

Destroy to Save

Geoffroy de Clippel
Brown University
Department of Economics
Providence, Rhode Island
declippel@brown.edu

Victor Naroditskiy
Brown University
Department of Computer
Science
Providence, Rhode Island
vnarodit@cs.brown.edu

Amy Greenwald
Brown University
Department of Computer
Science
Providence, Rhode Island
amy@brown.edu

ABSTRACT

We study the problem of how to allocate m identical items among $n > m$ agents, assuming each agent desires exactly one item and has a private value for consuming it. We assume the items are jointly owned by the agents, not by one uninformed center, so an auction cannot be used to solve our problem. Instead, the agents who receive items compensate those who do not.

This problem has been studied by others recently, and their solutions have modified the classic VCG mechanism. This approach guarantees strategy-proofness and allocative efficiency. Further, in an auction setting, VCG guarantees budget balance, because payments are absorbed by the auctioneer. In our setting, however, where payments are redistributed to the agents, some money must be burned in order to retain strategy-proofness.

While strategy-proofness is necessary for truthful implementation, allocative efficiency (allocating the m items to those that desire them most), is not always an appropriate goal in our setting. Rather, we contend that maximizing social surplus is. In service of this goal, we study a class of mechanisms that may burn not only money but destroy items as well. Our key finding is that destroying items can save money, and hence lead to greater social surplus.

More specifically, our first observation is that a mechanism is strategy-proof iff it admits a threshold representation. Given this observation, we restrict attention to specific threshold and payment functions for which we can numerically solve for an optimal mechanism. Whereas the worst-case ratio of the realized social surplus to the maximum possible is close to 1 when $m = 1$ and 0 when $m = n - 1$ under the VCG mechanism, the best mechanism we find coincides with VCG when $m = 1$ but has a ratio approaching 1 when $m = n - 1$ as n increases.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Algorithms

Keywords

Mechanism design, VCG, payment redistribution, inefficient allocation

1. INTRODUCTION

Suppose six roommates jointly own a car that seats five people. They decide to take a weekend trip to the countryside. While they all would like to go, there is not room for all of them. Some really need the fresh country air while others would not mind staying home. The roommates don't necessarily know one another's desires, but each of them knows her own true value of getting out of the city. How should they decide who gets to go?

More generally, we study a class of resource allocation problems, in which n agents commonly own $m < n$ identical items that they wish to distribute among themselves, assuming each agent wants exactly one item, and has a *value* for that item which is known to her alone. As further examples, one could think of the allocation of free tickets for a sport event among club members, or seats on an overbooked plane. This kind of problem is often discussed in the literature on social choice where the goal is fairness, often under the assumption that the agents' values are commonly known. Yet this assumption is rarely satisfied, and self-interested agents will misreport their private values if doing so would be profitable. This is why our primary focus is incentives, instead of fairness.

Significant progress has been made in the field of mechanism design on the general topic of incentives since the seventies. Most research efforts have been devoted to understanding what is achievable in the presence of informational constraints (e.g., revelation principles in mechanism design). Professors Hurwicz, Maskin and Myerson received the 2007 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel for their groundbreaking contributions in this direction. Much less is known, however, about how to select a mechanism that is socially optimal, among those that are incentive compatible. In other words, the extension of social choice theory to problems characterized

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

by asymmetric information remains an important avenue of further exploration.

In this work we will focus on *strategy-proof* mechanisms, which require that it is a dominant strategy for each agent to report her value truthfully. This requirement is less permissive, but more robust than Bayesian implementation. In particular, agents are more likely to play a dominant strategy than a strategy that is optimal only when other agents play their part of the truthful equilibrium. Perhaps even more importantly, dominant-strategy implementation does not require any assumptions about the distribution of the agents' values (or their beliefs), nor their attitude towards risk.

As for measuring the appeal of various strategy-proof mechanisms, we will apply a worst-case measure. More specifically, if one fixes the agents' profile of values, one can compute the ratio of the total social surplus realized by the mechanism over the maximal total social surplus that could be achieved, should these values be publicly known. Since these values are not known, nor is their probabilistic distribution, the appeal of a strategy-proof mechanism will be measured by the minimum of this ratio over all possible value profiles. We call this ratio a *social surplus index*, and we use it to determine a mechanism's worst-case (i.e., guaranteed) level of social surplus: reaching a level $\alpha \in [0, 1]$ means that a mechanism realizes at least a proportion α of the maximal social surplus for *every* possible profile of the agents' values.

We assume that the agents can make monetary payments, and further that they have *quasi-linear* utilities. Our problem is different from those studied in auction theory. There, any monetary payments go to the auctioneer, who is usually assumed to have no private information. The presence of such "residual claimants" makes it easy to achieve budget balance. In our problem, however, the objective is to redistribute as much as possible of the payments among the agents themselves without negatively impacting the incentives (agents may have an incentive to misreport their values to receive higher compensation).

