Strategy-proof Participatory Budgeting via a VCG-like Mechanism

Abstract

We study a divisible Participatory Budgeting model where agents determine both the total volume of expanses and the specific allocation. Our proposed mechanism is constructed as a modification of VCG to a less typical environment, namely where we do not assume quasi-linear utilities nor direct revelation. We show that it satisfies strategy-proofness in strictly dominant strategies. While incentives alignment in our mechanism, as in classic VCG, is achieved via individual payments we charge from agents, in a PB context that seems unreasonable. Our second main result thus provides that, under further specifications relevant in that context, these payments will vanish in large populations. In the last section we expand the mechanism's definition to a class of mechanisms in which the designer can prioritize certain outcomes she sees as desirable.

1 Introduction

With growing popularity worldwide, Participatory Budgeting[1] (PB) lets residents of a municipal authority to collectively decide on the allocation of public resources [1]. In accordance with most real-life PB implementations, much of the technical PB literature is dedicated to budget allocations among public projects with indivisible costs ([29, 2], among many others). Here, we study a 'divisible' PB model [17, 15, 27] where costs are completely flexible. Intuitively, that model describes more of a high level division of the budget among several public goods or city departments, e.g. education, transportation, public parks and so on, rather than concrete projects.[2] Moreover, in our model agents also determine the total volume of expanses they wish to fund through taxation.

Formally, a population of n agents is facing the following budgeting problem. There exist an initial budget of $B_0 \geq 0$ available for the funding of $m \geq 1$ public goods. The *budget decision* we seek is a pair (x, t) where $t \in [-\frac{B_0}{n}, \infty)$ is a monetary sum ("tax") that each agent adds to (or subtracts from) the budget, and $x \in \Delta^m := \{x \in R^n | x_j \geq 0 \; \forall j, \; \sum_j x_j = 1\}$ represents the allocation of the resulting budget $B_t = B_0 + nt$ among the m goods. Note that we allow $t < 0$, meaning that some of the initial budget B_0 will be distributed equally among agents rather than fund public expenditures. We are interested in constructing a collective decision mechanism to which every agent submits her preferred budget decision, and one that is incentive compatible (IC) particularly. Table 1 demonstrates how agents might report to such mechanism their preferences for budget allocation among three municipal services. Note e.g. that agent a suggests an individual payment of $t = 20$ each, and so to allocate a total of $B_0 + 3t = 100 + 3 \cdot 20 = 160$. Also note that agents b and c propose the same normalized allocation $x_b = x_c \in \Delta^3$ but differ in their preferred tax.

1.1 Contribution and Paper Structure

The Inroduction will proceed with a discussion on our modeling choices and its relation to existing literature. After introducing the full model in Sections 2, in Section 3 we introduce

[1] 'The Participatory Budget Project': `https://www.participatorybudgeting.org/`

[2] Although examples of specific projects that can be implemented fractionally do also exist - `https://pbstanford.org/boston16internal/knapsack`.

	Tax (t)	Education (x_1)	Parks (x_2)	Transport (x_3)
agent a	0	60 (0.6)	30 (0.3)	10 (0.1)
agent b	20	32 (0.2)	48 (0.3)	80 (0.5)
agent c	-20	8 (0.2)	12 (0.3)	20 (0.5)

Table 1: An example voting profile for three alternatives and three agents, with an initial budget $B_0 = 100$.

a welfare maximizing, strategy-proof in *strictly* dominant strategies (SDSIC) mechanism for the above problem. As an immediate corollary of SDSIC, it implements the social optimum as a Coalition-Proof Nash Equilibrium [7]. Conceptually, we see two main advantages in including the tax policy in the collective decision: granting the community with higher control over economic policies that shape their reality; and providing valuable feedback to the social planner on the real intensity to which the community values public expanses. Without it, classic PB only informs us on the relative preference among projects. Essentially, the mechanism we introduce is a VCG mechanism adapted to our PB environment, an adaptation that is not straight-forward for mainly two reasons. Most importantly, because our utility model does not satisfy quasi-linearity. Section 3.1 shows how that can be tackled using the information entailed in agents' reports, that in our case also includes their preferences regarding monetary transfers. Beyond that, VCG involves monetary transfers among agents, that seems inappropriate in a PB context. Thus, in Section 4 we show that these modified "VCG payments" we construct vanish in large populations, under the further assumption that agents utility is a function of per capita spending. It is important to note that in our setting the collectively chosen outcome itself includes homogeneous tax payments that finance public expenditures, whereas the (heterogeneous) modified VCG payments are charged on top of that, only to serve the purpose of aligning incentives. Our desire is to diminish the payments of the latter type. Then in section 5, we expand the mechanism to a class of mechanisms that insert a bias towards outcomes that the designer may see as favourable.

1.2 Our Modelling Approach

For its most part, the divisible PB literature adopts or generalizes the same assumptions used for PB with indivisible projects, which in turn stem from standard modelling assumptions in voting and computational social choice. In particular, strategy-proof (by its full meaning) mechanisms have so far been proposed for two types of utility models, namely, the spatial model [19, 17] and the binary additive model [10, 3]. However, we argue that divisible budget allocation is much closer in nature to problems typically treated within microeconomic framework [28, 34, 6]. Hence, it makes more sense to adopt conventional economic assumptions regarding demand and utility, as found in [15] and as we do here.[3] That includes:

- *Additive concave utilities model* [15, 21, 31], that is, that agents' utility is expressed as $U_i(X) = \sum_j \alpha_{i,j}\theta_j(X_j)$ where X_j where X_j is the amount spent on public good j and θ_j is an increasing concave function for all j, (smoothly) expressing the plausible assumption of decreasing marginal gains. ($\alpha_{i,j}$ are scalars that vary between agents). As we assume that part of the budget is collected via a tax-per-agent t, our model adds on the above the disutility of a voter from the tax payment.
- *Optimal points characterized by the MRS conditions*, that follows from the concavity and some additional conventions on utilities [22].

[3] For a thorough discussion on existing models' limitations see the full version of the paper [32].

- *Utility depends on public investment per capita.* (that we add to the model in Section 4). On a large scale, it is reasonable that the quality of public goods depends more on spending per capita rather than on the nominal amount.[4]

In contrast, *elicitation* is an issue that has received much more attention in the literature of mechanism design and computational social choice than in microeconomics. For example there is a live discussion in the context of indivisible PB on the tradeoff between expressiveness of the ballot and effort required for elicitation [16, 5]. Similarly, we argue that it does not make sense to assume that we have direct access to voters' preferences, and here we adopt from computational social choice the assumptions that voters simply report their most preferred allocation, as in [17].

In terms of applicability, however, the obvious shortcoming of our utility model is that it requires us to explicitly specify the concave functions $\{\theta_j\}_j$, as well as the disutility from monetary losses we add, which are fairly abstract. Importantly, we **do not** assume that agents are 'aware' of their assumed utility function, but, conventionally, only know their preferences regarding the decision space, that presumably can be interpreted as derived from an underlying utility model [22]. Of course, any such model would be an approximation at best. Nevertheless, it is fair to assume that any choice of monotonically increasing concave functions probably better approximates individuals' preferences—and thereby incentives—than the spatial model or a linear additive model [11].

1.3 Further related literature

The Economic literature on public goods markets, equilibria and optimal taxation is abundant. ([28, 31, 6], to name just a few). While our work adopts a similar approach to modelling and also optimizes social welfare, this branch of the literature rarely discusses mechanisms. One exception that we know of is found in [23], in which the socially optimal outcome is implemented in strictly dominant strategies using a method very similar to ours, however for quite a different utility model. To the best of our knowledge, the only existing PB mechanism that included tax in the collective decision previously to ours was studied by Garg et al. [18] in the context of experimenting *'iterative voting'* mechanisms. Interestingly, it may suggest some supporting evidence in favour of the additive concave utility model over spatial models in that context. Two other previous works [1, 11] incorporated private funding into a PB model, albeit in the form of voluntary donations that every agent can choose freely and not as a collectively decided sum that is collected from (or paid to) everyone, as we consider here.

The literature on divisible PB is relatively narrow. In terms of incentive compatibility, [17] and [10] presented the soundest results, under a spatial and binary utility models, respectively. For additive concave utilities, [15] and [19] introduce mechanisms that satisfy some relaxations of strategy-proofness. Aziz et al. [3] presented IC mechanisms for an additive linear model similar to that in [10], although their model is primarily motivated by randomized approval mechanisms. A similar utility model is also found in [29]. Overall, in relation to the divisible PB field, this work offers an SDSIC mechanism under concave additive utilities, to the best of our knowledge for the first time.

In the VCG area, [24] and [25] also deal with generalizations of VCG to non quasi-linear utility models, however quite different than the one we consider here. Moreover, our desire for diminishing the (modified) VCG payments resembles the idea of *redistribution* in mechanism design [12, 20]. Such methods are especially relevant in a discrete decision space and can eliminate surplus only partially, while in our model the complete (asymptotic) vanishing is much thanks to the continuity of the decision space.

[4] See for example `https://data.oecd.org/gga/general-government-spending.htm`, and [28].

Much of the PB literature deals with the concept of *fair allocations* [15, 27, 2]. While not a primary goal of our model, we show that the designer can bias the allocation closer to a favorable allocation—including one they see as fair.

