Optimizing Payments in Dominant-Strategy Mechanisms for Multi-Parameter Domains

Lachlan Dufton
Cheriton School of Computer Science
University of Waterloo, Canada
ltdufton@cs.uwaterloo.ca

Victor Naroditskiy
School of Electronics and Computer Science
University of Southampton, UK
vn@ecs.soton.ac.uk

Maria Polukarov
School of Electronics and Computer Science
University of Southampton, UK
mp3@ecs.soton.ac.uk

Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton, UK
nrj@ecs.soton.ac.uk

Abstract

In AI research, mechanism design is typically used to allocate tasks and resources to agents holding private information about their values for possible allocations. In this context, optimizing payments within the Groves class has recently received much attention, mostly under the assumption that agent's private information is single-dimensional. Our work tackles this problem in multi-parameter domains. Specifically, we develop a generic technique to look for a best Groves mechanism for any given mechanism design problem. Our method is based on partitioning the spaces of agent values and payment functions into regions, on each of which we are able to define a feasible linear payment function. Under certain geometric conditions on partitions of the two spaces this function is optimal. We illustrate our method by applying it to the problem of allocating heterogeneous items.

Introduction

Mechanism design is concerned with problems that involve multiple self-interested participants having private information about their types. In particular, it has been widely applied to task and resource allocation scenarios, which are fundamental to AI. A central result in mechanism design is the Groves class (Vickrey 1961; Clarke 1971; Groves 1973), which contains efficient[1], dominant-strategy mechanisms. Mechanisms within this class are parameterized by the *rebate function* that specifies the part of an agent's payment which is independent of his type. The most well-known mechanism of this kind is the Vickrey-Clarke-Groves (VCG)[2] mechanism that charges each agent his externality—i.e., the amount by which his presence affects the total value of the other agents.

However, while VCG is perhaps the most natural of the Groves mechanisms, it is not always the best choice. In particular, in settings with no auctioneer, VCG payments provide a dominant-strategy implementation but noone benefits from the collected payments, which represent a loss in social welfare. In such cases, redistribution of VCG payments back to the agents is desirable. To this end, the Groves mechanism that redistributes the highest possible fraction of social welfare has been proposed in (Moulin 2009; Guo and Conitzer 2009) for the sale of identical items. In other contexts, the question of choosing the "best" Groves mechanism has been studied in (Bailey 1997; Porter, Shoham, and Tennenholtz 2004; Cavallo 2006; Guo and Conitzer 2010b; Gujar and Narahari 2011; Guo 2011; 2012) for various models and objectives.

Generally, a mechanism design problem is specified by an objective function (e.g., welfare-maximization, fairness, or revenue-maximization) and constraints (e.g., individual rationality[3] or weak budget balance[4]). Our goal is to find a mechanism within the Groves class that optimizes the given objective under the given constraints. Let $v \in V = \Theta^n$ be a type profile, where Θ denotes the space of types for each of n agents. The vector $v_{-i} \in \Theta^{n-1}$ stands for types of the agents other than i, and the space $W = \Theta^{n-1}$ of all vectors v_{-i} is the same for each i. Mechanisms within the Groves class differ from one another by the part of an agent's payment that is independent of his type, which is specified by a rebate function $h : W \to \mathbb{R}$. Thereby, the problem of finding a best Groves mechanism for a given mechanism design problem can be expressed as an optimization problem:[5,6]

$$\text{maximize}_{h:W\to\mathbb{R}} \text{ objective value} \quad \text{s.t.} \quad \forall\, v \in V \qquad (1)$$
$$\text{objective value is achieved}$$
$$\text{constraints hold}$$

To date, however, most of the literature on optimal Groves mechanisms was devoted to single-parameter domains

[1] An mechanism is *efficient* if it chooses the allocation that maximizes the sum of the values of allocated agents.

[2] Also referred to as the pivotal or the Clarke mechanism.

[3] A mechanism is *individually rational* if each agent is never worse off after participating in the mechanism.

[4] A mechanism is *weakly budget-balanced* if the total amount of payments made by the agents is non-negative.

[5] We refer the reader to Section 3 of (Naroditskiy, Polukarov, and Jennings 2012) for a single-parameter example of a mechanism design problem modeled this way.

[6] Some combinations of constraints (e.g., weak budget balance and 2-fairness) may be impossible to implement: this is identified by the lack of a feasible solution to the optimization problem.

where an agent's private information is characterized by a single number (i.e., $\Theta \subseteq \mathbb{R}$). While these models are applicable to many important domains (e.g., public good and allocation of homogeneous items), they do not capture many other scenarios of practical interest such as allocation of heterogeneous items or public choice with several alternatives. Therefore, optimization within the Groves class for *multi-parameter* types is an important task.

In this direction, there has been comparatively little work. In particular, Gujar and Narahari (2011) present a conjecture about an optimal mechanism for allocation of heterogeneous items to agents with unit demand, which Guo (2012) proves correct; Cavallo (2006) gives a widely applicable, but not necessarily optimal mechanism, and Guo and Conitzer (2009) find an optimal mechanism for the allocation of homogeneous items among agents desiring multiple copies. More recently, Guo (2011) studied welfare-maximizing Groves mechanisms in a combinatorial auction setting and derived a number of results for allocation of heterogeneous items. In all these examples, agent types are represented by m-dimensional vectors, that is $\Theta \subseteq \mathbb{R}^m$ for some $m \in \mathbb{N}$. However, each of these papers studies a specific scenario.

