Worst-Case Optimal Redistribution of VCG Payments

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

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

ABSTRACT

For allocation problems with one or more items, the well-known Vickrey-Clarke-Groves (VCG) mechanism is efficient, strategy-proof, individually rational, and does not incur a deficit. However, the VCG mechanism is not (strongly) budget balanced: generally, the agents' payments will sum to more than 0. If there is an auctioneer who is selling the items, this may be desirable, because the surplus payment corresponds to revenue for the auctioneer. However, if the items do not have an owner and the agents are merely interested in allocating the items efficiently among themselves, any surplus payment is undesirable, because it will have to flow out of the system of agents. In 2006, Cavallo [3] proposed a mechanism that redistributes some of the VCG payment back to the agents, while maintaining efficiency, strategy-proofness, individual rationality, and the non-deficit property. In this paper, we extend this result in a restricted setting. We study allocation settings where there are multiple indistinguishable units of a single good, and agents have unit demand. (For this specific setting, Cavallo's mechanism coincides with a mechanism proposed by Bailey in 1997 [2].) Here we propose a family of mechanisms that redistribute some of the VCG payment back to the agents. All mechanisms in the family are efficient, strategy-proof, individually rational, and never incur a deficit. The family includes the Bailey-Cavallo mechanism as a special case. We then provide an optimization model for finding the optimal mechanism—that is, the mechanism that maximizes redistribution in the worst case—inside the family, and show how to cast this model as a linear program. We give both numerical and analytical solutions of this linear program, and the (unique) resulting mechanism shows significant improvement over the Bailey-Cavallo mechanism (in the worst case). Finally, we prove that the obtained mechanism is optimal among *all* anonymous deterministic mechanisms that satisfy the above properties.

EC'07, June 13–16, 2007, San Diego, California, USA.

Categories and Subject Descriptors

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

General Terms

Algorithms, Economics, Theory

Keywords

Mechanism Design, Vickrey-Clarke-Groves Mechanism, Payment Redistribution

1. INTRODUCTION

Many important problems in computer science and electronic commerce can be modeled as resource allocation problems. In such problems, we want to allocate the resources (or *items*) to the agents that value them the most. Unfortunately, agents' valuations are private knowledge, and self-interested agents will lie about their valuations if this is to their benefit. One solution is to *auction* off the items, possibly in a *combinatorial* auction where agents can bid on bundles of items. There exist ways of determining the payments that the agents make in such an auction that incentivizes the agents to report their true valuations—that is, the payments make the auction *strategy-proof*. One very general way of doing so is to use the VCG mechanism [23, 4, 12]. (The VCG mechanism is also known as the Clarke mechanism or, in the specific context of auctions, the Generalized Vickrey Auction.)

Besides strategy-proofness, the VCG mechanism has several other nice properties in the context of resource allocation problems. It is *efficient*: the chosen allocation always maximizes the sum of the agents' valuations. It is also *(ex-post) individually rational*: participating in the mechanism never makes an agent worse off than not participating. Finally, it has a *no-deficit* property: the sum of the agents' payments is always nonnegative.

In many settings, another property that would be desirable is *(strong) budget balance*, meaning that the payments sum to exactly 0. Suppose the agents are trying to distribute some resources among themselves that do not have a previous owner. For example, the agents may be trying to allocate the right to use a shared good on a given day. Or, the agents may be trying to allocate a resource that they have collectively constructed, discovered, or otherwise obtained. If the agents use an auction to allocate these resources, and the sum of the agents' payments in the auction is positive, then this surplus payment must leave the system of the agents (for example, the agents must give the money to an outside party, or burn it). Naïve redistribution of the surplus payment (*e.g.* each of the n agents receives $1/n$ of the surplus) will generally result in a mechanism that is not strategy-proof (*e.g.* in a Vickrey auction, the second-highest bidder would want to increase her bid to obtain a larger redistribution payment). Unfortunately, the VCG mechanism is not budget balanced: typically, there is surplus payment. Unfortunately, in general settings, it is in fact impossible to design mechanisms that satisfy budget balance in addition to the other desirable properties [16, 11, 21].

In light of this impossibility result, several authors have obtained budget balance by sacrificing some of the other desirable properties [2, 6, 22, 5]. Another approach that is perhaps preferable is to use a mechanism that is "more" budget balanced than the VCG mechanism, and maintains all the other desirable properties. One way of trying to design such a mechanism is to redistribute some of the VCG payment back to the agents in a way that will not affect the agents' incentives (so that strategy-proofness is maintained), and that will maintain the other properties. In 2006, Cavallo [3] pursued exactly this idea, and designed a mechanism that redistributes a large amount of the total VCG payment while maintaining all of the other desirable properties of the VCG mechanism. For example, in a single-item auction (where the VCG mechanism coincides with the second-price sealed-bid auction), the amount redistributed to bidder i by Cavallo's mechanism is $1/n$ times the second-highest bid among bids *other than* i's bid. The total redistributed is at most the second-highest bid overall, and the redistribution to agent i does not affect i's incentives because it does not depend on i's own bid.