2 Model and Preliminaries

We denote by Δ^m the set of distributions over m elements, and use $[m]$ as a shortcut for $\{1, \ldots, m\}$. A set of n agents (voters) need to collectively reach a *budget decision* (x, t) described as follows. $t \in \mathbb{R}$ is a lump-sum tax collected from every agent. $x \in \Delta^m$ is an allocation of the total available budget $B_t := B_0 + nt$ among some m pre-given public goods, where B_0 is an external, non tax-funded source of public funds. t is restricted only by $t > -\frac{B_0}{n}$, meaning that voters can decide either to finance a budget larger than B_0 through positive taxation, or allocate some of it to themselves directly (negative taxation). The collective decision is taken through some voting mechanism to which every agent submits her most preferred budget decision $(x^{(i)}, t^{(i)})$.

2.1 Preferences

We introduce here a concise description of the utility model we use. For a broader discussion on its justification and possible relaxations see the full version of the paper [32]. Every agent i is associated with a *type* $\alpha_i = (\alpha_{i,1}, \ldots, \alpha_{i,m}, \alpha_{i,f}) \in \Delta^m \times \mathbb{R}_{++}$ where $\alpha_{i,j} \geq 0 \, \forall j$, $\sum_j \alpha_{i,j} = 1$ and $\alpha_{i,f} > 0$. We use two different notions of utility functions. First, the *valuation* of agent i for budget decision (x, t) is given by

$$v_i(x, t) := \sum_{j=1}^{m} \alpha_{i,j} \theta_j \big(x_j \cdot B_t \big) - \alpha_{i,f} f(t)$$

In some places we use $X_j := x_j \cdot B_t$ (and X accordingly). Since our model involves additional payments beyond the collected tax t, we write the overall utility of agent i when the budget decision is (x, t) and she is charged with a payment p_i as

$$u_i(x, t, p_i) = \sum_{j=1}^{m} \alpha_{i,j} \theta_j (x_j \cdot B_t) - \alpha_{i,f} f(t + p_i)$$

In a more general sense, we also write u_α and v_α for functions of a hypothetical "type α" agent. We assume the following on the functions $\{\theta\}_{j=1}^m$ and f.

Assumption 1. For all $1 \leq j \leq m$, $\theta_j : \mathcal{D} \to \mathbb{R}, \mathcal{D} \in \{\mathbb{R}_+, \mathbb{R}_{++}\}$ (where $\mathbb{R}_+ := \{x \in \mathbb{R} | x \geq 0\}$, $\mathbb{R}_{++} := \{x \in \mathbb{R} | x > 0\}$) is increasing and strictly concave, and $\lim_{X_j \to 0} \theta(X_j) \leq 0$.

Assumption 2. $f : \mathbb{R} \to \mathbb{R}$ is increasing, strictly convex in $(-\infty, 0]$ and strictly concave in $[0, \infty)$, and $f(0) = 0$. We assume that f is either continuously differentiable in all $\mathbb{R}$ or anywhere but the origin, in which case $\lim_{z \to 0^+_-} f'(z) = \infty$. [5]

Our formulation of the disutility function f follows Kahneman and Tversky's Prospect Theory [30, 9], that assumes *loss aversion* in relation to monetary transfers, that is risk aversion with respect to gains while risk seeking with respect to losses.

We define a mapping g that takes any type α and outputs the optimal budget decision w.r.t. v_α, in other words that maps types to votes.

[5] In principle, our analysis requires differentiable utility functions. However, the most natural examples of elementary increasing concave functions, the logarithmic and power functions, are not differentiable at zero, albeit still enables our results. Thus, for the sake of giving more intuitive and simple examples, we will allow a diverging derivative at zero.

Definition 1. For all $\alpha \in \Delta^m \times \mathbb{R}_{++}$, define $g(\alpha) : \Delta^m \times \mathbb{R}_{++} \to \Delta^m \times [-\frac{B_0}{n}, \infty)$ such that

$$g(\alpha) = (x^{(\alpha)}, t^{(\alpha)}) \in \arg\max_{(x,t)} v_\alpha(x, t)$$

If that optimum is not unique, g chooses one arbitrarily. In some places we use this notation somewhat abusively, ignoring that indecisiveness in the specific choice of $g(\alpha)$. Since every agent reports $(x^{(i)}, t^{(i)}) = g(\alpha_i)$, we want $g(\alpha)$ to entail the full description of the underlying type α. Moreover, our SDSIC result will also depend on that decisiveness. This is provided by our next assumption.

Assumption 3 (MRS characterization)**.** For every type $\alpha \in \Delta^m \times \mathbb{R}_{++}$, there exists an optimal point w.r.t. v_α, $(x^{(\alpha)}, t^{(\alpha)}) \in \arg\max_{(x,t)} v_\alpha$,[6] and moreover, every such optimum satisfies the following ***MRS conditions***:

$$\frac{n\theta'_j(x_j^{(\alpha)} \cdot B_{t^{(\alpha)}})}{f'(t^{(\alpha)})} = \frac{\alpha_f}{\alpha_j} \quad \forall \alpha_j > 0, \text{ and } \alpha_j = 0 \implies x_j^{(\alpha)} = 0.$$

That characterization of optimal points, which is merely the first order conditions derived from the optimization of v_α, is a widely adopted convention in Economic literature [22, 28, 21, 8], and that will hold under a vast range of plausible circumstances in our model, e.g. every internal optimal point satisfies the MRS equations. We refer the reader to the full version [32] for a full discussion on the foundations of that assumption and possible relaxations.

Hence, a *budgeting instance* $\mathcal{I} = \{m, n, B_0, \vec{\alpha}, \{\theta_j\}_{j\in[m]}, f\}$ is defined by the number of public goods m, number of agents n, initial budget B_0, the *type profile* $\vec{\alpha} = (\alpha_i, \dots, \alpha_n)$ and functions $\{\theta_j\}_{j\in[m]}, f$ that respect Assumptions 1,2,3. For example,

Example 1 (Running example)**.** Consider an instance with $B_0 = 0, m = 2, n = 3$, where for every agent $i \in [3]$, the valuation is:

$$v_i(x, t) := \sum_{j=1}^{2} \alpha_{i,j} \ln(x_j \cdot 3t) - \alpha_{i,f} \sqrt{t} \tag{1}$$

2.2 Mechanisms

The mechanism we introduce in the next section is designed to maximize the social welfare $\sum_i v_i(x, t)$. However, agents do not report their full valuations $v_i(\cdot, \cdot)$ explicitly but rather their votes $g(\alpha_i) = (x^{(i)}, t^{(i)})$ $\forall i \in [n]$. Assumption 3, however, provides that this is enough to fully recover the underlying types, as $g(\alpha)$ defines α uniquely.

Corollary 1. For any given optimal point (x, t) there exists a unique type $\alpha \in \Delta^m \times \mathbb{R}_{++}$ such that $(x, t) = g(\alpha) \in \arg\max_{(x',t')} v_\alpha(x', t')$.

That is not difficult to see, since the m linear MRS equations above (Assumption 3), along with $\sum_j \alpha_j = 1$, has a unique solution.

Example 2 (Running example, cont.)**.** Every agent submits the optimal budget decision $g(\alpha_i) = (x^{(i)}, t^{(i)}) \in \Delta^m \times [-\frac{B_0}{n}, \infty)$ that maximizes (1). Using the above Corollary, we infer every agent's underlying type α_i from $\frac{(x_j^{(i)} \cdot 3t^{(i)})^{-1}}{(2\sqrt{t^{(i)}})^{-1}} = \frac{\alpha_{i,f}}{\alpha_{i,j}}$ if $x_j^{(i)} > 0$, and $\alpha_{i,j} = 0$ otherwise.

[6] Note that the 'existence' here means that $t^{(\alpha)} < \infty$ for all types, meaning that no agent would ever wish to fund an infinite budget through infinite taxation.

For example, the voting profile on the left can only be induced by the type profile α^{RE} on the right (RE for Running Example).

(votes)	t	$X_1^{(i)}$ $(x_1^{(i)})$	$X_2^{(i)}$ $(x_2^{(i)})$
voter 1	6.25	13.125 (0.7)	5.625 (0.3)
voter 2	2.36	0 (0)	7.1 (1)
voter 3	4	6 (0.5)	6 (0.5)

(types)	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,f}$
$i=1$	0.7	0.3	0.8
$i=2$	0	1	1.3
$i=3$	0.5	0.5	1

Thanks to Corollary 1, we may consider w.l.o.g. only *Direct Revelation* [26] mechanisms that take the explicit type profile as input, although that is not in fact what we assume. A ***mechanism*** is thus a pair $M(\vec{\alpha}) = (\phi(\vec{\alpha}), p(\vec{\alpha}))$ where $\vec{\alpha}$ is the type profile, ϕ is a social choice function that outputs a budget decision and p assigns a payment p_i for each agent i.

3 Utility-Sensitive VCG

In this section section we present our proposed mechanism (Def. 4) and discuss its properties. We start with the payment function p.