In contrast, our work is in the spirit of automated mechanism design (Conitzer and Sandholm 2002; Guo and Conitzer 2010a): we develop a generic technique that takes a mechanism design problem as an input and searches for an optimal Groves mechanism for the problem. The same technique can be applied to any of the problems often avoiding the need for custom-made solutions for each problem. Note that this generic optimization problem (1) cannot be solved directly: optimization is over functions and there is an infinite number of constraints. Nonetheless, the technique we propose makes it possible to tackle such problems effectively.

Specifically, we design an algorithm for multi-parameter domains, that allows optimization over Groves mechanisms by considering various subclasses of rebate functions. In more detail, each of these subclasses is associated with a subdivision[7] of the rebate space W (i.e., the domain of the rebate function) such that all rebate functions in the subclass are linear on each region of the subdivision. Given such a subdivision, we construct a linear program that finds the rebate function, which is optimal within its corresponding subclass. These rebates satisfy the constraints of a given mechanism design problem (1), and thus provide a *lower bound* on its objective. Importantly, the quality of the solution achieved depends on the choice of the subclass. That is, by initializing the process with different subdivisions, one can be guided in the direction of finding the subdivision that yields an optimal rebate function. The distance between the lower bound and the optimal solution can be measured by comparing the current objective value to an *upper bound*, which can be computed with another algorithm that we provide. If the bounds coincide, we have found the optimal rebate function.

[7] A *subdivision* (or, *partition*) of a space is a collection of its disjoint subsets (or, *regions*), whose union covers the entire space.

Our results build upon and significantly extend a technique for finding optimal rebate functions in single-parameter domains (Naroditskiy, Polukarov, and Jennings 2012). This previous technique reduces the problem of finding optimal rebates to identifying *consistent partitions* of the rebate space and the space of agent values. Unfortunately, consistent partitions are sometimes difficult to identify. To this end, this paper provides a heuristic for finding (not necessarily optimal) rebate functions when such partitions are not available. We would like to remark that the results we present do not require finding consistent partitions. However, if such partitions are available, our method will produce an optimal rebate function.

We illustrate our method on the problem of welfare-maximizing allocation of heterogeneous items under the constraints of weak budget balance, and individual rationality. Our analysis explains existing mechanisms in terms of subdivisions of the rebate space they induce. We describe the subdivision on which VCG rebates are linear, and then show how it is refined to define the Bailey/Cavallo mechanism (Bailey 1997; Cavallo 2006), and then further refined for the `HETERO` (Gujar and Narahari 2011) mechanism. Finally, the negative result in (Gujar and Narahari 2011) follows easily from this geometric perspective: using a single region in the rebate space, the best linear rebate function yields zero welfare in the worst case.

Concisely, our work extends the state of the art as follows:

- We propose a methodology to optimize payments in Groves mechanisms that is not problem-specific and can be applied to any problem. This is the first general technique for multi-parameter domains.
- For both single- and multi-parameter domains, this is the first technique that provides *feasible* (i.e., satisfying all of the given constraints) mechanisms, regardless of the availability of consistent partitions. In cases where such partitions are known, optimal solutions are guaranteed.
- Our method gives a unifying geometric perspective for understanding results developed by other authors for specific problems.

The paper unfolds as follows. We first present the formal model of multi-parameter domains, and cover preliminaries. Specifically, we provide all necessary definitions from polyhedral geometry and generalize the concept of consistent subdivisions to multi-parameter domains. In the following section, we introduce our main contribution, which is a heuristic technique for finding payment functions for cases with no (known) consistent partitions. We then apply the technique to allocation of heterogeneous items, before our conclusions and future work.

Multi-Parameter Domains

We consider multi-parameter domains where n agents participate in a mechanism choosing a social outcome $k \in K$. The type of agent i is given by an m-dimensional vector $v_i \in [0,1]^m$, where v_i^j denotes the value agent i gets in a publicly known subset of outcomes $K_i^j \subseteq K$.[8] We use $V = [0,1]^{nm}$ to denote the *value space*. In the case of allocating m heterogeneous items to agents with unit-demand, v_i^j denotes the value agent i derives from owning item j, and K_i^j is the set of all outcomes where agent i is allocated item j. The model also applies to non-allocation scenarios such as public project with multiple alternatives, where participants need to decide which one of m possible projects to undertake. There v_i^j denotes how much agent i values project j and K_i^j is the outcome undertaking project j (in this case, $K_1^j = \ldots = K_n^j$). Obviously, single-parameter domains are a special case obtained by setting $m = 1$.

