

Social Decision with Minimal Efficiency Loss: an Automated Mechanism Design Approach

Mingyu Guo, Hong Shen
School of Computer Science
University of Adelaide, Australia
{mingyu.guo, hong.shen}
@adelaide.edu.au

Taiki Todo, Yuko Sakurai, Makoto Yokoo
Graduate School of ISEE
Kyushu University, Japan
{todo,ysakurai,yokoo}
@inf.kyushu-u.ac.jp

ABSTRACT

We study the problem where a group of agents need to choose from a finite set of social outcomes. We assume every agent’s valuation for every outcome is bounded and the bounds are public information. For our model, no mechanism simultaneously satisfies strategy-proofness, individual rationality, non-deficit, and efficiency. In light of this, we aim to design mechanisms that are strategy-proof, individually rational, non-deficit, and minimize the worst-case efficiency loss.

We propose a family of mechanisms called the shifted Groves mechanisms, which are generalizations of the Groves mechanisms. We first show that if there exist mechanisms that are strategy-proof, individually rational, and non-deficit, then there exist shifted Groves mechanisms with these properties. Our main result is an Automated Mechanism Design (AMD) approach for identifying the (unique) optimal shifted Groves mechanism, which minimizes the worst-case efficiency loss among all shifted Groves mechanisms. Finally, we prove that the optimal shifted Groves mechanism is globally optimal among all deterministic mechanisms that are strategy-proof, individually rational, and non-deficit.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Algorithms, Economics, Theory

Keywords

Groves Mechanisms, Automated Mechanism Design, Social Decision

1. INTRODUCTION

Social decision-making is a fundamental problem in multiagent systems. This paper studies how to design economic mechanisms for implementing good social outcomes. We study the problem where a group of agents need to choose

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

from a finite set of social outcomes. The agents’ valuations for the outcomes are private information. We assume every agent’s valuation for every outcome is bounded and the bounds are public information (e.g., the bounds are prescribed by the nature of the outcomes). One example setting is the public project problem:

Example 1. Public Project Problem

n agents need to decide whether or not to build a public project (e.g., a bridge accessible by everyone). The cost of the project is 1. There are two outcomes:

- Outcome 1: not build
Every agent’s valuation for this outcome equals 0.
- Outcome 2: build, and every agent is responsible for an equal share of the cost, which equals $\frac{1}{n}$
Different agents value the project differently. We assume an agent’s valuation for the project is between 0 (not appreciating at all) and 1 (as high as the cost of the whole project). That is, taking the cost share into consideration, every agent’s valuation for this outcome is between $-\frac{1}{n}$ and $\frac{n-1}{n}$.

For social decision problems like the above, ideally, we prefer mechanisms that satisfy the following properties:

- *Strategy-proofness*: Every agent’s dominant strategy is to report truthfully, regardless of the other agents’ reports. It is without loss of generality to require strategy-proofness due to the revelation principle.
- *(Ex post) individual rationality*: Every agent’s utility (valuation of the chosen outcome minus payment) is non-negative. Individual rationality is necessary when participation is voluntary.
- *Non-deficit property*: The agents’ total payment is non-negative.¹ That is, no external subsidy is ever needed. Another related but stronger property is *budget balance*, which requires that the agents’ total payment is exactly 0. We require non-deficit, otherwise the mechanism is not sustainable.
- *Efficiency*: The mechanism always chooses the outcome that maximizes the agents’ total valuation. (For the public project problem, it means the project is built if and only if the agents’ total valuation for outcome 2 is higher.)

¹An individual agent may still pay negative payment.

For resource allocation settings, the above properties can be simultaneously achieved, *e.g.*, by the VCG redistribution mechanisms [1, 5, 13]. However, when it comes to social decision settings, sometimes an agent has to face an undesirable outcome that is forced onto him by the others (*e.g.*, an agent who does not appreciate the public project may be forced to chip in to build the project). The VCG redistribution mechanisms, when applied to the public project problem, are not individually rational (unless we change the definition of individual rationality) [7, 8, 14]. Actually, as we will discuss in more details later on, *for social decision settings, no mechanisms can be simultaneously strategy-proof, individual rational, non-deficit, and efficient.*

In light of the above, we choose to sacrifice efficiency.² Our aim is to design mechanisms that are strategy-proof, individually rational, non-deficit, and as close to efficiency as possible. Specifically, we try to minimize the *worst-case efficiency loss* (the loss of efficiency in the worst case, as the result of using an inefficient mechanism).

We adopt a (Computationally Feasible) Automated Mechanism Design approach [6, 2]. We propose a family of mechanisms called the shifted Groves mechanisms, which are generalizations of the Groves mechanisms. We first show that if there exist mechanisms that are strategy-proof, individually rational, and non-deficit, then there exist shifted Groves mechanisms with these properties. We use AMD to identify the (unique) optimal shifted Groves mechanism, which minimizes the worst-case efficiency loss among all shifted Groves mechanisms. Finally, we prove that the optimal shifted Groves mechanism is globally optimal among all deterministic mechanisms that are strategy-proof, individually rational, and non-deficit. Besides proposing optimal mechanisms for a fundamental mechanism design setting, our approach also demonstrates the effectiveness of using AMD to deliver new economic results.

2. FORMAL MODEL DESCRIPTION

There are n agents who need to choose from k outcomes. We use $\theta_i = (v_{i1}, v_{i2}, \dots, v_{ik})$ to denote agent i 's type, where v_{ij} represents agent i 's valuation for outcome j .

We make the following assumptions:

- v_{ij} is bounded below and above by *known constants* L_{ij} and U_{ij} . It should be noted that both L_{ij} and U_{ij} may be negative.
- For every i , agent i 's type space Θ_i is the whole of $[L_{i1}, U_{i1}] \times [L_{i2}, U_{i2}] \times \dots \times [L_{ik}, U_{ik}]$.
- The set of type profiles is the whole of $\Theta_1 \times \Theta_2 \times \dots \times \Theta_n$.
- There do not exist two outcomes, so that the agents always prefer one outcome than the other for all type profiles. That is, for every pair of j and j' , we have $\sum_i U_{ij} \geq \sum_i L_{ij'}$.

