}_n^{-i})] + \bar{u} > \mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i}))\right]. \end{aligned} \tag{A4}$$

From the definitions of Cases I and II, no G can satisfy Case I for some type realizations and Case II for other type realizations. Moreover, one can observe that (A3) and (A4) cannot be satisfied simultaneously. Therefore, there is no G satisfying Case II and the only way to ensure a positive TPP is by selecting a G that satisfies Case I. This proves parts (a) and (b).

To prove (c), note that

$$\begin{aligned} &T(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \leqslant 1, \text{ for all } \boldsymbol{\theta}_n \in \Theta^I \\ \underset{(*)}{\iff} &\frac{(1-\delta)\left[G(\boldsymbol{\theta}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i}))\right]}{\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - G(\tilde{\boldsymbol{\theta}}_n^{-i})\right] - \bar{u}\right]} \leqslant 1, \text{ for all } \boldsymbol{\theta}_n \in \Theta^I \\ \iff &G(\boldsymbol{\theta}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})) \\ &\leqslant \delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - G(\tilde{\boldsymbol{\theta}}_n^{-i})\right] + G(\boldsymbol{\theta}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})) - \bar{u}\right] \\ \iff &\delta \geqslant \underline{\delta}, \text{ for all } \boldsymbol{\theta}_n \in \Theta^I, \end{aligned}$$

where ($*$) is by defining formulae (3) and (5), and where

$$\underline{\delta} = \frac{G(\boldsymbol{\theta}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i}))}{\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[\sum_{i \in \mathcal{I}} u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - G(\tilde{\boldsymbol{\theta}}_n^{-i})\right] + G(\boldsymbol{\theta}_n^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta_n^j, a^*(\theta_n^i, \boldsymbol{\theta}_n^{-i})) - \bar{u}}. \tag{A5}$$

One can easily substitute the conditions in (A1) and (A2) to show that $\underline{\delta} \leqslant 1$, for all $\boldsymbol{\theta}_n \in \Theta^I$. This means that there will always be some $\delta \in [0,1]$ such that $0 \leqslant T(\theta_n^i, \boldsymbol{\theta}_n^{-i}) \leqslant 1$ for all $\boldsymbol{\theta}_n \in \Theta^I$ when (A1) and (A2) are satisfied. Moreover, substituting defining formula (3) into equation (A5), we get:

$$\underline{\delta} = \frac{t(\theta_n^i, \boldsymbol{\theta}_n^{-i})}{\mathbb{E}_{\tilde{\boldsymbol{\theta}}_n}\left[u(\tilde{\theta}_n^i, a^*(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})) - t(\tilde{\theta}_n^i, \tilde{\boldsymbol{\theta}}_n^{-i})\right] + t(\theta_n^i, \boldsymbol{\theta}_n^{-i}) - \bar{u}}, \tag{A6}$$

which is the expression in (8). □

Proof of Theorem 1

a) Let us consider a Tax on the Probability of Participation of the following form:

$$\hat{T}(\hat{\theta}^i, \hat{\boldsymbol{\theta}}^{-i}) = W\, t(\hat{\theta}^i, \hat{\boldsymbol{\theta}}^{-i}) \quad \forall\, i \in \mathcal{I}, \tag{A7}$$

where $W \in \mathbb{R}$. Moreover, let us set voter i's strategy in all future decisions to truthful revelation i.e. $s_m^i(\theta_m^i) = \theta_m^i$, for all $m > n$ where n is today's decision. We shall test whether deviating from truthful revelation in today's decision is profitable when the TPP is set according to (5). If it is not, then this will be so for all future decisions as well due to recursive expected utility function in (10).