We make the standard assumption of utilities linear in money and denote payments by $t : V \to \mathbb{R}^n$, where $t_i(v) \in \mathbb{R}$ is the payment collected *from* agent i given value profile v. We focus on Groves mechanisms, thus fixing the allocation function to choose the efficient[1] allocation, denoted by $f^*(v)$. In a slight abuse of notation, we use $v_i(k)$ to denote the value of agent i under the outcome k. Payments under a Groves mechanism are set according to the rule $t_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(f^*(v))$ where $h_i : W \to \mathbb{R}$ is an arbitrary function that only depends on values of the other agents $v_{-i} \in W = [0,1]^{m(n-1)}$, with W dubbed the *rebate space*. We will refer to the expression $\sum_{j \neq i} v_j(f^*(v))$ in the payment function as the *Groves payment* and to h_i as the *rebate function*. Apt *et al.* (2008) show that the restriction to anonymous rebate functions $h : W \to \mathbb{R}$ is without loss of generality. Thus, mechanisms within the Groves class differ only in the function h.

In this work we design a generic method to optimize the rebate function. That is, we take as an input a mechanism design problem in the form of equation (1) and search for the best Groves mechanism for it. Recall that the optimization problem (1) cannot be solved directly: we need to optimize over functions, with an infinite number of constraints. Therefore, to overcome these difficulties, we partition the value and the rebate spaces in a certain way and consider subclasses of rebate functions associated with these partitions. In the next section we cover the necessary preliminary results, before introducing consistent partitions, which give the basis for our method.

Preliminaries

In this section, we introduce the notation and preliminaries necessary for presentation of our main results in following sections. In particular, we provide relevant concepts and definitions from polyhedral geometry that are needed to extend the results on consistent partitions from (Naroditskiy, Polukarov, and Jennings 2012) to multi-parameter domains.

In particular, we define a d-dimensional polytope, p, as a convex hull of a finite set of points in the Euclidean space. Equivalently, we can define it as a finite intersection of halfspaces: $p = \{x \in \mathbb{R}^d \mid Ax \geq b\}$, where $A \in \mathbb{R}^{s \times d}$, $b \in \mathbb{R}^s$, and s is the number of halfspaces. Thus, a pair (A, b) denotes the corresponding polytope. For any p, let the *relative interior*, *relint*(p), denote the polytope without its facets.

Definition 1 *A set P_X of polytopes is a* subdivision *(equivalently, a* partition*) of the polytope X if the polytopes P_X do not overlap:* $\text{relint}(p) \cap \text{relint}(q) = \varnothing,\ \forall\, p, q \in P_X$*, and cover exactly the polytope X:* $\bigcup_{p \in P_X} p = X$.

Each polytope $p \in P_X$ defines a *region* of the partitioned space X. We will subdivide the value space V and the rebate space W so that the allocation function is constant on each region of P_V and the rebate function is linear on each region of P_W. In order to be able to define such piecewise linear rebates, which will be feasible for a given mechanism design problem, we will require these subdivisions to satisfy some consistency conditions (to be defined) w.r.t. the regions of both subdivisions and their *extreme points* (or, *vertices*).

We proceed to introduce some additional concepts necessary for defining the consistency conditions for P_V and P_W.

Definition 2 *A subdivision P_X* refines *a subdivision P'_X if for each $p \in P_X$ there is a $p' \in P'_X \mid p \subseteq p'$.*

The above definition extends to *refinements* of sets of polytopes that are not necessarily subdivisions, as follows.

Definition 3 *A subdivision P_X of the polytope X* refines a polytope q *if for all $p \in P_X$ the intersection with q is either empty or p:* $\text{relint}(p) \cap \text{relint}(q) = \varnothing \vee p \cap q = p$*. A subdivision P_X* refines a set of polytopes Q *if P_X refines all polytopes $q \in Q$.*

Recall that Groves mechanisms allocate the items efficiently. That is, the sum of values of the allocated agents for the items is maximal among all possible allocations. Thus, a region of the value space on which the allocation function is constant (e.g., agent 1 is allocated item 1, agent 2—item 2, . . . , and agent m—item m) is given by the intersection of halfspaces as defined by the inequalities indicating that the value of a particular allocation (e.g., $v_1^1 + v_2^2 + \ldots + v_m^m$) is greater than the value of any other allocation. Thereby, allocation regions induce a subdivision of the value space. We denote this *initial subdivision* by P_V^I.

Note that all typical constraints (e.g., individual rational and weak budget balance) are linear in Groves payments and rebates throughout each of these regions,[9] and thus can be represented with coefficients $\alpha \in \mathbb{R}^{nm}$, $\beta \in \mathbb{R}^n$ and $\gamma_0 \in \mathbb{R}$ as follows:

$$\sum_{i=1}^{n} \sum_{j=1}^{m} \alpha_i^j v_i^j + \sum_{i=1}^{n} \beta_i h(v_{-i}) \geq \gamma_0 \qquad (2)$$

We are going to exploit linearity of constraints throughout each region in the value space, referred to as a *value region*. To this end, we require the final subdivision P_V of the value space to refine P_V^I. In fact, given a subdivision of the rebate space, on which the rebate function is linear, we may need to refine the initial subdivision in order to guarantee that each agent's rebate is linear throughout a value region. This will allow the constraints in (2) to be represented by coefficients

[8] This notation is a natural generalization of the single-parameter domain definition (Definition 9.33) in (Nisan 2007).

[9] See Section 3 in (Naroditskiy, Polukarov, and Jennings 2012) for an example illustrating this linearity.

