

Optimal-in-Expectation Redistribution Mechanisms

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

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

Abstract

Many important problems in multiagent systems involve the allocation of multiple resources among the 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 strategy-proof (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.

Previous research on redistribution mechanisms has resulted in a worst-case optimal redistribution mechanism, that is, a mechanism that maximizes the fraction of VCG payments redistributed in the worst case. In contrast, in this paper, we assume that a prior distribution over the agents' valuations is available, and our goal is to maximize the expected total redistribution.

In the first part of this paper, we study multi-unit auctions with unit demand. We analytically solve for a mechanism that is optimal among *linear* redistribution mechanisms. We also propose *discretized* redistribution mechanisms. We show how to automatically solve for the optimal discretized 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 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 discretized redistribution mechanism with a very small step size.

In the second part of this paper, we study multi-unit auctions with nonincreasing marginal values. We extend the notion of linear redistribution mechanisms, previously defined only in the unit demand setting, to this more general setting. We introduce a linear program for finding the optimal linear redistribution mechanism. This linear program is unwieldy, so we also introduce one simplified linear program that produces relatively good linear redistribution mechanisms. We conjecture an analytical solution for the simplified linear program.

1 Introduction

Many important problems in multiagent systems can be seen as resource allocation problems. One natural way of allocating resources among agents is to auction off the items. An allocation mechanism (or *auction*) takes as input the agents' reported valuations for the items, and as output produces an allocation of the items to the agents, as well as payments to be made by or to the agents. We assume that agents are *self-interested*: an agent will reveal her true valuation function only if doing so maximizes her utility. A mechanism is *strategy-proof* 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 (aka. the *social welfare*).

The well-known *VCG* (*Vickrey-Clarke-Groves*) mechanism [24, 6, 13] is both strategy-proof 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 [12, 18]. The VCG mechanism has an additional nice property, which is that it satisfies the *non-deficit* property (in allocation settings): 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. This property is not satisfied by the VCG mechanism. In the context of auctions, often, this is not seen as a problem for the sake of maximizing the agents' welfare: the idea is that the payments are collected by the seller of the items, who is just another agent, so that nothing goes to waste. However, this reasoning does not apply to many multiagent settings; in particular, it does not apply to settings in which *there is no seller* who is separate from the agents. For example, consider the problem of dissolving a partnership: suppose there is a group of agents who have started a company together, but due to personal disagreements can no longer work together, so that it becomes essential to allocate the (currently jointly owned) company to just one of the agents. While it makes sense to auction off the company among the agents, ideally, the revenue of this auction is then distributed among the agents themselves—if the revenue leaves the system of the agents, their welfare is reduced. Similarly, the agents may be deciding how to allocate a resource that is not claimed by anyone—for example, the agents may have jointly discovered a valuable commodity (say, an oil field) in unclaimed territory, which they now need to allocate to the one of them that can make the best use of it. Finally, the agents may have a jointly owned resource (say, a powerful computer) that can only be used by one agent on any given day, and may wish to use an auction to determine which agent gets to use it today. In all these cases, any payment that is not redistributed to the agents truly goes to waste. Hence, 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

¹We 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.

achieve budget balance together with strategy-proofness and efficiency [19, 12, 21].² Incidentally, while these types of setting are perhaps not what one typically has in mind when considering “auctions” in the common sense of the word, the fact that we use auctions does not significantly limit the generality of our approach. Effectively, we just use “auctions” as a convenient word to describe resource allocation mechanisms that use payments.

Previous research has sacrificed either strategy-proofness or efficiency to achieve budget balance [11, 22, 10]. Another 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 [2, 23, 3]. For example, in the Bailey mechanism [2], 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 [3], 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 setting of multi-unit auctions with unit demand (that is, each agent wants at most one of the indistinguishable units)—the setting that we study in most of this paper—a mechanism that maximizes the worst-case redistribution fraction has been analytically characterized [16, 20] (one of these papers [16] also generalizes beyond the unit-demand case, to nonincreasing marginal values). 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 [17, 5], 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.) In this paper, we restrict ourselves to VCG redistribution mechanisms, so that the allocation is always efficient; other work has studied what can be done when this constraint is relaxed [10, 20, 14, 9] (all the resulting mechanisms are characterized analytically). We also restrict ourselves to static mechanisms; good redistribution mechanism has also been analytically characterized in a dynamic context [4].

The rest of this paper is presented as follows. From Section 2 to Section 5, we focus on multi-unit auctions with unit demand. In Section 2, we cover the necessary background and introduce our notation. In Section 3, we recall the definition of linear redistribution mechanisms and we solve for optimal-in-expectation linear (OEL)

²The dAGVA mechanism [8] 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 (strategy-proofness) in this paper.

redistribution mechanisms in our setting. We focus on deriving an analytical characterization of these OEL mechanisms. In Section 4, we show how to automatically (using linear programming) solve for (possibly nonlinear) mechanisms that are close to optimal, based on a discretization of the valuation space.³ This technique is only effective for cases with small number of agents. That is, it does not scale very well. Fortunately, the experimental results in Section 5 show that for auctions with many bidders, the optimal linear mechanism redistributes almost everything, whereas for auctions with few bidders, we can solve for the optimal discretized redistribution mechanism with a very small step size. That is, the two approaches are in some sense complementary. Finally, in Section 6, we study the more general setting of multi-unit auctions with nonincreasing marginal values. We extend the notion of linear redistribution mechanisms to this more general setting, and propose several models for finding optimal linear redistribution mechanisms. It is more difficult to work in this more general setting, since we also need to consider a type of ordering information; we discuss these difficulties in that section.

2 Background

From this section to Section 5, we focus on multi-unit auctions with unit demand.

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 *strategy-proof* 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 strategy-proof 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 $m < n$.⁴ Let the set of agents be $I = \{1, \dots, n\}$, 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, $\infty > U \geq v_1 \geq v_2 \geq \dots \geq v_n \geq L \geq 0$. 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

³This falls under the general research agenda of *automated mechanism design* [7], where we have an algorithm search through a space of possible mechanisms for an optimal one. However, here we use a formulation that is specifically tailored to this context. In fact, the linear programs elsewhere in this paper could in principle also be used for the purpose of automated mechanism design, but of course there is little purpose to doing so for the cases where we also provide an analytical solution.

⁴We remove this restriction in Section 6 where we consider settings without unit demand.

distribution by $f(v_1, \dots, v_n)$. We emphasize that we require neither that the agents' values are drawn from identical distributions, nor that they are independent.

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 $1, \dots, m$) 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 strategy-proof, efficient, and never incurs a deficit.

A redistribution mechanism works as follows: after collecting a vector of bids $v_1 \geq v_2 \geq \dots \geq v_n$, we first run the VCG mechanism ($(m + 1)$ th price auction). The resulting allocation is efficient (agents $1 \dots m$ 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 strategy-proofness, 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 strategy-proofness: 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 strategy-proof. In general, because our allocation is efficient, the requirement that r_i does not depend on v_i is also necessary for strategy-proofness [12, 18].) Hence, we can write i 's redistribution as $r(v_{-i})$, where v_{-i} is the multiset of bids other than v_i ; functions r determines i 's redistribution. In this paper, unless otherwise specified, we consider only *anonymous* redistribution mechanisms, for which the redistribution function is the same for all agents (denoted by r). 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.⁵ In our setting, this means that $\sum_{i=1}^n r(v_{-i}) \leq mv_{m+1}$.

Finally, one property of mechanisms that we have not discussed so far is *individual rationality* (aka. *voluntary participation*): participating in the mechanism should not make agents worse off. Since our objective is to maximize social welfare, if the prior distribution is symmetric across agents, then under any redistribution mechanism that redistributes a nonnegative amount of payment in expectation, every agent benefits from participating in the mechanism (the agent receives nonnegative expected utility). That is, *ex-interim* individual rationality is not a binding constraint. The techniques in this paper can also be used to design mechanisms that are *ex-interim* individually rational when the prior is not symmetric across agents, or mechanisms that satisfy the even stronger *ex-post* individual rationality. However, this would require additional constraints and make the analytical characterization in Section 3 too complex. For the above reasons, we omit individual rationality constraints in this paper.