3.1 Payments

Essentially, the payments we define are VCG payments adjusted to our non quasi-linear setup.[13] In general, the social choice function in a VCG mechanism ϕ^{VCG} outputs the socially optimal outcome $\Omega^* := argmax_\Omega \sum_i v_i(\Omega)$ and payments are defined as follows.

Definition 2 (VCG payments)**.**

$$p_i^{VCG} = -\sum_{k\neq i} v_k(\Omega^*) + h(\vec{\alpha}_{-i}) \quad \forall i \in [n]$$

where $h(\vec{\alpha}_{-i})$ could be any function of $\vec{\alpha}_{-i}$, the partial type profile submitted by all agents excluding i. The VCG model assumes *Quasi-linear* utility functions, meaning that

$$(*) \quad u_i\big(VCG(\vec{\alpha})\big) = v_i(\Omega^*) - p_i^{VCG} = \sum_k v_k(\Omega^*) - h(\vec{\alpha}_{-i})$$

The above expression is the key property on which truthfulness of a VCG mechanism relies.

However, in our model the effect of the payment is not linear, as it is mediated through f. If we naïvely charge agents with the VCG payments, their utility would be:

$$u_i = \sum_j \alpha_{i,j}\theta_j\big(x_j^* B_{t^*}\big) - \alpha_{i,f} f\big(t^* + p_i^{VCG}\big) \neq v_i(x^*, t^*) - p_i^{VCG},$$

and thus we cannot apply (*) to show truthfulness. However, we can adjust the payments appropriately so that the key property $(*)$ is maintained.

Definition 3 (Utility-Sensitive VCG Payments)**.** Let $\Omega^* = (x^*, t^*) \in \arg\max_{(x,t)} \sum_{i\in[n]} v_i(x,t)$ be the socially optimal budget decision. For every agent i, we define the **utility-sensitive VCG payment** p^{US} as

$$p_i^{US} = -t^* + f^{-1}\Big(f(t^*) + \frac{1}{\alpha_{i,f}} p_i^{VCG}\Big) \quad \forall i. \tag{2}$$

Lemma 1. *Let $\Omega^* = (x^*, t^*)$ be the social optimum and p_i^{US} the utility sensitive VCG payment given in 3 above. Then $u_i(\Omega^*, p_i^{US}) = v_i(x^*, t^*) - p_i^{VCG} = \sum_k v_k(\Omega^*) - h(\vec{\alpha}_{-i})$.*

The proof is straight-forward.

Remark 1. Here is where the inclusion of t in the decision space becomes crucial for our analysis. If we had only asked agents for their preferred allocation $x \in \Delta^m$, we would have no information on the α_f coefficients, and thus could not have use Definition 3 to set the payments appropriately.

Definition 4 (Utility-Sensitive VCG). The *Utility-Sensitive VCG* (US-VCG) mechanism $\mathcal{M}$ is defined by $\mathcal{M}(\vec{\alpha}) := \left(\Omega^*, p^{US}\right)$ for any input profile $\vec{\alpha}$.

3.2 Incentive Compatibility

Definition 5. A mechanism M is ***dominant strategy incentive compatible (DSIC)*** if for all $i \in [n]$, $\alpha_i \in \Delta^m \times \mathbb{R}_{++}$, $\vec{\alpha}_{-i} \in (\Delta^m \times \mathbb{R}_{++})^{n-1}$ and every $\alpha' \in \Delta^m \times \mathbb{R}_{++}$ s.t. $\alpha' \neq \alpha_i$,

$$u_i(M(\alpha_i, \vec{\alpha}_{-i})) \geq u_i(M(\alpha', \vec{\alpha}_{-i}))$$

If that inequality is strict for all i, α_i, $\vec{\alpha}_{-i}$ and α' we say that M is ***strictly dominant strategy incentive compatible (SDSIC)***.

Theorem 1. *The US-VCG mechanism is SDSIC.*

Remark 2. In general, VCG mechanisms only satisfy the weaker DSIC. Beyond being technically stronger, the SDSIC property brings a significant improvement in terms of robustness against manipulating coalitions, namely that truthfully reporting the types forms a Coalition-Proof Nash Equilibrium (CPNE). See appendix B for a detailed discussion and proof.

Before moving on to the proof, we point out here to some very useful characteristic of our model that plays a major role in both of our main results (Theorem 1 below and Theorem 2 in Section 4). That is that the outcome Ω^* depends solely on the average of types reported by all agents. Let $\bar{\alpha} := \frac{1}{n}\sum_{i\in[n]} \alpha_i$ denote the types mean. Now, simply changing the order of summation in the social welfare

$$\begin{aligned}\sum_i v_i(x,t) &= \sum_i \Big[\sum_j \alpha_{i,j}\theta_j(x_j B_t) - \alpha_{i,f} f(t)\Big] = \sum_j \Big[\sum_i \alpha_{i,j}\theta_j(x_j B_t) - \alpha_{i,f} f(t)\Big] \\ &= \sum_j n\bar{\alpha}_j \theta_j(x_j B_t) - n\bar{\alpha}_f f(t) = n \cdot v_{\bar{\alpha}}(x,t)\end{aligned}$$

shows that:

Observation 1. In any budget decision (x,t), the social welfare is given by $\sum_i v_i(x,t) = n \cdot v_{\bar{\alpha}}(x,t)$. Consequently, the social optimum $\Omega^* = (x^*, t^*)$ maximizes $\sum_i v_i(x,t)$ if and only if it maximizes $v_{\bar{\alpha}}(x,t)$, in other words $\Omega^* = g(\bar{\alpha})$.

Example 3 (Running example, cont.)**.** For the utility functions from Eq. (1) we have $g(\alpha) = \left((\alpha_1, \alpha_2), \left(\frac{2}{\alpha_f}\right)^2\right)$. The average type in the type profile α^{RE} from Example 2 is $\bar{\alpha} = \left(0.4, 0.6, 1.03\right)$, and thus the budget decision chosen by $\mathcal{M}$ is $g(\bar{\alpha}) = \left((0.4, 0.6), 3.74\right)$.

Note that in addition, by Observation 1 the social optimum Ω^* is computed in linear time. As for payments, the typical choice for $h(\alpha_{-i})$ is the "Clarke Pivot Rule" $h(\alpha_{-i}) = \sum_{k\neq i} v_k(\mathcal{M}(\alpha_{-i})) = \sum_{k\neq i} v_k(g(\bar{\alpha}_{-i}))$ [13], meaning, charging every agent with the social welfare of others in her absence. With that choice, computing every agent's payment is as hard (or as easy) as computing the outcome Ω^*. We end this section with the proof of Theorem 1 below.

Proof. Fix i, α_i and $\bar{\alpha}_{-i}$. Note that by Observation 1 the US-VCG mechanism outputs $g(\bar{\alpha})$, and, by Lemma 1,

$$u_i(\mathcal{M}(\vec{\alpha})) = \sum_k v_k(g(\bar{\alpha})) - h(\vec{\alpha}_{-i}) = n \cdot v_{\bar{\alpha}}(g(\bar{\alpha})) - h(\vec{\alpha}_{-i})$$

Now, assume that i falsely reports $\alpha_i' \neq \alpha_i$. Inevitably, that shifts the mean preferences to some $\bar{\alpha}' \neq \bar{\alpha}$, and the social optimum that $\mathcal{M}$ outputs to $g(\bar{\alpha}')$. Due to Corollary 1, $g(\bar{\alpha}')$ is certainly **not** an optimum of $v_{\bar{\alpha}}$, and therefore:

$$\begin{aligned} u_i(\mathcal{M}(\alpha_i', \vec{\alpha}_{-i})) &= n \cdot v_{\bar{\alpha}}(g(\bar{\alpha}')) - h(\vec{\alpha}_{-i}) \\ &< n \cdot v_{\bar{\alpha}}(g(\bar{\alpha})) - h(\vec{\alpha}_{-i}) = u_i(\mathcal{M}(\alpha_i, \vec{\alpha}_{-i})) \end{aligned}$$

□

4 Vanishing Payments Under Per-Capita Utilities

In this section we show that payments become negligible in large populations. While these payments are essential for aligning incentives, charging additional money from voters is not desired in a PB context. For the technical proofs, we will have to further specify the utility model so that it captures some important feature of divisible PB, which has been (justifiably) overlooked in past literature as well as in this work up to this point. That is, that the utility achieved from a given spending X_j on some public good j must also depend on the number of people that enjoy it, n. The reason we only have to address that now is that in this section we analyse the asymptotic behavior of payments w.r.t. n, and in our model the overall budget $B_t = B_0 + nt$ depends on it directly. The following example illustrates the problem.