$\gamma \in \mathbb{R}^{nm+1}$ throughout each value region:

$$\sum_{i=1}^{n}\sum_{j=1}^{m}\gamma_i^j v_i^j \geq \gamma_0 \quad (3)$$

The following definitions help us describe such regions.

Definition 4 Lifting *of a subdivision P_W from W to V is a set of polytopes in V*

$$\mathrm{lift}(P_W) = \bigcup_{(A,b)\in P_W}\bigcup_{i=1}^{n}(Av_{-i} \geq b) \cap V$$

Definition 5 (Region consistency) *Subdivisions P_V, P_W are* region-consistent *if P_V refines the polytopes* $\mathrm{lift}(P_W)$.

Region consistency is a geometric encoding of the property that for each region $q \in P_V$ and for each agent i, the rebate for this agent is given by the same rebate function h_p for all $v \in q$, where $p \in P_W$ (see Lemma 2 in (Naroditskiy, Polukarov, and Jennings 2012)).

Next we turn to the vertex consistency condition. The set of the extreme points of the polytopes in subdivision P_X is denoted by $\hat{P}_X$.

Definition 6 *Given a subdivision P_V, the projection of its extreme points $\hat{P}_V$ on W is*

$$\Pi_W(\hat{P}_V) = \bigcup_{v\in\hat{P}_V}\bigcup_{i=1}^{n} v_{-i}$$

Definition 7 (Vertex consistency) *Subdivisions P_V and P_W are* vertex-consistent *if the projection of the extreme points of P_V is the extreme points of P_W*

$$\Pi_W(\hat{P}_V) = \hat{P}_W$$

Armed with the definitions above, we are ready to present our results in the next two sections.

Consistent Partitions

In (Naroditskiy, Polukarov, and Jennings 2012), the authors showed for single-parameter domains that region- and vertex-consistency, coupled with some additional conditions on P_V and P_W collectively called *consistent partitions*, imply the existence of rebates that are linear on each region of P_W and are optimal for the mechanism design problem being solved.

Definition 8 *For a given initial partition P_V^I, partitions P_V and P_W are* consistent *if:* (i) *P_V refines P_V^I;* (ii) *P_V and P_W are region- and vertex-consistent; and* (iii) *each polytope in P_W has $(n-1)m+1$ extreme points.*

This definition generalizes consistent partitions from (Naroditskiy, Polukarov, and Jennings 2012) to multi-parameter domains. The only difference in definitions is in the number of extreme points in condition (iii): in the $(n-1)m$-dimensional W space, $(n-1)m+1$ points define a linear function.

Before showing how consistent partitions can be used to find optimal rebate functions, we need additional notation. Suppose the set of values is not the infinite set V, but a finite subset $\hat{V} \subset V$. When we look for optimal rebates assuming the set of value profiles is finite, we speak of the *restricted problem*. In fact, the restricted problem is a linear program, which can be solved to obtain an optimal solution as we formally define in Figure 1. Notice that in a restricted problem we are looking for a finite number of rebates (one for each element in $\Pi_W(\hat{V})$) rather than for rebate functions.

As Theorem 1 below shows, we can use consistent subdivisions with the restricted problem to find optimal rebates.

Algorithm `RestrictedProblem`
Input: set of profiles $\hat{V}$
Output: payment values $\hat{h}$ on $\Pi_W(\hat{V})$ and upper bound obj
/** solve a linear program **/
$\max_{obj\in\mathbb{R},\ \hat{h}\in\mathbb{R}^{|\hat{V}|}} obj$ s.t. $\forall\, v \in \hat{P}_V$
$constraints(v, \hat{h}(v_{-1}), \ldots, \hat{h}(v_{-n}))$
return $(\hat{w}, obj)$

Figure 1: LP for an upper bound.

Theorem 1 *Let P_V and P_W denote consistent subdivisions for a mechanism design problem with an allocation function inducing the initial partition P_V^I. Let $\{\hat{h}(w) \mid w \in \hat{P}_W\}$ denote the set of rebates from an optimal solution to the restricted problem, which only considers profiles $\hat{P}_V$. Further, let $\hat{p}$ denote the set of $(n-1)m+1$ extreme points of a polytope $p \in P_W$. For each polytope, define a linear rebate function $h_p(w) = \sum_{i=1}^{(n-1)m} a_i^p w_i + b^p$ with coefficients $a^p \in \mathbb{R}^{(n-1)m}, b^p \in \mathbb{R}$ given by a solution to the system of linear equations $\{\hat{h}(w) = \sum_{i=1}^{(n-1)m} a_i^p w_i + b^p \mid w \in \hat{p}\}$. Then, the following rebate function is* optimal *for the mechanism design problem: for $w \in p$, $h(w) = h_p(w)$.*

