

Redistribution of VCG Payments in Public Project Problems^{*}

Victor Naroditskiy¹ Mingyu Guo² Lachlan Dufton³ Maria Polukarov¹
Nicholas R. Jennings¹

¹ School of Electronics and Computer Science, University of Southampton, UK
`{vn,mp3,nrj}@ecs.soton.ac.uk`

² Computer Science Department, University of Liverpool, UK
`Mingyu.Guo@liverpool.ac.uk`

³ Cheriton School of Computer Science
University of Waterloo, Canada
`ltdufton@cs.uwaterloo.ca`

Abstract. Redistribution of VCG payments has been mostly studied in the context of resource allocation. This paper focuses on another fundamental model—the public project problem. In this scenario, the VCG mechanism collects in payments up to $\frac{n-1}{n}$ of the total value of the agents. This collected revenue represents a loss of social welfare. Given this, we study how to redistribute most of the VCG revenue back to the agents. Our first result is a bound on the best possible efficiency ratio, which we conjecture to be tight based on numerical simulations. Furthermore, the upper bound is confirmed on the case with 3 agents, for which we derive an optimal redistribution function. For more than 3 agents, we turn to heuristic solutions and propose a new approach to designing redistribution mechanisms.

1 Introduction

Public good or public project problems refer to situations where a group of agents need to decide whether or not to undertake a project or to procure a good. The project is “public” in the sense that everyone will enjoy the benefits of it. A typical example is a community deciding to build a bridge. If the bridge is built, everyone will be able to cross it. The challenge in deciding whether or not the bridge should be built, lies in learning how much the people need the bridge. Each person has a value for the bridge, but this value is known to him alone. The *efficient* outcome is to build the bridge if and only if the total value exceeds the cost of the bridge. Public project problems have been studied extensively in both economics and computer science literature (see, e.g., [9–11, 7, 1]).

In this context, we are interested in mechanisms that satisfy dominant-strategy incentive compatibility (DSIC), and maximize social welfare. The social

^{*} Naroditskiy, Polukarov, and Jennings gratefully acknowledge funding from the UK Research Council for project ORCHID, grant EP/I011587/1.

welfare is measured as the sum of the utilities of the agents. Specifically, any payments collected from the agents reduce social welfare. Some payments however are required by DSIC. The objective of social welfare is natural in public good provision problems: after all, the benefits are to be enjoyed by all non-exclusively, and public projects are normally undertaken in the interest of the participants. This is in contrast to private goods that are often sold to generate profit for the auctioneer (there is no auctioneer or residual claimant in the public good problems we consider).

Our focus here is on mechanisms that are efficient and weakly budget-balanced (i.e., do not require an external subsidy). The latter restriction is necessary, as otherwise one can achieve infinite social welfare by providing an infinite subsidy to the agents. To this end, mechanisms from the Groves class align the incentives of the agents with the objective of choosing the efficient outcome. Specifically, under a Groves mechanism, each agent prefers reporting her value truthfully regardless of the reports of the other agents. In fact, Groves mechanisms are the only mechanisms that are dominant-strategy incentive compatible (or, *truthful*) and efficient for public project problems [8]. The mechanisms within the Groves class differ in the amount of payment collected from the agents. An easy way to describe this class of mechanisms is through the most prominent Groves mechanism—the VCG mechanism: payment made by agent i under any Groves mechanism can be represented as the payment collected by the VCG mechanism minus a redistribution $h_i(v_{-i})$, which is a function of other agents’ values. For efficient mechanisms without an auctioneer, the objective of maximizing social welfare is equivalent to the objective of minimizing the revenue collected. Under this objective, the VCG mechanism has a very poor performance (i.e., collects a lot of revenue) as we detail next. Therefore, the question we study in this paper is how to design the *redistribution* functions that maximize social welfare.

We do not assume any prior on agent valuations and we evaluate mechanisms based on the worst-case performance over all possible value profiles. Following previous work on redistribution in resource allocation settings (e.g., [12, 6]), we make the performance metric unit free by measuring the performance as a percentage of the value of the efficient outcome achieved. We will refer to this metric as the *ratio*. Since there are no external subsidies, the value of the efficient outcome is the highest welfare that can be achieved, had all values been publicly known. Thus, the highest possible ratio is one.