⁵Without 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.

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:

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

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 [3, 2, 23, 16, 20] 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 n , m , 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, strategy-proof, and satisfy the non-deficit property.*

Let the o_i be defined as follows:

$$o_0 = U - Ev_1, \quad o_i = Ev_i - Ev_{i+1} \quad (i = 1, 2, \dots, n-1), \quad \text{and} \quad o_n = Ev_n - L.$$

The o_i are determined by the given prior distribution.

Let k be any integer satisfying

$$k \in \arg \min_i \{o_i / \binom{n}{i} \mid i - m \text{ odd}, i = 0, \dots, n\}$$

Let function G be defined as follows:

$$G(n, m, i) = \binom{n-i-1}{n-m-1} / \binom{m-1}{i-1} = \frac{(n-i-1)!(i-1)!}{(n-m-1)!(m-1)!}$$

- *If $0 < k < m$, then*

$$c_i = (-1)^{m-i} G(n, m, i) \text{ for } i = k+1, \dots, m,$$

$$c_k = m/n - \sum_{i=k+1}^m (-1)^{m-i} G(n, m, i),$$

and $c_i = 0$ for other i .

- *If $k = 0$, then*

$$c_i = (-1)^{m-i} G(n, m, i) \text{ for } i = 1, \dots, m,$$

$$c_0 = Um/n - U \sum_{i=1}^m (-1)^{m-i} G(n, m, i),$$

and $c_i = 0$ for other i .

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

The key property of the mechanisms in the theorem is that the waste is always a multiple of: 1) the expected difference between two adjacent (in terms of size) bids, or 2) the expected difference between the upper bound and the largest bid, or 3) the expected 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 [16, 20], the waste is a linear combination of all of the bids (except for the highest m).

We now present a special case that may 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. In our setting of multi-unit auctions with unit demand, this is exactly the Bailey-Cavallo mechanism. This observation is formally stated in the following corollary.

Corollary 1. *Given n , m , and a prior distribution over agents' valuations, we define the o_i as follows:*

$$o_0 = U - Ev_1, o_i = Ev_i - Ev_{i+1} \ (i = 1, 2, \dots, n - 1), \text{ and } o_n = Ev_n - L.$$

If the following condition holds:

$$o_{m+1} \leq o_i \binom{n}{m+1} / \binom{n}{i} \text{ for all } 0 \leq i \leq n \text{ with } i - m \text{ odd,}$$

then the Bailey-Cavallo mechanism maximizes expected redistribution, under the constraints that the mechanism must be a linear redistribution mechanism, efficient, strategy-proof, and satisfy the non-deficit property.

Next, we present two example OEL mechanisms.

Example 1. Consider the case where $n = 3$ and $m = 1$, and the bids are all drawn independently and uniformly from $[0, 1]$. In this case, $Ev_i = \frac{4-i}{4}$ for $i = 1, \dots, 3$. So, $U = 1, L = 0, o_i = \frac{1}{4}$ for $i = 0, \dots, 3$. (We recall that $o_0 = U - Ev_1, o_n = Ev_n - L$, and $o_i = Ev_i - Ev_{i+1}$ otherwise.) $\arg \min_i \{o_i / \binom{n}{i} | i - m \text{ odd}, i = 0, \dots, n\}$ is then $\{m+1\} = \{2\}$. The expected amount that fails to be redistributed is $o_2 m \binom{n-1}{m} / \binom{n}{2} = \frac{1}{6}$. (The expected total VCG payment is $\frac{1}{2}$.) The optimal solution is given by $c_2 = \frac{1}{3}$, and $c_i = 0$ for other i . Hence, this optimal-in-expectation linear redistribution mechanism is defined by $r_i = \frac{1}{3}v_{-i,2}$, which is actually the Bailey-Cavallo mechanism. The total redistribution is $\sum_{i=1}^n r_i = \frac{1}{3}v_2 + \frac{2}{3}v_3$. The expected amount that fails to be redistributed is $E(v_2 - \frac{1}{3}v_2 - \frac{2}{3}v_3) = \frac{2}{3}E(v_2 - v_3) = \frac{1}{6}$.

Example 2. 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 $i = 0, \dots, 8$. $\arg \min_i \{o_i / \binom{n}{i} | i - m \text{ odd}, i = 0, \dots, n\}$ is then $\{3, 5\}$. The expected amount that fails 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 $r_i = \frac{1}{4}v_{-i,3}$ (Bailey-Cavallo mechanism). The total redistribution is $\sum_{i=1}^n r_i = \frac{5}{4}v_3 + \frac{3}{4}v_4$. The expected amount that fails 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 $\sum_{i=1}^n r_i = 2v_3 - \frac{3}{4}v_5 + \frac{3}{4}v_6$. The expected amount that fails to be redistributed is $E(\frac{3}{4}(v_5 - v_6)) = \frac{1}{12}$.

3.2 Properties of the OEL mechanism

In the remainder of this section, we present some properties of the OEL mechanism. First, we have that there cannot be another redistribution mechanism that always redistributes at least as much *in total* as OEL, and strictly more in at least one case. That is, the OEL mechanism is *welfare undominated* [1].⁶

Proposition 1 ([1]). *For any m, n and any prior distribution, there does not exist any redistribution mechanism that, for every multiset of bids, redistributes at least as much in total as OEL, and redistributes strictly more in at least one case.*

The above proposition was shown in [1]. More precisely, that paper shows that the OEL mechanisms characterized in Theorem 1 are the only welfare undominated redistribution mechanisms that are anonymous and linear in multi-unit auctions with unit demand.

It should be noted that Proposition 1 only applies to the OEL mechanism, as defined in Theorem 1. Under certain circumstances (as detailed in Theorem 1), this mechanism

⁶This immediately implies that there cannot be another redistribution mechanism that always redistributes at least as much *for every agent* as OEL. That is, the OEL mechanism is also *undominated* [15].

is not uniquely optimal; and the other optimal mechanisms do not always have the property of Proposition 1.

The next proposition 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).

Proposition 2. *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. The original VCG mechanism (redistributing nothing) is also a linear redistribution mechanism (corresponding to $c_i = 0$ for all i). Hence, the OEL mechanism will always redistribute a nonnegative amount in expectation. That is, $E(\sum_{i=1}^n r_i) \geq 0$. If the distribution is symmetric across agents, $E(r_i) = E(r_j)$ for any i and j ($E(r_i)$ is the expected redistribution received by agent i , which is independent of her own report). So $E(r_i) \geq 0$ for all i . However, the VCG mechanism is well-known to be ex-interim (in fact, ex-post) individually rational in this setting, so that even if $E(r_i) = 0$, agents' expected utility from participating in the mechanism is nonnegative. It follows that OEL must also be ex-interim individually rational. \square

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 constraint⁷) into our optimization model. This still results in a linear program (but 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 proposition, we give an upper bound on this that does not depend on the prior.

Proposition 3. *For any prior, the OEL mechanism fails to redistribute at most*

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

in expectation. This bound is tight.

Proof. Given a prior distribution (and therefore, given the o_i), the expected amount that fails to be redistributed is $o_k m \binom{n-1}{m} / \binom{n}{k}$ for any $k \in \arg \min_i \{o_i / \binom{n}{i} | i - m \text{ odd}, i = 0, \dots, n\}$. If a distribution is constructed such that $o_i = (U - L) \binom{n}{i} / \sum_{i=0,\dots,n;i-m \text{ odd}} \binom{n}{i}$ 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 —we can just use a degenerate distribution where the agents' valuations are not independent), then $\arg \min_i \{o_i / \binom{n}{i} | i - m \text{ odd}, i = 0, \dots, n\} = \{i | 0 \leq i \leq n, i - m \text{ odd}\}$. So k can be any i as long as $i - m$ is odd. In this case, the expected amount not redistributed is exactly $(U - L)m \binom{n-1}{m} / \sum_{i=0,\dots,n;i-m \text{ odd}} \binom{n}{i}$.