Proof The proof is a straightforward generalization of the proof of Theorem 2 in (Naroditskiy, Polukarov, and Jennings 2012). Property (i) of Definition 8 ensures that the constraints on a value region $q \in P_V$ are of the form given in Equation (2). Further, property (ii) guarantees that if h is linear on P_W, then the constraints can be represented by linear coefficients as shown in Equation (3). We can construct a linear function h_p by interpolating the extreme points as in $(n-1)m$ dimensions, $(n-1)m+1$ linearly independent points define a linear function (if some of them are dependent, they define a family of functions). Thus, for a region p with $(n-1)m+1$ extreme points (property iii), there exists at least one linear function $h_p(w) = \sum_{i=1}^{(n-1)m} a_i^p w_i + b^p$ with $a^p \in \mathbb{R}^{(n-1)m}, b^p \in \mathbb{R}$ given by a solution to the system of linear equations $\{\hat{h}(w) = \sum_{i=1}^{(n-1)m} a_i^p w_i + b^p \mid w \in \hat{p}\}$. By construction, $h(w)$ is linear on P_W and the constraints hold at the extreme points of each $q \in P_V$. It is an easy observation that a linear constraint holds at the extreme points of a polytope if and only if it holds on the entire polytope, and thus $h(v_{-i})$ satisfies the constraints at all points $v \in q$ for each q. That is, a solution to the restricted problem that only includes constraints for extreme points of P_V is feasible in the original problem, i.e. the constraints are satisfied for each $v \in V$. As the objective function in (1) is represented by a constraint, and the objective value of the restricted problem (i.e., the upper bound) is achieved for all $v \in V$, the rebate function $h(w)$ is optimal. □

Note that Theorem 1 provides a constructive proof of the existence of a piecewise linear optimal solution for any mechanism design problem that admits consistent partitions. Given such partitions P_V and P_W, one just needs to solve the restricted problem over the set of extreme points $\hat{P}_V$ and then define a linear rebate function for each rebate region by linearly interpolating optimal rebates at its extreme points. It is therefore important to identify settings that admit such partitions. A special case of consistent partitions for multi-parameter domains has been studied in (Guo and Conitzer 2010b). There, "order-consistent" classes correspond to consistent subdivisions where the value and rebate spaces are subdivided with hyperplanes of the form $x_i = x_j$ for each pair of coordinates (i, j).[10]

However, for many problems it is not clear how to find consistent partitions or, indeed, whether they exist at all. In this paper, we provide a procedure that computes heuristic solutions to mechanism design problems where consistent partitions are not available.

The Heuristic Technique

In this section, we present our main result: a technique for finding and evaluating solutions by solving linear programs. All of the results in this section are novel for both single and multi-parameter domains. Our approach builds upon the following observation. While it may be difficult to find partitions that are both region- and vertex-consistent, we notice that region-consistency alone holds for any P_W and P_V that refines *lift*(P_W). As before, to preserve linearity of constraints on each region of a partition P_V, we require P_V to refine the initial partition P_V^I as defined by the allocation function. To define such a partition, we need the following notation.

Definition 9 *Given subdivisions P_X and P'_X of a polytope X, their* intersection *is the set of polyhedra*

$$\{p \cap p'\}_{p \in P_X,\ p' \in P'_X,\ \mathrm{relint}(p) \cap \mathrm{relint}(p') \neq \varnothing}$$

Intuitively, an intersection is obtained by placing one subdivision on top of the other. Importantly, as we show below, intersecting subdivisions produce a new subdivision.

Lemma 1 *The intersection of subdivisions is a subdivision.*

Proof By construction, the resulting set of polyhedra defines non-overlapping regions that cover the entire polytope. All that we need to show is that these regions are convex (i.e., polytopes). This follows immediately from the fact that each region is given by the intersection of convex regions, which itself is convex. □

A *minimal partition* that refines both P_V^I and *lift*(P_W) is their intersection. Note however that *lift*(P_W) is not necessarily a subdivision of the value space but a collection of possibly overlapping polytopes covering V, and so definition 9 is not directly applicable. A natural way to overcome this problem is to replace *lift*(P_W) with the collection of (disjoint) polyhedra that is induced by the facets of the polytopes *lift*(P_W). Yet, for this collection to define a subdivision of V, the polyhedra must be convex. The next lemma shows this is indeed the case.

Lemma 2 *For a subdivision P_W of the rebate space, the linear constraints of the polytopes* lift(P_W) *define a subdivision of the value space.*

Proof First note that the polytopes *lift*(P_W) cover the polytope V: any point $v \in V$ belongs to the lifted polytopes $\{p \in P_W \mid v_{-i} \in p\}_{i=1}^n$. Thus, the constraints of the polytopes *lift*(P_W) subdivide V into polyhedra. We need to prove that each polyhedron is a polytope (i.e., convex).

By construction, each boundary between adjacent regions (polyhedra) from *lift*(P_W) corresponds to a boundary in P_W for some i. Crossing a region boundary in V corresponds to crossing a boundary in P_W. Suppose that V contains a non-convex region q. Then one can cross a boundary while following an interval that starts and ends in q: for some $\lambda \in (0, 1)$, there must exist $v', v^* \in q \mid v = \lambda v' + (1 - \lambda) v^* \notin q$. Denote the region[11] and the lifting dimension that generated the boundary by $p \in P_W$ and i. By construction of q, the rebate for agent i comes from the same rebate region throughout the q region: $v_{-i} \mid v \in \mathit{relint}(q) \subseteq p$ for some unique $p \in P_W$. In particular, v'_{-i} and v^*_{-i} belong to the same region $p \in P_W$. Now, since p is convex, v_{-i} must also belong to p, which is impossible since we crossed the boundary. □

By Lemmas 1 and 2, given a subdivision of the rebate space P_W, we can define a subdivision P_V of the value space as the intersection of the initial partition P_V^I and the set of polytopes as induced by the constraints of *lift*(P_W). Moreover, as the following lemma shows, we can further intersect region-consistent partitions, while preserving region-consistency.