The ratio of the VCG mechanism is $\frac{1}{n}$, where n is the number of agents [7]. In this paper, we derive an upper bound on the optimal ratio. Unlike the ratio of VCG, which decreases with n , the upper bound increases with n . We conjecture the bound to be tight based on numerical simulations. Further, for the case of $n = 3$, we find an optimal mechanism which guarantees the upper bound ratio of $\frac{2}{3}$. Finally, we propose a general heuristic-based approach for deriving redistribution mechanisms. Using a simple sampling-based heuristic, we obtain a mechanism whose ratio is higher than that of VCG for $n = 4, 5, 6$.

Our work is related to, and builds upon, some recent research on redistribution mechanisms. The public good model and, in particular, the valuation

function of the agents are the same as in [7]. There, non-efficient but strongly budget-balanced mechanisms are considered. The authors discuss a randomized allocation function that guarantees a high expected ratio, while restricting the payments to add up to zero. In contrast, here we study *deterministic* mechanisms optimizing only over the payment functions, while the allocation rule is fixed to choose the efficient allocation. Our upper bound results suggest that full social welfare may be achievable asymptotically without resorting to randomized mechanisms.

Other work in various allocation settings has studied the problem of finding payments for Groves mechanisms that are optimal in terms of social welfare. In particular, Moulin [12] and Guo and Conitzer [6] simultaneously derived the optimal redistribution for allocating identical items to agents with unit demand. The results were further extended to multi-unit demand in [6]. An optimal Groves mechanism for allocating heterogeneous items was derived in [5]. General techniques have also been proposed for optimizing payments according to the mechanism designer’s objectives, for single-parameter and multi-parameter domains [13, 4]. In fact, we make use of a heuristic technique from [4] to derive an optimal solution for $n = 3$.

There are also other redistribution mechanisms aiming to minimize payments that can be applied to the public good setting. Bailey [2] proposed a redistribution mechanism for public good problems, but under the worst-case analysis it is not weakly budget-balanced. While the mechanism proposed by Cavallo [3] is efficient and weakly budget-balanced, it cannot redistribute any VCG revenue in public good problems [7]. In this paper, we propose weakly budget-balanced mechanisms that do redistribute some of the VCG revenue, which increases social welfare without requiring external subsidy.

The rest of the paper is organized as follows. We present the model in Section 2. A conjecture about the optimal ratio is derived analytically in Section 3. The optimal solution to the case with $n = 3$ is presented in Section 4. We then propose a heuristic-based approach for deriving redistribution mechanisms and analyze the resulting mechanism’s performance in Section 5. Section 6 relaxes the assumption that allowed us to restrict the value space while deriving prior results. We conclude and discuss directions for future work in Section 7.

2 The Model

There are n agents deciding whether or not to undertake a project, such as building a bridge. The cost of the bridge is C , which is commonly known. Each agent has a private type $\theta_i \geq 0$ denoting how much he will benefit if the bridge is built. We will assume $\theta_i \in [0, C]$, and will demonstrate in Section 6 that it is without loss of generality to consider types that are bounded from above by C . Also, without loss of generality, we can assign labels to the agents so that agent 1 is the agent with the highest value, agent 2—with the second highest, etc. Thus, $C \geq \theta_1 \geq \theta_2 \geq \dots \geq \theta_n \geq 0$, and we denote the space of agent values by $\Theta = \{\theta \in [0, C]^n \mid C \geq \theta_1 \geq \theta_2 \geq \dots \geq \theta_n \geq 0\}$. A mechanism for this problem

consists of the outcome and the payment functions. The outcome is denoted by $k(\theta) \in \{0, 1\}$ with $k(\theta) = 1$ if the bridge is built, and $t_i(\theta) \in \mathbb{R}$ are payments made by each agent i . We fix k to be the efficient rule: $k(\theta) = 1$ iff $\sum_i \theta_i > C$.

The value of each agent depends on his type and whether or not the bridge is built. Following [7], we define the value of the efficient outcome as follows:

$$s(\theta) = \max\left(\sum_i \theta_i, C\right)$$

This definition corresponds to the interpretation that if the bridge is not built, the agents get to distribute C among themselves (or, equivalently, they do not spend C on the bridge). This is reflected in the valuation function, which lets each agent keep $\frac{C}{n}$ if the bridge is not built:

$$v_i(k(\theta), \theta_i) = \begin{cases} \theta_i & \text{if } k(\theta) = 1 \\ \frac{C}{n} & \text{otherwise} \end{cases}$$