Our analysis throughout the paper relies on the above assumptions (*e.g.*, we need the above to prove that our AMD approach produces optimal mechanisms). However,

²The Faltings' mechanism [3] also sacrifices efficiency. The mechanism works for both resource allocation settings and social decision settings, but again, it is not individually rational for our problem.

it should be noted that only the first assumption is needed to *apply* the AMD approach, which takes the bounds as input. For example, it could be that certain type values do not exist (*e.g.*, v_{ij} 's possible values do not form an interval or form a much narrower interval than $[L_{ij}, U_{ij}]$) and/or certain types/type profiles do not exist (*e.g.*, due to interdependence). For these cases, the AMD approach still lead to a feasible mechanism. Of course, the mechanism is generally not optimal, because we fail to recognize that certain profiles are impossible.

As mentioned earlier, it is impossible to simultaneously have strategy-proofness, individually rationality, non-deficit, and efficiency. Below we present the details. According to Holmström [9], for convex type space (our setting), the *Groves mechanisms* [4] are the only mechanisms that are efficient and strategy-proof. The Groves mechanisms are defined as follows:

- The mechanism chooses an outcome from $\arg \max_j \sum_i v_{ij}$. That is, the mechanism chooses an outcome that maximizes the agents' total valuation.
- Let j^* be the chosen outcome. For every i , agent i receives $\sum_{i' \neq i} v_{i'j^*}$. That is, the amount an agent receives is equal to the other agents' total valuation for j^* .
- For every i , agent i then pays an amount that is independent of i 's own type, which is denoted by $h_i(\theta_{-i})$.³ Different sets of h_i functions correspond to different Groves mechanisms.

Every Groves mechanism is efficient and strategy-proof, but no Groves mechanism is both individually rational and non-deficit. We need the $h_i(\theta_{-i})$ term to be small enough, so that agent i 's utility is non-negative. We also need the $h_i(\theta_{-i})$ term to be large enough, so that the agents' total payment is non-negative. However, we cannot satisfy both. We use $E(\theta_i, \theta_{-i})$ to denote the agents' total valuation for the efficient outcome. Under the Groves mechanisms, agent i 's utility is simply $E(\theta_i, \theta_{-i}) - h_i(\theta_{-i})$, which needs to be non-negative for ensuring individual rationality. Hence, we have $h_i(\theta_{-i}) \leq E(\theta_i, \theta_{-i})$ for all θ_i and all θ_{-i} . The left-hand side does not depend on θ_i , and the right-hand side is minimized when $\theta_i = (L_{i1}, L_{i2}, \dots, L_{ik})$ (we use θ_i to denote this type). Hence, individual rationality is equivalent to for all i and all type profiles,

$$h_i(\theta_{-i}) \leq E(\theta_i, \theta_{-i}) = \max_j (L_{ij} + \sum_{i' \neq i} v_{i'j})$$

It does not hurt to set $h_i(\theta_{-i})$ to exactly $\max_j (L_{ij} + \sum_{i' \neq i} v_{i'j})$. As mentioned earlier, we want this term to be large enough for ensuring non-deficit. If we set it any larger, individual rationality is violated. Once we fix the $h_i(\theta_{-i})$ term, there is one Groves mechanism left. If this mechanism is not non-deficit, then no Groves mechanisms are both individually rational and non-deficit. We use the following example to show this.

Example 2. Let us consider a public project problem (described in Example 1) with 2 agents. We have $L_{i1} = U_{i1} = 0$,

³When we run the mechanism, we simply combine the payment terms in step 2 and 3. Throughout this paper, they are presented separately as we analyze them separately.

$L_{i2} = -\frac{1}{2}$, and $U_{i2} = \frac{1}{2}$. Let $\theta_1 = \theta_2 = (0, \frac{1}{2})$. That is, $v_{i1} = 0$ and $v_{i2} = \frac{1}{2}$. For this type profile, the decision is to build, both agents receive $\frac{1}{2}$ (here, $h_i(\theta_{-i}) = 0$). The non-deficit property is violated here.

In conclusion, in light of the above impossibility result, our aim is to design social decision mechanisms that are strategy-proof, individually rational, non-deficit, and minimize the *worst-case efficiency loss*.⁴

Definition 1. Let M be a mechanism. Let $M(\theta_i, \theta_{-i})$ be the agents' total valuation under M for type profile (θ_i, θ_{-i}) . We recall that $E(\theta_i, \theta_{-i})$ represents the agents' total valuation for the efficient outcome. Mechanism M 's **worst-case efficiency loss** is defined as:

$$\max_{\theta_i, \theta_{-i}} (E(\theta_i, \theta_{-i}) - M(\theta_i, \theta_{-i}))$$

3. SHIFTED GROVES MECHANISMS

We revisit the public project example. As we have seen, if outcome 2 (build) is chosen, deficit may occur. On the other hand, if outcome 1 is chosen, deficit never occurs.⁵ One idea is then to "discourage" outcome 2 by asking the agents to pay additional payments when outcome 2 is chosen (essentially changing the agents' valuations for outcome 2). For example, we may *shift* the outcome "build" to "build, and every agent pays x in addition to mechanism payments and building cost". This helps reducing deficit for two reasons. First, after introducing these additional payments, less type profiles are mapped to outcome 2, so the maximal deficit when outcome 2 is chosen may potentially decrease. Second, the additional payments also help offset the deficit. If outcome 2 is chosen, we know that we have nx additional payment to offset the deficit.

The above idea leads to the *shifted Groves mechanisms*:

- For every agent i , every outcome j , we introduce an additional payment t_{ij} . That is, if outcome j is chosen, agent i needs to pay t_{ij} in addition to her mechanism payment. The t_{ij} are constants set by the mechanism designer. The total additional payment for outcome j is $P_j = \sum_i t_{ij}$.
- Apply the Groves mechanisms on the modified outcomes.

- The mechanism chooses an outcome from

$$\arg \max_j \sum_i (v_{ij} - t_{ij}) = \arg \max_j \{-P_j + \sum_i v_{ij}\}$$

That is, the mechanism chooses an outcome that maximizes the agents' total valuation, considering the additional payments.

- Let j^* be the chosen outcome. For every i , agent i receives $\sum_{i' \neq i} (v_{i'j^*} - t_{i'j^*})$. That is, the amount an agent receives is equal to the other agents' total valuation for the chosen outcome j^* , considering

⁴For mechanisms that need to deal with tie-breaking, we assume tie-breaking is done randomly, and the worst-case efficiency loss considers all possible tie-breaking scenarios.