Example 4. Consider a budgeting instance where $m = 1, B_0 = 0, \theta(nt) = (nt)^p$ and $f(t) = t^q$ for some $0 < p < q < 1$. A type α_f agent thus maximizes her utility $u_{\alpha_f} = (nt)^p - \alpha_f t^q$ at

$$t = \Big[\frac{\alpha_f q}{p}\Big]^{\frac{1-q}{1-p}} \cdot n^{1-q} \xrightarrow[n \to \infty]{} \infty$$

That is a very unlikely result, whatever that sole public good may be. There is no reason to expect that larger populations would wish to pay infinitely larger taxes, nor is it the situation found in reality. The model allows that because every tax unit payed by an individual is presumably "matched" $n - 1$ times by others, thus making the substitution rate grow proportionally with n. However, while larger societies probably do have larger available resources, they are also likely to have greater needs. Just as a country's economic state is conventionally measured by its GDP index, on the large scale the quality of public goods should be associated with spending per capita rather than with nominal spending.[7] Hence, we now narrow down the definition of $v_i(x, t)$ to

Definition 6 (per capita valuations of public expenditures)**.**

$$v_i(x, t, n) = \sum_{j=1}^{m} \alpha_{i,j} \theta_j\Big(\frac{x_j \cdot B_t}{n}\Big) - \alpha_{i,f} f(t) = \sum_{j=1}^{m} \alpha_{i,j} \theta_j(x_j(b_0 + t)) - \alpha_{i,f} f(t) \text{ where } b_0 := \frac{B_0}{n}.$$

Note that all of our previous results follow through since for any fixed n, $\theta_j(X_j/n)$ is a particular case of $\theta_j(X_j)$. Realistically, the dependency on n might not be necessarily that we assumed, and we take $\theta_j(X_j/n)$ as a benchmark and relatively simple case of a more general class of functions of the form $\theta_j(X_j, n)$. While for some public goods $\frac{X_j}{n}$

[7]See for example `https://data.oecd.org/gga/general-government-spending.htm`, and [28].

may capture the relation adequately—for example, the quality of an education system surely depends on its resources per child—for others it may serve more as a large scale approximation—e.g., if the city offers free cloud services that allocates space equally among users, the total number of users affects each of them only to the point where the provided space exceeds their needs. Still, the benefit for users must be somehow connected the to the available space per user (which is determined by spending).

Some comments on the notations before we proceed. First, in this section we define $h(\vec{\alpha}_{-i})$ in the VCG payments as the conventional Clarke pivot-rule [13] function that charges a voter with the social welfare of all others in her absence $h(\vec{\alpha}_{-i}) = \sum_{k\neq i} v_k(g(\bar{\alpha}_{-i}))$, making the VCG payments

$$p_i^{VCG} = -\sum_{k\neq i} v_k(g(\bar{\alpha})) + \sum_{k\neq i} v_k(g(\bar{\alpha}_{-i})) = (n-1)\Big(v_{\bar{\alpha}_{-i}}(g(\bar{\alpha}_{-i})) - v_{\bar{\alpha}_{-i}}(g(\bar{\alpha}))\Big)$$

where the second equality is by Observation 1. Next, we give an alternative representation for valuation functions that would ease the technical analysis significantly.

Definition 7 (Alternative representation of valuation functions)**.** Define the vector valued function $V : \Delta^m \times \mathbb{R} \to \mathbb{R}^{m+1}$

$$V(x,t) = (\theta_1(x_1(b_0+t)), \dots, \theta_m(x_m(b_0+t)), -f(t))$$

For every $\alpha \in \Delta^m \times \mathbb{R}$, we write the valuation function v_α as the dot product of α and V:

$$v_\alpha(x,t) = \sum_j \alpha_j \theta_j(x_j(b_0+t)) - \alpha_f f(t) = \alpha \cdot V(x,t)$$

In these notations, the VCG payments are written as

$$p_i^{VCG} = (n-1)\bar{\alpha}_{-i} \cdot \Big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\Big)$$

Our main result in this section, Theorem 2, basically says that the expression above converges to zero as $n \to \infty$. For that we must have that $g(\bar{\alpha}_{-i}) \to g(\bar{\alpha})$ as $\bar{\alpha}_{-i} \to \bar{\alpha}$ with n, and fast enough. Thus, it relies on the guarantees we can provide for g's smoothness around the solution $g(\bar{\alpha})$. In our running example, for instance, it is easy to check that g is as smooth as you can wish for.

Example 5 (Running example, cont.)**.** Consider again Example 1, this time with per capita valuations: $v_\alpha(x,t) = \sum_{j=1}^{2} \alpha_j \ln(x_j \cdot t) - \alpha_f \sqrt{t}$. Then $g(\alpha) = ((\alpha_1, \alpha_2), (\frac{2}{\alpha_f})^2)$ is continuously differentiable for all $\alpha \in \Delta^2 \times \mathbb{R}_{++}$ (and in particular at $\bar{\alpha}$). [8]

The next lemma establishes the continuity of g at the solution $g(\bar{\alpha})$ when that is uniquely defined, which is sufficient for Theorem 2 that follows. We defer its proof to the appendix.

Lemma 2. *For any given $\alpha \in \Delta^m \times \mathbb{R}$, if v_α has a unique global maximum then*

$$\lim_{\beta\to\alpha} g(\beta) = g(\alpha).$$

Note that (a) this statement is not obvious because we did not assume that g is continuous, and (b) it holds for any function g that follows Definition 1, i.e. the specific arbitrary choice of $g(\beta)$ in case v_β has multiple optima is irrelevant.

One last preliminary definition, and we are ready.

[8]Note that in this case, moving to per capita valuations does not change $g(\alpha)$ as taking $\ln\Big(\frac{x_j \cdot 3t}{3}\Big)$ instead of $\ln(xj \cdot 3t)$ merely adds a negative factor to v_α.

Definition 8. The *characteristic triplet* in a budgeting instance $\mathcal{I}$ is $\sigma = (b_0, \mu, \bar{\alpha})$ where

- $b_0 := \frac{B_0}{n} \geq 0$ is the non tax funded budget source per capita.
- $\bar{\alpha} := \frac{1}{n}\sum_k \alpha_k \in \Delta^m \times \mathbb{R}$ is the mean preferences vector of all agents.
- $1/\mu < \alpha_{i,f} < \mu \; \forall i \in [n]$.

Theorem 2. *Let $\sigma = (b_0, \mu, \bar{\alpha})$ such that $v_{\bar{\alpha}}$ has a unique global maximum at $g(\bar{\alpha})$. Then for every $\epsilon > 0$ there exists $n_\epsilon(\sigma)$ such that in every budgeting instance with characteristic triplet σ and $n > n_\epsilon(\sigma)$,*

$$|p_i^{US}| < \epsilon \quad \forall i \in [n].$$

As stated, Theorem 2 means that prices vanish if the population is sufficiently large while not taking into account the likely possibility that in reality, new members that join a community might change it's characteristic parameters b_0, μ and $\bar{\alpha}$. That is, we are saying that in any given community with known parameters $(b_0, \mu, \bar{\alpha})$, prices will be arbitrarily small if the population is large enough. Thus, as there is no reason to assume some correlation between these parameters and the population's size, the theorem essentially implies that prices are likely to be small, even negligible, in larger societies.

Proof. The p^{VCG} payments are defined as the loss an agent imposes on all others by participating, and are therefore always non-negative. Now,

$$\begin{aligned} 0 \leq p_i^{VCG} &= (n-1)\bar{\alpha}_{-i}\Big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\Big) \\ &= (n-1)\bar{\alpha}\Big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\Big) + (n-1)(\bar{\alpha}_{-i} - \bar{\alpha})\Big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\Big) \\ &\leq (n-1)(\bar{\alpha}_{-i} - \bar{\alpha})\Big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\Big) \\ &\leq \frac{n-1}{n}|\bar{\alpha}_{-i} - \alpha_i||V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))| \xrightarrow[\bar{\alpha}_{-i} \to \bar{\alpha}]{} 0 \end{aligned} \tag{3}$$

Where the first inequality is by definition of $g(\bar{\alpha})$, and in the second we used $\bar{\alpha} = \frac{n-1}{n}\bar{\alpha}_{-i} + \frac{1}{n}\alpha_i$ and Cauchy–Schwarz. Since $|\bar{\alpha}_{-i} - \alpha_i|$ is bounded, and by Lemma 2 $V \circ g$ is continuous, we get the convergence at the end. Now, as $\bar{\alpha}_{-i} \to \bar{\alpha}$ with $n \to \infty$,

$$p_i^{US} = -t(\bar{\alpha}) + f^{-1}\Big(f(t(\bar{\alpha})) + \frac{1}{\alpha_{i,f}}p_i^{VCG}\Big) \xrightarrow[n \to \infty]{} -t(\bar{\alpha}) + t(\bar{\alpha}) = 0$$

(note that f^{-1} is continuous). Thus for any arbitrary small ϵ we can find $n_\epsilon(\sigma)$ that yields the result. $\square$

Note that the mean-dependency (Observation 1) is crucial for that result too, since $\bar{\alpha}_{-i}$ approaches $\bar{\alpha}$ at a $1/n$ rate, precisely canceling the increase in the number of agents n.

Non-Positive Payments. In appendix C, we further show a method for constructing non-positive payments, meaning that no agent would add any payment on top of the tax t, even negligible. Instead, they might be paid a "negative payment" that we can view as a bonus or a "tax discount" for their participation. Just to outline here the main idea, we show that if g is moreover differential at the solution we can, by linear approximation, bound individual payments by $p_i^{US} \leq \frac{\mathcal{B}}{n}$ for some constant factor $\mathcal{B}$. Thus, we can redefine the payments so that we pay back each agent at least the amount she was originally charged with, while guaranteeing that the total amount we pay to all agents will not diverge.