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 $o'_i m \binom{n-1}{m} / \binom{n}{k} > m \binom{n-1}{m} / \sum_{i=0,\dots,n;i-m \text{ odd}} \binom{n}{i} = o_i m \binom{n-1}{m} / \binom{n}{k}$ for any i with $i - m$

⁷A mechanism is ex-post individually rational if every agent receives nonnegative utility for *any* bids.

odd. It follows that $o'_i > o_i$ for any i with $i - m$ odd. Since $\sum_{i=0, \dots, n; i-m \text{ odd}} o_i = U - L$ and $o'_i \geq 0$ for any i with $i - m$ even, we have $\sum_{i=0, \dots, n} o'_i > U - L$, which is a contradiction. \square

For Example 1, Proposition 3 gives an upper bound on the expected amount that fails to be redistributed of 0.5 (we recall that the actual amount is $\frac{1}{6}$). For Example 2, Proposition 3 gives an upper bound on the expected amount that fails to be redistributed of 0.3281 (we recall that the actual amount is $\frac{1}{12}$).

The next proposition shows that for fixed m , as n goes to infinity, the expected amount that fails to be redistributed goes to 0; hence OEL is asymptotically optimal for fixed number of units.

Proposition 4. *For fixed m , as n goes to infinity, the expected amount that fails to be redistributed by OEL goes to 0.*

Proof. By Proposition 3, we only need to show that for fixed m , as n goes to infinity, $(U - L)m \binom{n-1}{m} / \sum_{i=0, 1, \dots, n; i-m \text{ odd}} \binom{n}{i}$ goes to 0.

We have that $(U - L)m \binom{n-1}{m} / \sum_{i=0, 1, \dots, n; i-m \text{ odd}} \binom{n}{i} \leq (U - L)m \binom{n-1}{m} / \binom{n}{m+1} = (U - L) \frac{m(n-1)!(m+1)!(n-m-1)!}{m!(n-m-1)!n!} = (U - L)(m+1)m/n$. The right-hand side goes to 0 as n goes to infinity. \square

On the other hand, if we increase both n and m , and keep their difference within constant C , then the expected amount fails to be redistributed by OEL also goes to 0: for large n , the expected amount fails to be redistributed by OEL is at most

$$(U - L)m \binom{n-1}{m} / \sum_{i=0, 1, \dots, n; i-m \text{ odd}} \binom{n}{i} \leq (U - L)n(n-1)^{C-1} / \sum_{i=0, 1, \dots, n; i-m \text{ odd}} \binom{n}{i} \\ = (U - L)n(n-1)^{C-1} / \left(\sum_{i=1, 2, \dots, n-1; i-m \text{ odd}} \left(\binom{n-1}{i-1} + \binom{n-1}{i} \right) + \sum_{i=0, n; i-m \text{ odd}} \binom{n}{i} \right).$$

Basically, the denominator is exponential in n , while the numerator is polynomial in n . Therefore, as n increases, the amount fails to be redistributed by OEL approaches 0.

So far, we have only considered anonymous redistribution mechanisms (that is, mechanisms with the same redistribution function $r(\cdot)$ for each agent).⁸ 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 proposition, which shows that the OEL mechanism remains optimal even among nonanonymous linear redistribution mechanisms, if the prior is symmetric.

Proposition 5. *If the prior distribution is symmetric across agents (for example, the agents' values are independent and identically distributed), then no nonanonymous*

⁸An exception is Proposition 1, which shows that there is not even a nonanonymous mechanism that always redistributes at least as much in total as OEL, and strictly more in at least one case.

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. \square

4 Discretized 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 *discretized redistribution mechanisms*, which can be automatically designed [7] 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, \dots, 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. For simplicity, we will assume that f is continuous. The optimization model is the following:

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

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: $I(k) = [L + kh, L + kh + h]$ ($k = 0, 1, \dots, N - 1$). Define $r^h : \mathbf{R}^{n-1} \rightarrow \mathbf{R}$ as follows: for all $U \geq x_1 \geq x_2 \geq \dots \geq x_{n-1} \geq L$, $r^h(x_1, x_2, \dots, x_{n-1}) = z^h[k_1, k_2, \dots, k_{n-1}]$ where $k_i = \lfloor (x_i - L)/h \rfloor$ (except that $k_i = N - 1$ if $x_i = U$). Here, the $z^h[k_1, k_2, \dots, k_{n-1}]$ are variables. We call such a mechanism a discretized redistribution mechanism of step size h .

Proposition 6. *A discretized redistribution mechanism satisfies the non-deficit constraint if and only if*

$$\sum_{i=1}^n z^h[k_1, k_2, \dots, k_{i-1}, k_{i+1}, \dots, k_n] \leq m(L + k_{m+1}h)$$

for every $N - 1 \geq k_1 \geq k_2 \geq \dots \geq k_n \geq 0$.

Proof. For every $N - 1 \geq k_1 \geq k_2 \geq \dots \geq k_n \geq 0$, if $v_i = L + k_i h$ for all i , then $\sum_{i=1}^n z^h[k_1, k_2, \dots, k_{i-1}, k_{i+1}, \dots, k_n]$ is the total redistribution and $m(L + k_{m+1}h)$ is the total VCG payment. It follows that if the mechanism satisfies the non-deficit property, the inequalities in the proposition must hold. Conversely, if all the inequalities in the proposition hold, then the total redistribution of the mechanism is never more than $m(L + k_{m+1}h)$, which is less than equal to the total VCG payment mv_{m+1} . So the mechanism never incurs a deficit if all the inequalities in the proposition hold. \square

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

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

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

Proposition 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 $U \geq x_1 \geq x_2 \geq \dots \geq x_{n-1} \geq L$, let $d(x_1, \dots, x_{n-1})$ be the minimal distance from any $x_i - L$ to the nearest multiple of h . For any $\delta > 0$, let $r_\delta(x_1, \dots, x_{n-1}) = r^{*h}(x_1, \dots, x_{n-1})$ if $d(x_1, \dots, x_{n-1}) > \delta$, and $r_\delta(x_1, \dots, x_{n-1}) = r^{*h}(x_1, \dots, x_{n-1}) - (\delta - d(x_1, \dots, x_{n-1}))(r^{*h}(x_1, \dots, x_{n-1}) - B)/\delta$ otherwise.

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

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

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

Proof. For all $2N - 1 \geq k_1 \geq k_2 \geq \dots \geq k_{n-1} \geq 0$, let $z^{h/2}[k_1, k_2, \dots, k_{n-1}] = z^{*h}[\lfloor k_1/2 \rfloor, \lfloor k_2/2 \rfloor, \dots, \lfloor k_{n-1}/2 \rfloor]$. The discretized redistribution mechanism corresponding to $z^{h/2}[\dots]$ is exactly r^{*h} . The discretized redistribution mechanism r^{*h} satisfies the non-deficit property. Hence the variables $z^{h/2}[\dots]$ form a feasible solution of the linear program corresponding to step size $h/2$, so its expected redistribution must be less than or equal to that of the optimal solution of the linear program corresponding to step size $h/2$. That is, $R^{*h} \leq R^{*h/2}$. \square

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

Proposition 9. $\lim_{h \rightarrow 0} R^{*h} = R^*$.