⁵If the chosen outcome is 1 (not build), then both agents receive 0 and $h_i(\theta_{-i}) = 0$ (e.g., $h_1(\theta_{-1}) = \max\{0, -\frac{1}{2} + v_{22}\} = 0$).

the additional payments. We also keep in mind that agent i needs to pay the additional payment t_{ij^*} . Therefore, agent i 's net income at this point should be

$$\sum_{i' \neq i} (v_{i'j^*} - t_{i'j^*}) - t_{ij^*} = -P_{j^*} + \sum_{i' \neq i} v_{i'j^*}$$

- For every i , agent i then pays an amount that is independent of i 's own type, which is denoted by $h_i(\theta_{-i})$.

Similar to the case without additional payments, we can find an upper bound on $h_i(\theta_{-i})$. Agent i 's utility equals $\max_j \{-P_j + \sum_i v_{ij}\} - h_i(\theta_{-i})$, which must be non-negative for ensuring individual rationality. Therefore, we have $h_i(\theta_{-i}) \leq \max_j \{-P_j + \sum_i v_{ij}\}$ for all θ_i . By setting $\theta_i = \underline{\theta}_i = (L_{i1}, L_{i2}, \dots, L_{ik})$, we have

$$h_i(\theta_{-i}) \leq \max_j \{-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}\}$$

Again, it does not hurt to set $h_i(\theta_{-i})$ to be exactly the right-hand side. (This value does not affect strategy-proofness and allocative efficiency. We need it large for achieving the non-deficit property, but any larger value violates individual rationality.)

Once we fix the $h_i(\theta_{-i})$ terms, we observe that the shifted Groves mechanisms are completely characterized by the total additional payments for the outcomes (the P_j). Both the outcome and the payments do not depend on the individual additional payments (the t_{ij}). This leads to a conciser definition of shifted Groves mechanisms.

Definition 2. A **shifted Groves mechanism** is characterized by k constants P_1, P_2, \dots, P_k .

- The mechanism chooses an outcome from

$$\arg \max_j \{-P_j + \sum_i v_{ij}\}$$

- Let j^* be the chosen outcome. For every i , agent i pays

$$\max_j \{-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}\} + P_{j^*} - \sum_{i' \neq i} v_{i'j^*}$$

We use $M(P_1, P_2, \dots, P_k)$ to denote the shifted Groves mechanism characterized by the P_j .

It should be noted that the shifted Groves mechanisms are special cases of *Affine Maximizer Auctions (AMA)* [10, 11]. There are many works on characterizing strategy-proof mechanisms in social choice settings. Roberts [15] shows that when the type space of agents is unrestricted (and there exist three or more choices), the only strategy-proof mechanisms are affine maximizers (of which the shifted Groves mechanisms are special cases). In our setting, the type space of agents is bounded so the results in [15] do not apply. There also exist works on characterizing strategy-proof mechanisms in a restricted domain. For example, [12] showed that all strategy-proof mechanisms are affine maximizers if the agents' valuations are drawn from intervals. However, [12] requires that the intervals be open (we use close intervals in this paper), and more importantly, it also requires that an agent's interval for all outcomes are identical

(we do not require this, *e.g.*, for the public project problem, the agents obviously have different intervals for not build and build). Therefore, [12] also does not apply. For our setting, it remains to see whether all strategy-proof mechanisms are affine maximizers. Nevertheless, in this paper, we show that the optimal shifted Groves mechanism is optimal among all strategy-proof mechanisms, without relying on such characterization results.⁶

Based on how they are defined, all shifted Groves mechanisms are strategy-proof and individual rational. Unlike the Groves mechanisms, some shifted Groves mechanisms can be non-deficit (while being individually rational).

THEOREM 1. *Given a setting, if there exist mechanisms that are strategy-proof, individually rational, and non-deficit, then there must exist shifted Groves mechanisms with these properties.*

To prove the above theorem, we introduce the following lemma. We recall that $\sum_i L_{ij^*}$ is the agents' lowest possible total valuation for outcome j^* .

LEMMA 1. *If $\sum_i L_{ij^*} \geq 0$, then under all shifted Groves mechanisms (for all values of the P_j), when outcome j^* is chosen, deficit never occurs.*

PROOF. When outcome j^* is chosen, agent i 's payment equals

$$\begin{aligned} & \max_j \{-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}\} + P_{j^*} - \sum_{i' \neq i} v_{i'j^*} \\ & \geq -P_{j^*} + L_{ij^*} + \sum_{i' \neq i} v_{i'j^*} + P_{j^*} - \sum_{i' \neq i} v_{i'j^*} = L_{ij^*} \end{aligned}$$

Therefore, the agents' total payment is at least $\sum_i L_{ij^*}$, which is non-negative. \square

Without loss of generality, we assume $\sum_i L_{i1} \geq \sum_i L_{i2} \geq \dots \geq \sum_i L_{ik}$. Now we are ready to prove the theorem.

PROOF. If $\sum_i L_{i1} < 0$, then for some type profiles, the agents' total valuation for every outcome is negative. For these type profiles, individual rationality and non-deficit cannot coexist.

If $\sum_i L_{i1} \geq 0$, then under any shifted Groves mechanism, outcome 1 never incurs deficit. We may simply "shift down" every outcome except for 1 (for all $j > 1$, set P_j to be really large), so that only outcome 1 can be possibly chosen (*e.g.*, ensure that for all $j > 1$, $-P_j + \sum_i U_{ij} < \sum_i L_{i1}$). The resulting mechanism simply always picks outcome 1, which guarantees the non-deficit property according to Lemma 1. We already know that all shifted Groves mechanisms are strategy-proof (they are part of the AMA family) and individually rational (the way we set the $h_i(\theta_{-i})$ term guarantees individual rationality).

In conclusion, if $\sum_i L_{i1} < 0$, no mechanisms can be both individually rational and non-deficit. If $\sum_i L_{i1} \geq 0$, there exist shifted Groves mechanisms that are strategy-proof, individually rational, and non-deficit. \square

⁶It should be noted that even if all strategy-proof mechanisms are affine maximizers for our setting, we still need to prove that only "shifts" are needed to reach the optimal mechanism, since affine maximizers also rely on "weights" (multiplication).

