Optimal-in-Expectation Redistribution Mechanisms

Mingyu Guo
Duke University
Department of Computer Science
Durham, NC, USA
mingyu@cs.duke.edu

Vincent Conitzer
Duke University
Department of Computer Science
Durham, NC, USA
conitzer@cs.duke.edu

ABSTRACT

Many important problems in multiagent systems involve the allocation of multiple resources to multiple agents. If agents are self-interested, they will lie about their valuations for the resources if they perceive this to be in their interest. The well-known VCG mechanism allocates the items efficiently, is incentive compatible (agents have no incentive to lie), and never runs a deficit. Nevertheless, the agents may have to make large payments to a party outside the system of agents, leading to decreased utility for the agents. Recent work has investigated the possibility of redistributing some of the payments back to the agents, without violating the other desirable properties of the VCG mechanism.

We study multi-unit auctions with unit demand, for which previously a mechanism has been found that maximizes the worst-case redistribution percentage. In contrast, we assume that a prior distribution over the agents' valuations is available, and try to maximize the expected total redistribution. We analytically solve for a mechanism that is optimal among *linear* redistribution mechanisms. The optimal linear mechanism is asymptotically optimal. We also propose *discretization* redistribution mechanisms. We show how to automatically solve for the optimal discretization redistribution mechanism for a given discretization step size, and show that the resulting mechanisms converge to optimality as the step size goes to zero. We also present experimental results showing that for auctions with many bidders, the optimal linear redistribution mechanism redistributes almost everything, whereas for auctions with few bidders, we can solve for the optimal discretization redistribution mechanism with a very small step size.

Categories and Subject Descriptors

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

General Terms

Algorithms, Economics, Theory

Keywords

Mechanism design, Vickrey-Clarke-Groves mechanism, payment redistribution, prior distributions over preferences

Cite as: Optimal-in-Expectation Redistribution Mechanisms, Mingyu Guo and Vincent Conitzer, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.

1. INTRODUCTION

Many important problems in multiagent systems can be seen as resource allocation problems. In such an allocation problem, there are one or more items that must be allocated to the agents. We assume that each agent has a privately held *valuation* function that indicates how much she values the items. Moreover, we assume that agents are *self-interested*: an agent will reveal her true valuation function only if doing so maximizes her utility. An allocation mechanism (or *auction*) takes as input the agents' reported valuations, and as output produces an allocation of the items to the agents, as well as payments to be made by or to the agents. A mechanism is *incentive compatible* if it is a dominant strategy for the agents to report their true valuations—that is, regardless of what the other agents do, an agent is best off reporting her true valuation. A mechanism is *efficient* if it always chooses an allocation that maximizes the sum of the agents' valuations.

The well-known *VCG (Vickrey-Clarke-Groves)* mechanism [20, 4, 10] is both incentive compatible and efficient.[1] In fact, in sufficiently general settings, the wider but closely related class of Groves mechanisms coincides exactly with the class of mechanisms that satisfy both properties [9, 14]. The VCG mechanism has an additional nice property, which is that it satisfies the *non-deficit* property: the sum of the payments from the agents is nonnegative, which means that the mechanism does not need to be subsidized by an outside party. A stronger property than the non-deficit property is that of *(strong) budget balance*, which requires that the sum of the payments from the agents is zero—so that no value flows out of the system of agents. To maximize social welfare (taking payments into account), we would prefer a budget balanced mechanism to one that merely achieves the non-deficit property (assuming both are efficient). Unfortunately, it is impossible to achieve budget balance together with incentive compatibility and efficiency [15, 9, 17]. [2] Previous research has sacrificed either incentive compatibility or efficiency to achieve budget balance [8, 18, 7]. Another

[1]We use the term "VCG mechanism" to refer to the Clarke mechanism. Sometimes people refer to the wider class of Groves mechanisms as "VCG mechanisms," but we will avoid this usage in this paper. In fact, the mechanisms proposed in this paper fall within the class of Groves mechanisms.

[2]The dAGVA mechanism [6] is efficient, (strongly) budget balanced, and *Bayes-Nash* incentive compatible, which means that if each agent's belief over the other agents' valuations is the distribution that results from conditioning the (common) prior distribution over valuations on the agent's own valuation, and other agents bid truthfully, then the agent is best off (in expectation) bidding truthfully. In practice, it is somewhat unreasonable to assume that agents' beliefs are so consistent with each other and with the mechanism designer's belief, so we use the much stronger and more common notion of dominant-strategies incentive compatibility in this paper.

approach is to allocate the items according to the VCG mechanism, and then to redistribute as much of the total VCG payment as possible back to the agents, in a way that does not affect the desirable properties of the VCG mechanism. Several papers have pursued this idea and proposed some natural redistribution mechanisms [1, 19, 2]. For example, in the Bailey mechanism [1], each agent receives a redistribution payment that equals $1/n$ times the VCG revenue that would result if this agent were removed from the auction. In the Cavallo mechanism [2], each agent receives a redistribution payment that equals $1/n$ times the minimal VCG revenue that can be obtained by changing this agent's own bid. For revenue monotonic settings, Bailey's and Cavallo's mechanisms coincide; in this case we refer to this mechanism as the Bailey-Cavallo mechanism. More recently, there has been some research on finding *optimal* redistribution mechanisms. For the case of a multi-unit auction with unit demand (that is, each agent wants at most one of the indistinguishable units), a mechanism that maximizes the worst-case redistribution percentage has been characterized [11, 16]. In this paper, we continue the search for optimal redistribution mechanisms. Unlike the worst-case work, we assume that a prior distribution over the agents' valuations is available, and we aim to maximize the *expected* total redistribution. (There are two related papers [13, 3], in which the authors propose mechanisms that maximize the sum of the agents' utilities (taking payments into account) in expectation. However, these papers operate under the constraint that every agent's total payment must be nonnegative, which results in very different mechanisms.)

The rest of this paper is layed out as follows. In Section 2, we cover the necessary background. In Section 3, we define linear redistribution mechanisms, and solve for the optimal linear redistribution mechanism. In Section 4, we show how to automatically (using linear programming) solve for mechanisms that are close to optimal based on a discretization of the valuation space. In Section 5, we compare the linear and discretization mechanisms experimentally.

2. BACKGROUND

We will focus on multi-unit auctions with unit demand in this paper. In a multi-unit auction, multiple indistinguishable units of the same good are for sale. In a multi-unit auction with unit demand, each agent wishes to obtain at most one unit—that is, if the agent receives more than one unit, her utility is the same as if she receives one unit. We note that an (unrestricted) single-item auction is a special case of multi-unit auctions with unit demand.