In addition to strategy-proofness, we also impose the following natural constraints 1) *feasibility*: no more than m items can be allocated, and monetary deficits are not allowed (i.e., no external subsidy), 2) *individual rationality*: each agent's total utility is nonnegative, and 3) *anonymity*: the allocation and payment decision applied to each agent does not depend on her identity. The question we are interested in can now be stated formally:

Find a mechanism that maximizes the worst-case social surplus index among all those that are strategy-proof, feasible, individually rational, and anonymous.

Recently, two sets of authors (Moulin [11] and Guo/Conitzer [6]) solved the above question under the additional assumption that the items be allocated to the m agents who value them most, at each possible profile of values. Both their solutions (derived independently) involved a class of mechanisms called VCG¹ mechanisms, which has received special attention in the economic literature because they admit a simple functional form (cf. Green and Laffont's [5] characterization). Using a VCG mechanism guarantees an efficient

¹VCG stands for Vickrey, Clarke, and Groves, who independently defined and studied some of these mechanisms in various contexts.

allocation of the m items available, but not necessarily a good level of overall efficiency (as measured for instance by the worst-case social surplus index), because allocative efficiency may come at the cost of "burning" quite a bit of money to meet the incentive constraints (when $m \geq 2$). So it may be better, in terms of overall efficiency, to destroy some items in order to save money. Indeed it is. It is not difficult to check that it is impossible to guarantee a strictly positive ratio using a VCG mechanism when $m = n - 1$. On the other hand, applying the best VCG mechanism after destroying one item would secure a strictly positive ratio.

Still, applying a VCG mechanism after destroying some fixed number of items is not the best strategy for optimizing overall efficiency. As a first step towards solving the general question, we offer a characterization of all strategy-proof mechanisms in terms of *threshold mechanisms*: i.e., an agent receives an item if and only if her reported value is larger than a threshold value *that may depend on other agents' reports*.² Although we don't believe that this result has been stated explicitly in papers discussing the very same model as ours, it is reminiscent of previous characterizations of VCG mechanisms (see, e.g., Green and Laffont [5]) and other strategy-proof mechanisms in related models ([4]).

Though helpful in understanding the question, this characterization result does not immediately allow us to solve it, because the feasibility and individual rationality constraints are nontrivial. The allocation function of the VCG mechanism is constant once the agents' values are ordered decreasingly (the agents with the m highest values receive the items). In that context Guo and Conitzer [6] and Moulin [11] manage to find the optimal payment function. The allocation function of the more general mechanisms that we study is not constant making the search for the optimal mechanism more difficult. We haven't found yet a way of solving the general question. But we have managed to identify specific classes of threshold and compensation functions that allow us to partition the set of value profiles into regions for which the resulting constrained optimization problem is linear in values. We then solve this LP problem numerically.

Our approach is designed to achieve the right balance between tractability, and showing that one can obtain a significant improvement of overall efficiency if one does not rely on the technical convenience of VCG mechanisms. Perhaps most striking is the case where $m = n - 1$. As already pointed out, VCG mechanisms cannot guarantee any strictly positive ratio in this case. Further, applying the best VCG mechanism after destroying a fixed number of items does not guarantee a ratio larger than $1/2$ (see numerical computations in Guo and Conitzer [7]). Our method of "contingent destruction" will identify a mechanism that guarantees a ratio $1 - \frac{2}{n^2 - n}$, which rapidly approaches 1 as n increases.

We conclude this introduction by discussing some related literature. Enhancing VCG mechanisms with payment redistribution has been studied in various settings. Bailey [1] proposes a way to redistribute some of the VCG tax in a public good domain. Cavallo [2] designs a redistribution mechanism for single-item allocation problems, and provides a characterization of redistribution mechanisms for more general allocation problems. As already mentioned, Guo and Conitzer [6] and Moulin [11] independently discover the

²This is a key distinction between our work and [7], where destroying the same number of items regardless of reported values was considered.

optimal VCG redistribution mechanism for the allocation domain studied here. In [8], Guo and Conitzer derive a linear redistribution VCG mechanism to maximize the expected social surplus when the distribution of agents' values is known. Porter et al. [12] study the problem of allocating undesirable goods (e.g., tasks) to agents in a fair manner.

Most related to our paper is Guo and Conitzer [7]. Starting from the same observation as ours that applying a VCG mechanism after destroying a fixed number of items may increase the worst-case social surplus index, they study mechanisms where the number of items destroyed may be a random variable. Introducing lotteries implies that one must take into account the agents' attitude towards risk. Guo and Conitzer's analysis requires the agents to be risk neutral. Also, the feasibility and individual rationality constraints hold only in expectation. Perhaps most importantly, the lottery that determines how many items to destroy does not vary with the players' reports. The key insight we offer in the present paper is that one can improve upon the optimal VCG redistribution mechanism without using lotteries, if one applies *contingent destruction* rules. If one is willing to use lotteries, then it may be of interest to combine the insights from our two papers, making Guo and Conitzer's random variables vary with reported values.

Other directions have also been followed when allowing for lotteries. Faltings [3], for instance, proposes a mechanism that picks an agent at random, and makes him the recipient of the VCG payments. The mechanism, which applies to domains more general than our allocation domain, achieves budget balance. However, if one applies this mechanism to our allocation domain, one sees that the resulting allocation is not efficient (unless the chosen recipient happens to value the item less than those who are allocated an item).