Utility of agent i is quasi-linear in the payment $t_i \in \mathbb{R}$ collected from him:

$$u_i(\theta) = v_i(k(\theta), \theta_i) - t_i(\theta)$$

Without loss of generality, for efficient and dominant-strategy incentive compatible mechanisms, we focus on the Groves class. Furthermore, we focus on Groves mechanisms that are anonymous, which, for our objective of maximizing worst-case performance (see Equations 2 and 4), is without loss of generality [1]. These mechanisms implement the efficient outcome, $k(\theta) = 1$ iff $\sum_i \theta_i > C$. Note that $\sum_i v_i(k(\theta), \theta_i) = s(\theta)$ for the efficient $k(\theta)$. DSIC is achieved by selecting t_i that aligns an agent's utility with the goal of selecting the efficient outcome:

$$t_i(\theta) = v_i(k(\theta), \theta_i) - s(\theta) + h(\theta_{-i})$$

which yields

$$u_i(\theta) = s(\theta) - h(\theta_{-i}) \tag{1}$$

where $h : W \rightarrow \mathbb{R}$ is an arbitrary function of the values of the agents other than the agent whose redistribution (or rather, *rebate*) is computed.⁴ Here, domain $W = \{w \in [0, C]^{n-1} \mid C \geq w_1 \geq w_2 \geq \dots \geq w_{n-1} \geq 0\}$ of rebate function h (which we will also term the *rebate space*) refers to the space of values of $n - 1$ agents (other than i). Importantly, the second term of utility, $h(\theta_{-i})$, characterizes all mechanisms within the Groves class. Our goal is to choose function h that maximizes social welfare subject to the constraint of weak budget balance.

Weak budget balance constraint means that the sum of payments made by the agents must be non-negative:

$$\sum_i t_i(\theta) = \sum_i (v_i(k(\theta), \theta_i) - s(\theta) + h(\theta_{-i})) = \sum_i h(\theta_{-i}) - (n - 1)s(\theta) \geq 0 \quad \forall \theta$$

⁴ Under Equation 1, h is the function that determines how much of the value of the efficient outcome the agent should pay back. This is hardly a redistribution/rebate, but we keep this terminology to be consistent with prior literature.

Next, we describe how the performance of a mechanism is measured. A mechanism guarantees the ratio r if the following holds:

$$\sum_i u_i(\theta) = ns(\theta) - \sum_i h(\theta_{-i}) \geq rs(\theta) \quad \forall \theta$$

Stated formally, we seek to solve the following optimization problem:

$$\max_{h:W \rightarrow \mathbb{R}, r \in \mathbb{R}} r \tag{2}$$

$$\sum_i h(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \in \Theta \tag{3}$$

$$ns(\theta) - \sum_i h(\theta_{-i}) \geq rs(\theta) \quad \forall \theta \in \Theta \tag{4}$$

In words, we are looking for a mechanism with the highest ratio (Equations 2 and 4) that is weakly budget-balanced (Equation 3). Note that both constraints can be written in one line as

$$(n-r)s(\theta) \geq \sum_i h(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \tag{5}$$

3 Optimal Ratio (conjecture)

In this section, we describe an interesting structure of the optimization problem (2)-(4). The problem has an infinite number of constraints, but our numerical results showed that it is sufficient to consider only $n+1$ constraints to obtain an upper bound on the ratio, such that this ratio does not change when we add additional constraints (of course, we were only able to check finite sets of constraints). This provides numerical evidence that the upper bound is tight. Furthermore, we derive this upper bound in closed form, which we show in the rest of this section.

First, we discuss how the ratio can be upper bounded computationally using the technique `RestrictedProblem` from [4]. The idea is to solve the problem while only enforcing a finite subset of constraints. The solution may violate some of the excluded constraints, thus providing an upper bound on the objective value (we are considering a maximization problem). In more detail, the optimization problem (2)-(4) has an infinite number of constraints (one for each $\theta \in \Theta$) and optimizes over functions (equivalently, there is an infinite number of variables—a rebate $h(w)$ for each $w \in W$). To make the problem more manageable, we limit the space of value profiles to a finite subset $\hat{\Theta} \subset \Theta$. Notice that once the set of profiles is finite, the set of rebates that appear in the constraints is also finite. It can be obtained by “projecting” each value profile into n profiles by removing one of the elements while keeping the rest. For example, when we restrict the value space to the set of profiles $\hat{\Theta} = \{(a, b, c), (d, e, f)\}$, the relevant rebates are defined for each profile in $\hat{W} = \{(b, c), (a, c), (a, b), (e, f), (d, f), (d, e)\}$. The constraints (3) and (4) appear once for each value profile, and the number of

variables is $|\hat{W}|$. With these restrictions, the optimization problem in (2)-(4) becomes a linear program, which we implemented and solved using CPLEX.