In this paper, we restrict our attention to a limited setting, and in this setting we extend Cavallo's result. We study allocation settings where there are multiple indistinguishable units of a single good, and all agents have *unit demand*, *i.e.* they want only a single unit. For this specific setting, Cavallo's mechanism coincides with a mechanism proposed by Bailey in 1997 [2]. Here we propose the family of *linear* VCG redistribution mechanisms. All mechanisms in this family are efficient, strategy-proof, individually rational, and never incur a deficit. The family includes the Bailey-Cavallo mechanism as a special case (with the caveat that we only study allocation settings with multiple indistinguishable units of a single good and unit demand, while Bailey's and Cavallo's mechanisms can be applied outside these settings as well). We then provide an optimization model for finding the optimal mechanism inside the family, based on worst-case analysis. Both numerical and analytical solutions of this model are provided, and the resulting mechanism shows significant improvement over the Bailey-Cavallo mechanism (in the worst case). For example, for the problem of allocating a single unit, when the number of agents is 10, our mechanism always redistributes more than 98% of the total VCG payment back to the agents (whereas the Bailey-Cavallo mechanism redistributes only 80% in the worst case). Finally, we prove that our mechanism is in fact optimal among *all* anonymous deterministic mechanisms (even nonlinear ones) that satisfy the desirable properties.

Around the same time, the same mechanism has been independently derived by Moulin [19].[1] Moulin actually pursues a different objective (also based on worst-case analysis): whereas our objective is to maximize the percentage of VCG payments that are redistributed, Moulin tries to minimize the overall payments from agents as a percentage of efficiency. It turns out that the resulting mechanisms are the same. Towards the end of this paper, we consider dropping the individual rationality requirement, and show that this does not change the optimal mechanism for our objective. For Moulin's objective, dropping individual rationality does change the optimal mechanism (but only if there are multiple units).

2. PROBLEM DESCRIPTION

Let n denote the number of agents, and let m denote the number of units. We only consider the case where $m < n$ (otherwise the problem becomes trivial). We also assume that m and n are always known. (This assumption is not harmful: in environments where anyone can join the auction, running a redistribution mechanism is typically not a good idea anyway, because everyone would want to join to collect part of the redistribution.)

Let the set of agents be $\{a_1, a_2, \ldots, a_n\}$, where a_i is the agent with ith highest report value $\hat{v}_i$—that is, we have $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_n \geq 0$. Let v_i denote the true value of a_i. Given that the mechanism is strategy-proof, we can assume $v_i = \hat{v}_i$.

Under the VCG mechanism, each agent among $a_1, \ldots, a_m$ wins a unit, and pays $\hat{v}_{m+1}$ for this unit. Thus, the total VCG payment equals $m\hat{v}_{m+1}$. When $m = 1$, this is the second-price or Vickrey auction.

We modify the mechanism as follows. After running the original VCG mechanism, the center returns to each agent a_i some amount z_i, agent a_i's *redistribution payment*. We do not allow z_i to depend on $\hat{v}_i$; because of this, a_i's incentives are unaffected by this redistribution payment, and the mechanism remains strategy-proof.

3. LINEAR VCG REDISTRIBUTION MECHANISMS

We are now ready to introduce the family of linear VCG redistribution mechanisms. Such a mechanism is defined by a vector of constants $c_0, c_1, \ldots, c_{n-1}$. The amount that the mechanism returns to agent a_i is $z_i = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + \ldots + c_{i-1}\hat{v}_{i-1} + c_i\hat{v}_{i+1} + \ldots + c_{n-1}\hat{v}_n$. That is, an agent receives c_0, plus c_1 times the highest bid *other* than the agent's own bid, plus c_2 times the second-highest other bid, *etc.* The mechanism is strategy-proof, because for all i, z_i is independent of $\hat{v}_i$. Also, the mechanism is anonymous.

It is helpful to see the entire list of redistribution payments:

$z_1 = c_0 + c_1\hat{v}_2 + c_2\hat{v}_3 + c_3\hat{v}_4 + \ldots + c_{n-2}\hat{v}_{n-1} + c_{n-1}\hat{v}_n$
$z_2 = c_0 + c_1\hat{v}_1 + c_2\hat{v}_3 + c_3\hat{v}_4 + \ldots + c_{n-2}\hat{v}_{n-1} + c_{n-1}\hat{v}_n$
$z_3 = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_4 + \ldots + c_{n-2}\hat{v}_{n-1} + c_{n-1}\hat{v}_n$
$z_4 = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-1} + c_{n-1}\hat{v}_n$
$\ldots$
$z_i = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + \ldots + c_{i-1}\hat{v}_{i-1} + c_i\hat{v}_{i+1} + \ldots + c_{n-1}\hat{v}_n$
$\ldots$
$z_{n-2} = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-1} + c_{n-1}\hat{v}_n$
$z_{n-1} = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-2} + c_{n-1}\hat{v}_n$
$z_n = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-2} + c_{n-1}\hat{v}_{n-1}$

[1]We thank Rakesh Vohra for pointing us to Moulin's working paper.

Not all choices of the constants $c_0, \ldots, c_{n-1}$ produce a mechanism that is individually rational, and not all choices of the constants produce a mechanism that never incurs a deficit. Hence, to obtain these properties, we need to place some constraints on the constants.

To satisfy the individual rationality criterion, each agent's utility should always be non-negative. An agent that does not win a unit obtains a utility that is equal to the agent's redistribution payment. An agent that wins a unit obtains a utility that is equal to the agent's valuation for the unit, minus the VCG payment $\hat{v}_{m+1}$, plus the agent's redistribution payment.