In this setting, each agent has a privately held true value for receiving (at least) one unit. If an agent wins one unit, her utility is her true value minus her payment; otherwise, her utility is the negative of her payment. In a *(sealed-bid) mechanism*, every agent reports her value (her *bid*), and the mechanism determines which agents win a unit, as well as how much each agent pays, as a function of these bids. A mechanism is *(dominant-strategies) incentive compatible* if it is a dominant strategy for each agent to bid her true valuation—that is, bidding truthfully is optimal regardless of what the other agents bid. Since we only study incentive compatible mechanisms in this paper, we do not need to make a clear distinction in our notation between the true values and the bids.

We assume that we know the number of agents n and the number of indistinguishable units m. If $m \geq n$, then we can give every agent a unit without charging any payments. Thus, we only consider the case $m < n$. Let the set of agents be $I = \{1, \ldots, n\}$, where agent i has the i-th highest value v_i. Let constants L and U be the lower bound and upper bound of the possible values. Hence, $\infty > U \geq v_1 \geq v_2 \geq \ldots \geq v_n \geq L \geq 0$. We also assume that we have a prior joint probability distribution over the agents' values v_i. We denote the probability density function of this joint distribution by $f(v_1, \ldots, v_n)$. We emphasize that we require neither that the agents' values are drawn from identical distributions, nor that they are independent. However, for the special case where agents' values are independently drawn from the same distribution $g(x)$ ($U \geq x \geq L$), we know from the theory of order statistics that $f(v_1, \ldots, v_n) = n! g(v_1) g(v_2) \ldots g(v_n)$ for all $U \geq v_1 \geq v_2 \geq \ldots \geq v_n \geq L$. If the agents' values are not drawn independently or are not drawn from the same distribution, then we do not always have an elegant analytical form for the joint distribution f. However, we will see later that optimal-in-expectation linear redistribution mechanisms depend only on the expectations of $v_1, \ldots, v_n$, which can usually be obtained by sampling.

In a multi-unit auction with unit demand, the VCG mechanism coincides with the $(m+1)$-th price auction. In this auction, the bidders with the highest m bids (bidders $1, \ldots, m$) each win one unit, and each pay at the price of the $(m+1)$-th bid (v_{m+1}). (When $m = 1$, this is the well-known second-price auction.) Because it is a special case of the VCG mechanism, the $(m+1)$-th price auction is incentive compatible, efficient, and never incurs a deficit.

A redistribution mechanism works as follows: after collecting a vector of bids $v_1 \geq v_2 \geq \ldots \geq v_n$, we first run the VCG mechanism ($(m+1)$-th price auction). The resulting allocation is efficient (agents $1 \ldots m$ each win a unit). However, because each winner has to pay v_{m+1}, a total VCG payment of mv_{m+1} leaves the system of agents. In order to achieve higher social welfare (taking payments into account), we try to redistribute a large portion of the total VCG payment back to the bidders, while maintaining the desirable properties of the VCG mechanism. Let r_i be the redistribution received by bidder i. To maintain incentive compatibility, r_i must be independent of bidder i's own bid v_i. (It is not difficult to see that this is sufficient for maintaining incentive compatibility: if an agent cannot affect her own redistribution payment, then she may as well ignore it when she determines her strategy; hence, the incentives for bidding are the same as in the VCG mechanism, which is incentive compatible. In general, because our allocation is efficient, the requirement that r_i does not depend on v_i is also necessary for incentive compatibility [9, 14].) Hence, we can write i's redistribution as $r_i(v_{-i})$ (sometimes short for r_i), where v_{-i} is the multiset of bids other than v_i; these functions r_i determine the redistribution mechanism. In this paper, unless otherwise specified, we consider only *anonymous* redistribution mechanisms, for which $r_i(\cdot) = r_j(\cdot) = r(\cdot)$ for all i, j. That is, the redistribution *function* is the same for all agents. This may still result in different redistribution payments for the agents, because the input to the function, v_{-i}, can be different for different i.

Another property of the VCG mechanism that we want to maintain is the *non-deficit* property: the payments collected by the mechanism are at least the payments redistributed by it. This is crucial if no external subsidy for the mechanism is available.[3] In our setting, this means that $\sum_{i=1}^{n} r_i(v_{-i}) \leq mv_{m+1}$.

3. LINEAR REDISTRIBUTION MECHANISMS

We first restrict our attention to the family of *linear* redistribution mechanisms. A linear redistribution mechanism is characterized by a linear redistribution function of the following form:

$$r_i(v_{-i}) = c_0 + c_1 v_{-i,1} + c_2 v_{-i,2} + \ldots + c_{n-1} v_{-i,n-1}$$

[3]Without the non-deficit constraint, we can simply redistribute $1/n$ of the expected total VCG payment to every agent, which leaves no waste in expectation.

where $v_{-i,j}$ is the j-th highest bid among v_{-i} (the set of bids other than v_i). The coefficients c_j completely characterize the redistribution mechanism. All previously proposed redistribution mechanisms for this setting [2, 1, 19, 11, 16] are in fact linear redistribution mechanisms.

3.1 Optimal-in-expectation linear redistribution mechanisms

We will prove the following result, which characterizes a linear redistribution mechanism that maximizes the expected total redistribution (among linear redistribution mechanisms). We call this mechanism OEL (optimal-in-expectation, linear).

THEOREM 1. *Given m, n, and a prior distribution over agents' valuations, the following c_i define a redistribution mechanism that maximizes expected redistribution, under the constraints that the mechanism must be a linear redistribution mechanism, efficient, incentive compatible, and satisfy the non-deficit property.*

Let the o_i be defined as follows:

$$o_0 = U - Ev_1,\ o_i = Ev_i - Ev_{i+1}, \text{ and } o_n = Ev_n - L.$$

The o_i are determined by the given prior distribution.

Let k be any integer satisfying

$$k \in argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i - m\ odd, i = 0, \dots, n\}$$

Let function G and H be defined as follows:

$$G(n,m,i) = \binom{n-i-1}{n-m-1}/\binom{m-1}{i-1}\ (i \le m)$$
$$H(n,m,i) = \binom{i-1}{m-1}/\binom{n-m-1}{n-i-1}\ (i \ge m)$$

- *If $0 < k \le m$, then*
  $c_i = (-1)^{m-i} G(n,m,i)$ *for* $i = k+1, \dots, m$,
  $c_k = m/n - \sum_{i=k+1}^{m} (-1)^{m-i} G(n,m,i)$,
  and $c_i = 0$ *for other* i.

- *If $k = 0$, then*
  $c_i = (-1)^{m-i} G(n,m,i)$ *for* $i = 1, \dots, m$,
  $c_0 = Um/n - U\sum_{i=1}^{m} (-1)^{m-i} G(n,m,i)$,
  and $c_i = 0$ *for other* i.

- *If $m + 1 \le k < n$, then*
  $c_i = (-1)^{m-i-1} H(n,m,i)$ *for* $i = m+1, \dots, k-1$,
  $c_k = m/n - \sum_{i=m+1}^{k-1} (-1)^{m-i-1} H(n,m,i)$,
  and $c_i = 0$ *for other* i.