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

$|r^*(x_1, x_2, \dots, x_{n-1}) - r^*(x'_1, x'_2, \dots, x'_{n-1})| \leq \epsilon$ as long as $\max_i \{|x_i - x'_i|\} \leq \delta$. Choose $h \leq \delta$, and define $z^h[k_1, k_2, \dots, k_{n-1}]$ by $r^*(L+k_1h, L+k_2h, \dots, L+k_{n-1}h)$ for all $N-1 \geq k_1 \geq k_2 \geq \dots \geq k_{n-1} \geq 0$. $z^h[\dots]$ corresponds to a feasible discretized mechanism r^h . In addition, $r^h \geq r^* - \epsilon$. Hence, the expected redistribution of the optimal discretized mechanism with step size (at most) h is $R^{*h} \geq R^* - \gamma - n\epsilon$. Since γ and ϵ are both arbitrarily small, $\lim_{h \rightarrow 0} R^{*h} \geq R^*$. By Proposition 7, $\lim_{h \rightarrow 0} R^{*h} \leq R^*$. \square

We note that a discretized redistribution mechanism r^h is defined by a finite number of real-valued variables: namely, one variable $z^h[k_1, k_2, \dots, k_{n-1}]$ for every $N-1 \geq k_1 \geq k_2 \geq \dots \geq k_{n-1} \geq 0$. Because of this, we can use a standard LP solver to solve for the optimal discretized redistribution mechanism r^h (for given m, n, h and prior). In general, this linear program involves exponential number of variables and does not scale. However, at least for small problem instances, we can set h to very small values, and by Proposition 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^* .

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

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

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

Proof. Let r be any feasible continuous (anonymous) redistribution mechanism. Now, consider the conditional expectation of a bidder's redistribution payment under r , given that, for each $i \in \{1, \dots, n-1\}$, the i th highest bid among other bidders is in $I(k_i) = [L + k_ih, L + k_ih + h]$. Let $z^h[k_1, k_2, \dots, k_{n-1}]$ denote this conditional

expectation. (We emphasize that this does not depend on which agent we choose, due to the i.i.d. assumption.)

Now, these $z^h[\dots]$ constitute a feasible solution of the above linear program, for the following reason. The left-hand side of the constraint in the above linear program is now the expected total redistribution of r , given that for each $i \in \{1, \dots, n\}$, the i th highest bid is in $I(k_i)$; and the right-hand side is the expected total VCG payment, given that for each $i \in \{1, \dots, n\}$, the i th highest bid is in $I(k_i)$. Because r satisfies non-deficit by assumption, the constraint must be met by the $z^h[\dots]$.

Moreover, the objective value of the feasible solution defined by the $z^h[\dots]$ is identical to the expected total amount redistributed by r . Hence, for every expected total amount redistributed by a feasible continuous mechanism, there exists a feasible solution to the above linear program that attains that value. It follows that $\hat{R}^h \geq R^*$. \square

So, we have that R^{*h} is a lower bound on R^* , and \hat{R}^h is an upper bound. The next proposition considers how close these two bounds are, in terms of the step size h .

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

Proof. Consider the right-hand side of the constraints of the above linear program. We have $mE(v_{m+1} | v_1 \in I(k_1), v_2 \in I(k_2), \dots, v_n \in I(k_n)) \leq m(L + k_{m+1}h + h)$, since $v_{m+1} \in I(k_{m+1})$ implies that $v_{m+1} \leq L + k_{m+1}h + h$. Consider an optimal solution of the linear program for determining \hat{R}^h . Now, from every variable $z^h[k_1, k_2, \dots, k_{n-1}]$, subtract mh/n . This results in a feasible solution of the linear program for determining R^{*h} , and the decrease in the objective value is $nmh/n = mh$. Hence, $\hat{R}^h \leq R^{*h} + mh$. \square

Hence, by solving the linear program for determining R^{*h} , we get a lower bound on R^* and a discretized 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 in expectation despite not being able to welfare dominate the OEL mechanism. On the other hand, the discretized 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, the experimental results below show that for auctions with many bidders, the OEL mechanism redistributes almost everything, whereas for auctions with few bidders, we can solve for the optimal discretized redistribution mechanism with a very

small step size. That is, the two types of redistribution mechanisms are in some sense complementary.

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 discretized mechanism (for specific step sizes). The bids are independently drawn from the uniform $[0, 1]$ distribution.

n,m	VCG	BC	OEL	R^{*h}	\hat{R}^h	%
3,1	0.5000	0.3333	0.3333	0.4218 (N=100)	0.4269	84.4
4,1	0.6000	0.5000	0.5000	0.5498 (N=40)	0.5625	91.6
5,1	0.6667	0.6000	0.6000	0.6248 (N=25)	0.6452	93.7
6,1	0.7143	0.6667	0.6667	0.6701 (N=15)	0.7040	93.8
3,2	0.5000	0.0000	0.3333	0.4169 (N=100)	0.4269	83.4
4,2	0.8000	0.5000	0.5000	0.6848 (N=40)	0.7103	85.6
5,2	1.0000	0.8000	0.8000	0.8944 (N=25)	0.9355	89.4
6,2	1.1429	1.0000	1.0000	1.0280 (N=15)	1.0978	89.9

Expected redistribution by VCG, BC, OEL, and discretized mechanisms, for small numbers of agents.

In the above table, the column “VCG” gives the expected total VCG payment; the column “BC” gives the expected redistribution by the Bailey-Cavallo mechanism; the column “OEL” gives the expected redistribution by the OEL mechanism; the column “ R^{*h} ” gives the expected redistribution by the optimal discretized 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}). The last column gives the percentages of the VCG payment that are redistributed by the optimal discretized redistribution mechanisms (rounding to the nearest tenth).

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

n,m	VCG	BC	OEL	%	n,m	VCG	BC	OEL	%
10,1	0.8182	0.8000	0.8143	99.5	20,1	0.9048	0.9000	0.9048	100.0
10,3	1.9091	1.8000	1.8000	94.3	20,5	3.5714	3.5000	3.5564	99.6
10,5	2.2727	2.0000	2.0000	88.0	20,10	4.7619	4.5000	4.5000	94.5
10,7	1.9091	1.4000	1.8000	94.3	20,15	3.5714	3.0000	3.5564	99.6
10,9	0.8182	0.0000	0.8143	99.5	20,19	0.9048	0.0000	0.9048	100.0

Expected redistribution by VCG, BC, and OEL for large numbers of agents.

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

6 Multi-Unit Auctions with Nonincreasing Marginal Values

In this section, we consider a more general setting in which agents do not necessarily have unit demand, that is, they may value receiving units in addition to the first. However, we assume that the marginal values are nonincreasing, that is, they value the earlier units (weakly) more. (Units remain indistinguishable.) We still use n and m to

denote the number of agents and the number of available units, but we no longer require that $m < n$. An agent's bid is now a nonincreasing sequence of m elements. We denote agent i 's bid by $B_i = \langle b_{i1}, b_{i2}, \dots, b_{im} \rangle$, where b_{ij} is agent i 's marginal value for getting her j th unit (so that $b_{ij} \geq b_{i(j+1)}$). That is, agent i 's valuation for receiving j units is $\sum_{k=1}^j b_{ik}$. A bid profile now consists of n vectors B_i , with $1 \leq i \leq n$, or equivalently mn elements b_{ij} , with $1 \leq i \leq n$ and $1 \leq j \leq m$. We represent the b_{ij} in matrix form as follows:

$$\begin{bmatrix} b_{1m} & b_{2m} & \dots & b_{nm} \\ \dots & \dots & \dots & \dots \\ b_{12} & b_{22} & \dots & b_{n2} \\ b_{11} & b_{21} & \dots & b_{n1} \end{bmatrix}$$

Without loss of generality, we assume that $b_{11} \geq b_{21} \geq \dots \geq b_{n1}$. That is, the agents are ordered according to their marginal values for winning their first unit. We denote the k th highest element among all the b_{ij} by v_k ($1 \leq k \leq mn$).