Consider agent a_n, the agent with the lowest bid. Since this agent does not win an item ($m < n$), her utility is just her redistribution payment z_n. Hence, for the mechanism to be individually rational, the c_i must be such that z_n is always nonnegative. If the c_i have this property, then it actually follows that z_i is nonnegative for *every* i, for the following reason. Suppose there exists some $i < n$ and some vector of bids $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_n \geq 0$ such that $z_i < 0$. Then, consider the bid vector that results from replacing $\hat{v}_j$ by $\hat{v}_{j+1}$ for all $j \geq i$, and letting $\hat{v}_n = 0$. If we omit $\hat{v}_n$ from this vector, the same vector results that results from omitting $\hat{v}_i$ from the original vector. Therefore, a_n's redistribution payment under the new vector should be the same as a_i's redistribution payment under the old vector—but this payment is negative.

If all redistribution payments are always nonnegative, then the mechanism must be individually rational (because the VCG mechanism is individually rational, and the redistribution payment only increases an agent's utility). Therefore, the mechanism is individually rational if and only if for any bid vector, $z_n \geq 0$.

To satisfy the non-deficit criterion, the sum of the redistribution payments should be less than or equal to the total VCG payment. So for any bid vector $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_n \geq 0$, the constants c_i should make $z_1 + z_2 + \ldots + z_n \leq m\hat{v}_{m+1}$.

We define the family of linear VCG redistribution mechanisms to be the set of all redistribution mechanisms corresponding to constants c_i that satisfy the above constraints (so that the mechanisms will be individually rational and have the no-deficit property). We now give two examples of mechanisms in this family.

Example 1 (Bailey-Cavallo mechanism): Consider the mechanism corresponding to $c_{m+1} = \frac{m}{n}$ and $c_i = 0$ for all other i. Under this mechanism, each agent receives a redistribution payment of $\frac{m}{n}$ times the $(m+1)$th highest bid from another agent. Hence, $a_1, \ldots, a_{m+1}$ receive a redistribution payment of $\frac{m}{n}\hat{v}_{m+2}$, and the others receive $\frac{m}{n}\hat{v}_{m+1}$. Thus, the total redistribution payment is $(m+1)\frac{m}{n}\hat{v}_{m+2} + (n-m-1)\frac{m}{n}\hat{v}_{m+1}$. This redistribution mechanism is individually rational, because all the redistribution payments are nonnegative, and never incurs a deficit, because $(m+1)\frac{m}{n}\hat{v}_{m+2} + (n-m-1)\frac{m}{n}\hat{v}_{m+1} \leq n\frac{m}{n}\hat{v}_{m+1} = m\hat{v}_{m+1}$. (We note that for this mechanism to make sense, we need $n \geq m+2$.)

Example 2: Consider the mechanism corresponding to $c_{m+1} = \frac{m}{n-m-1}$, $c_{m+2} = -\frac{m(m+1)}{(n-m-1)(n-m-2)}$, and $c_i = 0$ for all other i. In this mechanism, each agent receives a redistribution payment of $\frac{m}{n-m-1}$ times the $(m+1)$th highest reported value from other agents, minus $\frac{m(m+1)}{(n-m-1)(n-m-2)}$ times the $(m+2)$th highest reported value from other agents. Thus, the total redistribution payment is $m\hat{v}_{m+1} - \frac{m(m+1)(m+2)}{(n-m-1)(n-m-2)}\hat{v}_{m+3}$. If $n \geq 2m+3$ (which is equivalent to $\frac{m}{n-m-1} \geq \frac{m(m+1)}{(n-m-1)(n-m-2)}$), then each agent always receives a nonnegative redistribution payment, thus the mechanism is individually rational. Also, the mechanism never incurs a deficit, because the total VCG payment is $m\hat{v}_{m+1}$, which is greater than the amount $m\hat{v}_{m+1} - \frac{m(m+1)(m+2)}{(n-m-1)(n-m-2)}\hat{v}_{m+3}$ that is redistributed.

Which of these two mechanisms is better? Is there another mechanism that is even better? This is what we study in the next section.

4. OPTIMAL REDISTRIBUTION MECHANISMS

Among all linear VCG redistribution mechanisms, we would like to be able to identify the one that redistributes the greatest percentage of the total VCG payment.[2] This is not a well-defined notion: it may be that one mechanism redistributes more on some bid vectors, and another more on other bid vectors. We emphasize that we do not assume that a prior distribution over bidders' valuations is available, so we cannot compare them based on expected redistribution. Below, we study three well-defined ways of comparing redistribution mechanisms: best-case performance, dominance, and worst-case performance.

Best-case performance. One way of evaluating a mechanism is by considering the highest redistribution percentage that it achieves. Consider the previous two examples. For the first example, the total redistribution payment is $(m+1)\frac{m}{n}\hat{v}_{m+2} + (n-m-1)\frac{m}{n}\hat{v}_{m+1}$. When $\hat{v}_{m+2} = \hat{v}_{m+1}$, this is equal to the total VCG payment $m\hat{v}_{m+1}$. Thus, this mechanism redistributes 100% of the total VCG payment in the best case. For the second example, the total redistribution payment is $m\hat{v}_{m+1} - \frac{m(m+1)(m+2)}{(n-m-1)(n-m-2)}\hat{v}_{m+3}$. When $\hat{v}_{m+3} = 0$, this is equal to the total VCG payment $m\hat{v}_{m+1}$. Thus, this mechanism also redistributes 100% of the total VCG payment in the best case.