Clearly, the choice of the enforced constraints as governed by $\hat{\Theta}$ determines the quality of the upper bound. Adding more constraints can only improve the bound. Interestingly, we find that considering only $n + 1$ “important” profiles gave the best upper bound we could find among all sets of $\hat{\Theta}$ that we tried. In more detail, for a given n , we obtained the profiles $\hat{\Theta}$ by discretizing the space of values an agent may have. For example, discretizing into $z + 1$ possible values we get $\theta_i \in \{j \frac{C}{z}\}_{j=0}^z$. Without loss of generality we set $C = 1$, and focus on $\theta_i \in \{j \frac{1}{z}\}_{j=0}^z$. Looking deeper into the patterns, we observed an interesting structure, that let us characterize the upper bound analytically.

The best upper bound we observed numerically across n was obtained when solving the restricted problem with the following $n + 1$ value profiles: the zero profile and the profiles $(\underbrace{\frac{1}{b}, \dots, \frac{1}{b}}_k, \underbrace{0, \dots, 0}_{n-k})$ where b is the integer part of $\frac{n}{2}$ and

$1 \leq k \leq n$. For example, for $n = 5$ we have $b = 2$, and the profiles $(0, 0, 0, 0, 0)$, $(\frac{1}{2}, 0, 0, 0, 0)$, \dots , $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. We refer to these $n + 1$ profiles as *important profiles*. Next we provide an optimal solution to the restricted problem analytically.

Theorem 1. *No mechanism can achieve a ratio above r .*

$$r = 1 - \left(2 + \frac{2(\frac{n}{2}!)^2}{n} \sum_{j=0}^{\frac{n-4}{4}} \frac{(3n-4j)}{(2j)!(n-2j)!} \right)^{-1} \quad n = 4, 8, 12, \dots \quad (6)$$

$$r = 1 - \left(2 + \frac{2(\frac{n}{2}!)^2}{n} \sum_{j=0}^{\frac{n-2}{4}} \frac{(3n-4j-2)}{(2j+1)!(n-2j-1)!} \right)^{-1} \quad n = 6, 10, 14, \dots \quad (7)$$

$$r = 1 - \left(\frac{n(\frac{n-1}{2}!)^2}{n-1} \sum_{j=0}^{\frac{n-1}{4}} \frac{(n+4j-1)}{(2j)!(n-2j)!} \right)^{-1} \quad n = 5, 9, 13, \dots \quad (8)$$

$$r = 1 - \left(\frac{n(\frac{n-1}{2}!)^2}{n-1} \sum_{j=0}^{\frac{n-3}{4}} \frac{(n+4j+1)}{(2j+1)!(n-2j-1)!} \right)^{-1} \quad n = 3, 7, 11, \dots \quad (9)$$

Proof. The proof is available in the full version of the paper. □

Considering much larger sets of value profiles never improved the bound. This leads us to believe that the bound is tight. Furthermore, performing sensitivity analysis revealed that only the constraints used to derive the bound were tight in optimal solutions to restricted problems that included supersets of the important profiles. If the ratio is indeed tight, then we also have optimal rebates for the n rebate profiles used in deriving the bound: these values are unique, and thus, they cannot change in a solution that achieves the bound.

Observing the behavior of this upper bound (see Figure 1), we see that it approaches 1 as the number of agents increases. Thus, if this bound on the ratio is tight, then an optimal mechanism for the public project problem will have a

loss of social welfare approaching zero with additional agents. This is in contrast to the VCG mechanism, which has an overall social welfare of $\frac{1}{n}$ that approaches zero as the number of agents increase [7].




Fig. 1. The loss, $1 - r$, approaches zero as the number of agents increases.

4 Optimal Redistribution for $n = 3$

For the case of $n = 3$, we obtain an optimal redistribution function. It was derived using techniques described in [4]. We provide the details next.