Based on the proof of Theorem 1, from this point on, we assume $\sum_i L_{i1} \geq 0$. Otherwise, no feasible mechanisms exist. For many settings, there exists an outcome which corresponds to the *status quo* (*e.g.*, not build in the public project problem). The agents' valuations for the status quo are all equal to 0. If the status quo belongs to the set of outcomes, then according to Theorem 1, there exist mechanisms that are strategy-proof, individually rational, and non-deficit.

Given $\sum_i L_{i1} \geq 0$, we also have that for all j , $\sum_i U_{ij} \geq 0$, because our assumption is that there do not exist two outcomes, so that the agents always prefer one to the other for all type profiles.

We do not have to study all shifted Groves mechanisms. First of all, we only consider those that are non-deficit (we call these shifted Groves mechanisms *feasible*). We also ignore those that are not *onto*.

Definition 3. A shifted Groves mechanism $M(P_1, P_2, \dots, P_k)$ is **onto** iff under it, for every outcome j , there exists at least one type profile for which j is chosen by the mechanism (or may be chosen by the mechanism assuming random tie-breaking).

Formally, this means that for every pair of outcomes a and b , we have

$$-P_a + \sum_i U_{ia} \geq -P_b + \sum_i L_{ib}$$

LEMMA 2. *Let $M(P_1, P_2, \dots, P_k)$ be a shifted Groves mechanism that is feasible but not onto. There must exist another feasible shifted Groves mechanism that is onto, and has the same or higher allocative efficiency for all type profiles.*

PROOF. If $M(P_1, P_2, \dots, P_k)$ is feasible but not onto, then there exist two outcomes a and b , so that

$$-P_a + \sum_i U_{ia} < -P_b + \sum_i L_{ib}$$

We can simply decrease P_a so that

$$-P_a + \sum_i U_{ia} = -P_b + \sum_i L_{ib}$$

This modification does not change the allocative efficiency of any type profile, except for maybe when $v_{ia} = U_{ia}$ and $v_{ib} = L_{ib}$ for all i . For this case, a may be chosen instead of b . We have that $\sum_i U_{ij} > \sum_i L_{ij'}$ for every pair of j and j' . That is, the above modification never decreases the allocative efficiency. The modification also does not change the payments of any agents when outcome a is not chosen. So for outcomes other than a , the modified mechanism is still non-deficit. Finally, when a is indeed chosen, agent i 's payment equals

$$\begin{aligned} & \max_j \{-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}\} + P_a - \sum_{i' \neq i} v_{i'a} \\ & \geq -P_b + L_{ib} + \sum_{i' \neq i} v_{i'b} + P_a - \sum_{i' \neq i} v_{i'a} \\ & \geq -P_a + \sum_i U_{ia} + P_a - \sum_{i' \neq i} v_{i'a} \geq U_{ia} \end{aligned}$$

The total payment is then at least $\sum_i U_{ia}$, which is at least 0 (at least $\sum_i L_{i1}$). In summary, the modified mechanism

remains feasible and has the same or higher allocative efficiency. We can repeat the above process until the mechanism becomes onto. \square

THEOREM 2. *Let $M(P_1, P_2, \dots, P_k)$ be a shifted Groves mechanism that is both feasible and onto. $M(P_1, P_2, \dots, P_k)$'s worst-case efficiency loss is $\max_j P_j - \min_j P_j$.*

Theorem 2 directly follows from the following two lemmas:

LEMMA 3. *Let $M(P_1, P_2, \dots, P_k)$ be a shifted Groves mechanism. $M(P_1, P_2, \dots, P_k)$ and $M(P_1 - \Delta, P_2 - \Delta, \dots, P_k - \Delta)$ are equivalent (Δ is an arbitrary constant).*

The above lemma directly follows from Definition 2. It implies that $M(P_1, P_2, \dots, P_k)$ and $M(P_1 - \min_j P_j, P_2 - \min_j P_j, \dots, P_k - \min_j P_j)$ are equivalent. That is, it is without loss of generality to focus on cases where the P_j are non-negative, and at least one of the P_j is 0.

LEMMA 4. *Let $M(P_1, P_2, \dots, P_k)$ be a shifted Groves mechanism that is feasible and onto. The P_j are non-negative and $\min_j P_j = 0$. $M(P_1, P_2, \dots, P_k)$'s worst-case efficiency loss is $\max_j P_j$.*

PROOF. If $\max_j P_j = 0$, then the mechanism is simply a Groves mechanism, which is efficient (worst-case efficiency loss is 0). We then analyze cases with $\max_j P_j > 0$.

First of all, it is easy to see that the worst-case efficiency loss is at most $\max_j P_j$. We only need to prove that it is also at least this much.

Let $P_a = 0$ and $P_b = \max_j P_j$. a and b are different. We construct the following type profile: $v_{ij} = L_{ij}$ for all i and all $j \neq a, b$. Due to the onto property, we have $\sum_i U_{ia} = -P_a + \sum_i U_{ia} \geq \max_j(-P_j + \sum_i L_{ij})$. We certainly have $\sum_i L_{ia} = -P_a + \sum_i L_{ia} \leq \max_j(-P_j + \sum_i L_{ij})$. Therefore, we can set the v_{ia} so that $\sum_i v_{ia} = \max_j(-P_j + \sum_i L_{ij})$. This ensures that a is at least tied with all $j \neq b$. Again, due to the onto property, we have $-P_b + \sum_i U_{ib} \geq \max_j(-P_j + \sum_i L_{ij}) \geq -P_b + \sum_i L_{ib}$. We set the v_{ib} so that outcome b is tied with a . For this type profile, since we assume random tie-breaking, a may be chosen ahead of b , which corresponds to an efficiency loss of $\max_j P_j$. \square

Next we try to find the optimal shifted Groves mechanism. As mentioned earlier, we only need to consider shifted Groves mechanisms that are feasible and onto. Also, we only need to consider cases where the P_j are non-negative and at least one of the P_j is 0.

Definition 4. A shifted Groves mechanism $M(P_1, P_2, \dots, P_k)$ is **optimal**, if there does not exist another shifted Groves mechanism $M(P'_1, P'_2, \dots, P'_k)$, so that $P'_i \leq P_i$ for all i .

Our aim is to minimize the worst-case efficiency loss. Based on Lemma 4, the optimal shifted Groves mechanism should minimize $\max_j P_j$. We will prove later on that this is indeed the case, and the optimal mechanism is unique. Before that, we first propose a linear program based AMD approach for identifying the optimal mechanism.