This paper unfolds as follows. Section 2 formally states the problem we are studying. We characterize strategy-proof mechanisms for the allocation domain in Section 3. A computational method of searching for an optimal mechanism in a restricted setting is proposed in Section 4. Numerical results in this setting are presented in Section 5.

2. DEFINITIONS

An *allocation problem* is a triple $\langle n, m, v \rangle$, where n is the number of agents, $m < n$ is the number of (identical) items available to allocate, and $v \in \mathbb{R}_+^n$ represents the agents' satisfaction from consuming one item (agents do not care for consuming multiple units). Monetary compensations are possible, and utilities are quasi-linear. An *allocation* is a pair $(a, t) \in \{0, 1\}^n \times \mathbb{R}^n$, where $a_i = 1$ if and only if agent i gets one item, and t_i represents the amount of money that agent i receives (this number can be negative, of course, in which case agent i pays that amount). Hence the total utility of agent i when implementing the allocation (a, t) is $a_i v_i + t_i$, if her value for the item is v_i . A *mechanism* is a pair of functions $f : \mathbb{R}_+^n \rightarrow \{0, 1\}^n$ and $t : \mathbb{R}_+^n \rightarrow \mathbb{R}^n$. Thus it determines an allocation for each possible report from the agents regarding their value for the item. We focus on mechanisms that satisfy the following constraints:

- *Feasibility*: no more than m items should be allocated, and the sum of payments to the agents should be less than or equal to zero, for all value vectors v . In other words,

$$\sum_{i=1}^n f_i(v) \leq m \text{ and } \sum_{i=1}^n t_i(v) \leq 0 \quad \forall v \in \mathbb{R}_+^n$$

- *Strategy-proofness*: It is a dominant strategy for each agent to report her value truthfully. Formally, $\forall v \in \mathbb{R}_+^n, i, v'_i$

$$f_i(v_i, v_{-i})v_i + t_i(v_i, v_{-i}) \geq f_i(v'_i, v_{-i})v_i + t_i(v'_i, v_{-i})$$

- *Individual Rationality*: It is in each agent's interest to participate in the mechanism, for all value vectors v , i.e.

$$f_i(v)v_i + t_i(v) \geq 0 \quad \forall v \in \mathbb{R}_+^n, i$$

We now define the index that we will use to measure the overall efficiency of a mechanism (f, t) that is implemented truthfully (an equivalent index was used in [11, 6, 7]). If the true value vector is v , then the (utilitarian) surplus realized by the mechanism is equal to $\sum_{i=1}^n [v_i f_i(v) + t_i(v)]$. This absolute number is less interesting than knowing how far it is from the first-best solution, i.e. the maximal surplus one could achieve if the agents' values were known. In order to have an index that is unit-free (i.e. homogenous of degree zero), it is natural to consider a ratio. Finally, since the agents' values are not known, nor their probabilistic distribution, it is natural to consider the worst-case index. To summarize, the index that we will use to measure the performance of a mechanism (f, t) that is truthfully implemented is given by the following number:

$$\min_{v \in \mathbb{R}_+^n} \frac{\sum_{i=1}^n [f_i(v)v_i + t_i(v)]}{\max_{a \in \mathcal{F}(m)} \sum_{i=1}^n a_i v_i},$$

where $\mathcal{F}(m) = \{a \in \{0, 1\}^n \mid \sum_{i=1}^n a_i \leq m\}$. Finding a mechanism whose index is α means that a proportion α of the maximal total surplus is achieved, *independently* of what the true values are.

The formal content of the question stated in the Introduction can thus be summarized by the following optimization problem:

$$\begin{aligned} \max_{(f,t)} \quad & \min_{v \in \mathbb{R}_+^n} \frac{\sum_{i=1}^n [f_i(v)v_i + t_i(v)]}{\max_{a \in \mathcal{F}(m)} \sum_{i=1}^n a_i v_i} \quad (1) \\ \sum_{i=1}^n f_i(v) \leq m \quad & \forall v \in \mathbb{R}_+^n \quad (2) \\ \sum_{i=1}^n t_i(v) \leq 0 \quad & \forall v \in \mathbb{R}_+^n \quad (3) \\ f_i(v_i, v_{-i})v_i + t_i(v_i, v_{-i}) \geq & f_i(v'_i, v_{-i})v_i + t_i(v'_i, v_{-i}) \quad (4) \\ & \forall v \in \mathbb{R}_+^n, i, v'_i \\ f_i(v)v_i + t_i(v) \geq 0 \quad & \forall v \in \mathbb{R}_+^n, i \quad (5) \end{aligned}$$

3. CHARACTERIZATION OF STRATEGY-PROOFNESS: THE THRESHOLD MECHANISMS

The allocation domain places strong restrictions on value functions of the agents. Specifically, an agent's value is zero in all outcomes where the agent is not allocated an item and the private value $v_i > 0$ in all outcomes where the agent is allocated an item. We use this restriction on the values to characterize strategy-proof mechanisms in the following proposition.³