The upper bound linear program described in Section 3 can be modified to produce a heuristic redistribution function using another technique from [4]. The idea is to optimize over the space of rebate functions that are piecewise linear within a specified set of regions. The algorithm `LinearRebates` described in [4] takes a subdivision of the rebate space into regions and produces a redistribution function (and the ratio it achieves) that is optimal over all rebate functions that are linear within these regions. We use `LinearRebates` with the subdivision shown in Figure 2 to obtain a redistribution function. This piecewise linear function is composed of linear functions for each of the 4 regions

$$h(w) = \frac{2}{3}C + \begin{cases} 0 & \text{if } w \in \text{region 0} \\ \frac{2}{3}w_1 + \frac{2}{3}w_2 - \frac{C}{3} & \text{if } w \in \text{region 1} \\ \frac{1}{3}w_1 + \frac{2}{3}w_2 - \frac{C}{6} & \text{if } w \in \text{region 2} \\ \frac{7}{6}w_1 + \frac{3}{2}w_2 - C & \text{if } w \in \text{region 3} \end{cases}$$

This function can be represented more compactly. Let $s(\theta, C) = \max(\sum_i \theta_i, C)$, denote the value of the efficient outcome for agents defined by value profile θ and some total cost C . The optimal piecewise linear redistribution function is

$$h(\theta_{-i}) = \frac{5}{6}s(\theta_{-i}, C) + \frac{2}{3}s(\theta_{-i}, \frac{C}{2}) - \frac{1}{3}s(\theta_{-i}^1, \frac{C}{2}) - \frac{C}{3} \quad (10)$$

where θ_{-i}^1 refers to the first element of the vector θ_{-i} .

The ratio obtained by this function is $\frac{2}{3}$. However, $\frac{2}{3}$ is also the upper bound on the ratio as computed in Equation 9. This means that the rebate function we found is optimal.

We next provide an interpretation of the rebate function, which may help generalize it to more than 3 agents. In the analytical form used in Equation 10 to express the function, each region boundary of the subdivision is encoded in a single $s(\cdot)$ term. Note that, without the coefficient, the first term is the rebate agent i would receive in a normal VCG mechanism. The second term is the VCG rebate for a project with cost $\frac{C}{2}$. The first two terms are piecewise linear, with boundaries at $\sum_{j \neq i} \theta_j = C$ and $\sum_{j \neq i} \theta_j = \frac{C}{2}$, respectively. In Figure 2, these are the region-2-3 boundary and region-0-1 boundary, respectively. Finally, since we assume agents are sorted, the max-valued agent in the third term is always agent w_1 , and this third term is piecewise linear, with a boundary at $\max_{j \neq i} \theta_j = \frac{C}{2}$, i.e. the region-1-2 boundary.

The next step is to generalize the rebate function above to problems with more than 3 agents. One way to do this is through finding a subdivision of the rebate space such that an optimal mechanism for this subdivision improves over the VCG mechanism. However, generalizing the subdivision in Figure 2 to 3- or higher dimensional rebate spaces proved elusive, and the question remains open.

5 Heuristic-based Redistribution

In the previous section, we have solved for an optimal mechanism for 3 agents. However, when there are more agents, we do not yet know how to solve for the optimal mechanisms. Given this, we propose a new heuristic-based approach for designing weakly budget-balanced mechanisms with high social welfare. By using a simple sampling-based heuristic, we derive the *sampling-based redistribution (SBR)* mechanism. We show that SBR's ratio is higher than that of VCG for $n = 4, 5, 6$, and conjecture that this is still the case for $n > 6$. Both the heuristic-based approach and the SBR mechanism are general enough that they may potentially be used in settings other than public project problems.

Our approach builds on the Cavallo mechanism [3], which works as follows: We first run VCG. Besides paying the VCG payment, agent i also receives

$$\frac{1}{n} \min_{\theta'_i} VCG(\theta'_i, \theta_{-i})$$

Here, $VCG(\theta'_i, \theta_{-i})$ represents the total VCG payment for the profile under which agent i reports θ'_i , and the other agents report θ_{-i} . In words, agent i




Fig. 2. Subdivision of the space θ_{-i} for 3 agents. The rebate function $h(\theta_{-i}) = h(w_1, w_2)$ is linear within each of the 4 regions.

receives $\frac{1}{n}$ times the *minimal possible* total VCG payment given that the other agents report θ_{-i} . Since the additional amount agent i receives is independent of her own type, the Cavallo mechanism is dominant-strategy incentive compatible. Then, since every agent at most receives $\frac{1}{n}$ times the actual total VCG payment, the Cavallo mechanism is weakly budget-balanced. In many settings (*e.g.*, resource allocation with free disposal and public good provision), VCG is pay-only. In these settings, the additional amount an agent receives is non-negative. Unfortunately, for our model, the additional amount an agent receives is always 0.⁵ That is, the Cavallo mechanism always coincides with VCG.