Given a shifted Groves mechanism $M(P_1, P_2, \dots, P_k)$, we can calculate the maximal deficit for outcome j^* using the following linear program:

<p>Constants: L_{ij}, U_{ij}, P_j, j^* Variables: v_{ij} Dummy Variables: d_i (d_i represents agent i's payment) Maximize: $-\sum_i d_i$ (maximize deficit) Subject to: For all $j \neq j^*$,</p> $-P_{j^*} + \sum_i v_{ij^*} \geq -P_j + \sum_i v_{ij} \quad (1)$ <p>For all j,</p> $d_i \geq (-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}) + P_{j^*} - \sum_{i' \neq i} v_{i'j^*} \quad (2)$ <p>For all i and j,</p> $L_{ij} \leq v_{ij} \leq U_{ij}$
--

Constraint 1 ensures that we are dealing with type profiles where outcome j^* is chosen. Constraint 2 ensures that d_i represents agent i 's payment. That is,

$$d_i = \max_j \{-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}\} + P_{j^*} - \sum_{i' \neq i} v_{i'j^*}$$

This is true because the linear program will minimize the d_i in order to maximize the objective.

The above linear program can be significantly simplified. We may divide Constraint 2 into the following two:

For all $j \neq j^*$,

$$d_i \geq (-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}) + P_{j^*} - \sum_{i' \neq i} v_{i'j^*} \quad (3)$$

For $j = j^*$,

$$d_i \geq (-P_j + L_{ij} + \sum_{i' \neq i} v_{i'j}) + P_{j^*} - \sum_{i' \neq i} v_{i'j^*} = L_{ij^*} \quad (4)$$

Constraint 4 has nothing to do with the v_{ij} . For Constraint 1 and 3, they are the most relaxed if we set $v_{ij} = L_{ij}$ for all i and all $j \neq j^*$, and we set $v_{ij^*} = U_{ij^*}$ for all i . After simplification, we have

<p>Constants: L_{ij}, U_{ij}, P_j, j^* Dummy Variables: d_i (d_i represents agent i's payment) Maximize: $-\sum_i d_i$ (maximize deficit) Subject to: For all $j \neq j^*$,</p> $-P_{j^*} + \sum_i U_{ij^*} \geq -P_j + \sum_i L_{ij} \quad (5)$ <p>For all $j \neq j^*$ and all i,</p> $d_i \geq (-P_j + \sum_{i'} L_{i'j}) + P_{j^*} - \sum_{i' \neq i} U_{i'j^*} \quad (6)$ $d_i \geq L_{ij^*} \quad (7)$
--

Constraint 5 is already guaranteed by the onto property. The objective value of the above linear program is simply

$$-\sum_i \max_{j \neq j^*} \{\max\{-P_j + \sum_{i'} L_{i'j}\} + P_{j^*} - \sum_{i' \neq i} U_{i'j^*}, L_{ij^*}\}$$

We use $D_{j^*}(P_1, P_2, \dots, P_k)$ to denote the above objective value. That is, $D_{j^*}(P_1, P_2, \dots, P_k)$ represents the maximal deficit under the shifted Groves mechanism $M(P_1, P_2, \dots, P_k)$, when outcome j^* is chosen,

We notice that $D_{j^*}(P_1, P_2, \dots, P_k)$ is non-increasing in P_{j^*} and non-decreasing in P_j for $j \neq j^*$. This leads to the following AMD approach.

- Start from $P_1 = P_2 = \dots = P_k = 0$. That is, we start from the original Groves mechanism, which is onto but not necessarily feasible.
- For every outcome j^* , if $D_{j^*}(P_1, P_2, \dots, P_k)$ is positive, then increase P_{j^*} until $D_{j^*}(P_1, P_2, \dots, P_k)$ reaches 0. This move maintains the onto property. If after the modification, it is no longer onto, then we have $-P_{j^*} + \sum_i U_{ij^*} < \max_{j \neq j^*} \{-P_j + \sum_i L_{ij}\}$. For $D_{j^*}(P_1, P_2, \dots, P_k)$, we then have

$$\begin{aligned} & -\sum_i \max\{\max_{j \neq j^*} \{-P_j + \sum_{i'} L_{i'j}\} + P_{j^*} - \sum_{i' \neq i} U_{i'j^*}, L_{ij^*}\} \\ & \leq -\sum_i (\max_{j \neq j^*} \{-P_j + \sum_{i'} L_{i'j}\} + P_{j^*} - \sum_{i' \neq i} U_{i'j^*}) \\ & < -\sum_i (\sum_i U_{ij^*} - \sum_{i' \neq i} U_{i'j^*}) = -\sum_i U_{ij^*} \leq 0 \end{aligned}$$

That is, $D_{j^*}(P_1, P_2, \dots, P_k)$ is negative, which contradicts with the algorithm description.

THEOREM 3. *The AMD process produces a unique optimal shifted Groves mechanism that is feasible, onto, and has minimal worst-case efficiency loss among all shifted Groves mechanisms.*

PROOF. If we follow the AMD steps, then P_1 must remain 0, because according to Lemma 1, $D_1(P_1, P_2, \dots, P_k)$ is never positive. This implies that in all calculations, for $j > 1$,

$$-P_1 + \sum_i L_{i1} = \sum_i L_{i1} \geq \sum_i L_{ij} \geq -P_j + \sum_i L_{ij}$$

This fact further simplifies $D_{j^*}(P_1, P_2, \dots, P_k)$ for $j^* > 1$:

$$\begin{aligned} & -\sum_i \max\{\max_{j \neq j^*} \{-P_j + \sum_{i'} L_{i'j}\} + P_{j^*} - \sum_{i' \neq i} U_{i'j^*}, L_{ij^*}\} \\ & = -\sum_i \max\{\sum_{i'} L_{i'i} + P_{j^*} - \sum_{i' \neq i} U_{i'j^*}, L_{ij^*}\} \end{aligned}$$

That is, the maximal deficit for outcome j^* depends only on P_{j^*} (the other P_j values are irrelevant). Therefore, the AMD algorithm does produce a feasible mechanism that is onto.