PROPOSITION 1. *An allocation mechanism (f, t) is strategy-proof if and only if it is a "threshold mechanism," meaning that, for each $i = 1, \dots, n$, there exist a threshold function $\tau_i : \mathbb{R}_+^{n-1} \rightarrow \mathbb{R} \cup \{\infty\}$ and a compensation function $c_i : \mathbb{R}_+^{n-1} \rightarrow \mathbb{R}$ such that*

$$\begin{cases} f_i(v) = 1 \text{ and } t_i(v) = c_i(v_{-i}) - \tau_i(v_{-i}) & \text{if } v_i \geq \tau_i(v_{-i}) \\ f_i(v) = 0 \text{ and } t_i(v) = c_i(v_{-i}) & \text{if } v_i < \tau_i(v_{-i}), \end{cases}$$

or

$$\begin{cases} f_i(v) = 1 \text{ and } t_i(v) = c_i(v_{-i}) - \tau_i(v_{-i}) & \text{if } v_i > \tau_i(v_{-i}) \\ f_i(v) = 0 \text{ and } t_i(v) = c_i(v_{-i}) & \text{if } v_i \leq \tau_i(v_{-i}). \end{cases}$$

Remark The threshold mechanisms are easy to interpret. Each agent faces a personalized price (the threshold) that is determined by the reports of the other agents. She gets the good if and only if her reported value is (strictly) larger than this price, and must pay it in exchange. The collected money can be redistributed to some extent to the agents via the compensation function. The VCG mechanisms form a special case, where i 's threshold is the m^{th} largest component of v_{-i} .

PROOF. The sufficient condition is straightforward to check. So we provide an argument only for the necessary part. Fix i and the reports v_{-i} from the other agents. Strategy-proofness implies that

$$f_i(v) = f_i(v'_i, v_{-i}) \Rightarrow t_i(v) = t_i(v'_i, v_{-i}), \quad (6)$$

for all v_i, v'_i . It is easy to write (f, t) as a threshold mechanism if $f_i(v) = 0$, for all v_i , or $f_i(v) = 1$, for all v_i . Suppose thus that there exists v_i, v'_i such that $f_i(v) = 1$ and $f_i(v'_i, v_{-i}) = 0$. For any such pair, strategy-proofness implies that $v_i + t_i(v) \geq t_i(v'_i, v_{-i}) \geq v'_i + t_i(v)$. Hence $v_i \geq v'_i$. The space of agent's values $v_i \in \mathbb{R}_+$ can be partitioned in two intervals based on the mapping to either $f_i(v_i, v_{-i}) = 0$ or $f_i(v_i, v_{-i}) = 1$: there exists a threshold, denoted $\tau_i(v_{-i})$, such that $f_i(v) = 1$ if and only if $v_i \geq \tau_i(v_{-i})$ (or $f_i(v) = 1$ if and only if $v_i > \tau_i(v_{-i})$). Given (6), let $t_i^1(v_{-i})$ (resp. $t_i^0(v_{-i})$) be the payment made by i when she receives (resp. does not receive) the item. Strategy-proofness implies that $\tau(v_{-i}) + \epsilon + t_i^1(v_{-i}) \geq t_i^0(v_{-i}) \geq \tau(v_{-i}) - \epsilon + t_i^1(v_{-i})$, for each $\epsilon > 0$. Making ϵ tend to zero, we conclude that $t_i^1(v_{-i}) = t_i^0(v_{-i}) - \tau(v_{-i})$, and the result follows by taking $c = t_i^0$. \square

If we add anonymity to strategy-proofness in Proposition 1, the mechanism will change only in dropping indexes i from τ and c . For notational convenience from now on we will restrict our attention to generic profiles v where all

³We thank Yves Sprumont for pointing out this simple result to us.

components are distinct. This restriction is introduced without loss of generality as we can extend the mechanism to all value profiles (including vectors with equal components) by using uniform lotteries⁴ to break ties, as is usually done in papers on auctions.

This characterization of strategy-proofness is reminiscent of other well-known results for VCG and other more general strategy-proof mechanisms (see [10, 9]). A result similar to Proposition 1 appears in [4] in a slightly different context.

We restrict our attention to the first class of mechanisms identified in Proposition 1 (the one with $v_i \geq \tau_i(v_{-i})$) and restate the constrained optimization problem (1)-(4) using the threshold characterization:

$$\begin{aligned} \max_{(c, \tau)} \quad & \min_{v \in \mathbb{R}_+^n} \frac{\sum_{i|v_i \geq \tau(v_{-i})} (v_i - \tau(v_{-i})) + \sum_{i=1}^n c(v_{-i})}{\max_{a \in \mathcal{F}(m)} \sum_{i=1}^n a_i v_i} \\ \# \{i|v_i \geq \tau(v_{-i})\} & \leq m \quad \forall v \in \mathbb{R}_+^n \\ \sum_{i=1}^n c(v_{-i}) & \leq \sum_{i|v_i \geq \tau(v_{-i})} \tau(v_{-i}) \quad \forall v \in \mathbb{R}_+^n \\ c(v_{-i}) & \geq 0 \quad \forall v \in \mathbb{R}_+^n, i \end{aligned}$$

The first constraint is the feasibility constraint with respect to the items being allocated, while the second constraint is the feasibility constraint with respect to money (the sum of all compensations or rebates should be no more than the sum of the money collected from the agents that get an item). The third constraint is the individual rationality constraint (remember that an agent's value must be larger than the threshold when she gets an item, and so the IR constraint is trivially satisfied for her as well).