5 Biased Mechanisms

In this section, we expand the US-VCG definition to a class of mechanisms that insert a bias towards an arbitrary desired outcome or a set of outcomes. Indeed, the designer may have some goal in mind that she may want to balance with welfare, for example a particular project she wants to promote, a legacy allocation from previous years, or an allocation she sees as fair. The general form of mechanisms in that class follows the familiar *affine-maximizer* generalization of a VCG mechanism. (See [26], p. 228). We start with choosing a *bias function* $\mathcal{C}(x,t) : \Delta^m \times \mathbb{R} \mapsto \mathbb{R}$, that in one way or another favours—that is, assigns higher values to—outcomes we see as desirable. Note that $\mathcal{C}(x,t)$ must be independent of the realization of preferences $\vec{\alpha} = (\alpha_1, \ldots, \alpha_n)$ to maintain incentive compatibility. Then, we generalize the US-VCG definition as follows.

Definition 9 (Biased Utility-Sensitive VCG). For any bias function $\mathcal{C}(x,t) : \Delta^m \times \mathbb{R} \mapsto \mathbb{R}$, we define the *Biased Utility-Sensitive VCG* (BUS-VCG) mechanism:

$$\hat{\mathcal{M}}(\vec{\alpha}) := (\hat{g}(\bar{\alpha}), \hat{p}^{US})$$

where

- $\forall \alpha \in \Delta^m \times \mathbb{R}_{++},\ \hat{g}(\alpha) := (x,t) \in \arg\max_{(x,t)} v_\alpha(x,t) + \mathcal{C}(x,t)$ [9]
- $\hat{p}^{US}$ is the payment assignment

$$\hat{p}_i^{US} := -\hat{t} + f^{-1}\Big(f(\hat{t}) + \frac{1}{\alpha_{i,f}}(p_i^{VCG} + n\mathcal{C}(\hat{g}(\bar{\alpha}_{-i})) - n\mathcal{C}(\hat{g}(\bar{\alpha}))\Big)$$

Example 6. ***Equitable/Egalitarian allocations.*** The designer might favor allocations of the budget B_t (once it is determined by t) that promote some familiar concepts of fair division. The *Egalitarian* allocation, for example, maximizes the minimum possible utility of a voter given t, i.e.,

$$\hat{e}^t := \arg\max_{x \in \Delta^m} \min_{\alpha \in \Delta^m} \sum_j \alpha_j \theta_j(x_j B_t)$$

Since that utility is a convex combination of $\{\theta_j(x_j B_t)\}_j$, it is not difficult to see that $\hat{e}^t = \arg\max_{x \in \Delta^m} \min_j \theta_j(x_j B_t)$. Moreover, that choice of $\hat{e}^t$ also minimizes the difference

$$\Big(\max_{\alpha \in \Delta^m} \sum_j \alpha_j \theta_j(x_j B_t) - \min_{\alpha \in \Delta^m} \sum_j \alpha_j \theta_j(x_j B_t)\Big)$$

In particular, if that difference is zero then $\hat{e}^t$ is an *equitable* allocation where $\sum_j \alpha_{i,j} \theta_j(\hat{e}_j^t B_t)$ has the same value for every agent i. That case is in fact quite generic, for example if $\theta_j(0)$ has the same value for all j, for any given t we can allocate the budget B_t so that $\theta_j(\hat{e}_j^t B_t) = \theta_k(\hat{e}_k^t B_t)\, \forall j,k$.

5.1 Properties inherited by $\hat{\mathcal{M}}$

We would obviously like to preserve useful properties of the US-VCG mechanism when generalizing to BUS-VCG. Some of them carry over quite easily. First, note that mean-dependency (Observation 1) holds for $\hat{g}$ too. (We in fact rely on that in the definition of $\hat{\mathcal{M}}$). It is not difficult to see now that DSIC extends under any choice for $\mathcal{C}$, as the definition of $\hat{p}^{US}$

[9]We assume that $\mathcal{C}$ is comstructed so that indeed such a maximum exists. Note that $\hat{g}$ Obviously depends on our choice for $\mathcal{C}$, but we do not refer to that explicitly as it should be clear in the context.

imitates the situation where $v_i(x,t) \to v_i(x,t) + \mathcal{C}(x,t)$ and $h(\alpha_{-i}) \to h(\alpha_{-i}) + n\mathcal{C}(\hat{g}(\bar{\alpha}_{-i}))$ for every agent i.[10] One also easily checks that payments vanish if $\hat{g}$ is continuous at $\bar{\alpha}$, the same condition we needed earlier for g. (Albeit, for $\hat{g}$ we will have to explicitly demand that while for g we had other terms that implied its smoothness). [11] SDSIC, however, will not carry over that easily. The condition we need for that is $\alpha \neq \alpha' \implies \hat{g}(\alpha) \neq \hat{g}(\alpha')$, which is not satisfied by any arbitrary choice of $\mathcal{C}(x,t)$. In Appendix D we describe a class of bias functions that do preserve SDSIC.

6 Concluding Remarks

We presented a divisible participatory budgeting mechanism, the US-VCG mechanism, that concerns both the allocation and total volume of expenses. It is essentially a VCG mechanism adjusted to our setting, in which we had to tackle some issues. Mainly, we had to reformulate the payments to suit our preference model of non quasi-linear utilities. The US-VCG mechanism is welfare-maximizing and strategy-proof in strictly dominant strategies, and consequently implements the social optimum as a Coalition-Proof Nash Equilibrium. In Section 4, we showed that our modified VCG payments the mechanism charges become negligible in large populations, if we accept the plausible assumption of per-capita utilities. Finally, we showed a generalization of the US-VCG mechanism that inserts a bias towards any set of outcomes of one's choice.

Future Directions. In the introduction, we discussed the theoretic advantages of an additive concave utility model over other examples from the literature. The obvious downside is, when considering a mechanism that aggregates preferences, is the difficulty in assessing the concrete functions we should assume. While we can nevertheless argue that any such functions are probably a better approximation for the true underlying preferences than previous suggestions, future experimental research attempting to evaluate them, similar to those performed in relation to the disutility monetary function f [9], could make a valuable contribution to the field.

[10] See [26] for a rigorous proof.

[11] Similarly, Theorem 3 in the appendix extends if we have differentiability at $\bar{\alpha}$.

References

[1] Haris Aziz and Aditya Ganguly. Participatory funding coordination: Model, axioms and rules. In *Algorithmic Decision Theory: 7th International Conference, ADT 2021, Toulouse, France, November 3–5, 2021, Proceedings 7*, pages 409–423. Springer, 2021.

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

[3] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 753–781, 2019.

[4] Yoram Bachrach. Honor among thieves: collusion in multi-unit auctions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 617–624. Citeseer, 2010.

[5] Gerdus Benade, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. *Management Science*, 67(5):2813–2827, 2021.

[6] Alain Bernard. Optimal taxation and public production with budget constraints. In *The Economics of Public Services: Proceedings of a Conference held by the International Economic Association at Turin, Italy*, pages 361–389. Springer, 1977.

[7] B Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. Coalition-proof nash equilibria i. concepts. *Journal of economic theory*, 42(1):1–12, 1987.

[8] Kjetil Bjorvatn and Guttorm Schjelderup. Tax competition and international public goods. *International tax and public finance*, 9:111–120, 2002.

[9] Adam S Booij, Bernard MS Van Praag, and Gijs Van De Kuilen. A parametric analysis of prospect theory's functionals for the general population. *Theory and Decision*, 68: 115–148, 2010.

[10] Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. Distribution rules under dichotomous preferences: Two out of three ain't bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, page 158–179, 2021. ISBN 9781450385541.

[11] Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. Funding public projects: A case for the nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.

[12] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of vcg payments. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 882–889, 2006.

[13] Edward H Clarke. Multipart pricing of public goods. *Public choice*, pages 17–33, 1971.

[14] Vincent Conitzer and Tuomas Sandholm. Failures of the vcg mechanism in combinatorial auctions and exchanges. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 521–528, 2006.

[15] Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *Web and Internet Economics: 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings 12*, pages 384–399. Springer, 2016.

[16] Roy Fairstein, Adam Lauz, Reshef Meir, and Kobi Gal. Modeling people's voting behavior with poll information. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1422–1430, 2019.

[17] Rupert Freeman, David M Pennock, Dominik Peters, and Jennifer Wortman Vaughan. Truthful aggregation of budget proposals. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 751–752, 2019.

[18] Nikhil Garg, Vijay Kamble, Ashish Goel, David Marn, and Kamesh Munagala. Iterative local voting for collective decision-making in continuous spaces. *Journal of Artificial Intelligence Research*, 64:315–355, 2019.

[19] Ashish Goel, Anilesh K Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)*, 7(2):1–27, 2019.

[20] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of vcg payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009.

[21] Kamal Jain and Vijay V Vazirani. Eisenberg-gale markets: Algorithms and structural properties. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 364–373, 2007.

[22] Paul Krugman and Robin Wells. *Microeconomics.* Macmillan, 2008.

[23] Ratul Lahkar and Saptarshi Mukherjee. Dominant strategy implementation in a large population public goods game. *Economics Letters*, 197:109616, 2020.

[24] Hongyao Ma, Reshef Meir, and David C Parkes. Social choice with non quasi-linear utilities. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 49–49, 2018.