We assume that we know the joint distribution of the b_{ij} (and hence we also know the joint distribution of the v_k). We continue to use U to denote the known upper bound on the values that the b_{ij} can take (U is also the upper bound on the v_k). In this part of the paper, we will not consider the case where there is a lower bound $L > 0$ on all the b_{ij} (v_k); that is, we assume the lower bound is 0. (In fact, if there is a lower bound $L > 0$, we can simply require the agents to bid how far above L their marginal values are, that is, require them to submit $b'_{ij} = b_{ij} - L$, in which case we arrive at the case that we study below. The VCG payments under these modified bids will always be mL less than under the original bids, but we can easily redistribute this additional mL . Hence, the restriction that $L = 0$ comes without loss of generality.)

Let B be a bid profile. We denote the set of bids other than B_i (agent i 's own bid) by B_{-i} . B_{-i} consists of $mn - m$ elements. We can write B_{-i} in matrix form as follows:

$$\begin{bmatrix} b_{1m} & \dots & b_{i-1,m} & b_{i+1,m} & \dots & b_{nm} \\ \dots & \dots & \dots & \dots & \dots & \dots \\ b_{12} & \dots & b_{i-1,2} & b_{i+1,2} & \dots & b_{n2} \\ b_{11} & \dots & b_{i-1,1} & b_{i+1,1} & \dots & b_{n1} \end{bmatrix}$$

We denote the k th highest element in B_{-i} by $v_{-i,k}$ ($1 \leq k \leq mn - m$).

Our definition for VCG redistribution mechanisms in this setting is similar to our earlier definition. Namely, in a VCG redistribution mechanism, we first allocate the units efficiently, according to the VCG mechanism; then, each agent receives a redistribution payment that is independent of her own bid. An efficient allocation is obtained by accepting the m highest marginal values (v_1, v_2, \dots, v_m). That is, if x elements among v_1, v_2, \dots, v_m come from agent i 's bid, then agent i wins x units. Agent i 's redistribution equals $r(B_{-i})$, where r is the function that characterizes the redistribution rule.

We now need a definition of *linear* redistribution mechanisms in this setting. We could define linear redistribution mechanisms as follows:

$$r(B_{-i}) = c_0 + c_1 v_{-i,1} + c_2 v_{-i,2} + \dots + c_{mn-m} v_{-i,mn-m}$$

We will study this particular definition later in the paper; however, it should immediately be noted that this definition ignores some potentially valuable information in B_{-i} , as shown by the following example.

Example 3. Let $n = 3$ and $m = 2$.

- Case 1: Let B_{-i} be $\begin{bmatrix} 0 & 0 \\ U & U \end{bmatrix}$.
- Case 2: Let B_{-i} be $\begin{bmatrix} U & 0 \\ U & 0 \end{bmatrix}$.

In both cases, we have $v_{-i,1} = U$, $v_{-i,2} = U$, $v_{-i,3} = 0$, and $v_{-i,4} = 0$. Hence, if we define the linear redistribution mechanisms as above, then the redistribution payment must be the same in both cases.

We can see that the above definition loses some information about the ordering of the elements in the matrix. We will show later that this information loss can in fact come at a cost (less payments can be redistributed). It would be good if we can incorporate the information about the order of the b_{ij} in B_{-i} in the definition of linear redistribution mechanisms. This is what we will do next.

Let B and B' be two bid profiles. The elements in B and B' are denoted by b_{ij} and b'_{ij} , respectively, for $1 \leq i \leq n$ and $1 \leq j \leq m$. We say B and B' are **order consistent**, denoted by $B \simeq B'$, if for any i_1, j_1, i_2, j_2 , we have that $b_{i_1 j_1} > b_{i_2 j_2}$ implies $b'_{i_1 j_1} \geq b'_{i_2 j_2}$, and $b'_{i_1 j_1} > b'_{i_2 j_2}$ implies $b_{i_1 j_1} \geq b_{i_2 j_2}$. An *order consistent class* of bid profiles consists of bid profiles that are all pairwise order consistent. The set of all allowable bid profiles can be divided into a finite number of maximal order consistent classes (that is, order consistent classes that are not proper subsets of other order consistent classes). (Specifically, we have one such class for every strict ordering $<$ on the ordered pairs (i, j) ($1 \leq i \leq m$ and $1 \leq j \leq n$) such that $(i, j + 1) < (i, j)$ and $(i + 1, 1) < (i, 1)$ everywhere. We note that some bid profiles are part of more than one of these maximal order consistent classes: for example, the bid profile with all 0 elements belongs to all the classes.) We can apply the same definitions of order consistency and (maximal) order consistent classes to the profiles of *other* bids, the B_{-i} . Let $I(B_{-i})$ denote the maximal order consistent class that contains B_{-i} .⁹

The following definition of linear redistribution mechanisms successfully captures the ordering information of B_{-i} , by having separate coefficients for every maximal order consistent class.

$$r(B_{-i}) = c_{I(B_{-i}),0} + c_{I(B_{-i}),1}v_{-i,1} + \dots + c_{I(B_{-i}),mn-m}v_{-i,mn-m}$$

Since $\begin{bmatrix} 0 & 0 \\ U & U \end{bmatrix}$ and $\begin{bmatrix} U & 0 \\ U & 0 \end{bmatrix}$ are not order consistent, they can result in different redistribution payments in this class of redistribution mechanisms.

Of course, this set of coefficients is unwieldy. As it turns out, we can simplify the representation of these mechanisms if we assume that they are continuous.

⁹If B_{-i} belongs to multiple maximal order consistent classes, then $I(B_{-i})$ is the class with the smallest index in any predetermined order of all the classes. If we assume continuity of the redistribution function, as we will do below, then in fact it does not matter which maximal order consistent class we choose for B_{-i} .

Let r be a linear redistribution mechanism (as just defined). Let $T(B_{-i}, k)$ be the result of changing the largest k elements of B_{-i} into U , and changing the remaining elements of B_{-i} into 0. (We assume that ties for the top k values are broken in a consistent way.) We note that $T(B_{-i}, k) \simeq B_{-i}$ for all $0 \leq k \leq mn - m$. For example, $T\left(\begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix}, 1\right) = \begin{bmatrix} 0 & 0 \\ U & 0 \end{bmatrix}$ and $T\left(\begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix}, 2\right) = \begin{bmatrix} 0 & 0 \\ U & U \end{bmatrix}$.

We define the following function r' :

$$r'(B_{-i}) = r(T(B_{-i}, 0)) + \frac{r(T(B_{-i}, 1)) - r(T(B_{-i}, 0))}{U} v_{-i,1} + \dots + \frac{r(T(B_{-i}, mn - m)) - r(T(B_{-i}, mn - m - 1))}{U} v_{-i, mn - m}$$

Proposition 12. *If r is continuous, then $r = r'$.*

Proof. We first restrict our attention to profiles B_{-i} in a specific (but arbitrary) maximal order consistent class; moreover, we only consider profiles B_{-i} in which no two elements are equal. For any B_{-i} in this class, we use the same $mn - m + 1$ coefficients of r , and $T(B_{-i}, k)$ (and hence $r(T(B_{-i}, k))$) depends only on k . That is, both the coefficients and $T(B_{-i}, k)$ are constant in B_{-i} .

If r is continuous, then when B_{-i} approaches $T(B_{-i}, k)$, we have that $r(B_{-i})$ approaches $r(T(B_{-i}, k))$. By the definition of r' , we also have that when B_{-i} approaches $T(B_{-i}, k)$, that is, when the first k elements of B_{-i} approach U and the remaining elements of B_{-i} approach 0, we have that $r'(B_{-i})$ approaches $r(T(B_{-i}, k))$. That is, $r(T(B_{-i}, k)) = r'(T(B_{-i}, k))$ for $0 \leq k \leq mn - m$; that is, the functions agree in $mn - m + 1$ different places. Since r and r' are both linear functions with $mn - m + 1$ constant coefficients, r and r' must be the same function when B_{-i} is restricted to one class. Since the choice of class was arbitrary, we have that $r = r'$. \square

From now on, we only consider continuous r . Hence, we can characterize r by the values it attains at all possible $T(B_{-i}, k)$. $T(B_{-i}, k)$ consists of only the numbers U and 0. We represent $T(B_{-i}, k)$ by an integer vector of length n , where the i th coordinate of the vector is the number of U s in the i th column of $T(B_{-i}, k)$.