- *If $k = n$, then*
  $c_i = (-1)^{m-i-1} H(n,m,i)$ *for* $i = m+1, \dots, n-1$,
  $c_0 = Lm/n - L\sum_{i=m+1}^{n-1} (-1)^{m-i-1} H(n,m,i)$,
  and $c_i = 0$ *for other* i.

In expectation, this mechanism fails to redistribute

$$o_k m\binom{n-1}{m}/\binom{n}{k}$$

This mechanism is uniquely optimal among all linear redistribution mechanisms if and only if the choice of k is unique and there does not exist an even i and an odd j such that $o_i = o_j = 0$.

The mechanism is complicated, and is perhaps easier to understand using the auxiliary variables that we define in the derivation of this mechanism (in the next subsection).

The key property of the mechanisms in the theorem is that the waste is always a multiple of: 1) the difference between two adjacent (in terms of size) bids, or 2) the difference between the upper bound and the largest bid, or 3) the difference between the lowest bid and the lower bound. Moreover, the multiplication coefficient is determined by m and n. Then, the OEL mechanism simply chooses the best of these options. In contrast, under the worst-case optimal mechanism [11, 16], the waste is a linear combination of all of the bids (except for the highest m).

The following special case and example should give some further intuition.

The case where $k = m + 1$ in Theorem 1 corresponds to the redistribution mechanism in which each agent receives a redistribution payment that is equal to m/n times the $(m+1)$-th highest bid from the other agents. This is exactly the Bailey-Cavallo mechanism in our setting (multi-unit auctions with unit demand).

Example 1. Consider the case where $n = 8$ and $m = 2$, and the bids are all drawn independently and uniformly from $[0, 1]$. In this case, $Ev_i = \frac{9-i}{9}$ for $i = 1, \dots, 8$. So $U = 1, L = 0, o_i = \frac{1}{9}$ for $i = 0, \dots, 8$. (We recall that $o_0 = U - Ev_1, o_n = Ev_n - L$, and $o_i = Ev_i - Ev_{i+1}$ otherwise.) $argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i - m\ odd, i = 0, \dots, n\}$ is then $\{3, 5\}$. The expected amount failed to be redistributed is $o_3 m\binom{n-1}{m}/\binom{n}{3} = \frac{1}{12}$. (The expected total VCG payment is $\frac{4}{3}$.)

One optimal solution is given by $c_3 = \frac{1}{4}$, and $c_i = 0$ for other i. Hence this expectation optimal linear redistribution mechanism is defined by $r_i = \frac{1}{4}v_{-i,3}$, which is actually the Bailey-Cavallo mechanism[1, 2]. The total redistribution is $\sum_{i=0}^{n} r_i = \frac{5}{4}v_3 + \frac{3}{4}v_4$. The expected amount failed to be redistributed is $E(2v_3 - \frac{5}{4}v_3 - \frac{3}{4}v_4) = \frac{3}{4}E(v_3 - v_4) = \frac{1}{12}$.

The other optimal solution is given by $c_3 = \frac{2}{5}, c_4 = -\frac{3}{10}, c_5 = \frac{3}{20}$, and $c_i = 0$ for other i. Hence this expectation optimal linear redistribution mechanism is defined by $r_i = \frac{2}{5}v_{-i,3} - \frac{3}{10}v_{-i,4} + \frac{3}{20}v_{-i,5}$. The total redistribution is $\sum_{i=0}^{n} r_i = 2v_3 - \frac{3}{4}v_5 + \frac{3}{4}v_6$. The expected amount failed to be redistributed is $E(\frac{3}{4}(v_5 - v_6)) = \frac{3}{4}E(v_5 - v_6) = \frac{1}{12}$.

3.2 Deriving an optimal linear redistribution mechanism

In this subsection, we derive the OEL mechanism and prove its optimality. Our objective is to find an linear redistribution mechanism that redistributes the most in expectation. To optimize among the family of linear redistribution mechanisms, we must solve for the optimal values of the c_i. We want the resulting redistribution mechanism to be incentive compatible and efficient, and we want it to satisfy the non-deficit property. The first two properties are satisfied by all the mechanisms inside the linear family, so the only constraint is the non-deficit property. The following optimization model can be used to find the linear redistribution mechanism (the c_i) that redistributes the most in expectation, while satisfying the non-deficit property.

Variables: $c_0, c_1, \dots, c_{n-1}$
Maximize $E(\sum_{i=1}^{n} r_i)$
Subject to:
For every bid vector $U \ge v_1 \ge v_2 \ge \dots \ge v_n \ge L$
$\sum_{i=1}^{n} r_i \le m v_{m+1}$
$r_i = c_0 + c_1 v_{-i,1} + c_2 v_{-i,2} + \dots + c_{n-1} v_{-i,n-1}$

Given the prior distribution, $E(mv_{m+1})$ is a constant, so we may rewrite the objective of the above model as

Minimize $E(mv_{m+1} - \sum_{i=1}^{n} r_i)$

Since $r_i = c_0 + c_1 v_{-i,1} + c_2 v_{-i,2} + \ldots + c_{n-1} v_{-i,n-1}$, where $v_{-i,j}$ is the j-th highest bid among bids other than i's own bid, we have the following:

$r_1 = c_0 + c_1 v_2 + c_2 v_3 + c_3 v_4 \ldots + c_{n-2} v_{n-1} + c_{n-1} v_n$
$r_2 = c_0 + c_1 v_1 + c_2 v_3 + c_3 v_4 \ldots + c_{n-2} v_{n-1} + c_{n-1} v_n$
$r_3 = c_0 + c_1 v_1 + c_2 v_2 + c_3 v_4 \ldots + c_{n-2} v_{n-1} + c_{n-1} v_n$
$\ldots$
$r_{n-1} = c_0 + c_1 v_1 + c_2 v_2 + c_3 v_3 \ldots + c_{n-2} v_{n-2} + c_{n-1} v_n$
$r_n = c_0 + c_1 v_1 + c_2 v_2 + c_3 v_3 \ldots + c_{n-2} v_{n-2} + c_{n-1} v_{n-1}$

We can write $mv_{m+1} - \sum_{i=1}^{n} r_i$ as $q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n$, where the coefficients q_i are listed below:

$$q_0 = -nc_0$$
$$q_i = -(i-1)c_{i-1} - (n-i)c_i \text{ for } i = 1, 2, \ldots, m, m+2, \ldots, n$$
$$q_{m+1} = m - mc_m - (n-m-1)c_{m+1}$$