Now if we set truthful revelation in future decisions, then according to (10), a voter i's expected utility will be of the form:

$$\begin{aligned}
\bar{U}^i(\theta^i, \boldsymbol{\theta}^{-i}, s^i) &= u(\theta^i, a^*(s^i(\theta^i), \boldsymbol{\theta}^{-i})) + (1 - \hat{T}(s^i(\theta^i), \boldsymbol{\theta}^{-i}))\delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] \\
&\quad + \hat{T}(s^i(\theta^i), \boldsymbol{\theta}^{-i}) \frac{\delta}{1-\delta} \mathbb{E}_{(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}})}[u(\tilde{\theta}^i, a^*(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}}^{-i}))] \\
&\underset{(*)}{=} u(\theta^i, a^*(s^i(\theta^i), \boldsymbol{\theta}^{-i})) + \delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] \\
&\quad - W\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] - \frac{1}{1-\delta}\mathbb{E}_{(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}})}[u(\tilde{\theta}^i, a^*(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}}^{-i}))]\right] \\
&\quad \left[G(\boldsymbol{\theta}^{-i}) - \sum_{j \neq i,\, j \in \mathcal{I}} u(\theta^j, a^*(s^i(\theta^i), \boldsymbol{\theta}^{-i}))\right],
\end{aligned} \tag{A8}$$

where (*) follows from (A7) and (3). One can observe that if we set:

$$W = \frac{1}{\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] - \frac{1}{1-\delta}\mathbb{E}_{(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}})}[u(\tilde{\theta}^i, a^*(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}}^{-i}))]\right]},$$

we get:

$$\bar{U}^i(\theta^i, \boldsymbol{\theta}^{-i}, s^i) = \sum_{j \in \mathcal{I}} u(\theta^j, a^*(s^i(\theta^i), \boldsymbol{\theta}^{-i})) - G(\boldsymbol{\theta}^{-i}) + \delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})]. \tag{A9}$$

One can observe from (A9) that a voter's expected utility function is the social welfare function maximized by the outcome (2) plus a constant. Therefore, it is a voter's optimal strategy to reveal their true type i.e. to set $s^i(\theta^i) = \theta^i \ \forall \ i \in \mathcal{I}$.

Moreover for any voter, we can further compute W from (A9) in the following manner:

$$\begin{aligned}
\bar{U}^i(\theta^i, \boldsymbol{\theta}^{-i}) &= \sum_{j \in \mathcal{I}} u(\theta^j, a^*(\theta^i, \boldsymbol{\theta}^{-i})) - G(\boldsymbol{\theta}^{-i}) + \delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] \\
\Rightarrow \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] &= \mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[\sum_{j \in \mathcal{I}} u(\tilde{\theta}^j, a^*(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})) - G(\tilde{\boldsymbol{\theta}}^{-i})\right] + \delta \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i] \\
\Rightarrow \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] &= \frac{\mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[\sum_{j \in \mathcal{I}} u(\tilde{\theta}^j, a^*(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})) - G(\tilde{\boldsymbol{\theta}}^{-i})\right]}{1-\delta}.
\end{aligned} \tag{A10}$$

By substituting (A10) into the expression for W we get:

$$\begin{aligned}
\Rightarrow W &= \frac{1}{\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\bar{U}^i(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})] - \frac{1}{1-\delta}\mathbb{E}_{(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}})}[u(\tilde{\theta}^i, a^*(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}}^{-i}))]\right]} \\
&= \frac{1-\delta}{\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[\sum_{j \in \mathcal{I}} u(\tilde{\theta}^j, a^*(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})) - G(\tilde{\boldsymbol{\theta}}^{-i})\right] - \mathbb{E}_{(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}})}[u(\tilde{\theta}^i, a^*(\tilde{\theta}^r, \tilde{\boldsymbol{\theta}}^{-i}))]\right]} \\
&= \frac{(1-\delta)}{\delta\left[\mathbb{E}_{\tilde{\boldsymbol{\theta}}}\left[u(\tilde{\theta}^i, a^*(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})) - t(\tilde{\theta}^i, \tilde{\boldsymbol{\theta}}^{-i})\right] - \bar{u}\right]},
\end{aligned} \tag{A11}$$

where (A11) follows from (3) and (4). Finally, substituting (A11) in (A7), we see that our TPP in (5) achieves Bayesian incentive-compatibility when the conditions in Lemma 1 are satisfied. This is because a voter will not have an incentive to deviate from truthful revelation in any decision.

b) Since $\hat{\theta}^i = \theta^i$ for all $i \in \mathcal{I}$ at a Bayesian Nash equilibrium, we see that the outcome according to (2) will maximize social welfare at a Bayesian Nash equilibrium.