For example,

$$T\left(\begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix}, 2\right) = \begin{bmatrix} U & 0 \\ U & 0 \end{bmatrix} \rightarrow \langle 2, 0 \rangle$$

$$T\left(\begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix}, 3\right) = \begin{bmatrix} 0 & U \\ U & U \end{bmatrix} \rightarrow \langle 1, 2 \rangle$$

Using this, $r(T(B_{-i}, k))$ can be rewritten as $r[x_1, x_2, \dots, x_{n-1}]$, where $\langle x_1, x_2, \dots, x_{n-1} \rangle$ is the vector representing $T(B_{-i}, k)$ (with for each i , $0 \leq x_i \leq m$, and $\sum x_i = k$). Moreover, because we have, for example, that $r\left(\begin{bmatrix} 0 & U \\ U & U \end{bmatrix}\right) = r\left(\begin{bmatrix} U & 0 \\ U & U \end{bmatrix}\right)$, we can assume without loss of generality that $x_1 \geq x_2 \geq \dots \geq x_{n-1}$.

The following is an example of how to compute an agent's redistribution payment based on the values of $r[x_1, x_2, \dots, x_{n-1}]$.

Example 4. Let $n = 3$ and $m = 2$. Let $B_{-i} = \begin{bmatrix} 3 & 1 \\ 4 & 2 \end{bmatrix}$.

$$\begin{aligned}
r(B_{-i}) &= r(T(B_{-i}, 0)) + \frac{r(T(B_{-i}, 1)) - r(T(B_{-i}, 0))}{U} v_{-i,1} + \dots + \\
&\quad \frac{r(T(B_{-i}, mn - m)) - r(T(B_{-i}, mn - m - 1))}{U} v_{-i, mn - m} \\
&= r\left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}\right) + \frac{r\left(\begin{bmatrix} 0 & 0 \\ U & 0 \end{bmatrix}\right) - r\left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}\right)}{U} \cdot 4 + \frac{r\left(\begin{bmatrix} U & 0 \\ U & 0 \end{bmatrix}\right) - r\left(\begin{bmatrix} 0 & 0 \\ U & 0 \end{bmatrix}\right)}{U} \cdot 3 \\
&\quad + \frac{r\left(\begin{bmatrix} U & 0 \\ U & U \end{bmatrix}\right) - r\left(\begin{bmatrix} U & 0 \\ U & 0 \end{bmatrix}\right)}{U} \cdot 2 + \frac{r\left(\begin{bmatrix} U & U \\ U & U \end{bmatrix}\right) - r\left(\begin{bmatrix} U & 0 \\ U & U \end{bmatrix}\right)}{U} \cdot 1 \\
&= r[0, 0] + \frac{r[1, 0] - r[0, 0]}{U} \cdot 4 + \frac{r[2, 0] - r[1, 0]}{U} \cdot 3 + \frac{r[2, 1] - r[2, 0]}{U} \cdot 2 + \frac{r[2, 2] - r[2, 1]}{U} \cdot 1
\end{aligned}$$

Since the values of the $r[x_1, x_2, \dots, x_{n-1}]$ completely characterize the continuous linear redistribution mechanism, we can solve for values of the $r[x_1, x_2, \dots, x_{n-1}]$ for which the corresponding linear redistribution mechanism satisfies the non-deficit property and produces the least waste in expectation under this constraint.

The following proposition characterizes the non-deficit linear redistribution mechanisms.

Proposition 13. *A linear redistribution mechanism satisfies the non-deficit property if and only if the corresponding $r[x_1, x_2, \dots, x_{n-1}]$ satisfy the following inequalities: For all $m \geq x_1 \geq x_2 \geq x_3 \geq \dots \geq x_n \geq 0$, $\sum_{i=1}^n r[x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n] \leq U \cdot (\sum_{i=1}^n \min\{(\sum_{j=1}^n x_j) - x_i, m\} - (n-1) \min\{\sum_{j=1}^n x_j, m\})$. (The right-hand side of the inequality corresponds to the total VCG payment for the profile $\langle x_1, x_2, \dots, x_n \rangle$.)*

Proof. To see why the right-hand side $U \cdot (\sum_{i=1}^n \min\{(\sum_{j=1}^n x_j) - x_i, m\} - (n-1) \min\{\sum_{j=1}^n x_j, m\})$ corresponds to the total VCG payment, we note that $U \cdot \min\{(\sum_{j=1}^n x_j) - x_i, m\}$ is the total efficiency when i is removed, so that $U \cdot \sum_{i=1}^n \min\{(\sum_{j=1}^n x_j) - x_i, m\}$ is the sum of all the terms corresponding to efficiencies when one agent is removed. $U \cdot (n-1) \min\{\sum_{j=1}^n x_j, m\}$ corresponds to the sum of the basic Groves terms in the payments from the agents: in this term, each agent receives the total efficiency obtained by the other agents (when the agent is not removed), and if we sum over all the agents, that means each agent is counted $n-1$ times.

Now we can prove the main part of the proposition. If the non-deficit property is satisfied for all bid profiles, then it should also be satisfied when the marginal values are restricted to be either U or 0 . This proves the “only if” direction. Now we prove the “if” direction. Let B be any bid profile from a fixed maximal order consistent class. This implies that the maximal order consistent class of B_{-i} is fixed as well, for every i . The total VCG payment equals the sum over all i of the m highest elements

in B_{-i} , minus $n - 1$ times the sum of the m highest elements in B . In either case, because we are restricting attention to a fixed class, the m highest elements are the same ones for any B in the class. Because of this, the VCG payments are linear in the v_i . Additionally, again because we are restricting attention to one particular class, the redistribution payments are also linear in the v_i .

Now, if the inequalities hold, that means that the total VCG payment minus the total redistribution is nonnegative when the marginal values are restricted to either U or 0. That is, the non-deficit constraints hold for these extreme cases. But by Lemma 1, if a non-deficit constraint is violated anywhere, then a non-deficit constraint must be violated for one of these extreme cases. It follows that the non-deficit constraints hold everywhere in the class that we were considering, and because this class was arbitrary, the non-deficit constraint must hold everywhere. \square

Let z be the total number of maximal order consistent classes. Let Z_j be an arbitrary bid profile that is (only) in the j th class. Let $P(B \in I(Z^j))$ be the probability that a bid profile is drawn that is (only) in the j th class, and let $E(v_{-i,k} | B \in I(Z^j))$ be the expectation of the k th-highest marginal value among B_{-i} , given that B is (only) in the j th class. We assume that the probability that we draw a bid vector that is in more than one class is zero (this would require that two values are exactly equal).

Now we are ready to introduce a linear program that solves for the optimal-in-expectation linear redistribution mechanism.¹⁰ This linear program is based on the alternative representation of linear redistribution mechanisms, whose correctness was established by Proposition 12, and on the characterization of the non-deficit constraints established for this representation by Proposition 13.

<p>Variables: $r[x_1, x_2, \dots, x_{n-1}]$ for all integer $m \geq x_1 \geq x_2 \geq \dots \geq x_{n-1} \geq 0$</p> <p>Maximize:</p> $\sum_j P(B \in I(Z^j)) \sum_i [r(T(Z_{-i}^j, 0)) + \frac{r(T(Z_{-i}^j, 1)) - r(T(Z_{-i}^j, 0))}{U} E(v_{-i,1} B \in I(Z^j)) + \dots + \frac{r(T(Z_{-i}^j, mn-m)) - r(T(Z_{-i}^j, mn-m-1))}{U} E(v_{-i, mn-m} B \in I(Z^j))]$ <p>Subject to:</p> <p>For all $m \geq x_1 \geq x_2 \geq x_3 \geq \dots \geq x_n \geq 0$,</p> $\sum_{i=1}^n r[x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n] \leq U \cdot (\sum_{i=1}^n \min\{(\sum_{j=1}^n x_j) - x_i, m\} - (n-1) \min\{\sum_{j=1}^n x_j, m\})$

We do not have an analytical solution to this linear program; all that we can do is solve for the optimal mechanism for specific values of m and n . More problematically, in general it is not easy to compute the constants $P(B \in I(Z^j))$ and $E(v_{-i,k} | B \in I(Z^j))$.