(We note that we introduced a dummy variable c_n in the above equations—since there are only $n-1$ other bids, c_n will always be multiplied by 0, but adding this variable makes the definition of the q_i more elegant.) Given n and m, $q_0, \ldots, q_n$ ($n+1$ values) are determined by $c_0, \ldots, c_{n-1}$ (n values). Conversely, if $q_0, \ldots, q_{n-1}$ are fixed, then we can completely solve for the values of $c_0, \ldots, c_{n-1}$ (and hence also for q_n). This results in the following relation among the q_i:

$$q_1 - \frac{n-1}{1!} q_2 + \frac{(n-1)(n-2)}{2!} q_3 - \frac{(n-1)(n-2)(n-3)}{3!} q_4 + \ldots + (-1)^{n-1} \frac{(n-1)(n-2)\ldots 2\cdot 1}{(n-1)!} q_n = (-1)^m m \frac{(n-1)(n-2)\ldots(n-m)}{m!}$$

After simplification we obtain:

$$\sum_{i=1}^{n} (-1)^{i-1} \binom{n-1}{i-1} q_i = (-1)^m m \binom{n-1}{m}$$

Now, we can use the q_i as the variables of the optimization model, since from them we will be able to infer the c_i. Because $mv_{m+1} - \sum_{i=1}^{n} r_i = q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n$, we can rewrite the non-deficit constraint by requiring that the latter summation is nonnegative. Also, the q_i must satisfy the previous inequality (otherwise there will be no corresponding c_i).

> **Variables:** $q_0, q_1, \ldots, q_n$
> **Minimize** $E(q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n)$
> **Subject to:**
> For every bid vector $U \geq v_1 \geq v_2 \geq \ldots \geq v_n \geq L$
> $q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n \geq 0$
> $\sum_{i=1}^{n} (-1)^{i-1} \binom{n-1}{i-1} q_i = (-1)^m m \binom{n-1}{m}$

In what follows, we will cast the above model into a linear program. We begin with the following lemma[11]:

LEMMA 1. *The following are equivalent:*
(1) $q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n \geq 0$ *for all* $U \geq v_1 \geq v_2 \geq \ldots \geq v_n \geq L$
(2) $q_0 + L\sum_{i=1}^{n} q_i + (U-L)\sum_{i=1}^{k} q_i \geq 0$ *for* $k = 0, \ldots, n$

PROOF. (1)⇒(2): (2) can be obtained from (1) by setting $v_1 = v_2 = \ldots = v_k = U$ and $v_{k+1} = v_{k+2} = \ldots = v_n = L$.

(2)⇒(1): Let us rewrite $T = q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n$ as $q_0 + L\sum_{i=1}^{n} q_i + (v_1 - v_2)\sum_{i=1}^{1} q_i + (v_2 - v_3)\sum_{i=1}^{2} q_i + \ldots + (v_{n-1} - v_n)\sum_{i=1}^{n-1} q_i + (v_n - L)\sum_{i=1}^{n} q_i$. If $\sum_{i=1}^{k} q_i \geq 0$ for every $k = 1, \ldots, n$, then $T \geq q_0 + L\sum_{i=1}^{n} q_i \geq 0$ (because $v_1 - v_2, v_2 - v_3, \ldots, v_n - L$ are all nonnegative). Otherwise, let k' be the index so that $\sum_{i=1}^{k'} q_i$ is minimal (hence negative). To make T minimal, we want $v_{k'} - v_{k'+1}$ (which is multiplied by $\sum_{i=1}^{k'} q_i$) to be maximal. So the minimal value for T is $q_0 + L\sum_{i=1}^{n} q_i + (U-L)\sum_{i=1}^{k'} q_i \geq 0$, which is attained when $v_1 = v_2 = \ldots = v_{k'} = U$ and $v_{k'+1} = v_{k'+2} = \ldots = v_n = L$. Hence T is always nonnegative. □

Let $x_k = (q_0 + L\sum_{i=1}^{n} q_i)/(U-L) + \sum_{i=1}^{k} q_i$ for $k = 0, \ldots, n$. The x_i correspond (one to one) to the q_i, so we can use the x_i as the variables in the optimization model. The first constraint of the optimization model now becomes $x_k \geq 0$ for every k. Since $x_k - x_{k-1} = q_k$ for $k = 1, \ldots, n$, the second constraint becomes

$$\sum_{i=1}^{n} (-1)^{i-1} \binom{n-1}{i-1} (x_i - x_{i-1}) = (-1)^m m \binom{n-1}{m}$$

After simplification we get:

$$\sum_{i=0}^{n} (-1)^i \binom{n}{i} x_i = (-1)^{m-1} m \binom{n-1}{m}$$

Let $o_0 = U - Ev_1$, $o_i = Ev_i - Ev_{i+1}$ ($i = 1, \ldots, n-1$) and $o_n = Ev_n - L$. The o_i are all nonnegative constants that we know from the prior distribution. The objective of the optimization model can be rewritten as follows:

$E(q_0 + q_1 v_1 + q_2 v_2 + \ldots + q_n v_n)$
$= q_0 + q_1 Ev_1 + q_2 Ev_2 + \ldots + q_n Ev_n$
$= x_0(U-L) + q_1(Ev_1 - L) + q_2(Ev_2 - L) + \ldots + q_n(Ev_n - L)$
$= x_0((U-L) - (Ev_1 - L)) + (x_0 + q_1)((Ev_1 - L) - (Ev_2 - L)) + (x_0 + q_1 + q_2)((Ev_2 - L) - (Ev_3 - L)) + \ldots + (x_0 + q_1 + \ldots + q_n)(Ev_n - L)$
$= o_0 x_0 + o_1 x_1 + \ldots + o_n x_n$

We finally obtain the following linear program:

> **Variables:** $x_0, x_1, \ldots, x_n$
> **Minimize** $o_0 x_0 + o_1 x_1 + \ldots + o_n x_n$
> **Subject to:**
> $x_i \geq 0$
> $\sum_{i=0}^{n} (-1)^i \binom{n}{i} x_i = (-1)^{m-1} m \binom{n-1}{m}$

At this point, for any given n and m, for any prior distribution, it is possible to solve this linear program using any LP solver; then, using the above, the resulting x_i can be transformed back to c_i to obtain an optimal-in-expectation linear redistribution mechanism. However, this will not be necessary. The following claim gives an analytical solution of this linear program.

CLAIM 1. *Let k be any integer satisfying*
$k \in argmin_i \{o_i m \binom{n-1}{m} / \binom{n}{i} | i - m\ odd, i = 0, \ldots, n\}$.

The above linear program has the following optimal solution: $x_k = m\binom{n-1}{m}/\binom{n}{k}$, *and* $x_i = 0$ *for* $i \neq k$.

The optimal objective value is $o_k m \binom{n-1}{m}/\binom{n}{k}$.

This solution is the unique optimal solution if and only if the choice of k is unique and there does not exist an even i and an odd j such that $o_i = o_j = 0$.