Moreover, there are actually infinitely many mechanisms that redistribute 100% of the total VCG payment in the best case—for example, any convex combination of the above two will redistribute 100% if both $\hat{v}_{m+2} = \hat{v}_{m+1}$ and $\hat{v}_{m+3} = 0$.

Dominance. Inside the family of linear VCG redistribution mechanisms, we say one mechanism *dominates* another mechanism if the first one redistributes at least as much as the other for *any* bid vector. For the previous two examples, neither dominates the other, because they each redistribute 100% in different cases. It turns out that there is no mechanism in the family that dominates all other mechanisms in the family. For suppose such a mechanism exists. Then, it should dominate both examples above. Consider the *remaining* VCG payment (the VCG payment failed to be redistributed). The remaining VCG payment of the dominant mechanism should be 0 whenever $\hat{v}_{m+2} = \hat{v}_{m+1}$ or $\hat{v}_{m+3} = 0$. Now, the remaining VCG payment is a linear function of the $\hat{v}_i$ (linear redistribution), and therefore also a polynomial function. The above implies that this function can be written as $(\hat{v}_{m+2} - \hat{v}_{m+1})(\hat{v}_{m+3})P(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n)$, where P is a

[2]The percentage redistributed seems the natural criterion to use, among other things because it is scale-invariant: if we multiply all bids by the same positive constant (for example, if we change the units by re-expressing the bids in euros instead of dollars), we would not want the behavior of our mechanism to change.

polynomial function. But since the function must be linear (has degree at most 1), it follows that $P = 0$. Thus, a dominant mechanism would always redistribute all of the VCG payment, which is not possible. (If it were possible, then our worst-case optimal redistribution mechanism would also always redistribute all of the VCG payment, and we will see later that it does not.)

Worst-case performance. Finally, we can evaluate a mechanism by considering the lowest redistribution percentage that it guarantees. For the first example, the total redistribution payment is $(m+1)\frac{m}{n}\hat{v}_{m+2}+(n-m-1)\frac{m}{n}\hat{v}_{m+1}$, which is greater than or equal to $(n-m-1)\frac{m}{n}\hat{v}_{m+1}$. So in the worst case, which is when $\hat{v}_{m+2} = 0$, the percentage redistributed is $\frac{n-m-1}{n}$. For the second example, the total redistribution payment is $m\hat{v}_{m+1} - \frac{m(m+1)(m+2)}{(n-m-1)(n-m-2)}\hat{v}_{m+3}$, which is greater than or equal to $m\hat{v}_{m+1}(1-\frac{(m+1)(m+2)}{(n-m-1)(n-m-2)})$. So in the worst case, which is when $\hat{v}_{m+3} = \hat{v}_{m+1}$, the percentage redistributed is $1 - \frac{(m+1)(m+2)}{(n-m-1)(n-m-2)}$. Since we assume that the number of agents n and the number of units m are known, we can determine which example mechanism has better worst-case performance by comparing the two quantities. When $n = 6$ and $m = 1$, for the first example (Bailey-Cavallo mechanism), the percentage redistributed in the worst case is $\frac{2}{3}$, and for the second example, this percentage is $\frac{1}{2}$, which implies that for this pair of n and m, the first mechanism has better worst-case performance. On the other hand, when $n = 12$ and $m = 1$, for the first example, the percentage redistributed in the worst case is $\frac{5}{6}$, and for the second example, this percentage is $\frac{14}{15}$, which implies that this time the second mechanism has better worst-case performance.

Thus, it seems most natural to compare mechanisms by the percentage of total VCG payment that they redistribute in the worst case. This percentage is undefined when the total VCG payment is 0. To deal with this, technically, we define the worst-case redistribution percentage as the largest k so that the total amount redistributed is at least k times the total VCG payment, for all bid vectors. (Hence, as long as the total amount redistributed is at least 0 when the total VCG payment is 0, these cases do not affect the worst-case percentage.) This corresponds to the following optimization problem:

> **Maximize** k (the percentage redistributed in the worst case)
> **Subject to:**
> For every bid vector $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_n \geq 0$
> $z_n \geq 0$ (individual rationality)
> $z_1 + z_2 + \ldots + z_n \leq m\hat{v}_{m+1}$ (non-deficit)
> $z_1 + z_2 + \ldots + z_n \geq km\hat{v}_{m+1}$ (worst-case constraint)
> We recall that $z_i = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + \ldots + c_{i-1}\hat{v}_{i-1} + c_i\hat{v}_{i+1} + \ldots + c_{n-1}\hat{v}_n$.

5. TRANSFORMATION TO LINEAR PROGRAMMING

The optimization problem given in the previous section can be rewritten as a linear program, based on the following observations.