Next, we prove the optimal mechanism is unique. Let $M(P_1^*, P_2^*, \dots, P_k^*)$ be an optimal mechanism. When we start the AMD process, we have $P_1 = P_2 = \dots = P_k = 0$. That is, we have $P_j \leq P_j^*$ for all j . During the AMD, if $D_j(P_1, P_2, \dots, P_k) > 0$. The AMD process increases P_j to $P_j + \Delta$ so that $D_j(P_1, P_2, \dots, P_k) = 0$. We recall that $D_j(P_1, P_2, \dots, P_k)$ is non-increasing in P_j and non-decreasing in $P_{j'}$ for $j' \neq j$. Therefore,

$$0 = D_j(P_1, \dots, P_{j-1}, P_j + \Delta, P_{j+1}, \dots, P_k) \leq$$

$$D_j(P_1^*, \dots, P_{j-1}^*, P_j + \Delta, P_{j+1}^*, \dots, P_k^*)$$

We also have

$$D_j(P_1^*, \dots, P_{j-1}^*, P_j^*, P_{j+1}^*, \dots, P_k^*) \leq 0$$

Therefore, we have $P_j^* \geq P_j + \Delta$. That is, after the AMD, we have $P_j \leq P_j^*$ for all j . Since we assumed $M(P_1^*, P_2^*, \dots, P_k^*)$ to be optimal, we should have $P_j = P_j^*$ for all j after the AMD. Since $M(P_1^*, P_2^*, \dots, P_k^*)$ is arbitrarily chosen, the optimal mechanism must be unique.

If there exists a non-optimal mechanism $M(P'_1, P'_2, \dots, P'_k)$ that minimize the worst-case efficiency loss. Then there must exist an optimal mechanism $M(P_1^*, P_2^*, \dots, P_k^*)$ satisfying $P_j^* \leq P'_j$ for all j . Then by Lemma 4, $M(P_1^*, P_2^*, \dots, P_k^*)$ should have the same or less worst-case efficiency loss. \square

Finally, based on Lemma 1, for all j with $\sum_i L_{ij} \geq 0$, we have $D_j(P_1, P_2, \dots, P_k) \leq 0$ for all the P_j . This implies that when conducting the AMD, if $\sum_i L_{ij} \geq 0$, we simply leave P_j to be 0.

4. GLOBAL OPTIMALITY

Theorem 3 only shows that the optimal shifted Groves mechanism is optimal *within* the family of shifted Groves mechanisms. Here, we show that it is also optimal among all deterministic mechanisms that are strategy-proof, individually rational, and non-deficit.

LEMMA 5. *When there are only two outcomes, the optimal shifted Groves mechanism has the lowest worst-case efficiency loss among all deterministic mechanisms that are strategy-proof, individually rational, and non-deficit.*

PROOF. We denote outcome 1 by a and outcome 2 by b . We have the usual assumptions that $\sum_i L_{ia} \geq \sum_i L_{ib}$ and $\sum_i L_{ia} \geq 0$. Let H be a deterministic mechanism that is strategy-proof, individually rational, and non-deficit. Let Δ be H 's worst-case efficiency loss. We show that there exists a shifted Groves mechanism with the same or lower worst-case efficiency loss.

If $\Delta \geq \sum_i U_{ib} - \sum_i L_{ia}$, then we can simply construct a trivial shifted Groves mechanism that always chooses a (choosing a never results in deficit), and this mechanism's worst-case efficiency loss is at most Δ . Therefore, we only need to consider $\Delta < \sum_i U_{ib} - \sum_i L_{ia}$.

We consider the shifted Groves mechanism $M(0, \Delta + \epsilon)$, where ϵ is a small positive value ($\Delta + \epsilon < \sum_i U_{ib} - \sum_i L_{ia}$). For simplicity, we simply call this mechanism M . M 's worst-case efficiency loss is at most $\Delta + \epsilon$. If we can show that M is non-deficit, then that means one feasible shifted Groves mechanism has a worst-case efficiency loss of at most $\Delta + \epsilon$. ϵ can be made arbitrary small, which means that there must exist a feasible shifted Groves mechanism with a worst-case efficiency loss of at most Δ . This proves that the optimal shifted Groves mechanism is globally optimal.

Now we prove that M is non-deficit. When outcome a is chosen, by Lemma 1, deficit will not occur. We only need to prove that when M chooses b , there will not be deficit.

Whenever M chooses b , we know that the agents prefer b to a by at least $\Delta + \epsilon$, which means that H must also choose b .

Let (θ_i, θ_{-i}) be a type profile for which M chooses b . For this type profile, H must also choose b . We analyze agent i 's payment under H and M .

Case 1: Under M , given θ_{-i} , no matter what agent i reports, the outcome is always b . If this is the case, then under H , given θ_{-i} , no matter what agent i reports, the outcome is always b . (Whenever b is chosen under M , H must also choose b .) H is deterministic. Agent i 's payment under H must be fixed (does not change with the report). Otherwise, H is not strategy-proof. Let x be agent i 's payment. x must be at most L_{ib} . Otherwise, H is not individually rational. Agent i 's payment under M is exactly L_{ib} according to the definition. That is, for Case 1, agent i pays the same or more under M .

Case 2: Under M , given θ_{-i} , both a and b may be chosen (depending on θ_i). When b is chosen, agent i pays

$$\max\{L_{ia} + \sum_{i' \neq i} (v_{i'a} - v_{i'b}) + \Delta + \epsilon, L_{ib}\}$$

We divide Case 2 into the following sub-cases.

Case 2i: If under H , given θ_{-i} , no matter what agent i reports, the outcome is always b , then as analyzed earlier, agent i pays at most L_{ib} . That is, for Case 2i, agent i pays the same or more under M .