PROOF. We can rewrite the second constraint as

$$\sum_{i=0}^{n} ((-1)^{i-m+1} \binom{n}{i}) / (m \binom{n-1}{m}) x_i = 1$$

This results in the program

> **Variables:** $x_0, x_1, \ldots, x_n$
> **Minimize** $o_0 x_0 + o_1 x_1 + \ldots + o_n x_n$
> **Subject to:**
> $x_i \geq 0$
> $\sum_{i=0 \ldots n; i-m\ odd} \binom{n}{i} / (m \binom{n-1}{m}) x_i = \sum_{i=0 \ldots n; i-m\ even} \binom{n}{i} / (m \binom{n-1}{m}) x_i + 1$

The o_i are nonnegative. To minimize the objective, we want all the x_i to be as small as possible. It is not hard to see that it does not hurt to set the x_i for which $i-m$ is even to zero: in fact, setting them to a larger value will only force the x_i for which $i-m$ is odd to take on larger values, by the last constraint. (It should be noted that if there exists an even i and an odd j such that $o_i = o_j = 0$, then we can increase the corresponding x_i and x_j at no cost to the objective without breaking the constraint, hence the solution is not unique in that case.) This results in the following linear program:

Variables: $x_0, x_1, \ldots, x_n$
Minimize $o_0x_0 + o_1x_1 + \ldots + o_nx_n$
Subject to:
$x_i \geq 0$
$\sum_{i=0 \ldots n; i-m \; odd} \binom{n}{i}/(m\binom{n-1}{m})x_i = 1$

We want the x_i to be as small as possible. However, the second constraint makes it impossible to set all the x_i to 0. For each x_i with $i-m$ odd, if we increase it by δ, the left side of the second constraint is increased by $\binom{n}{i}/(m\binom{n-1}{m})\delta$ and the objective value is increased by $o_i\delta$. We need the left side of the second constraint to increase to 1 (starting from 0), while minimizing the increase in the objective value. To do so, we want to find the x_i (with $i-m$ odd) that has the minimal cost-gain ratio (where the cost is $o_i\delta$, and the gain is $\binom{n}{i}/(m\binom{n-1}{m})\delta$). It follows that for any integer k satisfying $k \in argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i-m \; odd, i = 0, \ldots, n\}$, the linear program has the following optimal solution: $x_k = m\binom{n-1}{m}/\binom{n}{i}$ and $x_i = 0$ for $i \neq k$. The resulting optimal objective value is $o_k m\binom{n-1}{m}/\binom{n}{k}$.

In the above argument, there were only two conditions under which we made a choice that is not necessarily uniquely optimal: if (and only if) there exists an even i and an odd j such that $o_i = o_j = 0$, then, as we explained, there exist optimal solutions where some x_i with $m-i$ even is set to a positive value (in fact, it can be set to any value in this case); and if (and only if) $argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i-m \; odd, i = 0, \ldots, n\}$ is not a singleton set, then there exists another optimal solution with another x_k set to a positive value (in fact, in this case, multiple x_k may simultaneously be set to a positive value). □

By transforming the x_i from Claim 1 to the corresponding c_i, we obtain the OEL mechanism from Theorem 1.

3.3 Properties of the OEL mechanism

In the remainder of this section, we prove some properties of the OEL mechanism. First we have that there cannot be another redistribution mechanism that always redistributes at least as much to every agent as OEL. That is, the OEL mechanism is *undominated* [12]. (This does not immediately follow from Theorem 1, because that theorem only proved optimality among linear redistribution mechanisms, whereas this claim applies to all redistribution mechanisms.)

CLAIM 2. *For any m, n and any L, U, there does not exist any redistribution mechanism (other than OEL) that, for every multiset of bids, redistributes at least as much to every agent as OEL.*

PROOF. Omitted due to space constraints. Undominated redistribution mechanisms are characterized in [12]. □

It should be noted that Claim 2 only applies to the OEL mechanism, as defined in Theorem 1. Under certain circumstances (as detailed in Theorem 1), this mechanism is not uniquely optimal; and the other optimal mechanisms do not always have the property of Claim 2.

One property of mechanisms that we have not discussed so far is *individual rationality*: participating in the mechanism should not make agents worse off. The next claim shows that, if the prior distribution does not distinguish among agents, OEL is *ex-interim* individually rational—that is, in expectation, agents benefit from participating in the mechanism (they receive nonnegative expected utilities).

CLAIM 3. *If the prior distribution is symmetric across agents (for example, the agents' values are independent and identically distributed), then the OEL redistribution mechanism is ex-interim individually rational.*

PROOF. Omitted due to space constraints. □

As an aside, if the prior is not symmetric across agents, then we can explicitly add the ex-interim individual rationality constraint (or the stronger *ex-post* individual rationality constraint[4]) into our optimization model. This still results in a linear program. While it is possible to give a special-purpose algorithm for solving this linear program, it does not admit an elegant analytical solution.

In Theorem 1, we gave an expression for the expected amount that OEL fails to redistribute, which depended on the prior. In the next claim, we give an upper bound on this that does not depend on the prior.

CLAIM 4. *For any prior, the OEL mechanism fails to redistribute at most*

$$(U-L)m\binom{n-1}{m}/\sum_{i=0,1,\ldots,n; i-m \; odd} \binom{n}{i}$$

in expectation. This bound is tight.

PROOF. Given a prior distribution (and therefore, given the o_i), the expected amount failed to be redistributed is $o_k m\binom{n-1}{m}/\binom{n}{k}$ for any $k \in argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i-m \; odd, i = 0, \ldots, n\}$. If $o_i = (U-L)\binom{n}{i}/\sum_{i=0,\ldots,n; i-m \; odd}\binom{n}{i}$ for all i with $i-m$ odd, and $o_i = 0$ for all other i (this is in fact a feasible setting of the o_i), then $argmin_i\{o_i m\binom{n-1}{m}/\binom{n}{i} | i-m \; odd, i = 0, \ldots, n\} = \{i | 0 \leq i \leq n, i-m \; odd\}$. So k can be any i as long as $i-m$ is odd. In this case, the expected amount not redistributed is exactly $(U-L)m\binom{n-1}{m}/\sum_{i=0,\ldots,n; i-m \; odd}\binom{n}{i}$.

Now suppose that there is another distribution under which the mechanism fails to redistribute strictly more in expectation. Then, the new set of o'_i must satisfy $o'_i m\binom{n-1}{m}/\binom{n}{k} > m\binom{n-1}{m}/\sum_{i=0,\ldots,n; i-m \; odd}\binom{n}{i} = o_i m\binom{n-1}{m}/\binom{n}{k}$ for any i with $i-m$ odd. It follows that $o'_i > o_i$ for any i with $i-m$ odd. Since $\sum_{i=0,\ldots,n; i-m \; odd} o_i = U-L$ and $o'_i \geq 0$ for any i with $i-m$ even, we have $\sum_{i=0,\ldots,n} o'_i > U-L$, which is a contradiction. □

We note that as n goes to infinity (for fixed m), the expected amount that fails to be redistributed goes to 0; hence OEL is asymptotically optimal. For Example 1, Claim 4 gives an upper bound on the expected amount failed to be redistributed of 0.3281 (we recall that the actual amount is $1/12$, so the bound is not very close in this case).