CLAIM 1. *If $c_0, c_1, \ldots, c_{n-1}$ satisfy both the individual rationality and the non-deficit constraints, then $c_i = 0$ for $i = 0, \ldots, m$.*

PROOF. First, let us prove that $c_0 = 0$. Consider the bid vector in which $\hat{v}_i = 0$ for all i. To obtain individual rationality, we must have $c_0 \geq 0$. To satisfy the non-deficit constraint, we must have $c_0 \leq 0$. Thus we know $c_0 = 0$. Now, if $c_i = 0$ for all i, there is nothing to prove. Otherwise, let $j = min\{i|c_i \neq 0\}$. Assume that $j \leq m$. We recall that we can write the individual rationality constraint as follows: $z_n = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-2} + c_{n-1}\hat{v}_{n-1} \geq 0$ for any bid vector. Let us consider the bid vector in which $\hat{v}_i = 1$ for $i \leq j$ and $\hat{v}_i = 0$ for the rest. In this case $z_n = c_j$, so we must have $c_j \geq 0$. The non-deficit constraint can be written as follows: $z_1 + z_2 + \ldots + z_n \leq m\hat{v}_{m+1}$ for any bid vector. Consider the same bid vector as above. We have $z_i = 0$ for $i \leq j$, because for these bids, the jth highest other bid has value 0, so all the c_i that are nonzero are multiplied by 0. For $i > j$, we have $z_i = c_j$, because the jth highest other bid has value 1, and all lower bids have value 0. So the non-deficit constraint tells us that $c_j(n-j) \leq m\hat{v}_{m+1}$. Because $j \leq m$, $\hat{v}_{m+1} = 0$, so the right hand side is 0. We also have $n - j > 0$ because $j \leq m < n$. So $c_j \leq 0$. Because we have already established that $c_j \geq 0$, it follows that $c_j = 0$; but this is contrary to assumption. So $j > m$. □

Incidentally, this claim also shows that if $m = n - 1$, then $c_i = 0$ for all i. Thus, we are stuck with the VCG mechanism. From here on, we only consider the case where $m < n - 1$.

CLAIM 2. *The individual rationality constraint can be written as follows: $\sum_{i=m+1}^{j} c_i \geq 0$ for $j = m+1, \ldots, n-1$.*

Before proving this claim, we introduce the following lemma.

LEMMA 1. *Given a positive integer k and a set of real constants $s_1, s_2, \ldots, s_k$, ($s_1t_1 + s_2t_2 + \ldots + s_kt_k \geq 0$ for any $t_1 \geq t_2 \geq \ldots \geq t_k \geq 0$) if and only if ($\sum_{i=1}^{j} s_i \geq 0$ for $j = 1, 2, \ldots, k$).*

PROOF. Let $d_i = t_i - t_{i+1}$ for $i = 1, 2, \ldots, k-1$, and $d_k = t_k$. Then ($s_1t_1 + s_2t_2 + \ldots + s_kt_k \geq 0$ for any $t_1 \geq t_2 \geq \ldots \geq t_k \geq 0$) is equivalent to (($\sum_{i=1}^{1} s_i)d_1 + (\sum_{i=1}^{2} s_i)d_2 + \ldots + (\sum_{i=1}^{k} s_i)d_k \geq 0$ for any set of arbitrary non-negative d_j). When $\sum_{i=1}^{j} s_i \geq 0$ for $j = 1, 2, \ldots, k$, the above inequality is obviously true. If for some j, $\sum_{i=1}^{j} s_i < 0$, if we set $d_j > 0$ and $d_i = 0$ for all $i \neq j$, then the above inequality becomes false. So $\sum_{i=1}^{j} s_i \geq 0$ for $j = 1, 2, \ldots, k$ is both necessary and sufficient. □

We are now ready to present the proof of Claim 2.

PROOF. The individual rationality constraint can be written as $z_n = c_0 + c_1\hat{v}_1 + c_2\hat{v}_2 + c_3\hat{v}_3 + \ldots + c_{n-2}\hat{v}_{n-2} + c_{n-1}\hat{v}_{n-1} \geq 0$ for any bid vector $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_{n-1} \geq \hat{v}_n \geq 0$. We have already shown that $c_i = 0$ for $i \leq m$. Thus, the above can be simplified to $z_n = c_{m+1}\hat{v}_{m+1} + c_{m+2}\hat{v}_{m+2} + \ldots + c_{n-2}\hat{v}_{n-2} + c_{n-1}\hat{v}_{n-1} \geq 0$ for any bid vector. By the above lemma, this is equivalent to $\sum_{i=m+1}^{j} c_i \geq 0$ for $j = m+1, \ldots, n-1$. □

CLAIM 3. *The non-deficit constraint and the worst-case constraint can also be written as linear inequalities involving only the c_i and k.*

PROOF. The non-deficit constraint requires that for any bid vector, $z_1 + z_2 + \ldots + z_n \leq m\hat{v}_{m+1}$, where $z_i = c_0 + c_1\hat{v}_1 +$

$c_2\hat{v}_2+\ldots+c_{i-1}\hat{v}_{i-1}+c_i\hat{v}_{i+1}+\ldots+c_{n-1}\hat{v}_n$ for $i=1,2,\ldots,n$. Because $c_i=0$ for $i \le m$, we can simplify this inequality to

$q_{m+1}\hat{v}_{m+1}+q_{m+2}\hat{v}_{m+2}+\ldots+q_n\hat{v}_n \ge 0$
$q_{m+1}=m-(n-m-1)c_{m+1}$
$q_i=-(i-1)c_{i-1}-(n-i)c_i$, for $i=m+2,\ldots,n-1$ (when $m+2>n-1$, this set of equalities is empty)
$q_n=-(n-1)c_{n-1}$

By the above lemma, this is equivalent to $\sum_{i=m+1}^{j} q_i \ge 0$ for $j=m+1,\ldots,n$. So, we can simplify further as follows:

$q_{m+1} \ge 0 \iff (n-m-1)c_{m+1} \le m$
$q_{m+1}+\ldots+q_{m+i} \ge 0 \iff n\sum_{j=m+1}^{j=m+i-1} c_j+(n-m-i)c_{m+i} \le m$ for $i=2,\ldots,n-m-1$
$q_{m+1}+\ldots+q_n \ge 0 \iff n\sum_{j=m+1}^{j=n-1} c_j \le m$

So, the non-deficit constraint can be written as a set of linear inequalities involving only the c_i.