We now propose a last formulation of our optimization problem. We remove the minimization over v by introducing a variable $r \in \mathbb{R}$ denoting the best ratio that holds for any profile of values. The resulting optimization program is:

$$\begin{aligned} \max_{r, c, \tau} \quad & r \quad (7) \\ \sum_{i|v_i \geq \tau(v_{-i})} (v_i - \tau(v_{-i})) + \sum_{i=1}^n c(v_{-i}) & \geq r \max_{a \in \mathcal{F}(m)} \sum_{i=1}^n a_i v_i \quad \forall v \quad (8) \\ \# \{i|v_i \geq \tau(v_{-i})\} & \leq m \quad \forall v \in \mathbb{R}_+^n \quad (9) \\ \sum_{i=1}^n c(v_{-i}) & \leq \sum_{i|v_i \geq \tau(v_{-i})} \tau(v_{-i}) \quad \forall v \in \mathbb{R}_+^n \quad (10) \\ c(v_{-i}) & \geq 0 \quad \forall v \in \mathbb{R}_+^n, i \quad (11) \end{aligned}$$

⁴Suppose for instance that agent i should receive an item, and that more than m other agents have the same value as i . Anonymity would then come in conflict with feasibility. A uniform lottery will then be used to determine which subset of agents will receive the item, among all those that have the same value. Even so, the way agents react to risk is irrelevant because all the outcomes of the lottery are equivalent in terms of utility. Specifically, the lottery is between receiving the item worth v_i at the price p_i and receiving compensation c_i such that $c_i = v_i - p_i$.

A solution to the mathematical program above is the optimal mechanism for the allocation domain. However the program is hard to solve for different reasons. Firstly, maximization is over arbitrary functions c and τ , and there is little hope in optimizing over the space of arbitrary functions. Secondly, the program has an infinite number of constraints as the set of possible value vectors $v \in \mathbb{R}_+^n$ is infinite. We address these problems in the next section where we make assumptions about the form of the functions c and τ , show that it is sufficient to consider a finite number of constraints, and solve the resulting problem computationally.

4. A SIMPLER PROBLEM

There are infinitely many constraints in our optimization problem, since they are indexed by the value profiles v . Our main insight for dealing with this difficulty comes from the simple observation that linear constraints are satisfied by all the elements of a convex polytope if and only if they are satisfied by its extreme points. So we will restrict attention to the threshold and compensation functions that add some linearity to the general problem. The main difference with [11, 6] is that the number of items allocated, and thereby the constraints in the optimization problem, will vary with v when considering non-VCG mechanisms. So we will have to decompose the optimization problem into different regions of value profiles where the number of items allocated remains constant. To guarantee that these regions are convex (to be able to apply our insight), we will focus on threshold functions that are linear on regions partitioning \mathbb{R}_+^n . We chose a class of functions that is rich enough, but also that allows a simple characterization of these regions. Notice first that from now on we will restrict attention to value profiles v such that $v_1 \geq v_2 \geq \dots \geq v_n \geq 0$. This is without loss of generality since our problem involves only anonymous mechanisms. With this convention in mind, we are ready to define the class of threshold functions we will consider:

ASSUMPTION 1. *The threshold function is of the form $\tau(v_{-i}) = \max(kv_{-i}^p, v_{-i}^m)$, with $k \in [0, 1]$ and $p \in \{1, 2, \dots, (m-1)\}$.*

Adding v_{-i}^m -component in the max operator guarantees that no more than m items will be allocated, as required by the feasibility constraint. The p parameter determines how many items are guaranteed to be allocated, independently of the agents' reports. The k parameter controls how large lower values should be compared to larger values for the items $(p+1), \dots, m$ to be allocated. Taking $k = 0$ brings us back to VCG mechanisms, while setting of $k = 1$ means that items $p+1, \dots, m$ are always destroyed.