Case 2ii: Under H , given θ_{-i} , for some θ_i , outcome a is chosen, and for some other θ_i , outcome b is chosen. We recall that H is deterministic. When a is chosen, agent i 's payment must be a fixed amount that is independent of θ_i . Otherwise, H is not strategy-proof. We use p_a^H to denote this amount. Similarly, we use p_b^H to denote the fixed payment paid by i whenever b is chosen. Let $C^H = p_b^H - p_a^H$. This C^H value is a critical value: whenever agent i favors outcome b by more than C^H , outcome b will be chosen (otherwise, a is chosen). This critical value C^H must be less than $\sum_{i' \neq i} (v_{i'a} - v_{i'b}) + \Delta + \epsilon$. Otherwise, H may choose a when the agents favor b by more than Δ , which contradicts with the fact that H 's worst-case efficiency loss is Δ . H is individually rational, which means that we either have $p_a^H \leq L_{ia}$ or $p_b^H \leq L_{ib}$. Otherwise, agent i may end up not able to afford either outcome. So we have $p_b^H - C^H \leq L_{ia}$ or $p_b^H \leq L_{ib}$. That is,

$$\begin{aligned} p_b^H &\leq \max\{L_{ia} + C^H, L_{ib}\} \\ &\leq \max\{L_{ia} + \sum_{i' \neq i} (v_{i'a} - v_{i'b}) + \Delta + \epsilon, L_{ib}\} \end{aligned}$$

Therefore, for Case 2ii, agent i pays the same or more under M when b is chosen. \square

THEOREM 4. *The optimal shifted Groves mechanism has the lowest worst-case efficiency loss among all deterministic mechanisms that are strategy-proof, individually rational, and non-deficit.*

PROOF. If there is only one outcome, then the optimal shifted Groves mechanism is the original Groves mechanism, which obviously minimizes the worst-case efficiency loss (as it is efficient). Lemma 5 already proved the case with two outcomes. We then consider cases with more than two outcomes.

We first analyze the worst-case efficiency loss of the optimal shifted Groves mechanism $M(P_1^*, P_2^*, \dots, P_k^*)$. Reusing the analysis in the proof of Theorem 3, we must have $P_1^* = 0$. For $j > 1$, $P_j^* = 0$ if

$$-\sum_i \max\{\sum_{i'} L_{i'1} - \sum_{i' \neq i} U_{i'j}, L_{ij}\} \leq 0$$

Otherwise, P_j^* is the minimal value that makes

$$-\sum_i \max\{\sum_{i'} L_{i'1} + P_j^* - \sum_{i' \neq i} U_{i'j}, L_{ij}\} = 0$$

According to Lemma 4, the worst-case efficiency loss is simply $\max_j P_j^*$. Let $P_b^* = \max_j P_j^*$. If $P_b^* = 0$, then the shifted Groves mechanism is efficient, which obviously minimizes the worst-case efficiency loss. We only consider $P_b^* > 0$. When $P_b^* > 0$, P_b^* is the minimal value that makes

$$-\sum_i \max\{\sum_{i'} L_{i'1} + P_b^* - \sum_{i' \neq i} U_{i'b}, L_{ib}\} = 0$$

Let us then consider the following settings:

- Setting 1 is the original setting. That is, for all i and j , v_{ij} is in $[L_{ij}, U_{ij}]$.
- Setting 2 is a restricted version of setting 1.
 - Outcome 1: For all i , let $v_{i1} = L_{i1}$.
 - Outcome b : For all i , let v_{ib} 's range be the same as before, that is, in $[L_{ib}, U_{ib}]$.
 - Outcome j other than 1 and b : For all i , let $v_{ij} = C_{ij}$. C_{ij} is a constant chosen from $[L_{ij}, U_{ij}]$. We need $\sum_i C_{ij} = \sum_i L_{i1}$. The C_{ij} exist because $\sum_i L_{ij} \leq \sum_i L_{i1} \leq \sum_i U_{ij}$.
- Setting 3 is obtained by dropping all outcomes other than 1 and b .
 - Outcome 1: For all i , let $v_{i1} = L_{i1}$.
 - Outcome b : For all i , let v_{ib} 's range be $[L_{ib}, U_{ib}]$.

Let H_i be the strategy-proof, individually rational, and non-deficit mechanism that minimizes the worst-case efficiency loss in setting i . Let α_i be H_i 's worst-case efficiency loss.

$\alpha_1 \geq \alpha_2$: Setting 2 is a restricted version of setting 1, so H_1 applies in setting 2, which means that the optimal worst-case efficiency loss in setting 2 is at most α_1 .

$\alpha_2 \geq \alpha_3$: We first consider setting 2 and H_2 . It is without loss of generality to assume that H_2 only chooses outcome 1 and b , because given a type profile, if H_2 chooses outcome $j \neq 1, b$, then we may simply change H_2 so that it chooses 1 instead. Then for all i , ask agent i to pay an additional payment equal to $L_{i1} - C_{ij}$. This makes the agents "as if" they are choosing outcome j . This change does not affect the agents' utilities (strategy-proofness and individual rationality are not affected). This change does not affect the total payment, as $\sum_i (L_{i1} - C_{ij}) = 0$ (non-deficit is not affected). This change does not affect the allocative efficiency, as the agents' total valuations for outcome 1 and j are the same. In conclusion, we can apply H_2 to setting 3, which implies that the optimal worst-case efficiency loss in setting 3 is at most α_2 .

Above, we have proved that $\alpha_1 \geq \alpha_3$. According to Lemma 5, in setting 3, α_3 is achieved by the optimal shifted Groves mechanism. We run the AMD process on setting 3. We get that α_3 is the minimal value that makes

$$-\sum_i \max\{\sum_{i'} L_{i'1} + \alpha_3 - \sum_{i' \neq i} U_{i'b}, L_{ib}\} = 0$$

Now earlier we have shown that the optimal shifted Groves mechanism in setting 1 (the original setting) has the same

worst-case efficiency loss as α_3 . Given that $\alpha_1 \geq \alpha_3$, we must have that the optimal shifted Groves mechanism is optimal in the original setting, which concludes the proof. \square

COROLLARY 1. *There exists a mechanism that is strategy-proof, individually rational, non-deficit, and efficient, iff for all $j > 1$, we have*

$$-\sum_i \max\left\{\sum_{i'} L_{i'1} - \sum_{i' \neq i} U_{i'j}, L_{ij}\right\} \leq 0$$

PROOF. According to Theorem 4, there exists a mechanism that is strategy-proof, individually rational, non-deficit, and efficient, iff the optimal shifted Groves mechanism is efficient.

Now the optimal shifted Groves mechanism is efficient iff it is $M(0, 0, \dots, 0)$. That is, the optimal shifted Groves mechanism is efficient iff the AMD process carries out no updates at all, which is exactly when for all $j > 1$, we have

$$-\sum_i \max\left\{\sum_{i'} L_{i'1} - \sum_{i' \neq i} U_{i'j}, L_{ij}\right\} \leq 0$$

\square

If for all $j > 1$, $\sum_i L_{ij} \geq 0$, then the above corollary applies. In this case, the original Groves mechanism $M(0, 0, \dots, 0)$ (efficient) is the optimal shifted Groves mechanism.