The worst-case constraint can be also written as a set of linear inequalities, by the following reasoning. The worst-case constraint requires that for any bid input $z_1+z_2+\ldots+z_n \ge km\hat{v}_{m+1}$, where $z_i=c_0+c_1\hat{v}_1+c_2\hat{v}_2+\ldots+c_{i-1}\hat{v}_{i-1}+c_i\hat{v}_{i+1}+\ldots+c_{n-1}\hat{v}_n$ for $i=1,2,\ldots,n$. Because $c_i=0$ for $i \le m$, we can simplify this inequality to

$Q_{m+1}\hat{v}_{m+1}+Q_{m+2}\hat{v}_{m+2}+\ldots+Q_n\hat{v}_n \ge 0$
$Q_{m+1}=(n-m-1)c_{m+1}-km$
$Q_i=(i-1)c_{i-1}+(n-i)c_i$, for $i=m+2,\ldots,n-1$
$Q_n=(n-1)c_{n-1}$

By the above lemma, this is equivalent to $\sum_{i=m+1}^{j} Q_i \ge 0$ for $j=m+1,\ldots,n$. So, we can simplify further as follows:

$Q_{m+1} \ge 0 \iff (n-m-1)c_{m+1} \ge km$
$Q_{m+1}+\ldots+Q_{m+i} \ge 0 \iff n\sum_{j=m+1}^{j=m+i-1} c_j+(n-m-i)c_{m+i} \ge km$ for $i=2,\ldots,n-m-1$
$Q_{m+1}+\ldots+Q_n \ge 0 \iff n\sum_{j=m+1}^{j=n-1} c_j \ge km$

So, the worst-case constraint can also be written as a set of linear inequalities involving only the c_i and k. □

Combining all the claims, we see that the original optimization problem can be transformed into the following linear program.

Variables: $c_{m+1},c_{m+2},\ldots,c_{n-1},k$
Maximize k (the percentage redistributed in the worst case)
Subject to:
$\sum_{i=m+1}^{j} c_i \ge 0$ for $j=m+1,\ldots,n-1$
$km \le (n-m-1)c_{m+1} \le m$
$km \le n\sum_{j=m+1}^{j=m+i-1} c_j+(n-m-i)c_{m+i} \le m$ for $i=2,\ldots,n-m-1$
$km \le n\sum_{j=m+1}^{j=n-1} c_j \le m$

6. NUMERICAL RESULTS

For selected values of n and m, we solved the linear program using Glpk (GNU Linear Programming Kit). In the table below, we present the results for a single unit ($m=1$). We present $1-k$ (the percentage of the total VCG payment that is not redistributed by the worst-case optimal mechanism in the worst case) instead of k in the second column because writing k would require too many significant digits. Correspondingly, the third column displays the percentage of the total VCG payment that is not redistributed by the Bailey-Cavallo mechanism in the worst case (which is equal to $\frac{2}{n}$).

n	$1-k$	**Bailey − Cavallo Mechanism**
3	66.7%	66.7%
4	42.9%	50.0%
5	26.7%	40.0%
6	16.1%	33.3%
7	9.52%	28.6%
8	5.51%	25.0%
9	3.14%	22.2%
10	1.76%	20.0%
20	$3.62e-5$	10.0%
30	$5.40e-8$	$6.67e-2$
40	$7.09e-11$	$5.00e-2$

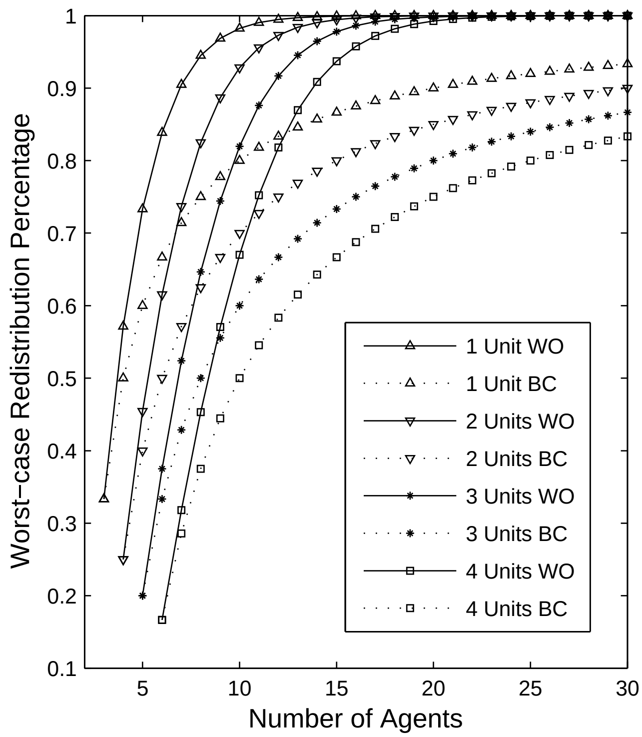

Figure 1: A comparison of the worst-case optimal mechanism (WO) and the Bailey-Cavallo mechanism (BC).