¹⁰Incidentally, we can give a similar linear program for finding the linear redistribution mechanism that is worst-case optimal, that is, it maximizes the fraction of total VCG payment redistributed in the worst case. In previous work [16], we have already identified a worst-case optimal linear mechanism for the nonincreasing marginal values case; however, that mechanism is only optimal under the requirement of ex-post individual rationality. The linear programming technique here can be used to find the worst-case optimal mechanism when individual rationality is not required. For the sake of coherence of this paper, we will not go into further detail on this here.

$I(Z^j)$). One way to work around this problem is to approximate the final result. That is, instead of computing an exact optimal linear redistribution mechanism, we can draw a few sample bid profiles, and solve for a linear redistribution mechanism that is optimal for the samples. This way, we do not need to compute any probabilities or conditional expectations; we simply sum over the profiles in the sample in the objective. (However, we still enforce the constraints everywhere, not just on the samples.) Because the linear redistribution mechanisms are continuous and we assume continuous and bounded prior distributions for the valuations, it follows that as the number of samples grows, we approach an optimal mechanism.

We now return to the original idea for the definition of linear redistribution mechanisms: what if we ignore the ordering information and just use coefficients c_k for $0 \leq k \leq mn - m$, which do not depend on the maximal order consistent class? This will be a more scalable approach, although it will come at a loss. To find an optimal mechanism in this class, we can take a similar approach as we did above for the more general definition of linear redistribution mechanisms (and this approach is correct for similar reasons). We consider the extreme bid vectors where all marginal values are U or 0, represented by vectors of integers x_1, x_2, \dots, x_n , as before. The fact that we ignore the ordering information now implies that we require that $r[x_1, x_2, \dots, x_{n-1}] = r[y_1, y_2, \dots, y_{n-1}]$ whenever $\sum_{i=1}^{n-1} x_i = \sum_{i=1}^{n-1} y_i$. So, we can rewrite $r[x_1, \dots, x_{n-1}]$ as $r[\sum_{i=1}^{n-1} x_i]$. That is, the variables now are $r[k]$ for $k = 0, 1, \dots, mn - m$. The redistribution function now becomes:

$$r(B_{-i}) = r[0] + \frac{r[1] - r[0]}{U} v_{-i,1} + \dots + \frac{r[mn - m] - r[mn - m - 1]}{U} v_{-i, mn - m}$$

The linear program for finding an optimal mechanism becomes:

Variables: $r[k]$ for integer $0 \leq k \leq mn - m$
Maximize:

$$\sum_i [r[0] + \frac{r[1] - r[0]}{U} E(v_{-i,1}) + \dots + \frac{r[mn - m] - r[mn - m - 1]}{U} E(v_{-i, mn - m})]$$

Subject to:
For all $m \geq x_1 \geq x_2 \geq x_3 \geq \dots \geq x_n \geq 0$,
 $\sum_{i=1}^n r[(\sum_{j=1}^n x_j) - x_i] \leq$
 $U \cdot (\sum_{i=1}^n \min\{(\sum_{j=1}^n x_j) - x_i, m\} - (n - 1) \min\{\sum_{j=1}^n x_j, m\})$

While this linear program is much more manageable, it may lead to worse results than the earlier linear program, which optimizes over the more general class of linear redistribution mechanisms that take the ordering information into account. We now study some example solutions to this linear program, and compare them to the Bailey-Cavallo redistribution mechanism [2, 3]. We recall that the Bailey-Cavallo mechanism redistributes to every agent $\frac{1}{n}$ times the VCG payment that would result if this agent were removed from the auction. If we only consider bid profiles from a specific maximal order consistent class, then for any i , the VCG payment that would result if i is removed is a linear combination of the $v_{-i,k}$. Therefore, the Bailey-Cavallo mechanism

belongs to the family of linear redistribution mechanisms that consider the ordering information (and hence, the optimal solution to the earlier linear program will do at least as well as the Bailey-Cavallo mechanism). The Bailey-Cavallo mechanism does not belong to the family of linear redistribution mechanisms that ignore the ordering information: in fact, we will see that it sometimes performs better than the optimal mechanism among linear redistribution mechanisms that ignore the ordering information. Hence, ignoring the ordering information in general comes at a cost.

For these examples, let us recall that agent i 's bid vector B_i consists of m elements $b_{i1}, b_{i2}, \dots, b_{im}$. In both examples, we assume that the values of $b_{i1}, b_{i2}, \dots, b_{im}$ are drawn independently from the uniform $[0, 1]$ distribution, with b_{ij} being the j th highest among the m drawn values. We also assume that B_1, B_2, \dots, B_n are independent.

Example 5. Suppose that $n = 3$ and $m = 2$. By solving the above linear program (the one that ignores the ordering information), we get the following linear redistribution mechanism that ignores ordering information: $r(B_{-i}) = \frac{2}{3}v_{-i,3}$. That is, an agent's redistribution is equal to two thirds of the third highest marginal value among the set of other bids. The expected waste of this mechanism is 0.2571. In contrast, the expected waste of the Bailey-Cavallo mechanism is 0.4571. (The expected total VCG payment is 1.0571.) So, for this example, the optimal linear redistribution mechanism that ignores the ordering information outperforms the Bailey-Cavallo mechanism.

Example 6. Suppose that $n = 7$ and $m = 2$. By solving the above linear program (the one that ignores the ordering information), we get the following linear redistribution mechanism that ignores ordering information: $r(B_{-i}) = \frac{1}{5}v_{-i,3} + \frac{3}{35}v_{-i,4}$. That is, an agent's redistribution is equal to $\frac{1}{5}$ times the third highest marginal value among the set of other bids, plus $\frac{3}{35}$ times the fourth highest marginal value among the set of other bids. The expected waste of this mechanism is 0.0923. In contrast, the expected waste of the Bailey-Cavallo mechanism is 0.0671. (The expected total VCG payment is 1.5846.) So, for this example, the Bailey-Cavallo mechanism outperforms the optimal linear redistribution mechanism that ignores the ordering information.

In both of these examples (as well as in other examples for which we solved the linear program, including examples with other distributions), the optimal linear redistribution mechanism that ignores the ordering information is a special case of the following more general mechanism.

Mechanism M^* is defined as follows, where $t = m + \lfloor \frac{m(n-2)}{n} \rfloor$.

- $r[k] = U^{\frac{k-m}{n-2}}$ for $k = m + 1, m + 2, \dots, t$
- $r[k] = U^{\frac{m}{n}}$ for $k > t$

The redistribution an agent receives is:

$$r(B_{-i}) = \sum_{m+1 \leq k \leq t} \frac{1}{n-2} v_{-i,k} + \left(\frac{m}{n} - \frac{t-m}{n-2} \right) v_{-i,t+1}$$

We conjecture that there are some more general conditions under which M^* is the optimal linear redistribution mechanism that ignores the ordering information.

7 Conclusion

The well-known VCG mechanism allocates the items efficiently, is strategy-proof, 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. Previous research on redistribution mechanisms has resulted in a worst-case optimal redistribution mechanism, that is, a mechanism that maximizes the fraction of VCG payments redistributed in the worst case. In contrast, in this paper, we assumed that a prior distribution over the agents' valuations is available, and studied the goal of maximizing the expected total redistribution.