[25] Ilan Nehama. Almost quasi-linear utilities in disguise: Positive-representation an extension of roberts' theorem. *arXiv preprint arXiv:1910.12131*, 2019.

[26] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory.* Cambridge university press, 2007.

[27] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities, 2022.

[28] Joseph E Stiglitz. The theory of local public goods. In *The Economics of Public Services: Proceedings of a Conference Held by the International Economic Association at Turin, Italy*, pages 274–333. Springer, 1977.

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

[30] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5:297–323, 1992.

[31] Vijay V Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)*, 58(3):1–25, 2011.

[32] Jonathan Wagner and Reshef Meir. Strategy-proof budgeting via a vcg-like mechanism. *arXiv preprint arXiv:2303.06923*, 2023.

[33] Kevin Wainwright et al. *Fundamental methods of mathematical economics*. Boston, Mass. McGraw-Hill/Irwin, 2005.

[34] David E Wildasin. Nash equilibria in models of fiscal competition. *Journal of public economics*, 35(2):229–240, 1988.

A Missing Proofs

Lemma 2. *For any given $\alpha \in \Delta^m \times \mathbb{R}$, if v_α has a unique global maximum then*

$$\lim_{\beta \to \alpha} g(\beta) = g(\alpha).$$

Proof. By assumption, $0 < \alpha \cdot (V(g(\alpha)) - V(g(\beta)))$. On the other hand, putting $\epsilon := (\alpha - \beta)$,

$$\begin{aligned} 0 < \alpha \cdot (V(g(\alpha)) - V(g(\beta))) &= \beta \cdot (V(g(\alpha)) - V(g(\beta))) + \epsilon \cdot (V(g(\alpha)) - V(g(\beta)) \\ &\leq \epsilon \cdot (V(g(\alpha)) - V(g(\beta))) \leq |\epsilon| \cdot |V(g(\alpha)) - V(g(\beta))|. \end{aligned}$$

where the second inequality is by definition of $g(\beta)$ and the last is by Cauchy-Shwartz. We now want to show that $|V(g(\alpha)) - V(g(\beta))|$ is bounded, that would imply the convergence we want. Indeed, for any type β the optimum $g(\beta) = (x^{(\beta)}, t^{(\beta)})$ satisfies the MRS conditions

$$\frac{\theta_j'(x_j^{(\beta)}(b_0 + t^{(\beta)}))}{f'(t^{(\beta)})} = \frac{\beta_f}{\beta_j} \quad \forall j \; s.t. \; \beta_j > 0$$

and there exists $j \in [m]$ such that $x_j^{(\beta)} \geq 1/m$. Since θ_j' is decreasing,

$$\frac{\theta_j'(t^{(\beta)}/m)}{f'(t^{(\beta)})} \geq \frac{\theta_j'(x_j^{(\beta)} \cdot t^{(\beta)})}{f'(t^{(\beta)})} \geq \frac{\theta_j'(x_j^{(\beta)}(b_0 + t^{(\beta)}))}{f'(t^{(\beta)})} = \frac{\beta_f}{\beta_j}$$

By Assumption **??**, the LHS of that equality vanishes as $t^{(\beta)} \to \infty$, however the ratio on the RHS is bounded away from zero because $\beta \to \alpha$. Thus $t^{(\beta)}$ cannot increase infinitely as $\beta \to \alpha$, meaning that $g(\beta)$ is bounded and consequently $V(g(\beta))$ is too. Thus $|\epsilon| \cdot |V(g(\alpha)) - V(g(\beta))| \xrightarrow[\epsilon \to 0]{} 0$, ergo,

$$\lim_{\beta \to \alpha} \alpha \cdot (V(g(\alpha)) - V(g(\beta))) = \lim_{\beta \to \alpha} v_\alpha(g(\alpha)) - v_\alpha(g(\beta)) = 0$$

Since v_α is continuous, $g(\alpha)$ is unique and $g(\beta)$ is bounded, it must be that $g(\beta) \to g(\alpha)$. $\square$

B Manipulations By Coalitions

In general, VCG mechanisms are known to be highly prone to group manipulations [4, 14]. While individuals cannot benefit from reporting false preferences when the reports of all others are fixed, a group of agents can sometimes coordinate their misreports in such way that each of them (or some at least) benefits due to the untruthful reports of others. The US-VCG is no different in that. However, the SDSIC property ensures that any such coalition would not be sustainable in the sense that the colluding agents cannot trust each other to follow the agreed scheme. Thus, it may suggest that such coalitions are not likely to form in the first place. That softer robustness demand where we allow for manipulating coalitions as long as they are unsustainable in the above sense is captured in the *Coalition Proof Nash Equilibrium* (CPNE) solution concept [7]. While the original term is quite involved, the application to our context is intuitive: in an SDSIC mechanism, no sustainable coalition could exist since the individual sole optimal action, under any circumstances, is for every agent to report her true preference. We thus formulate here a simpler term that the US-VCG satisfies, and that implies CPNE.[12] First, we define a manipulation by a coalition as a coordinated misreport by all its members, that benefits them.

[12]We refer the reader to [7] for the original CPNE definition.

Definition 10 (Coalition manipulation)**.** A **manipulation** by a **coalition** $S \subset [n]$ is a partial type profile $\alpha'_S = \{\alpha'_i\}_{i \in S}$ such that $\alpha'_i \neq \alpha_i \; \forall i \in S$,

$$u_i(M(\alpha'_S, \alpha_{-S}) \geq u_i(M(\alpha_S, \alpha_{-S}) \; \forall i \in S$$

and there exists $i \in S$ for which the inequality is strict.

Now, we demand that if such a manipulation exists then α'_i is suboptimal for at least one agent in the coalition.

Definition 11 (”One Step Coalition-Proof”)**.** We say that a mechanism is *One Step Coalition-Proof* (OSCP) if for any manipulation α'_S by a coalition S, there exists $i \in S$ and $\alpha''_i \neq \alpha'_i$ s.t.

$$u_i(M(\alpha''_i, \alpha'_{S_{-i}}, \alpha_{-S})) > u_i(M(\alpha'_i, \alpha'_{S_{-i}}, \alpha_{-S})) \; \forall i \in S$$

Corollary 2. The US-VCG mechanism is OSCP (and consequently implements the social optimum as a CPNE).

The claim follows trivially from SDSIC, by taking $\alpha''_i = \alpha_i$ for any $i \in S$ in any coalition S.

C Non-Positive Payments

We here give a method for constructing the payment assignment p^{US} so that no agent pays any amount beyond the tax t. It will require that g is not only continious but also differentiable near the solution $\bar{\alpha}$. That is, instead of charging a ”fee” for participating in the vote (and moreover, one that is not identical for everyone), agents will be paid by the mechanism, which you might see as a reward for their participation. The definition below describes a sufficient condition for that, as we show in Lemma 3 and Corollary 3 that follow.

Definition 12. Define $F : (\Delta^m \times \mathbb{R})^2 \to \mathbb{R}^{m+1}$ as:

$$F_j(\alpha, (x, t)) = \alpha_j \theta'_j(x_j(b_0 + t)) - \alpha_f f'(t) \quad \forall 1 \leq j \leq m$$
$$F_{m+1}(\alpha, (x, t)) = \sum_j x_j - 1$$

By Assumption 3, $F(\alpha, (x^{(\alpha)}, t^{(\alpha)})) = 0$ for all $\alpha \in \Delta^m \times R$. We say that $g(\alpha) = (x^{(\alpha)}, t^{(\alpha)})$ is a ***”regular maximum”*** of v_α if, moreover,

$$\det \left[\frac{\partial F}{\partial(x, t)} \big(\alpha, g(\alpha)\big) \right] \neq 0$$

where $\frac{\partial F}{\partial(x,t)}$ is the $(m+1) \times (m+1)$ Jacobi matrix of F with respect to the variables (x, t).

The next Lemma is a direct application of the Implicit Function Theorem [33].

Lemma 3. *Let $g(\alpha)$ be a regular maximum of v_α. Then there exist an open neighborhood of α $S \subset \Delta^m \times R$ and a unique mapping $s : S \to \Delta^m \times [-\frac{B_0}{n}, \infty)$ such that :*

- *$s(\alpha) = g(\alpha)$.*
- *s is continuously differentiable in S.*

- $\forall \beta \in S,\ F(\beta, s(\beta)) = 0$

The Lemma does not yet provide the smoothness we want for g, because it only promises that $\{s(\beta)\}_{\beta \in S}$ are critical points of v_β that satisfy $F(\beta, s(\beta)) = 0$, not necessarily maxima. In other words s and g are not the same function by definition. However, an immediate corollary is that if $g(\alpha)$ is also unique they indeed must coincide.

Corollary 3. Assume that $g(\alpha)$ is a regular unique global maximum of v_α. Then there exist a neighborhood of α, $S \subset \Delta^m \times R$, such that g is uniquely defined and continuously differentiable in S.