□

References

Anagnostides, I., Fotakis, D., Patsilinakos, P., 2022. Metric-Distortion Bounds under Limited Information. Journal of Artificial Intelligence Research 74, 1449–1483.

Arrow, K.J., 1979. The property rights doctrine and demand revelation under incomplete information, in: Boskin, M.J. (Ed.), Economics and human welfare. Elsevier, pp. 23–39.

Azrieli, Y., Kim, S., 2014. Pareto efficiency and weighted majority rules. International Economic Review 55, 1067–1088.

Ball, I., Kattwinkel, D., 2024. Quota Mechanisms: Robustness and Finite-Sample Performance. arXiv preprint arXiv:2309.07363 .

Barberà, S., Jackson, M., 2006. On the Weights of Nations: Assigning Voting Weights in a Heterogeneous Union. Journal of Political Economy 114, 317–339.

Benadè, G., Gölz, P., Procaccia, A.D., 2019. No Stratification Without Representation, in: Proceedings of the 2019 ACM Conference on Economics and Computation, pp. 281–314.

Benjamin, D., Heffetz, O., Kimball, M., Lougee, D., 2017. The relationship between the normalized gradient addition mechanism and quadratic voting. Public Choice 172, 233–263.

Benjamin, D.J., Heffetz, O., Kimball, M.S., Szembrot, N., 2013. Aggregating local preferences to guide marginal policy adjustments. American Economic Review 103, 605–610.

Caragiannis, I., Micha, E., Peters, J., 2024. Can a Few Decide for Many? The Metric Distortion of Sortition. arXiv preprint arXiv:2406.02400 .

Casella, A., 2005. Storable votes. Games and Economic Behavior 51, 391–419.

Casella, A., Macé, A., 2021. Does vote trading improve welfare? Annual Review of Economics 13, 57–86.

Casella, A., Sanchez, L., 2022. Democracy and Intensity of Preferences: A Test of Storable Votes and Quadratic Voting on Four California Propositions. The Journal of Politics 84, 607–612.

Clarke, E.H., 1971. Multipart pricing of public goods. Public Choice 11, 17–33.

d'Aspremont, C., Gérard-Varet, L.A., 1979. Incentives and incomplete information. Journal of Public Economics 11, 25–45.

Drexl, M., Kleiner, A., . Preference Intensities in Repeated Collective Decision-Making. Kiel und Hamburg: ZBW-Deutsche Zentralbibliothek .

Drexl, M., Kleiner, A., 2018. Why Voting? A Welfare Analysis. American Economic Journal: Microeconomics 10, 253–271.

Ebadian, S., Kehne, G., Micha, E., Procaccia, A.D., Shah, N., 2022. Is sortition both representative and fair? Advances in Neural Information Processing Systems 35, 3431–3443.

Ebadian, S., Micha, E., 2024. Boosting Sortition via Proportional Representation. arXiv preprint arXiv:2406.00913 .

Eguia, J.X., Immorlica, N., Lalley, S.P., Ligett, K., Weyl, G., Xefteris, D., 2023. Efficiency in Collective Decision-Making via Quadratic Transfers. arXiv preprint arXiv:2301.06206 .

Eguia, J.X., Xefteris, D., 2021. Implementation by Vote-Buying Mechanisms. The American economic review 111, 2811–2828.