Our heuristic-based approach works as follows:

- We start with a dominant-strategy incentive-compatible mechanism (*e.g.*, VCG). Let $P(\theta)$ be the total payment under this mechanism for profile θ .
- Besides paying the payment under the initial mechanism, agent i also receives

$$\frac{1}{n} EM(\theta_{-i})$$

Here, $EM(\theta_{-i})$ represents agent i 's estimation of the total payment under the initial mechanism, given that the others report θ_{-i} . Agent i 's estimation should not depend on her own report, which is to maintain dominant-strategy incentive compatibility. The estimation function EM can be based on any heuristic. (One naive choice would be $EM(\theta_{-i}) = P(0, \theta_{-i})$, which uses the total payment assuming $\theta_i = 0$ to be the estimation.) The goal of this step is to modify the initial mechanism, so that it becomes as close to strong budget

⁵ $\min_{\theta'_i} VCG(\theta'_i, \theta_{-i})$ is always 0 [1]: if $\sum_{j \neq i} \theta_j \geq \frac{n-1}{n}C$, then set θ'_i to be C ; otherwise, set θ'_i to be 0.

balance as possible. Generally, we cannot achieve perfect budget balance. That is, even if EM is based on a good heuristic, the mechanism at this point still incurs some small amount of waste or deficit, depending on the profile.

- To ensure weak budget balance, we finally collect from every agent $\frac{1}{n}$ times the *maximum possible* deficit, given the heuristic that we use (EM) and given the other agents' reports. Formally, we collect from agent i the following amount:

$$\frac{1}{n} \max_{\theta'_i} \left\{ \sum_j \frac{1}{n} EM(\hat{\theta}_{-j}) - P(\hat{\theta}) \right\}$$

Here, $\hat{\theta}$ represents the profile (θ'_i, θ_{-i}) . It should be noted that this step is based on exactly the idea behind the Cavallo mechanism. Dominant-strategy incentive compatibility is maintained because the amount we charge from an agent does not depend on her own report. Furthermore, since the total amount we charge is never less than the actual deficit, the resulting mechanism is weakly budget-balanced.

We start with VCG, by using a simple sampling-based heuristic, we obtain a specific mechanism, which we call the *sampling-based redistribution (SBR)* mechanism. In detail, to estimate the total VCG payment given the others' report θ_{-i} , we just assume that agent i 's type is drawn uniformly at random from θ_{-i} , and then use the expected total VCG payment as the estimation. Formally, EM is defined as follows:

$$EM(\theta_{-i}) = \frac{\sum_{j \neq i} VCG(\theta_j, \theta_{-i})}{n-1}$$

Next, we show how to derive a lower bound on the ratio of SBR. Without loss of generality, we let $C = 1$.

The social welfare under SBR is:

$$s(\theta) - VCG(\theta) + \sum_i \frac{1}{n} EM(\theta_{-i}) - \sum_i \frac{1}{n} \max_{\theta'_i} \left\{ \sum_j \frac{1}{n} EM(\hat{\theta}_{-j}) - VCG(\hat{\theta}) \right\}$$

We have:

$$-VCG(\theta) + \sum_i \frac{1}{n} EM(\theta_{-i}) \geq \min_{\theta} \left\{ \sum_i \frac{1}{n} EM(\theta_{-i}) - VCG(\theta) \right\}$$

Also,

$$\begin{aligned} \sum_i \frac{1}{n} \max_{\theta'_i} \left\{ \sum_j \frac{1}{n} EM(\hat{\theta}_{-j}) - VCG(\hat{\theta}) \right\} &\leq \sum_i \frac{1}{n} \max_{\theta} \left\{ \sum_j \frac{1}{n} EM(\theta_{-j}) - VCG(\theta) \right\} \\ &= \max_{\theta} \left\{ \sum_i \frac{1}{n} EM(\theta_{-i}) - VCG(\theta) \right\} \end{aligned}$$

We use $EMVCG(\theta)$ to denote $\sum_i \frac{1}{n} EM(\theta_{-i}) - VCG(\theta)$. The social welfare under SBR is then at least:

$$s(\theta) + \min_{\theta} EMVCG(\theta) - \max_{\theta} EMVCG(\theta)$$

The ratio of SBR is then:

$$\frac{s(\theta) + \min_{\theta} EMVCG(\theta) - \max_{\theta} EMVCG(\theta)}{s(\theta)} \geq 1 + \min_{\theta} EMVCG(\theta) - \max_{\theta} EMVCG(\theta)$$

(We recall that $s(\theta)$ is at least $C = 1$.)