So far, we have only considered anonymous redistribution mechanisms (that is, mechanisms with the same redistribution function

[4]A mechanism is ex-post individually rational if every agent receives nonnegative utility for *any* bids.

$r(\cdot)$ for each agent).[5] If we allow the redistribution mechanism to be nonanonymous, then we can use different c_i for different bidders. Moreover, even for the same bidder, we can use different c_i depending on the order of the other bidders (in terms of their bids), and there are $(n-1)!$ such orders. Thus, it is clear that to optimize among the class of nonanonymous linear redistribution mechanisms, we need significantly more variables, and analytical solution of the linear program no longer seems tractable. However, we do have the following claim, which shows that the OEL mechanism remains optimal even among nonanonymous linear redistribution mechanisms, if the prior is symmetric.

CLAIM 5. *If the prior distribution is symmetric across agents (for example, the agents' values are independent and identically distributed), then no nonanonymous linear redistribution mechanism can redistribute strictly more than the OEL mechanism (which is anonymous) in expectation.*

PROOF. Let us define the average of two (not necessarily anonymous) redistribution mechanisms as follows: for any multiset of bids, for any agent i, if one redistribution mechanism redistributes x to agent i, and the other redistribution mechanism redistributes y to i, then the average mechanism redistributes $(x+y)/2$ to i. It is not difficult to see that if two redistribution mechanisms both never incur a deficit, then the average of these two mechanisms also satisfies the non-deficit property. This averaging operation is easily generalized to averaging over three or more mechanisms.

Now let us assume that r is a nonanonymous linear redistribution mechanism, and that r redistributes strictly more than the OEL mechanism in expectation when the prior distribution is symmetric across agents. Let π be any permutation of n elements. We permute the way r treats the agents according to π, and denote the new mechanism by r^π. That is, r^π treats agent $\pi(i)$ the way r treats i. Since we assumed that the prior distribution is symmetric across agents, the expected total amount redistributed by r^π should be the same as that redistributed by r. Now, if we take the average of the r^π over all permutations π, we obtain an anonymous linear redistribution mechanism that redistributes as much in expectation as r (and hence more than the OEL mechanism). But this contradicts the optimality of the OEL mechanism among anonymous linear redistribution mechanisms. □

4. DISCRETIZATION REDISTRIBUTION MECHANISMS

In the previous section, we only considered linear redistribution mechanisms. This restriction allowed us to find the optimal linear redistribution mechanism by analytically solving a linear program. In this section, we consider a larger domain of eligible mechanisms, and propose *discretization redistribution mechanisms*, which can be automatically designed [5] and can outperform the OEL mechanism. (In this section, for simplicity and to be able to compare to the previous section, we only consider anonymous mechanisms, and we do not impose an individual rationality constraint. However, all of the below can be generalized to allow for nonanonymous mechanisms and an individual rationality constraint.)

We study the following problem: given a prior distribution f (the joint pdf of $v_1, v_2, \ldots, v_n$), we try to find a redistribution mechanism that redistributes the most in expectation among all redistribution mechanisms that can be characterized by continuous functions.

[5]An exception is Claim 2, which shows that there is not even a nonanonymous mechanism that always redistributes at least as much as OEL to every agent (besides OEL itself).

For simplicity, we will assume that f is continuous. The optimization model is the following:

Variable function: $r : \mathbf{R}^{n-1} \to \mathbf{R}$, r continuous
Maximize
$\int_{U \geq v_1 \geq \ldots \geq v_n \geq L} \sum_{i=1}^{n} r(v_{-i}) f(v_1, v_2, \ldots, v_n) dv_1 dv_2 \ldots dv_n$
Subject to:
For every bid vector $U \geq v_1 \geq v_2 \geq \ldots \geq v_n \geq L$
$\sum_{i=1}^{n} r(v_{-i}) \leq m v_{m+1}$

Let R^* be the optimal objective value for this model. (To be precise, we have not proved that an optimal solution exists for this model: it could be that the set of feasible solution values does not include its least upper bound. In this case, simply let R^* be the least upper bound.) Since we are not able to solve this model analytically, we try to solve it numerically.

We divide the interval $[L, U]$ (within which the bids lie) into N equal parts, with step size $h = (U-L)/N$. Let k denote the subinterval: $I(k) = [L + kh, L + kh + h]$ $(k = 0, 1, \ldots, N-1)$. Define $r^h : \mathbf{R}^{n-1} \to \mathbf{R}$ as follows: for all $U \geq x_1 \geq x_2 \geq \ldots \geq x_{n-1} \geq L$, $r^h(x_1, x_2, \ldots, x_{n-1}) = z^h[k_1, k_2, \ldots, k_{n-1}]$ where $k_i = \lfloor (x_i - L)/h \rfloor$ (except that $k_i = N-1$ if $x_i = U$). Here, the $z^h[k_1, k_2, \ldots, k_{n-1}]$ are variables. We call such a mechanism a discretization redistribution mechanism of step size h.

CLAIM 6. *A discretization redistribution mechanism satisfies the non-deficit constraint if and only if*

$$\sum_{i=1}^{n} z^h[k_1, k_2, \ldots, k_{i-1}, k_{i+1}, \ldots, k_n] \leq m(L + k_{m+1} h)$$
for every $N-1 \geq k_1 \geq k_2 \geq \ldots \geq k_n \geq 0$.

PROOF. Omitted due to space constraints. □

The following linear program finds the optimal discretization redistribution mechanism for step size h. The variables are $z^h[k_1, k_2, \ldots, k_{n-1}]$ for all integers k_i satisfying $N-1 \geq k_1 \geq k_2 \geq \ldots \geq k_{n-1} \geq 0$. The objective is the expected total redistribution, where $p[k_1, k_2, \ldots, k_n] = P(v_1 \in I(k_1), v_2 \in I(k_2), \ldots, v_n \in I(k_n))$ (we note that the $p[k_1, k_2, \ldots, k_n]$ are constants).

Variables: $z^h[\ldots]$
Maximize $\sum_{N-1 \geq k_1 \geq k_2 \geq \ldots \geq k_n \geq 0}$
$p[k_1, k_2, \ldots, k_n] \sum_{i=1}^{n} z[k_1, k_2, \ldots, k_{i-1}, k_{i+1}, \ldots, k_n]$
Subject to:
For every $N-1 \geq k_1 \geq k_2 \geq \ldots \geq k_n \geq 0$
$\sum_{i=1}^{n} z[k_1, k_2, \ldots, k_{i-1}, k_{i+1}, \ldots, k_n] \leq m(L + k_{m+1} h)$

Let $z^{*h}[\ldots]$ denote the optimal solution of the above linear program, and let r^{*h} denote the corresponding optimal discretization redistribution mechanism. Let R^{*h} denote the optimal objective value. The next claim shows that discretization redistribution mechanisms cannot outperform the best continuous redistribution mechanisms.