Flanigan, B., Gölz, P., Gupta, A., Procaccia, A.D., 2020. Neutralizing self-selection bias in sampling for sortition. Advances in Neural Information Processing Systems 33, 6528–6539.

Flanigan, B., Gölz, P., Gupta, A., Hennig, B., Procaccia, A.D., 2021a. Fair algorithms for selecting citizens' assemblies. Nature 596, 548–552.

Flanigan, B., Kehne, G., Procaccia, A.D., 2021b. Fair sortition made transparent. Advances in Neural Information Processing Systems 34, 25720–25731.

Georgescu, L., Fox, J., Gautier, A., Wooldridge, M., 2024. Fixed-budget and multiple-issue quadratic voting.

Gershkov, A., Moldovanu, B., Shi, X., 2016. Optimal Voting Rules. The Review of Economic Studies 84, 688–717.

Ghosh, R., Pivato, M., 2024. Quadratically normalized utilitarian voting .

Goeree, J.K., Zhang, J., 2017. One man, one bid. Games and economic behavior 101, 151–171.

Good, I.J., 1977. Justice in voting by demand revelation. Public Choice 29, 65–70.

Gorokh, A., Banerjee, S., Iyer, K., 2021. From monetary to nonmonetary mechanism design via artificial currencies. Mathematics of Operations Research 46, 835–855.

Groves, T., 1973. Incentives in Teams. Econometrica 41, 617.

Groves, T., Loeb, M., 1975. Incentives and public inputs. Journal of Public Economics 4, 211–226.

Grüner, H.P., Tröger, T., 2019. Linear Voting Rules. Econometrica 87, 2037–2077.

Hortala-Vallve, R., 2010. Inefficiencies on linking decisions. Social Choice and Welfare 34, 471–486.

Hortala-Vallve, R., 2012. Qualitative voting. Journal of Theoretical Politics 24, 526–554.

Hylland, A., Zeckhauser, R., 1980. A Mechanism for Selecting Public Goods When Preferences Must be Elicited .

Jackson, M.O., Sonnenschein, H.F., 2007. Overcoming Incentive Constraints by Linking Decisions. Econometrica 75, 241–257.

Laine, C.R., 1977. Strategy in Point Voting: A Note*. The Quarterly Journal of Economics 91, 505.

Lalley, S.P., Weyl, E.G., 2019. Nash Equilbria for Quadratic Voting. arXiv preprint arXiv:1409.0264 .

Laurence, B., Sher, I., 2017. Ethical considerations on quadratic voting. Public Choice 172, 195–222.

Macías, A., 2024. Storable votes with a "pay as you win" mechanism. Journal of Economic Interaction and Coordination 19, 121–150.

Meir, R., Sandomirskiy, F., Tennenholtz, M., 2021. Representative Committees of Peers. Journal of Artificial Intelligence Research 71, 401–429.

Patty, J.W., Penn, E.M., 2017. Uncertainty, polarization, and proposal incentives under quadratic voting. Public Choice 172, 109–124.

Posner, E.A., Weyl, E.G., 2017. Quadratic voting and the public good: introduction. Public Choice 172, 1–22.

Saran, R., Tumennasan, N., 2013. Whose opinion counts? Implementation by sortition. Games and Economic Behavior 78, 72–84.

Saran, R., Tumennasan, N., 2019. Eliciting and Aggregating Information By Sortition in Collective Choice. The Economic Journal 129, 1924–1952.

Schmitz, P.W., Tröger, T., 2012. The (sub-)optimality of the majority rule. Games and Economic Behavior 74, 651–665.

Tideman, N., Plassmann, F., 2017. Efficient collective decision-making, marginal cost pricing, and quadratic voting. Public Choice 172, 45–73.

Weyl, E.G., 2012. Quadratic Vote Buying. Unpublished, University of Chicago .

Weyl, E.G., 2017. The robustness of quadratic voting. Public Choice 172, 75–107.