Proof. Denote by $\tilde{S}$ the neighborhood of α from the statement of Lemma 3 and let $g(\tilde{S})$ be the image of $\tilde{S}$ under g. By Lemma 2, $g(\beta)$ approaches $g(\alpha) = s(\alpha)$ as $\beta \to \alpha$, therefore there exists a neighborhood $S \subseteq \tilde{S}$ of α in which $g(\beta) \in g(\tilde{S})\ \forall \beta \in S$. Now, every solution $g(\beta)$ must satisfy $F(\beta, g(\beta)) = 0$, but, by Lemma 3 s is the *unique* function that maps $\beta \in \tilde{S}$ to $s(\beta)$ such that $F(\beta, s(\beta)) = 0$, therefore $g\Big|_S = s$. Thus by Lemma 3, g is uniquely defined and continuously differentiable in S. $\square$

We now have the necessary for the stronger version of payments convergence stated in Theorem 3. Basically, this further result is due to the differentiabilty of g that allows for a linear approximation of the difference $\big(V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))\big)$, which essentially determines the VCG payment of an agent i. Thus, we can not only say that this difference vanishes as we did in Theorem 2, but also specify the convergence rate of $O(1/n)$.

Theorem 3. *Let $\sigma = (b_0, \mu, \bar{\alpha})$ such that $v_{\bar{\alpha}}$ has a regular unique global maximum $g(\bar{\alpha})$. Then there exist some $\mathcal{B} \in \mathbb{R}$ and $n(\sigma)$ such that in every population with characteristic triplet σ and size $n > n(\sigma)$,*

$$|p_i^{US}| \leq \frac{\mathcal{B}}{n} \quad \forall i \in [n]$$

Proof. In the proof of Theorem 2 we showed that:

$$0 \leq p_i^{VCG} \leq \frac{n-1}{n} |\bar{\alpha}_{-i} - \alpha_i| |V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))|$$

By Corollary 3, g is differentiable in some neighborhood of $\bar{\alpha}$ S, and so $V \circ g$ is too. Thus, if we take n_0 sufficiently large so that $\bar{\alpha}_{-i} \in S$, we have for all $n > n_0$ that

$$\begin{aligned} p_i^{VCG} &\leq \frac{n-1}{n} |\bar{\alpha}_{-i} - \alpha_i| \left| \mathcal{D}_{V \circ g}(\bar{\alpha}) \big(\bar{\alpha}_{-i} - \bar{\alpha}\big) + o|\bar{\alpha}_{-i} - \bar{\alpha}| \right| \\ &\leq \frac{n-1}{n} |\bar{\alpha}_{-i} - \alpha_i| \Big(\|\mathcal{D}_{V \circ g}(\bar{\alpha})\| |\bar{\alpha}_{-i} - \bar{\alpha}| + o|\bar{\alpha}_{-i} - \bar{\alpha}| \Big) \\ &= \frac{n-1}{n} |\bar{\alpha}_{-i} - \alpha_i| \Big(\|\mathcal{D}_{V \circ g}(\bar{\alpha})\| |\bar{\alpha}_{-i} - \alpha_i| \frac{1}{n} + o\big(|\bar{\alpha}_{-i} - \alpha_i| \frac{1}{n}\big) \Big) \end{aligned}$$

Where $\mathcal{D}_{V \circ g}$ is the Jacobian matrix of $V \circ g$ and $\|\mathcal{D}_{V \circ g}\|$ is its matrix-norm, and in the equality at the end we put $\bar{\alpha} = \frac{n-1}{n} \bar{\alpha}_{-i} + \frac{1}{n} \alpha_i$. Since the $\{\alpha_{i,f}\}$ are bounded there exists some $\gamma \in \mathbb{R}$ such that $|\bar{\alpha}_{-i} - \alpha_i| < \gamma$ for all i, thus

$$p_i^{VCG} \leq \frac{n-1}{n^2} \gamma^2 \Big(\|\mathcal{D}_{V \circ g}(\bar{\alpha})\| + o(1) \Big)$$

Now taking $n_1 \geq n_0$ such that the $o(1)$ term above is less than 1 and $\mathcal{B}_0 = \gamma^2 \Big(\|\mathcal{D}_{V \circ g}(\bar{\alpha})\| + 1 \Big)$ gives

$$p_i^{US} = -t(\bar{\alpha}) + f^{-1}\Big(f(t(\bar{\alpha})) + \frac{1}{\alpha_{i,f}} p_i^{VCG} \Big) < (f^{-1})'(t(\bar{\alpha})) \frac{\mu \mathcal{B}_0}{n} + o(\frac{\mu \mathcal{B}_0}{n})$$

(As mentioned earlier, in case f is not differentiable at zero then $t(\bar{\alpha}) \neq 0$ and thus $(f^{-1})'$ always exists.) And now taking $\mathcal{B} = \mu(f^{-1})'(t(\bar{\alpha}))\mathcal{C}_0 + 1$ with sufficiently large $n(\sigma) \geq n_1$ completes the proof. $\square$

An immediate result of Theorem 3 is that the total sum we charge from agents is bounded.

Corollary 4. In any budgeting instance with characteristic triplet σ and $n > n(\sigma)$, and such that $v_{\bar{\alpha}}$ has a regular unique global maximum,

$$\sum_{i \in [n]} p_i^{US} \leq \mathcal{B}$$

Theorem 3 and the corollary that follows set the ground for constructing payments that not only vanish asymptotically, but are also non-positive for all agents. The idea is quite simple—after charging an agent with her VCG payment, we "pay her back" no less than the maximum possible payment she could have been charged with given the preferences of her peers $\vec{\alpha}_{-i}$. This way no one is charged with a strictly positive payment. Thanks to Corollary 4, we know that the overall sum needed to implement that will not diverge with the number of agents.

Definition 13 (Non-positive utility-sensitive VCG payments)**.** Define the *non-positive VCG payment assignment* $\tilde{p}^{VCG}$ as

$$\tilde{p}_i^{VCG} := p_i^{VCG} - \frac{\gamma^2}{n}\Big(\|\mathcal{D}_{V\circ g}(\bar{\alpha}_{-i})\| + 1\Big) - \frac{r}{n}$$

where γ is the bound defined in the proof of Theorem 3 and $r \geq 0$ is any constant factor.

Note that this payment assignment does not violate SDSIC as it does not involve α_i. $\frac{\gamma^2}{n}\Big(\|\mathcal{D}_{V\circ g}(\bar{\alpha}_{-i})\| + 1\Big)$ is a bound we take for p_i^{VCG} (See clarification below) and the purpose of the $\frac{r}{n}$ factor is to allow for even higher payments to the agents, as much as the social planner wishes and can afford. We thus conclude that:

Corollary 5. In any budgeting instance with characteristic triplet σ and $n > n(\sigma)$ such that $v_{\bar{\alpha}}$ has a regular unique global maximum, the payment assignment $\tilde{p}_i^{US} := -t(\bar{\alpha}) + f^{-1}\Big(f(t(\bar{\alpha})) + \frac{1}{\alpha_{i,f}}\tilde{p}_i^{VCG}\Big)$ satisfies:

1. $\tilde{p}_i^{US} \leq 0 \ \forall i \in [n]$
2. $\sum_{i\in[n]} \tilde{p}_i^{US} \geq -\tilde{\mathcal{B}}$ for some $\tilde{\mathcal{B}} \in \mathbb{R}_+$.

Let us clear out these statements. First, $\tilde{p}_i^{US} \leq 0$ because $\tilde{p}_i^{VCG} \leq 0$ and f^{-1} is increasing. Theoretically, we could have just put $\Big[\max_{\alpha_i \in \Delta^m \times \mathbb{R}} p_i^{VCG}(\alpha_{-i}, \alpha_i)\Big]$ instead of $\frac{\gamma^2}{n}\Big(\|\mathcal{D}_{V\circ g}(\bar{\alpha}_{-i})\| + 1\Big)$ in the above definition, however we cannot guarantee that this maximum can be found efficiently. Note that $\frac{\gamma^2}{n}\Big(\|\mathcal{D}_{V\circ g}(\bar{\alpha}_{-i})\| + 1\Big)$ is almost exactly the upper bound we put on p_i^{VCG} in the proof of Theorem 3, the only difference is that we now use the derivative at $\bar{\alpha}_{-i}$ to evaluate the difference $V(g(\bar{\alpha}_{-i})) - V(g(\bar{\alpha}))$, instead of the derivative at $\bar{\alpha}$. Obviously, that is just as valid and we could have similarly reached that bound. The reason for this substitution is that the payments formula must be independent of α_i to maintain Incentive Compatibility. Now, While $\mathcal{D}_{V\circ g}(\bar{\alpha}_{-i})$ is different for every agent i, we know that these are bounded globally because $\bar{\alpha}_{-i} \to \bar{\alpha}$ and the Jacobian matrix is a continuous function. Thus $\tilde{p}_i^{VCG} = O(\tilde{p}_i^{US}) = O(1/n)$.