Lemma 3 *Let P_V^* be region-consistent with P_W^* and P'_V region-consistent with P'_W. The subdivisions P_V and P_W given by the intersection of P_V^* with P'_V and P_W^* with P'_W are region-consistent.*

Proof We need to argue that all region boundaries of *lift*(P_W) are used to define P_V. But this is immediate as the only boundaries of P_W are the ones from P_W^* and P'_W, and by consistency of P_V^*, P_W^* and P'_V, P'_W, the boundaries *lift*(P_W^*) and *lift*(P'_W) are present in P_V. □

We use the lemmas above in the derivation of our main result in Theorem 2. Intuitively, we choose a subdivision of the rebate space and find rebate functions that are optimal among linear rebates on this subdivision. The objective value achieved by this mechanism provides a lower bound

[10] An interested reader is referred to Chapter 6.3.2 in (De Loera, Rambau, and Santos 2010) for more details on such triangulations of hypercubes.

[11] There may be multiple such regions, in which case apply the proof to each of them, and there will be at least one that will lead us to a contradiction.

on the true objective value. The performance of the rebate function can be evaluated against an upper bound that we obtain by solving the restricted problem with a finite number of profiles presented in Figure 1. If these two bounds match, the mechanism is optimal, otherwise the gap indicates how suboptimal the mechanism may be. A large gap is an indication that one of the bounds should be improved, which can be done as we discuss next.

These upper and lower bound procedures provide a computational way of searching for good payment functions. The selected subdivision of the rebate space determines the efficacy of the rebate function produced in the lower bound LP, so we can try different subdivisions in search for the highest lower bound. Similarly, the choice of value profiles determines the tightness of the upper bound, so we can plug in different subsets of profiles in search for the lowest upper bound.

Heuristic Rebates (Lower Bound)

For any subdivision P_W, we present a technique for finding a feasible rebate function that is optimal among rebate functions linear on P_W. Specifically, we first find a subdivision of the value space P_V as the intersection of P_V^I and *lift*(P_W). Then we define a linear program that (like the restricted problem in Figure 1) includes constraints only for the extreme points of P_V, and also enforces linearity of rebates on the regions of P_W. The resulting linear program is stated in Figure 2. The variables are the rebates $\hat{h}$ for the extreme points $\hat{P}_V$ and the coefficients $a^p \in \mathbb{R}^{(n-1)m+1}$ defining a linear rebate function on each region $p \in P_W$. As we show, an optimal solution (in fact, any feasible solution) to this linear program is a feasible solution in the original problem.[12]

Algorithm `LinearRebates`
Input: initial subdivision P_V^I, subdivision P_W
Output: rebate function h and the objective value obj
P_V - intersection of *lift*(P_W) and P_V^I
/** solve a linear program **/
$\max_{obj \in \mathbb{R},\ \hat{h} \in \mathbb{R}^{|\Pi_W(\hat{P}_V)|},\ a \in \mathbb{R}^{|P_W| \times ((n-1)m+1)}} obj$
s.t. $constraints(v, \hat{h}(v_{-1}), \ldots, \hat{h}(v_{-n})) \quad \forall\, v \in \hat{P}_V$
$\hat{h}(w) = \sum_{i=1}^{(n-1)m} a_i^p w + a_{(n-1)m+1}^p \quad \forall\, p \in P_W,\ w \in \hat{p}$
/** coefficients a define a rebate function linear on P_W **/
$h(w) = \{\sum_{i=1}^{(n-1)m} a_i^p w + a_{(n-1)m+1}^p\}_{p \in P_W}$
return $(h(w), obj)$

Figure 2: LP for a lower bound.

Theorem 2 (Main Result) *Given a rebate space subdivision P_W and an initial subdivision P_V^I of the value space, the algorithm* `LinearRebates`(P_W, P_V^I) *finds a feasible rebate function.*

Proof Recall that P_V is constructed by placing the polytopes *lift*(P_W) on top of the subdivision P_V^I. By Lemmas 1 and 2, this results in a valid subdivision of V that refines both *lift*(P_W) and P_V^I. Hence, P_V and P_W are region consistent, and for any $q \in P_V$, the rebate for agent i is given by the same rebate function h_p for all $v \in q$. Again, by construction the rebate function is linear on each $p \in P_W$: the linearity constraints are enforced explicitly in Figure 2. Combining this with the fact that P_V refines P_V^I, we get that the constraints are linear on each region $v \in q$: i.e., can be represented with a set of coefficients α as in (3). The constraints are satisfied at the extreme points of each $q \in P_V$ and are linear throughout q. Thus, they are satisfied for at all $v \in q$ and the rebate function is feasible. $\square$

The rebates found by `LinearRebates`(P_W, P_V^I) are optimal for a restricted class of solutions, thus providing a lower bound on the objective value.