The worst-case optimal mechanism significantly outperforms the Bailey-Cavallo mechanism in the worst case. Perhaps more surprisingly, the worst-case optimal mechanism sometimes does better in the *worst* case than the Bailey-Cavallo mechanism does on *average*, as the following example shows.

Recall that the total redistribution payment of the Bailey-Cavallo mechanism is $(m+1)\frac{m}{n}\hat{v}_{m+2}+(n-m-1)\frac{m}{n}\hat{v}_{m+1}$. For the single-unit case, this simplifies to $\frac{2}{n}\hat{v}_3+\frac{n-2}{n}\hat{v}_2$. Hence the percentage of the total VCG payment that is not redistributed is $\frac{\hat{v}_2-\frac{2}{n}\hat{v}_3-\frac{n-2}{n}\hat{v}_2}{\hat{v}_2}=\frac{2}{n}-\frac{2}{n}\frac{\hat{v}_3}{\hat{v}_2}$, which has an expected value of $E(\frac{2}{n}-\frac{2}{n}\frac{\hat{v}_3}{\hat{v}_2})=\frac{2}{n}-\frac{2}{n}E\frac{\hat{v}_3}{\hat{v}_2}$. Suppose the bid values are drawn from a uniform distribution over $[0,1]$. The theory of order statistics tells us that the joint probability density function of $\hat{v}_2$ and $\hat{v}_3$ is $f(\hat{v}_3, \hat{v}_2) = n(n-1)(n-2)\hat{v}_3^{n-3}(1-\hat{v}_2)$ for $\hat{v}_2 \geq \hat{v}_3$. Now, $E\frac{\hat{v}_3}{\hat{v}_2} = \int_0^1 \int_0^{\hat{v}_2} \frac{\hat{v}_3}{\hat{v}_2} f(\hat{v}_3, \hat{v}_2) d\hat{v}_3 d\hat{v}_2 = \frac{n-2}{n-1}$. So, the expected value of the remaining percentage is $\frac{2}{n} - \frac{2}{n}\frac{n-2}{n-1} = \frac{2}{n(n-1)}$. For $n = 20$, this is $5.26e-3$, whereas the remaining percentage for the worst-case optimal mechanism is $3.62e-5$ in the *worst* case.

Let us present the optimal solution for the case $n = 5$ in detail. By solving the above linear program, we find that the optimal values for the c_i are $c_2 = \frac{11}{45}, c_3 = -\frac{1}{9}$, and $c_4 = \frac{1}{15}$. That is, the redistribution payment received by each agent is: $\frac{11}{45}$ times the second highest bid among the other agents, minus $\frac{1}{9}$ times the third highest bid among the other agents, plus $\frac{1}{15}$ times the fourth highest bid among the other agents.

The total amount redistributed is $\frac{11}{15}\hat{v}_2 + \frac{4}{15}\hat{v}_3 - \frac{4}{15}\hat{v}_4 + \frac{4}{15}\hat{v}_5$; in the worst case, $\frac{11}{15}\hat{v}_2$ is redistributed. Hence, the percentage of the total VCG payment that is not redistributed is never more than $\frac{4}{15} = 26.7\%$.

Finally, we compare the worst-case optimal mechanism to the Bailey-Cavallo mechanism for $m = 1, 2, 3, 4$, $n = m + 2, \ldots, 30$. These results are in Figure 1.

We see that for any m, when $n = m + 2$, the worst-case optimal mechanism has the same worst-case performance as the Bailey-Cavallo mechanism (actually, in this case, the worst-case optimal mechanism is identical to the Bailey-Cavallo mechanism). When $n > m + 2$, the worst-case optimal mechanism outperforms the Bailey-Cavallo mechanism (in the worst case).

7. ANALYTICAL CHARACTERIZATION OF THE WORST-CASE OPTIMAL MECHANISM

We recall that our linear program has the following form:

Variables: $c_{m+1}, c_{m+2}, \ldots, c_{n-1}, k$
Maximize k (the percentage redistributed in the worst case)
Subject to:
$\sum_{i=m+1}^{j} c_i \geq 0$ for $j = m+1, \ldots, n-1$
$km \leq (n-m-1)c_{m+1} \leq m$
$km \leq n\sum_{j=m+1}^{j=m+i-1} c_j + (n-m-i)c_{m+i} \leq m$ for $i = 2, \ldots, n-m-1$
$km \leq n\sum_{j=m+1}^{j=n-1} c_j \leq m$

A linear program has no solution if and only if either the objective is unbounded, or the constraints are contradictory (there is no feasible solution). It is easy to see that k is bounded above by 1 (redistributing more than 100% violates the non-deficit constraint). Also, a feasible solution always exists, for example, $k = 0$ and $c_i = 0$ for all i. So an optimal solution always exists. Observe that the linear program model depends only on the number of agents n and the number of units m. Hence the optimal solution is a function of n and m. It turns out that this optimal solution can be analytically characterized as follows.

THEOREM 1. *For any m and n with $n \geq m+2$, the worst-case optimal mechanism (among linear VCG redistribution mechanisms) is unique. For this mechanism, the percentage redistributed in the worst case is*

$$k^* = 1 - \frac{\binom{n-1}{m}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}$$

The worst-case optimal mechanism is characterized by the following values for the c_i:

$$c_i^* = \frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{i\sum_{j=m}^{n-1}\binom{n-1}{j}} \frac{1}{\binom{n-1}{i}} \sum_{j=i}^{n-1}\binom{n-1}{j}$$

for $i = m+1, \ldots, n-1$.