Similarly we will focus on compensation functions that are linear on convex subsets of the set of value profiles. Here too we face a trade-off between tractability and generality. A finer partition enlarges the set of functions being considered, but also increases the number of constraints in the linear program to be solved. The simplest choice would be to choose functions that are linear on the whole cone characterized by $v_1 \geq v_2 \geq \dots \geq v_n \geq 0$. Focusing on VCG mechanisms, [11, 6] proved that the optimal compensation function actually falls in that class. This is not necessarily the case anymore when considering non-VCG mechanisms, as we came to realize after running some simulations. Another simple choice would be to take compensation functions that are linear on the regions where the number of

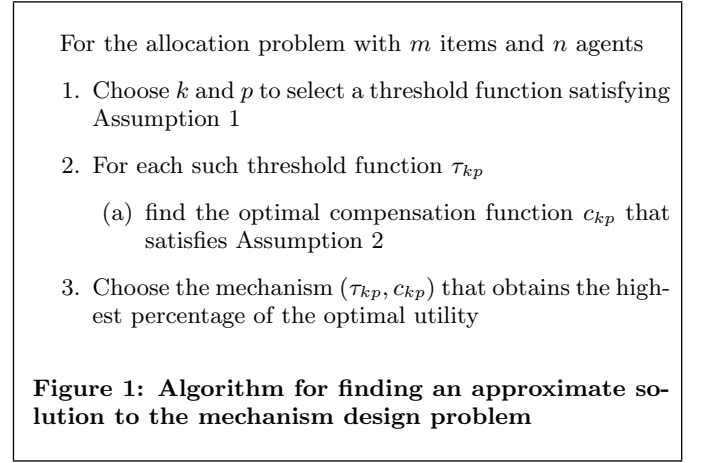
items being allocated is constant. This choice is not permitted, though, because these regions depend on the values of all the agents, while the compensation function can depend only on the values of the agents different from the one receiving the compensation. We decided to choose the closest match, imposing a condition that mimics the definition of these regions while using only the right values.

ASSUMPTION 2. *The compensation function c is linear in values on two regions:*

$$c(v_{-i}; a, b) = \begin{cases} av_{-i} & \text{if } kv_{-i}^p \geq v_{-i}^m \\ bv_{-i} & \text{otherwise} \end{cases}$$

where $a, b \in \mathbb{R}^{n-1}$.

The approach we will follow to find a mechanism that achieves a large index within the class of mechanisms that are feasible, strategy-proof, individually rational, and anonymous is summarized in Figure 1.



Our next goal is to characterize the regions where the compensation function is linear and the number of allocated items is constant. To do this, it is helpful to state explicitly the mechanisms satisfying Assumptions 1 and 2:

- $i \in \{1 \dots p\}$: $f_i = 1$
 $t_i = -\max(kv_{p+1}, v_{m+1}) + \begin{cases} av_{-i} & \text{if } kv_{p+1} \geq v_{m+1} \\ bv_{-i} & \text{otherwise} \end{cases}$
- $i \in \{(p+1) \dots m\}$
if $v_i \geq kv_p$: $f_i = 1$
 $t_i = -\max(kv_p, v_{m+1}) + \begin{cases} av_{-i} & \text{if } kv_p \geq v_{m+1} \\ bv_{-i} & \text{otherwise} \end{cases}$
otherwise: $f_i = 0, t_i = \begin{cases} av_{-i} & \text{if } kv_p \geq v_{m+1} \\ bv_{-i} & \text{otherwise} \end{cases}$
- $i \in \{(m+1) \dots n\}$: $f_i = 0, t_i = \begin{cases} av_{-i} & \text{if } kv_p \geq v_m \\ bv_{-i} & \text{otherwise} \end{cases}$

By the definition of the threshold function $\tau = \max(kv_{-i}^p, v_{-i}^m)$, there are $m-p+1$ possible allocations (the first p agents get the items, the first $p+1$ agents get the items, \dots , the first m agents get the items) determined by the position of kv_p

among $v_p \dots v_m$. The compensation function c is resolved to one of the two linear functions (av_{-i} or bv_{-i}) when the position of kv_p relative to v_m and v_{m+1} and the position of kv_{p+1} relative to v_{m+1} are determined. Each region below is defined to have a constant number of allocated items and a linear compensation function (i.e., resolved to either av_{-i} or bv_{-i}).

$$\forall j \in \{p \dots m\}, j' \in \{\max(p+1, j) \dots m\}$$

$$V_{j,j'} = \{v \in \mathbb{R}_+^n | v_1 \geq \dots \geq v_p \geq \dots \geq v_j \geq kv_p \geq v_{j+1} \geq \dots \geq v_{j'} \geq kv_{p+1} \geq v_{j'+1} \geq \dots \geq v_m \geq \dots \geq v_n\}$$

$$\forall j \in \{p \dots m\}$$

$$V_{j,m+1} = \{v \in \mathbb{R}_+^n | v_1 \geq \dots \geq v_p \geq \dots \geq v_j \geq kv_p \geq v_{j+1} \geq \dots \geq v_m \geq \dots \geq v_n \text{ AND } v_{m+1} \geq kv_{p+1}\}$$

$$V_{m+1,m+1} =$$

$$\{v \in \mathbb{R}_+^n | v_1 \geq \dots \geq v_n \text{ AND } v_{m+1} \geq kv_p \text{ AND } v_{m+1} \geq kv_{p+1}\}$$