CLAIM 7. $R^{*h} \leq R^*$.

PROOF. For any $\epsilon > 0$, we will show how to construct a continuous function r'_ϵ so that $r'_\epsilon \leq r^{*h}$ everywhere, and the measure of the set $\{r^{*h} \neq r'_\epsilon\}$ is less than ϵ.

Let B be the greatest lower bound of r^{*h} (r^{*h} is bounded below because it is a piecewise constant function with finitely many pieces). For given $U \geq x_1 \geq x_2 \geq \ldots \geq x_{n-1} \geq L$, let $d(x_1, \ldots, x_{n-1})$ be the minimal distance from any $x_i - L$ to the nearest multiple of h. For any $\delta > 0$, let $r_\delta(x_1, \ldots, x_{n-1}) = r^{*h}(x_1, \ldots, x_{n-1})$ if $d(x_1, \ldots, x_{n-1}) > \delta$, and $r_\delta(x_1, \ldots, x_{n-1})$

$= r^{*h}(x_1, \ldots, x_{n-1}) - (\delta - d(x_1, \ldots, x_{n-1}))(r^{*h}(x_1, \ldots, x_{n-1}) - B)/\delta$ otherwise.

It is easy to see that the function r_δ is continuous at any point where $d(x_1, \ldots, x_{n-1}) > \delta$, because at these points, r^{*h} is continuous. Furthermore, the function is continuous at any point where $\delta > d(x_1, \ldots, x_{n-1}) > 0$, because r^{*h} and d are both continuous at these points. Moreover, it is also continuous at any point where $d(x_1, \ldots, x_{n-1}) = \delta$, because at such a point $r^{*h}(x_1, \ldots, x_{n-1}) - (\delta - d(x_1, \ldots, x_{n-1}))(r^{*h}(x_1, \ldots, x_{n-1}) - B)/\delta = r^{*h}(x_1, \ldots, x_{n-1})$. Finally, at any point where $d(x_1, \ldots, x_{n-1}) = 0$, the function is continuous because on any point $x'_1, \ldots, x'_{n-1}$ at distance at most $\gamma > 0$ from $x_1, \ldots, x_{n-1}$, the function will take value at most $\gamma(H - B)/\delta$, where H is an upper bound on r^{*h} (H is finite).

As δ goes to 0, so does the measure of the set $\{r^{*h} \neq r_\delta\}$. Moreover, $r_\delta \leq r^{*h}$ everywhere. Hence we can obtain r'_ϵ with the desired property by letting it equal r_δ for sufficiently small δ.

Now, r'_ϵ is a feasible redistribution mechanism, because it always redistributes less than r^{*h}. Moreover, because f is a continuous pdf on a compact domain, as $\epsilon \to 0$, the difference in expected value between r'_ϵ and r^{*h} goes to 0. Hence, we can create continuous redistribution functions that come arbitrarily close to R^{*h} in terms of expected redistribution, and hence R^* (the least upper bound of the expected redistributions that can be obtained with continuous functions) is at least R^{*h}. □

The next claim shows that if we make the discretization finer, we will do no worse.

CLAIM 8. $R^{*h} \leq R^{*h/2}$.

PROOF. Omitted due to space constraints. □

The next claim shows that as we make the discretization finer and finer, we converge to the optimal value for continuous redistribution mechanisms.

CLAIM 9. $\lim_{h \to 0} R^{*h} = R^*$.

PROOF. For any $\gamma > 0$, there exists a continuous redistribution mechanism r^* such that its expected redistribution is at least $R^* - \gamma$. r^* is continuous on a closed and bounded domain, so r^* is uniformly continuous. Hence for any $\epsilon > 0$, there exists $\delta > 0$ so that $|r^*(x_1, x_2, \ldots, x_{n-1}) - r^*(x'_1, x'_2, \ldots, x'_{n-1})| \leq \epsilon$ as long as $\max_i\{|x_i - x'_i|\} \leq \delta$. Choose $h \leq \delta$, and define $z^h[k_1, k_2, \ldots, k_{n-1}]$ by $r^*(L + k_1 h, L + k_2 h, \ldots, L + k_{n-1} h)$ for all $N - 1 \geq k_1 \geq k_2 \geq \ldots \geq k_{n-1} \geq 0$. $z^h[\ldots]$ corresponds to a feasible discretization mechanism r^h. In addition, $r^h \geq r^* - \epsilon$. Hence, the expected redistribution of the optimal discretization mechanism with step size (at most) h is $R^{*h} \geq R^h \geq R^* - \gamma - n\epsilon$. Since γ and ϵ are both arbitrarily small, $\lim_{h \to 0} R^{*h} \geq R^*$. By Claim 7, $\lim_{h \to 0} R^{*h} \leq R^*$. □

We note that a discretization redistribution mechanism r^h is defined by a finite number of real-valued variables: namely, one variable $z^h[k_1, k_2, \ldots, k_{n-1}]$ for every $N - 1 \geq k_1 \geq k_2 \geq \ldots \geq k_{n-1} \geq 0$. Because of this, we can use a standard LP solver to solve for the optimal discretization redistribution mechanism r^h (for given m, n, h and prior). At least for small problem instances, we can set h to very small values, and by Claim 9, we expect the resulting mechanism to be close to optimal.

But how do we know how far from optimal we are? As it turns out, the discretization method can also be used to find upper bounds on R^*. Here, we will assume that agents' values are independent and identically distributed. The following linear program gives an upper bound on R^*.

Variables: $z^h[\ldots]$
Maximize $\sum_{N-1 \geq k_1 \geq k_2 \geq \ldots \geq k_n \geq 0}$
$p[k_1, k_2, \ldots, k_n] \sum_{i=1}^{n} z[k_1, k_2, \ldots, k_{i-1}, k_{i+1}, \ldots, k_n]$
Subject to:
For every $N - 1 \geq k_1 \geq k_2 \geq \ldots \geq k_n \geq 0$
$\sum_{i=1}^{n} z[k_1, k_2, \ldots, k_{i-1}, k_{i+1}, \ldots, k_n] \leq$
$mE(v_{m+1} | v_1 \in I(k_1), v_2 \in I(k_2), \ldots, v_n \in I(k_n))$

The intuition behind this linear program is the following. In the previous linear program, the non-deficit constraints were effectively set for the *lowest* values within each discretized block, which guaranteed that they would hold for every value in the block. In this linear program, however, we set the non-deficit constraints by taking the *expectation* over the values in each block. Generally, this will result in deficits for values inside the block, so this program does not produce feasible mechanisms.