Upper Bound on the Objective Value

The solution to `LinearRebates` provides a feasible rebate function and the objective value achieved with this function. This value provides a lower bound on the optimal objective value. Obviously, the quality of this bound depends on the choice of partition P_W of the rebate space. We can check the quality of a solution by comparing its lower bound to an upper bound, which we can obtain by finding optimal rebates for a finite set of profiles $\hat{V} \subset V$ using the `RestrictedProblem` algorithm in Figure 1. A natural choice for $\hat{V}$ is the set of extreme points of the subdivision P_V considered in `LinearRebates`.

We can consider different sets of profiles $\hat{V} \subset V$ in the search for the lowest upper bound. Similarly, we can try different P_W to search for the rebate function with the highest lower bound. The rebate function is optimal if the highest lower bound coincides with the lowest upper bound. Note, that whenever consistent partitions P_V and P_W are available, `LinearRebates` provides an optimal solution when supplied with P_W.

Example: Heterogeneous Item Allocation

We apply our technique to item allocation among n agents with unit demand, who compete for $m < n$ heterogeneous items. Each agent (weakly) desires the items, and evaluates each item independently of the other items and agents. Thus, an agent's valuation for the m items is represented by an m-dimensional, non-negative vector, where the jth component specifies the value the agent gets from being allocated item j. We evaluate the performance of a mechanism by the social welfare it guarantees, which is the minimum fraction of the value of the efficient allocation it provides, for all value profiles. The goal is to find a Groves mechanism that maximizes this worst-case social welfare, while satisfying constraints of weak budget balance (i.e., the total amount redistributed does not exceed the total price paid by the allocated agents), and individual rationality (i.e., utility of each agent is nonnegative).

[12] It is possible that the `LinearRebates` LP does not have a feasible solution. This may be caused by the linearity restriction, in which case trying a different P_W may result in a feasible solution. Or, the original problem may not have a feasible solution. The latter case can often be confirmed by observing that the `RestrictedProblem` LP has no feasible solution for some $\hat{V}$.

Theorem 2 states that in order to find a feasible rebate function using `LinearRebates`, we need to determine the initial subdivision P_V^I, defining regions where the allocation is fixed. This subdivision is given by a set of linear inequalities, comparing the values of all pairs of allocations. We also need a rebate space subdivision for the rebate function, which we base on allocation subdivisions for subsets of agents, as a natural extension from previous work. If s is a subset of agents, then define $P_V(s)$ as the allocation subdivision for agents in s. Further, let P_V^y be the intersection of $P_V(s)$ for all y-sized subsets $s \subseteq N$. Define $P_W(s)$ and P_W^y similarly.

Theorem 3 *Subdivisions P_V^y and P_W^y are region-consistent.*

Proof We need to show that the region boundaries of the polytopes *lift*(P_W^y) are in P_V^y. By construction, a polytope $(Aw \geq b) \in P_W(s)$ determines the allocation among agents s, $|s| = y$. Let S refer to the indexes of the agents included in this subset. Lifting the polytope for any agent i, we obtain $Av_{-i} \geq b$. But $Av_{-i} \geq b$ specifies an allocation region for y agents with the following indexes

$$S'(j) = \begin{cases} S(j) & \text{if } S(j) < i \\ S(j)+1 & \text{if } S(j) \geq i \end{cases}$$

This region is part of the subdivision P_V^y as, by construction, it includes allocation regions for all subsets of size y. □

Let $P_V^{j,k}$ denote the intersection of subdivisions P_V^y for $j \leq y \leq k$ and define $P_W^{j,k}$ analogously. By Lemma 3, the subdivisions $P_V^{j,k}$ and $P_W^{j,k}$ are region-consistent. Such subdivisions are induced by the existing mechanisms we discuss below.

Designing a mechanism by choosing the rebate space partition gives a geometric perspective on the process of finding an optimal mechanism. While searching for an appropriate subdivision, we can tie together previous results for this problem under this subdivision analysis.

First, VCG rebates involve the value of the efficient allocation $f^*(v_{-i})$, for each $i \in N$. Within each region of P_W^{n-1}, the allocation of the agents in $w = v_{-i}$ stays constant, so the subdivisions of P_W^{n-1} define the regions of linearity for VCG rebates. The Bailey/Cavallo mechanism defines the rebate function to be the average sum of the *VCG payments* made in the market with $n-1$ agents with values w. An agent's rebate depends on the value of the efficient allocation for all $n-1$ agents to determine each agent's Groves payment, which will be linear within regions of P_W^{n-1}. The VCG rebates of each of the $(n-1)$ agents depend on the efficient allocation for each subset of $(n-2)$ agents (excluding each agent from w in turn). The subdivision P_W^{n-2} ensures each subset of size $(n-2)$ has a fixed allocation within a region. Combining these, an agent's rebate function for the Bailey/Cavallo mechanism is linear in regions of $P_W^{n-2,n-1}$.