The collection of regions above partitions the space $\{v \in \mathbb{R}_+^n | v_1 \geq v_2 \geq \dots \geq v_n \geq 0\}$. We group constraints by the regions and state the optimization problem in Figure 2.⁵ Notice that on each region the constraints are of the form $dv \geq 0$ for some $d \in \mathbb{R}^n$, which means that they are satisfied at λv ($\forall \lambda > 0$) as soon as they are satisfied at v . Hence we can assume without loss of generality that $v_1 = 1$ and focus on polytopes of vectors $(v_2, \dots, v_n) \in \mathbb{R}^{n-1}$ characterized by $(n+2)$ inequalities:

$$\begin{aligned} 1 &\geq v_2 \\ v_2 &\geq v_3 \\ &\dots \\ v_j &\geq kv_p \\ kv_p &\geq v_{j+1} \\ &\dots \\ v_{j'} &\geq kv_{p+1} \\ kv_{p+1} &\geq v_{j'+1} \\ &\dots \\ v_{n-1} &\geq v_n \\ v_n &\geq 0 \end{aligned}$$

Extreme points of these polytopes have the property that $n-1$ of these inequalities are binding. It is easy to check that this is possible only if the variables v_2, \dots, v_n take the values $1, k, k^2, 0$. Therefore all of the extreme points are of the form $(1, \dots, 1, k, \dots, k, k^2, \dots, k^2, 0, \dots, 0)$. Making sure the constraints hold on all such vectors guarantees that the constraints hold everywhere on $V_{j,j'}$. Now the linear program in Figure 2 can be stated with a finite number of constraints.

Example As an example consider the allocation problem with $n = 3$, $m = 2$ and the threshold function with $k = .5$, $p = 1$: $\tau = \max(.5v_{-1}^1, v_{-2}^2)$. The threshold function for agent 1 is $\max(.5v_2, v_3) < v_1$. So agent 1 is always allocated an item. The threshold for agent 2 is $\max(.5v_1, v_3)$. Agent 2 is allocated an item only when $v_2 > .5v_1$. Agent 3 is never allocated an item as the threshold for agent 3 is $\max(.5v_1, v_2) > v_3$.

⁵Recall, that $v_1 \geq v_2 \geq \dots \geq v_n \geq 0$ and therefore $\max_{a \in \mathcal{F}(m)} \sum_{i=1}^n a_i v_i = \sum_{i=1}^m v_i$.



The compensation function is linear when in addition to the allocation the position of $.5v_1$ and $.5v_2$ relative to v_3 is determined. Taking $v_1 = 1$ we can represent this on a 2-dimensional graph (Figure 3). The space is divided into 5 regions, with each region having a linear compensation function and a fixed allocation. To make sure the constraints hold for all $\{v \in \mathbb{R}_+^n | v_1 \geq v_2 \geq v_3\}$, we just need to enforce each region's constraints on its extreme points. For example, the extreme points of the right bottom region after adding $v_1 = 1$ as the first component are $(1, .5, 0)$, $(1, .5, .25)$, $(1, 1, .5)$, $(1, 1, 0)$.

5. RESULTS

We find mechanisms for different values of n and m using the computational approach described in Figure 1. The class of threshold functions we consider is given by all pairs (k, p) where k takes values in $\{0, .025, .05, \dots, .975\}$ and p in $\{1, 2, \dots, m-1\}$. We used CPLEX 11.2.0 as a linear program solver.

Figure 4 illustrates the results we generate for each setting of n , m , and p . The value for the parameter k is varied along the horizontal axis. For each value of k , the corresponding threshold function is $\tau = \max(kv_{-i}^p, v_{-i}^m)$, and we can solve the linear program from Figure 2 to find an optimal compensation function c_{kp} . The ratio for each mechanism (τ_{kp}, c_{kp}) is plotted for the corresponding k value. We refer to the resulting graph as the *performance curve*.




Figure 3: ($v_1 = 1$) Regions where the number of allocated items remains constant and the compensation function is linear for 3 agents and 2 items. Each region is labeled with the number of items allocated and the coefficients used in the compensation function for each agent, e.g. (1:b,a,a) means that 1 item is allocated and the compensation functions for agents 1,2, and 3 are bv_{-1} , av_{-2} , and av_{-3} respectively.




Figure 4: Performance curve

We scan the values of k for the one that has the highest ratio. In Figure 4, the best ratio is for $k = .20$. Notice that the shape of the curve suggests that there is only one peak. We try other values of k around .175 to find the peak at $k = \frac{1}{6}$. In all of our results we noticed that the performance curve as a function of k is single-peaked.

The threshold function with $k = 0$ corresponds to the efficient allocation function and the mechanism we find for $k = 0$ is the best VCG mechanism. The ratio of the best VCG mechanism appears at $k = 0$ and as argued before is zero when $n = m + 1$.

For any fixed values of n and m we found that a mechanism with p set to $m - 1$ achieves the highest ratio. This setting of p means that at most one item is destroyed. This result is consistent with the one obtained by Guo and Conitzer [7] for randomized VCG mechanisms. They find that the best mechanism randomizes between destroying one item and not destroying any items. The performance curves for different values of p are shown in Figure 5. Notice that the highest ratio is obtained on the graph for $p = m - 1 = 8$ ($k = .1$).