5. NUMERICAL EXPERIMENTS

Given a setting, let the optimal shifted Groves mechanism be $M(P_1^*, P_2^*, \dots, P_k^*)$. We recall that we must have $P_1^* = 0$. For $j > 1$, $P_j^* = 0$ if

$$-\sum_i \max\left\{\sum_{i'} L_{i'1} - \sum_{i' \neq i} U_{i'j}, L_{ij}\right\} \leq 0$$

Otherwise, P_j^* is the minimal value that makes

$$-\sum_i \max\left\{\sum_{i'} L_{i'1} + P_j^* - \sum_{i' \neq i} U_{i'j}, L_{ij}\right\} = 0$$

At the end, the maximal P_j^* defines the optimal worst-case efficiency loss.

We revisit Example 1 (the public project problem). There are only two outcomes. The optimal shifted Groves mechanism is $M(0, P_2^*)$. For $j = 2$, we have

$$\begin{aligned} & -\sum_i \max\left\{\sum_{i'} L_{i'1} - \sum_{i' \neq i} U_{i'2}, L_{i2}\right\} \\ & = -\sum_i \max\left\{-\sum_{i' \neq i} \frac{n-1}{n}, -\frac{1}{n}\right\} > 0 \end{aligned}$$

We then need to find the minimal P_2^* that makes

$$-\sum_i \max\left\{P_2^* - \sum_{i' \neq i} \frac{n-1}{n}, -\frac{1}{n}\right\} = 0$$

We get that $P_2^* = \frac{(n-1)^2}{n}$, which is then the optimal worst-case efficiency loss of the public project problem.

Resource allocations with externalities can also be modelled as social decision problems, for which our AMD approach can be applied. Below we present a variant of the k -winner selection problem described in [16].

Example 3. There are n agents. We need to choose at most 2 winners (e.g. there are 2 Ad slots for n advertisers). There are $1 + n + \frac{n(n-1)}{2}$ outcomes. There is 1 outcome that selects 0 winners. There are n outcomes that select 1 winner. Finally, there are $\frac{n(n-1)}{2}$ outcomes that select 2 winners.

Agent i 's valuation for an outcome consists of two parts. One part is her valuation for winning, which is a value from 0 to 1. The other part is her "externality" toward the other agent selected by the outcome, if she is not the only winner. The externality value is from -1 to 1.

For outcome $\{\}$ (no winners), the agents' valuations are all 0. For outcome $\{a\}$ (the only winner is a), agent a 's valuation is from 0 to 1, and the other agents' valuations are all 0. For outcome $\{a, b\}$ (the winners are a and b), agent a 's valuation is from -1 to 2 (same for b), and the other agents' valuations are all 0.

As usual, we sort the outcomes according to the agents' minimal total valuations. So outcome 1 should be $\{\}$. We have $L_{i1} = U_{i1} = 0$. The next n outcomes should be $\{a\}$ for $a = 1, 2, \dots, n$. For outcome j among these outcomes, we have $\sum_i L_{ij} = 0$, so these outcomes never cause deficit and can be ignored. Finally, the next $\frac{n(n-1)}{2}$ outcomes should be $\{a, b\}$ for every pair of a and b from $1, 2, \dots, n$.

The worst-case efficiency loss is caused by one of the outcomes that selects 2 winners. Due to symmetry, we only consider the outcome $\{1, 2\}$. We denote this outcome by j . We have that the worst-case efficiency loss for this outcome is P_j , where P_j is the minimal value that makes

$$-\sum_i \max\left\{\sum_{i'} L_{i'1} + P_j - \sum_{i' \neq i} U_{i'j}, L_{ij}\right\} = 0$$

Plugging in the numbers, we have

$$-\sum_i \max\{P_j - \sum_{i' \neq i} U_{i'j}, L_{ij}\} = 0$$

Simplify it, we get

$$-2 \max\{P_j - 2, -1\} - (n-2) \max\{P_j - 4, 0\} = 0$$

Solving the above equation, we get that P_j should be 2. That is, the optimal worst-case efficiency loss is 2.

6. CONCLUSION

For the problem where a group of agents need to choose from a finite set of outcomes, we proposed the optimal shifted Groves mechanism, which is strategy-proof, individually rational, non-deficit, and minimizes the worst-case efficiency loss. One immediate future research direction is to study other objectives, such as minimizing the expected efficiency loss given a prior distribution. We may also consider settings where we drop certain assumptions (e.g., an agent's valuation space is not a closed interval, or the agents' valuations are interdependent).

7. ACKNOWLEDGMENTS

This work was partially supported by JSPS KAKENHI Grant Number 24220003 and 26730005, and JST PRESTO program.

REFERENCES

- [1] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.
- [2] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, Edmonton, Canada, 2002.
- [3] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce (AMEC), LNAI, 3435*, pages 30–43, 2005.
- [4] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [5] 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.
- [6] M. Guo and V. Conitzer. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1676–1679, Atlanta, GA, USA, 2010. NECTAR track.
- [7] M. Guo, E. Markakis, K. R. Apt, and V. Conitzer. Undominated groves mechanisms. *Journal of Artificial Intelligence Research*, 46:129–163, 2013.
- [8] M. Guo, V. Naroditskiy, V. Conitzer, A. Greenwald, and N. R. Jennings. Budget-balanced and nearly efficient randomized mechanisms: Public goods and beyond. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE)*, Singapore, 2011.
- [9] B. Holmström. Groves’ scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.
- [10] A. Likhodedov and T. Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 232–237, San Jose, CA, USA, 2004.
- [11] A. Likhodedov and T. Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, USA, 2005.
- [12] D. Mishra and A. Sen. Roberts’ theorem with neutrality: A social welfare ordering approach. *Games and Economic Behavior*, 75(1):283–298, 2012.
- [13] H. Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
- [14] V. Naroditskiy, M. Guo, L. Dufton, M. Polukarov, and N. R. Jennings. Redistribution of VCG payments in public project problems. In *Proceedings of the Eighth Workshop on Internet and Network Economics (WINE)*, Liverpool, 2012.
- [15] K. Roberts. The characterization of implementable social choice rules. In J.-J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland Publishing Company, 1979.
- [16] Y. Sakurai, T. Okimoto, M. Oka, and M. Yokoo. Strategy-proof mechanisms for the k-winner selection problem. In *Principles and Practice of Multi-Agent Systems (PRIMA)*, Dunedin, New Zealand, 2013.