The process of subdividing the rebate space can naturally continue to include subdivisions into allocation regions among $n-3$ agents $P_W^{n-3,n-1}$, $n-4$ agents $P_W^{n-4,n-1}$, etc. Adding extra subdivisions does not hurt the optimality of the mechanism, but it will lead to unnecessary complexity (both in the LPs of our technique and in the produced mechanism) without improving the objective value. The optimal mechanism for homogeneous items (Moulin 2009; Guo and Conitzer 2009) has rebates that are linear when the set and order of the lowest $(n-m-2)$ agents is fixed, and this occurs within regions of $P_W^{m,n-1}$. This subdivision also defines regions where the `HETERO` mechanism defined in (Gujar and Narahari 2011) has linear rebates.

We ran our heuristic technique to find the optimal mechanism for allocating 2 heterogeneous items to 4 agents, using rebates linear in $P_W^{(2,3)}$. Using this subdivision, our lower bound LP produced a mechanism with a worst-case ratio of $\frac{1}{4}$, which is the best achievable for the allocation of 2 *homogeneous* items among 4 agents (Moulin 2009; Guo and Conitzer 2009).

Conclusions

Groves mechanisms are a powerful tool in mechanism design as they are the class of efficient, dominant-strategy mechanisms. While there has been much work on finding optimal Groves mechanisms, it usually focused on single-parameter domains. In contrast, multi-parameter domains are not well understood, neither in terms of characterization of implementable mechanisms, nor in terms of the extant results on optimal payment functions. Our focus on Groves mechanisms stems from the fact that no other characterization exists for multi-parameter domains. Given this, our work is the first step in optimization of general dominant-strategy mechanisms for multi-parameter domains.

We presented a technique for finding optimal Groves mechanisms in multi-parameter domains. This built upon results on consistent partitions in single-parameter domains (Naroditskiy, Polukarov, and Jennings 2012). Consistent partitions of the agents' value space and the rebate space provide optimal rebates. However, consistent partitions may be hard or impossible to find. To this end, we introduced a heuristic technique to find the best rebate function when consistent partitions do not exist or are not known. This procedure consists of a lower bound and an upper bound linear programs. The lower bound LP produces optimal piecewise linear rebates for the given rebate space partitioning, which may not necessarily be globally optimal. However, the achievable global optimal objective value can be upper bounded using the upper bound LP. To demonstrate our procedure, we applied it to the problem of allocating heterogeneous items to agents with unit demand while maximizing social welfare.

One concern may be that the `LinearRebates` and `RestrictedProblem` linear programs can be quite computationally intensive to solve. Indeed, we observed that optimal piecewise linear functions require an exponential number of regions, making the complexity unavoidable if we are after an optimal solution. However, this computation needs to be performed only once to derive rebate functions. The output of the computation is an analytical solution: a rebate function that can be written down. This rebate function can then be used for any reports of agent values without any additional computation.

In future work, this technique could also be used to look for optimal rebates in combinatorial auctions. Moreover, social welfare is just one objective that can be maximized using the approach we presented. Alternatively, our approach can be applied to optimize fairness. Apart from the search for optimal solutions, the technique can be used to search for simple solutions with good performance. In particular, we can investigate the tradeoff between rebate function complexity and solution quality by varying the complexity of the partitioning of the rebate space when finding heuristic solutions. Since solving the linear programs can be computationally difficult, future work can experimentally evaluate the time required to find the rebate function. Finally, our technique can be used to optimize rebates in non-Groves mechanisms. Specifically, in single-parameter domains, one can apply the technique to non-efficient mechanisms that are dominant-strategy implementable.

Acknowledgments

The authors gratefully acknowledge funding from the UK Research Council for project ORCHID (`www.orchid.ac.uk`), grant EP/I011587/1.

References

Apt, K. R.; Conitzer, V.; Guo, M.; and Markakis, E. 2008. Welfare undominated groves mechanisms. 426–437.

Bailey, M. 1997. The demand revealing process: To distribute the surplus. *Public Choice* 91(2):107–26.

Cavallo, R. 2006. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS*, 882–889.

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

Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *UAI*, 103–110.

De Loera, J.; Rambau, J.; and Santos, F. 2010. *Triangulations*. Springer.

Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.

Gujar, S., and Narahari, Y. 2011. Redistribution mechanisms for assignment of heterogeneous objects. *Journal of Artificial Intelligence Research* 41:131–154.

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

Guo, M., and Conitzer, V. 2010a. Computationally feasible automated mechanism design: General approach and case studies. In *AAAI*, 1676–1679.

Guo, M., and Conitzer, V. 2010b. Optimal-in-expectation redistribution mechanisms. *Artif. Intell.* 174(5-6):363–381.

Guo, M. 2011. VCG redistribution with gross substitutes. In *AAAI*, 675–680.

Guo, M. 2012. Worst-case optimal redistribution of VCG payments in heterogeneous-item auctions with unit demand. In *AAMAS*.

Moulin, H. 2009. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory* 144(1):96–119.

Naroditskiy, V.; Polukarov, M.; and Jennings, N. 2012. Optimization of payments in dominant strategy mechanisms for single-parameter domains. *ACM Transactions on Economics and Computation (to appear)*.

Nisan, N. 2007. Introduction to mechanism design (for computer scientists). In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. 209–241.

Porter, R.; Shoham, Y.; and Tennenholtz, M. 2004. Fair imposition. *Journal of Economic Theory* 118(2):209 – 228.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16(1):8–37.