Let $\hat{z}^h[\ldots]$ denote the optimal solution of the above linear program, and let $\hat{r}^h$ denote the (not necessarily feasible) corresponding optimal discretization redistribution mechanism. Let $\hat{R}^h$ denote the optimal objective value. We have the following claims:

CLAIM 10. *If the bids are independent and identically distributed, then* $\hat{R}^h \geq R^*$.

PROOF. Omitted due to space constraints. □

CLAIM 11. *If the bids are independent and identically distributed, then* $\hat{R}^h \leq R^{*h} + mh$.

PROOF. Omitted due to space constraints. □

Hence, by solving the linear program for determining R^{*h}, we get a lower bound on R^* and a discretization redistribution mechanism that comes close to it. If we also have that the bids are independent and identically distributed, by solving the linear program for determining $\hat{R}^h$, we get an upper bound on R^* that is close to R^{*h}.

5. EXPERIMENTAL RESULTS

We now have two different types of redistribution mechanisms with which we can try to maximize the expected total redistributed. The OEL mechanism has the advantage that Theorem 1 gives a simple expression for it, so it is easy to scale to large auctions. In addition, it is optimal among all linear redistribution mechanisms, although nonlinear redistribution mechanisms may perform even better. On the other hand, the discretization mechanisms have the advantage that, as we decrease the step size h, we will converge to the maximum amount that can be redistributed by any continuous redistribution mechanism. The disadvantage of this approach is that it does not scale to large auctions. Fortunately, we will see that, as the auctions get larger, OEL redistributes almost the entire total VCG payment, so OEL is certainly very close to optimal. On the other hand, for smaller auctions, OEL is not that close to optimal, but for these auctions we are able to solve for the optimal discretization redistribution mechanism with very small step size, which we show is very close to optimal using the upper bounding technique. Thus, we can redistribute almost optimally for both small and large auctions.

In the following table, for different n (number of agents) and m (number of units), we list the expected amount of redistribution by both the OEL mechanism and the optimal discretization mechanism (for specific step sizes). The bids are independently drawn from the uniform $[0, 1]$ distribution.

n,m	VCG	OEL	R^{*h}	$\hat{R}^h$
3,1	0.5000	0.3333	0.4218 (N=100)	0.4269
4,1	0.6000	0.5000	0.5498 (N=40)	0.5625
5,1	0.6667	0.6000	0.6248 (N=25)	0.6452
6,1	0.7143	0.6667	0.6701 (N=15)	0.7040
3,2	0.5000	0.3333	0.4169 (N=100)	0.4269
4,2	0.8000	0.5000	0.6848 (N=40)	0.7103
5,2	1.0000	0.8000	0.8944 (N=25)	0.9355
6,2	1.1429	1.0000	1.0280 (N=15)	1.0978

In the above table, the column "VCG" gives the expected total VCG payment; the column "OEL" gives the expected redistribution by the OEL mechanism; the column "R^{*h}" gives the expected redistribution by the optimal discretization redistribution mechanism (step size $1/N$); the column "$\hat{R}^h$" gives the upper bound on the expected redistribution by any continuous redistribution mechanism (same step size as that of R^{*h}).

Finally, when the number of agents is large, the OEL mechanism is very close to optimal, as shown below:

n,m	VCG	OEL	%	n,m	VCG	OEL	%
10,1	0.8182	0.8143	99.5	20,1	0.9048	0.9048	100.0
10,3	1.9091	1.8000	94.3	20,5	3.5714	3.5564	99.6
10,5	2.2727	2.0000	88.0	20,10	4.7619	4.5000	94.5
10,7	1.9091	1.8000	94.3	20,15	3.5714	3.5564	99.6
10,9	0.8182	0.8143	99.5	20,19	0.9048	0.9048	100.0

The fourth and eighth columns give the percentages of the VCG payment that are redistributed by the OEL mechanisms (rounding to the nearest tenth).

6. CONCLUSION

The well-known VCG mechanism allocates the items efficiently, is incentive compatible (agents have no incentive to lie), and never runs a deficit. Nevertheless, the agents may have to make large payments to a party outside the system of agents, leading to decreased utility for the agents. Recent work has investigated the possibility of redistributing some of the payments back to the agents, without violating the other desirable properties of the VCG mechanism. We studied multi-unit auctions with unit demand, for which previously a mechanism had been found that maximizes the worst-case redistribution percentage. In contrast, in this paper, we assumed that a prior distribution over the agents' valuations is available, and tried to maximize the expected total redistribution.

We first considered *linear* redistribution mechanisms. We gave an analytical solution for a redistribution mechanism that, among linear redistribution mechanisms, maximizes the expected redistribution, and gave conditions under which it is unique. We also proved some other desirable properties of this mechanism—it is asymptotically optimal and undominated. We then proposed *discretization* redistribution mechanisms, which discretize the space of possible valuations, and determine redistributions solely based on the discretized values (however, the incentive compatibility and non-deficit constraints still hold over the non-discretized space). Given a discretization step size, we showed how to solve for the optimal discretization redistribution mechanism using a linear program. We also showed that as the step size goes to 0, the mechanism converges to the optimal value for all continuous mechanisms (and we proved a bound on how close to optimal we are). We presented experimental results showing that for auctions with many bidders, the optimal linear redistribution mechanism redistributes almost everything, whereas for auctions with few bidders, we can solve for the optimal discretization redistribution mechanism with a very small step size.

Future research on optimal-in-expectation redistribution mechanisms can take a number of directions. One can try to solve for an optimal-in-expectation redistribution mechanism that is not necessarily linear. One can also try to extend the results of this paper to more general settings, for example, settings without unit demand, or even combinatorial auctions. Finally, it would be interesting to see whether agents' expected welfare can be improved even further by allocating units inefficiently, and if so, by how much.

7. REFERENCES

[1] M. J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.

[2] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Japan, 2006.

[3] S. Chakravarty and T. Kaplan. Manna from heaven or forty years in the desert: Optimal allocation without transfer payments, October 2006. Working Paper.

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

[5] 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.

[6] C. d'Aspremont and L.-A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[7] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce (AMEC), LNAI, 3435*, pages 30–43, 2005.

[8] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001. Early version in STOC-00.

[9] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 1977.

[10] T. Groves. Incentives in teams. *Econometrica*, 1973.

[11] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 30–39, San Diego, CA, USA, 2007.

[12] M. Guo and V. Conitzer. Undominated VCG redistribution mechanisms. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Estoril, Portugal, 2008.

[13] J. Hartline and T. Roughgarden. Money burning and implementation, January 2007. Working Paper.

[14] B. Holmström. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[15] L. Hurwicz. On the existence of allocation systems whose manipulative Nash equilibria are Pareto optimal, 1975. Presented at the 3rd World Congress of the Econometric Society.

[16] H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment, March 2007. Working Paper.

[17] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[18] D. Parkes, J. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1161–1168, Seattle, WA, 2001.

[19] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118:209–228, 2004. Early version appeared in IJCAI-01.

[20] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.