Given n , $\min_{\theta} EMVCG(\theta)$ and $\max_{\theta} EMVCG(\theta)$ are constants. For small n , we can numerically solve for their values. Specifically, instead of minimizing/maximizing over all possible profiles, we only consider profiles where every agent's report is an integer multiple of $1/N$. Larger values of N generally correspond to more accurate results. We notice that as long as N is a multiple of $2n$ (e.g., $N = 2n, N = 4n, \dots, N = 100n$), we always end up with the same maximizing/minimizing profiles. To double check, for every maximizing/minimizing profile obtained, we generate 10,000 random vectors, and perturb the profile along these 10,000 directions. At the end, no perturbation ever leads to higher maximum or lower minimum. The results are presented in the following table. We only considered $n \leq 6$ due to the exponential complexity of this approach.

	$n = 3$	$n = 4$	$n = 5$	$n = 6$
$\max_{\theta} EMVCG(\theta)$	2/9	1/4	6/25	2/9
$\arg \max_{\theta} EMVCG(\theta)$	(1, 1, 0)	(1, 1, 0, 0)	(1, 1, 0, 0, 0)	(1, 1, 0, 0, 0, 0)
$\min_{\theta} EMVCG(\theta)$	-4/9	-19/48	-2/5	-23/60
$\arg \min_{\theta} EMVCG(\theta)$	(1, 0, 0)	(5/8, 3/8, 0, 0)	(3/5, 2/5, 0, 0, 0)	(7/12, 5/12, 0, 0, 0, 0)
$1 + \min_{\theta} EMVCG(\theta)$	1/3	17/48	9/25	71/180
$-\max_{\theta} EMVCG(\theta)$	≈ 0.333333	≈ 0.354167	≈ 0.360000	≈ 0.394444

There are two interesting observations. First, at least for $3 \leq n \leq 6$, the lower bound of the ratio of SBR increases as n increases. We conjecture that this trend remains when n is greater than 6. Second, when $n = 3$, the lower bound of the ratio of SBR is the same as VCG's ratio ($1/n$), and when $4 \leq n \leq 6$, the lower bound of the ratio of SBR is higher than VCG's ratio.

Finally, it should be noted that even though we do not know how to estimate the ratio of SBR when $n > 6$, we do know that SBR is always dominant-strategy incentive-compatible and weakly budget-balanced. Also, SBR's payments are computationally easy to calculate. Therefore, we can always apply it. It is just that for $n > 6$, we do not know how well it will perform. We tried to experimentally evaluate the ratio of SBR for larger values of n . For example, for $n = 10$, we randomly generated 1,000,000 profiles (every agent's type is drawn from *i.i.d.* uniform distribution from 0 to 1). For these profiles, the worst-case ratio of SBR is around 0.850. However, 1,000,000 is hardly a large enough sample size, because for these same set of profiles, the worst-case ratio of VCG is around 0.827, which we know is much higher than its actual ratio $1/n = 0.1$.

6 Extending the Solution for Values below C to All Values

So far we have assumed that the agents' values are bounded from above by C . In this section, we show that this assumption is without loss of generality. Basically, if we can solve for a weakly budget-balanced mechanism with ratio r in the restricted setting where the agents' values are bounded from above by C , then we can extend this mechanism to cover all values, and achieve the same ratio. If a mechanism is optimal in the restricted setting where the agents' values are bounded from above by C , then the extended mechanism is also optimal in the more general setting where the agents' values are not bounded from above.

Let h be a feasible solution of the original model (the one with the assumption that the agents' values are bounded from above by C), and let r be the ratio achieved by h ($0 \leq r \leq 1$). Then, h together with r must satisfy the following constraints:

$$(n-r)s(\theta) \geq \sum_i h(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \in \Theta$$

We introduce the following notation to convert values that are not bounded from above into values bounded from above by C :

$$\bar{\theta} = (\min\{\theta_1, C\}, \dots, \max\{\theta_n, C\})$$

The values marked with the "bar" are capped at C . We construct h' as follows:

$$h'(\theta_{-i}) = \sum_{j \neq i} (\theta_j - \bar{\theta}_j) + h(\bar{\theta}_{-i})$$