For the setting of multi-unit auctions with unit demand, 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 for fixed number of units and welfare undominated. We then proposed *discretized* redistribution mechanisms, which discretize the space of possible valuations, and determine redistributions solely based on the discretized values (however, the strategy-proofness and non-deficit constraints still hold over the non-discretized space). Given a discretization step size, we showed how to solve for the optimal discretized 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 discretized redistribution mechanism with a very small step size.

For the setting of multi-unit auctions with nonincreasing marginal values, we first generalized the definition of linear redistribution mechanisms. We then introduced a linear program for finding the optimal linear redistribution mechanism. Because this linear program is unwieldy, we also introduced a simplified linear program that produces relatively good (though not necessarily optimal) linear redistribution mechanisms. We also conjectured an analytical solution to the linear program, which we expect to be correct for most reasonable distributions.

Future research on optimal-in-expectation redistribution mechanisms can take a number of directions. For the setting of nonincreasing marginal utilities, one can try to find subclasses of the linear redistribution mechanisms that are more general than the subclass we considered but still lead to more tractable optimization problems. In general, one can also try to solve for an optimal-in-expectation redistribution mechanism that is not necessarily linear. Another direction is to extend the results of this paper to more general settings, for example, combinatorial auctions. Finally, inefficient mechanisms sometimes achieve higher social welfare than efficient mechanisms [14, 9] in the worst case. It would be interesting to see whether agents' expected social welfare can also be improved by allocating units inefficiently, and if so, by how much.

Acknowledgments

Parts of this work have been presented at the INFORMS-07, AAMAS-08, and GAMES-08 conferences. We thank Hervé Moulin and Giuseppe Lopomo for helpful discussions. This work has been supported by NSF under IIS-0812113 and an Alfred P. Sloan Research Fellowship.

References

- [1] K. Apt, V. Conitzer, M. Guo, and E. Markakis. Welfare undominated Groves mechanisms. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 426–437, Shanghai, China, 2008.
- [2] M. J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.
- [3] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.
- [4] R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 220–229, Chicago, IL, USA, 2008.
- [5] S. Chakravarty and T. Kaplan. Manna from heaven or forty years in the desert: Optimal allocation without transfer payments, October 2006. Working Paper.
- [6] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [7] 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.
- [8] C. d’Aspremont and L.-A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.
- [9] G. de Clippel, V. Naroditskiy, and A. Greenwald. Destroy to save. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 207–214, Stanford, CA, USA, 2009.
- [10] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce (AMEC), LNAI, 3435*, pages 30–43, 2005.
- [11] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001. Early version in STOC-00.

- [12] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [13] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [14] M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 210–219, Chicago, IL, USA, 2008.
- [15] M. Guo and V. Conitzer. Undominated VCG redistribution mechanisms. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1039–1046, Estoril, Portugal, 2008.
- [16] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009.
- [17] J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, 2008.
- [18] B. Holmström. Groves’ scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.
- [19] 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.
- [20] H. Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
- [21] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [22] 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, USA, 2001.
- [23] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118:209–228, 2004. Early version appeared in IJCAI-01.
- [24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

A Deriving an optimal linear redistribution mechanism

Here we derive the OEL mechanism and prove its optimality. Our objective is to find a 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 strategy-proof

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 $U \geq v_1 \geq v_2 \geq \dots \geq v_n \geq L$
 $\sum_{i=1}^n r_i \leq mv_{m+1}$
 $r_i = c_0 + c_1 v_{-i,1} + c_2 v_{-i,2} + \dots + c_{n-1} v_{-i,n-1}$

Given the prior distribution, $E(mv_{m+1})$ is a constant, so the objective of the above model may be rewritten as **Minimize** $E(mv_{m+1} - \sum_{i=1}^n r_i)$.

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

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

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

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

(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, \dots, q_n ($n+1$ values) are determined by c_0, \dots, c_{n-1} (n values). Conversely, if q_0, \dots, q_{n-1} are fixed, then we can completely solve for the values of c_0, \dots, c_{n-1} (and hence also for q_n). This results in the following relation among the q_i :

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

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 $mv_{m+1} - \sum_{i=1}^n r_i = q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n$, 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, \dots, q_n
Minimize $E(q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n)$
Subject to:
 For every bid vector $U \geq v_1 \geq v_2 \geq \dots \geq v_n \geq L$
 $q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n \geq 0$
 $\sum_{i=1}^n (-1)^{i-1} \binom{n-1}{i-1} q_i = (-1)^m m \binom{n-1}{m}$

In what follows, we will cast the above model into a linear program. We begin with the following lemma[16]:

Lemma 1. *The following are equivalent:*

- (1) $q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n \geq 0$ for all $U \geq v_1 \geq v_2 \geq \dots \geq v_n \geq L$
- (2) $q_0 + L \sum_{i=1}^n q_i + (U - L) \sum_{i=1}^k q_i \geq 0$ for $k = 0, \dots, n$

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

(2) \Rightarrow (1): Let us rewrite $T = q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n$ as $q_0 + L \sum_{i=1}^n q_i + (v_1 - v_2) \sum_{i=1}^1 q_i + (v_2 - v_3) \sum_{i=1}^2 q_i + \dots + (v_{n-1} - v_n) \sum_{i=1}^{n-1} q_i + (v_n - L) \sum_{i=1}^n q_i$. If $\sum_{i=1}^k q_i \geq 0$ for every $k = 1, \dots, n$, then $T \geq q_0 + L \sum_{i=1}^n q_i \geq 0$ (because $v_1 - v_2, v_2 - v_3, \dots, v_n - L$ 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 $v_1 = v_2 = \dots = v_{k'} = U$ and $v_{k'+1} = v_{k'+2} = \dots = v_n = L$. Hence T is always nonnegative. \square

Let $x_k = (q_0 + L \sum_{i=1}^n q_i) / (U - L) + \sum_{i=1}^k q_i$ for $k = 0, \dots, n$. 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 $k = 1, \dots, n$, 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 $o_0 = U - Ev_1$, $o_i = Ev_i - Ev_{i+1}$ ($i = 1, \dots, n-1$) 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:

$$\begin{aligned}
& E(q_0 + q_1v_1 + q_2v_2 + \dots + q_nv_n) \\
&= q_0 + q_1Ev_1 + q_2Ev_2 + \dots + q_nEv_n \\
&= x_0(U - L) + q_1(Ev_1 - L) + q_2(Ev_2 - L) + \dots + q_n(Ev_n - L) \\
&= x_0((U - L) - (Ev_1 - L)) + (x_0 + q_1)((Ev_1 - L) - (Ev_2 - L)) + (x_0 + q_1 + q_2)((Ev_2 - L) - (Ev_3 - L)) + \dots + (x_0 + q_1 + \dots + q_n)(Ev_n - L) \\
&= o_0x_0 + o_1x_1 + \dots + o_nx_n
\end{aligned}$$

We finally obtain the following linear program:

<p>Variables: x_0, x_1, \dots, x_n Minimize $o_0x_0 + o_1x_1 + \dots + o_nx_n$ Subject to: $x_i \geq 0$ $\sum_{i=0}^n (-1)^i \binom{n}{i} x_i = (-1)^{m-1} m \binom{n-1}{m}$</p>

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 proposition gives an analytical solution of this linear program.

Proposition 14. *Let k be any integer satisfying*

$$k \in \arg \min_i \{o_i / \binom{n}{i} \mid i - m \text{ odd}, i = 0, \dots, n\}$$

The above linear program has the following optimal solution:

$$x_k = m \binom{n-1}{m} / \binom{n}{k}, \text{ and } x_i = 0 \text{ for } i \neq k$$

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

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

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:

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

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 $k \in \arg \min_i \{o_i / \binom{n}{i} \mid i - m \text{ odd}, i = 0, \dots, n\}$, the 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 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) $\arg \min_i \{o_i / \binom{n}{i} \mid i - m \text{ odd}, i = 0, \dots, n\}$ 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). \square

By transforming the x_i from Proposition 14 to the corresponding c_i , we obtain the OEL mechanism from Theorem 1.