The mechanisms we find provide the most improvement




Figure 6: Performance of the mechanisms as a function of the number of items.

when the number of items is close to the number of agents. In the extreme case when $n = m + 1$ our mechanism achieves the ratio of $1 - \frac{2}{n^2 - n}$, while the VCG mechanisms have the ratio of 0. Our ratio becomes closer to the VCG ratio as the number of items becomes smaller and approximately around $m = \frac{n}{2}$ the ratios and the mechanisms coincide. Figure 6 shows this trend for 10 agents and varying number of items.

In our threshold algorithm the parameter p is set to $m - 1$ allocating at least $m - 1$ items. Also plotted are the ratios achieved by the best VCG mechanism as well as the ratio achieved by the mechanism that first destroys a fixed number of items and then applies an optimal VCG mechanism (see deterministic burning mechanism in [7]). All mechanisms coincide when the number of items is 4 or fewer.

We now illustrate the kind of mechanisms we find. Recall the problem where six roommates (i.e., agents) need to distribute five seats in the car (i.e., items). For six agents and five items our mechanism is given by the following parameters: $k = \frac{1}{6}$, $p = 4$, $a = b = (0, 0, 0, -\frac{1}{6}, 1)$. Under the mechanism, the first 4 roommates always go and each of them pays $\frac{1}{6}v_5$. Allocation and payment for agents 5 and 6 is determined as follows:

- if $v_5 \geq \frac{1}{6}v_4$
 - roommate 5 goes and pays $\frac{1}{6}v_4$
 - roommate 6 does not go and gets $(v_5 - \frac{1}{6}v_4)$
- if $v_5 < \frac{1}{6}v_4$
 - roommate 5 does not go and gets 0
 - roommate 6 does not go and gets 0

6. CONCLUSION

Finding the solution to the general optimization problem (1) remains an important open question. Taking a

⁶It is not difficult to check that this ratio is achieved by the following mechanism that meet the requirements of strategy-proofness, feasibility and anonymity: $\tau(v_{-i}) = \max(\frac{1}{n}v_{-i}^{m-1}, v_{-i}^m)$ and the coefficients $a = (0, \dots, 0)$ and $b = (0, \dots, -\frac{1}{n}, 1)$




Figure 5: Performance curves for different settings of p

step in that direction, we developed a practical methodology that improves upon previous contributions which were restricting attention to VCG mechanisms for technical convenience. Motivated by our characterization of strategy-proofness in terms of threshold and compensation functions (see Proposition 1), we imposed some restrictions on those functions which guarantee that the optimization problem can be solved via linear programming techniques. The key observation for this simplification is that linear inequalities hold at all points in a polytope if and only if they hold at its extreme points. Though it is possible that more intricate mechanisms would achieve an even greater social surplus, we observed that our approach already significantly improves upon the previous VCG analysis. The reason is that the combination of allocative efficiency, a characteristic feature of VCG mechanisms, and strategy-proofness may come at the cost of “burning” a lot of money. This insight is likely to prove helpful in other contexts as well.

The most striking illustration of the benefits of our approach in our problem is the allocation of $n - 1$ items among n agents. No redistribution of VCG payments is possible in that case, and for some value profiles the amount of VCG payments is as high as the sum of the $n - 1$ highest values. We find that destroying one item for some profiles of values significantly reduces the degree of payments. For example, the mechanism that destroys one item if the $(n - 1)^{\text{th}}$ highest value is less than $\frac{1}{n}$ of the $(n - 2)^{\text{th}}$ highest value guarantees that the amount of payment is less than $\frac{2}{n^2 - n}$ of the sum of $n - 1$ highest values.

At this time, we are experimenting with more flexible threshold and compensation functions in attempts to find a provably optimal mechanism. In the process, we are gathering intuition about the binding constraints for different threshold and compensation functions. In the future, we plan to investigate more general allocation settings characterized by allocation of non-identical items, agents desiring more than one item, agents with utilities that depend on whether other agents receive the items (externalities), and common-value models.

7. ACKNOWLEDGEMENTS

Insightful conversations with Hervé Moulin and Yves Sprumont are gratefully acknowledged.

8. REFERENCES

- [1] M. J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–26, April 1997.
- [2] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of vcg payments. In *Proc. of the 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'06)*, Hakodate, Japan, 2006.
- [3] B. Faltings. A Budget-balanced, Incentive-compatible Scheme for Social Choice. In P. Faratin and J. A. Rodriguez-Aguilar, editors, *Agent-Mediated Electronic Commerce VI*, volume 3435, pages 30–43. Springer, 2005.
- [4] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [5] J. R. Green and J.-J. Laffont. *Incentives in public decision-making*. North Holland, New York, 1979.
- [6] M. Guo and V. Conitzer. Worst-case optimal redistribution of vcg payments. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 30–39, New York, NY, USA, 2007. ACM.
- [7] M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 210–219, New York, NY, USA, 2008. ACM.
- [8] M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 1047–1054, Richland, SC, 2008.
- [9] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [10] H. Moulin. *Axioms of Cooperative Decision Making (Econometric Society Monographs)*. Cambridge University Press, July 1991.
- [11] H. Moulin. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory*, 2008.
- [12] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118(2):209 – 228, 2004.