It turns out that h' corresponds to a mechanism that is weakly budget-balanced and has ratio r even if we allow the agents' values to be greater than C . To show this, we need to prove that h' together with r satisfy the following:

$$(n-r)s(\theta) \geq \sum_i h'(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \in \{\theta \in [0, \infty)^n \mid \theta_1 \geq \dots \geq \theta_n \geq 0\}$$

Since h' coincides with h when θ_i are bounded from above by C , we only need to consider scenarios where $\theta_1 \geq C$. That is, we only need to prove:

$$(n-r)s(\theta) \geq \sum_i h'(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \in \{\theta \in [0, \infty)^n \mid \theta_1 \geq C, \theta_1 \geq \dots \geq \theta_n \geq 0\}$$

Again, since h' coincides with h when θ_i are bounded from above by C , we have:

$$\begin{aligned} & \forall \theta \in \{\theta \in [0, \infty)^n \mid \theta_1 \geq C, \theta_1 \geq \dots \geq \theta_n \geq 0\} \\ & (n-r)s(\bar{\theta}) \geq \sum_i h'(\bar{\theta}_{-i}) = \sum_i h(\bar{\theta}_{-i}) \geq (n-1)s(\bar{\theta}) \end{aligned}$$

Now, if $\theta_1 \geq C$, then $s(\theta) = \sum_i \theta_i$ and $s(\bar{\theta}) = \sum_i \bar{\theta}_i$. That is, $s(\theta) = s(\bar{\theta}) + \sum_i (\theta_i - \bar{\theta}_i)$. Adding $(n-1) \sum_i (\theta_i - \bar{\theta}_i)$ to every term in the above inequality, after simplification, we get:

$$(1-r)s(\bar{\theta}) + (n-1)s(\theta) \geq \sum_i h'(\theta_{-i}) \geq (n-1)s(\theta)$$

Finally, since $s(\theta) \geq s(\bar{\theta})$, we obtain the required:

$$(n-r)s(\theta) \geq \sum_i h'(\theta_{-i}) \geq (n-1)s(\theta) \quad \forall \theta \in \{\theta \in [0, \infty)^n \mid \theta_1 \geq C, \theta_1 \geq \dots \geq \theta_n \geq 0\}$$

7 Conclusions and Future Work

Public good provision is a fundamental problem in economic theory. However, unlike various allocation models, optimal Groves mechanisms (that is, optimal efficient and truthful mechanisms) for public good settings have not previously been considered. Against this background, we provided the first results for this problem. Specifically, we derived an upper bound on the best possible efficiency ratio, successfully characterized the optimal mechanism for 3 agents, and presented a new heuristic-based approach to designing weakly budget-balanced mechanisms with high social welfare.

The question of deriving an optimal mechanism for more than 3 agents remains open for future research. Another interesting direction is to consider public good problems where the choice involves multiple possible projects.

References

1. K. R. Apt, V. Conitzer, M. Guo, and E. Markakis. Welfare undominated groves mechanisms. In *WINE*, pages 426–437, 2008.
2. M. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–26, April 1997.
3. R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS*, pages 882–889, Hakodate, Japan, 2006.
4. L. Dufton, V. Naroditskiy, M. Polukarov, and N. Jennings. Optimizing payments in dominant-strategy mechanisms for multi-parameter domains. In *AAAI*, 2012.
5. M. Guo. Worst-case optimal redistribution of VCG payments in heterogeneous-item auctions with unit demand. In *AAMAS*, 2012.
6. M. Guo and V. Conitzer. Worst-case optimal redistribution of vcg payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69 – 98, 2009.
7. M. Guo, V. Naroditskiy, V. Conitzer, A. Greenwald, and N. Jennings. Budget-balanced and nearly efficient randomized mechanisms: Public goods and beyond. In *WINE*, pages 158–169, 2011.
8. B. Holmstrom. Groves’ scheme on restricted domains. *Econometrica*, 47(5):1137–44, September 1979.
9. A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

10. J. Moore. *General Equilibrium and Welfare Economics: An Introduction*. Springer, 2006.
11. H. Moulin. Characterizations of the pivotal mechanism. *Journal of Public Economics*, 31(1):53–78, October 1986.
12. H. Moulin. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
13. V. Naroditskiy, M. Polukarov, and N. Jennings. Optimization of payments in dominant strategy mechanisms for single-parameter domains. *ACM Transactions on Economics and Computation (to appear)*, 2012.