D Bias Functions That Preserve SDSIC

D.1 Phantom-Agents

We shall define our special bias function in the form of utility functions of fictitious agents that favour our targeted outcomes. In the simplest case where we have a sole budget decision in mind we can take $\mathcal{C}(x,t) = \lambda \cdot v_{\tilde{\alpha}}(x,t)$ where $\tilde{\alpha}$ is chosen so that $g(\tilde{\alpha})$ is the desired outcome and $\lambda > 0$ indicates the extent of bias we want. The SDSIC then extends trivially, as the our new biased mechanism is essentially the US-VCG mechanism for the n real agents plus λn fictitious ones with mean preference $\tilde{\alpha}$. However, diverting the mechanism towards a set of outcomes while maintaining SDSIC will not be as simple (Note that if we add fictitious agents of multiple types, their impact on the outcome is ultimately determined only by the mean of fictitious preferences). The bias function we introduce below (Def. 15) is composed as the sum of two functions, one that favours certain tax decisions and another that targets specific allocations $\hat{x}^t$ for any given $t \in [-\frac{B_0}{n}, \infty)$. That separation corresponds to the nature of our optimization problem, that can be solved in two steps accordingly. Nonetheless, as we do not demand that $\hat{x}^t$ necessarily exists for all t, i.e., it is possible that under some tax decisions no allocation is favourable, it should not harm its generality. Meaning, any arbitrary set of outcomes $W \subset \Delta^m \times [-\frac{B_0}{n}, \infty)$ can be targeted this way. We first define the following once $\{x^t\}_t$ are chosen.

Definition 14. For any choice of $\{\hat{x}^t\}_t$ and for every $t \in (-b_0, \infty)$ such that $\hat{x}^t$ exists, we define the *corresponding phantom type* $\hat{\alpha}^t = (\hat{\alpha}^t_1, \ldots, \hat{\alpha}^t_m)$ and *valuation function* $\hat{\Theta}^t$ such that

$$\arg\max_{x \in \Delta^m} \hat{\Theta}^t(xB_t) = \arg\max_{x \in \Delta^m} \sum_j \hat{\alpha}^t_j \theta_j(x_j B_t) = \hat{x}^t$$

and if $\hat{x}^t$ does not exist put $\hat{\alpha}^t = \vec{0}$.

For all t, once $\hat{x}^t$ is chosen then $\hat{\alpha}^t$ is defined by the first order conditions derived from the optimization that defines $\hat{x}^t$.

$$\frac{\theta'_j(\hat{x}_j B_t)}{\theta'_k(\hat{x}_k B_t)} = \frac{\hat{\alpha}^t_k}{\hat{\alpha}^t_j} \;\forall j \in [m]$$

Now, we define our bias function as follows.

Definition 15 (Phantoms Bias function)**.** For any choice of $\{\hat{x}^t\}_t$ and for every $\lambda > 0$, define

$$\mathcal{P}_\lambda(x,t) := \lambda\big(\hat{\Theta}^t(xB_t) - \hat{\Theta}^t(\hat{x}^t B_t)\big) + \psi(t)$$

where $\psi : (-\frac{B_0}{n}, \infty) \mapsto \mathbb{R}$ is continuous and $\lim_{t\to\infty} \psi(t) = 0$.

That is, for any given t we add the (non-positive) utility loss of λ fictitious agents that favour $\hat{x}^t$ over x (but have no preferences regarding the tax t), plus $\psi(t)$ that expresses the designer's preference on tax decisions. $\lim_{t\to\infty} \psi(t) = 0$ assures that $\hat{g}$ is well defined, i.e. that $v_\alpha(x,t) + \mathcal{P}_\lambda(x,t)$ has a global maximum for all $\alpha \in \Delta^m \times [-\frac{B_0}{n}, \infty)$.

Example 7 (Running example with bias to equitable allocation)**.** We show an example for executing the BUS-VCG mechanism with $\hat{x}^t$ taken as the equitable allocation

$$\hat{e}^t := \arg\min_x \Big(\max_{\alpha \in \Delta^m} \sum_j \alpha_j \theta_j(x_j B_t) - \min_{\alpha \in \Delta^m} \sum_j \alpha_j \theta_j(x_j B_t)\Big)$$

. In case that θ_j are identical for all j then $\hat{e}^t$ and $\hat{\alpha}^t$ are trivial, $\hat{e}^t_j = \hat{\alpha}^t = 1/m \ \forall j$. Thus,

$$\mathcal{P}_\lambda(x,t) = \lambda\Big(10\sum_j \frac{1}{m}\ln(x_j \cdot 3t) - 10\sum_j \frac{1}{m}\ln\Big(\frac{1}{m}\cdot 3t\Big)\Big)$$

To compute $\hat{g}(\hat{\alpha}) = (\hat{x},\hat{t})$, let us first find $\hat{x}$ w.r.t. any fixed t. When t is fixed, solving

$$\begin{cases}\max \ v_{\bar{\alpha}}(x,t) + 3\mathcal{P}_\lambda(x,t) \\ x \in \Delta^m\end{cases}$$

is just as solving the original problem, only for a modified preferences mean $\beta_j = \frac{\bar{\alpha}_j + \lambda/m}{1+\lambda/m} \forall j \in [m]$, and with logarithmic functions we know that $\hat{x}_j = \beta_j \ \forall j$, independently of t. Now we need to solve

$$\begin{cases}\max \ v_{\bar{\alpha}}(\hat{x},t) + 3\mathcal{P}_\lambda(\hat{x},t) \\ t \in [-\frac{B_0}{n},\infty)\end{cases}$$

but, note that $\mathcal{P}(\hat{x},t)$ is a constant function of t, thus not affecting $\hat{t}$. To conclude, the introduction of $\mathcal{P}_\lambda$ shifts the allocation from $x^*_j = \bar{\alpha}_j \forall j$ to $\hat{x}_j = \frac{\bar{\alpha}_j+\lambda/m}{1+\lambda/m}\forall j$ while not affecting the tax decision $t^* = \hat{t} = \big(\frac{2\cdot 10}{3\bar{\alpha}_f}\big)^2$.

As mentioned above, to maintain SDSIC we also have to see that $\hat{g}(\alpha)$ defines α uniquely. That requires further assumption on the smoothness of $\mathcal{P}_\lambda$.

Lemma 4. *Assume that $\mathcal{P}_\lambda(x,t)$ is differentiable in t. Then for any two distinct preferences vectors $\alpha \neq \beta \in \Delta^m \times \mathbb{R}$, $\hat{g}(\alpha) \neq \hat{g}(\beta)$.*

Proof. Fix any $\alpha \in \Delta^m \times [-\frac{B_0}{n},\infty)$, and let

$$\hat{g}(\alpha) = (x^*,t^*) \in \underset{(x,t)\in\Delta^m\times[-\frac{B_0}{n},\infty)}{\arg\max} \quad v_\alpha(x,t) + \mathcal{P}_\lambda(x,t).$$

Then in particular, x^* solves $\max_{x\in\Delta^m} v_\alpha(x,t^*) + \mathcal{P}_\lambda(x,t^*)$. Since

$$\begin{aligned} v_\alpha(x,t^*) + \mathcal{P}_\lambda(x,t^*) = \\ &\sum_j \alpha_j\theta_j(x_jB^*_t) - \alpha_f f(t^*) + \lambda\sum_j \hat{\alpha}^{t^*}_j \theta_j(x_j {B_t}^*) \\ &-\lambda\sum_j \hat{\alpha}^{t^*}_j\theta_j(\hat{x}^{t^*}_j B_{t^*}) + \psi(t^*), \end{aligned}$$

x^* also solves $\max_{x\in\Delta^m}\sum_j(\alpha_j + \lambda\hat{\alpha}^{t^*}_j)\theta_j(x_jB_{t^*})$ because once we fix t^* all the remaining terms are just constants. By the proof of Lemma **??**, that problem has a unique solution for every $(\alpha_1,\ldots,\alpha_m)\in\Delta^m$ and moreover, two distinct vectors cannot share the same solution. We thus conclude that $\hat{g}(\alpha) = \hat{g}(\beta) \implies \alpha_j = \beta_j \ \forall j \in [m]$.
Now, let $\alpha,\beta \in \Delta^m \times [-\frac{B_0}{n},\infty)$ where $\alpha_j = \beta_j \ \forall j \in [m]$ and assume that $\hat{g}(\alpha) = \hat{g}(\beta) = (x^*,t^*)$. Then

$$t^* \in \arg\max_t v_\alpha(x^*,t) + \mathcal{P}_\lambda(x^*,t)$$
$$\text{and } t^* \in \arg\max_t v_\beta(x^*,t) + \mathcal{P}_\lambda(x^*,t)$$

Since $\alpha_j = \beta_j \ \forall j \in [m]$, we can write

$$v_\alpha(x^*, t) + \mathcal{P}_\lambda(x^*, t) = \Gamma(t) - \alpha_f f(t);$$
$$v_\beta(x^*, t) + \mathcal{P}_\lambda(x^*, t) = \Gamma(t) - \beta_f f(t)$$

for some $\Gamma : (-b_0, \infty) \mapsto \mathbb{R}$ that is differentiable by our assumption on $\mathcal{P}$ and the initial assumptions on $\theta_j,\ j \in [m]$. Thus,

$$\Gamma'(t^*) - \alpha_f f(t^*) = \Gamma'(t^*) - \beta_f f(t^*) = 0 \implies \alpha_f = \beta_f$$

$\square$

Corollary 6. For all $\lambda \geq 0$ and for every bias function $\mathcal{P}_\lambda(x, t)$ that is differentiable in t, the BUS-VCG mechanism is SDSIC in every budgeting instance that the US-VCG is.