It should be noted that we have proved $c_i = 0$ for $i \leq m$ in Claim 1.

PROOF. We first rewrite the linear program as follows. We introduce new variables $x_{m+1}, x_{m+2}, \ldots, x_{n-1}$, defined by $x_j = \sum_{i=m+1}^{j} c_i$ for $j = m+1, \ldots, n-1$. The linear program then becomes:

Variables: $x_{m+1}, x_{m+2}, \ldots, x_{n-1}, k$
Maximize k
Subject to:
$km \leq (n-m-1)x_{m+1} \leq m$
$km \leq (m+i)x_{m+i-1} + (n-m-i)x_{m+i} \leq m$ for $i = 2, \ldots, n-m-1$
$km \leq nx_{n-1} \leq m$
$x_i \geq 0$ for $i = m+1, m+2, \ldots, n-1$

We will prove that for any optimal solution to this linear program, $k = k^*$. Moreover, we will prove that when $k = k^*$, $x_j = \sum_{i=m+1}^{j} c_i^*$ for $j = m+1, \ldots, n-1$. This will prove the theorem.

We first make the following observations:

$(n-m-1)c_{m+1}^*$
$= (n-m-1)\frac{(n-m)\binom{n-1}{m-1}}{(m+1)\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{m+1}}\sum_{j=m+1}^{n-1}\binom{n-1}{j}$
$= (n-m-1)\frac{(n-m)\binom{n-1}{m-1}}{(m+1)\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{m+1}}(\sum_{j=m}^{n-1}\binom{n-1}{j} - \binom{n-1}{m})$
$= (n-m-1)\frac{m}{n-m-1} - (n-m-1)\frac{m\binom{n-1}{m}}{(n-m-1)\sum_{j=m}^{n-1}\binom{n-1}{j}}$
$= m - (1-k^*)m = k^*m$

For $i = m+1, \ldots, n-2$,
$ic_i^* + (n-i-1)c_{i+1}^*$
$= i\frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{i\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{i}}\sum_{j=i}^{n-1}\binom{n-1}{j} +$
$(n-i-1)\frac{(-1)^{i+m}(n-m)\binom{n-1}{m-1}}{(i+1)\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{i+1}}\sum_{j=i+1}^{n-1}\binom{n-1}{j}$
$= \frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{i}}\sum_{j=i}^{n-1}\binom{n-1}{j} -$
$(n-i-1)\frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{(i+1)\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{i+1}{\binom{n-1}{i}(n-i-1)}\sum_{j=i+1}^{n-1}\binom{n-1}{j}$
$= \frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}$
$= (-1)^{i+m-1}m(1-k^*)$

Finally,
$(n-1)c_{n-1}^*$
$= (n-1)\frac{(-1)^{n+m}(n-m)\binom{n-1}{m-1}}{(n-1)\sum_{j=m}^{n-1}\binom{n-1}{j}}\frac{1}{\binom{n-1}{n-1}}\sum_{j=n-1}^{n-1}\binom{n-1}{j}$
$= (-1)^{m+n}m(1-k^*)$

Summarizing the above, we have:
$(n-m-1)c_{m+1}^* = k^*m$
$(m+1)c_{m+1}^* + (n-m-2)c_{m+2}^* = m(1-k^*)$
$(m+2)c_{m+2}^* + (n-m-3)c_{m+3}^* = -m(1-k^*)$
$(m+3)c_{m+3}^* + (n-m-4)c_{m+4}^* = m(1-k^*)$
$\vdots$

$(n-3)c^*_{n-3} + 2c^*_{n-2} = (-1)^{m+n-2}m(1-k^*)$
$(n-2)c^*_{n-2} + c^*_{n-1} = (-1)^{m+n-1}m(1-k^*)$
$(n-1)c^*_{n-1} = (-1)^{m+n}m(1-k^*)$

Let $x^*_j = \sum_{i=m+1}^{j} c^*_i$ for $j = m+1, m+2, \dots, n-1$, the first equation in the above tells us that
$(n-m-1)x^*_{m+1} = k^*m$.

By adding the first two equations, we get
$(m+2)x^*_{m+1} + (n-m-2)x^*_{m+2} = m$
By adding the first three equations, we get
$(m+3)x^*_{m+2} + (n-m-3)x^*_{m+3} = k^*m$
By adding the first i equations, where $i = 2, \dots, n-m-1$, we get
$(m+i)x^*_{m+i-1} + (n-m-i)x^*_{m+i} = m$ if i is even
$(m+i)x^*_{m+i-1} + (n-m-i)x^*_{m+i} = k^*m$ if i is odd
Finally by adding all the equations, we get
$nx^*_{n-1} = m$ if $n-m$ is even;
$nx^*_{n-1} = k^*m$ if $n-m$ is odd.

Thus, for all of the constraints other than the nonnegativity constraints, we have shown that they are satisfied by setting $x_j = x^*_j = \sum_{i=m+1}^{j} c^*_i$ and $k = k^*$. We next show that the nonnegativity constraints are satisfied by these settings as well.

For $m+1 \le i, i+1 \le n-1$, we have
$\frac{1}{i}\frac{\sum_{j=i}^{n-1}\binom{n-1}{j}}{\binom{n-1}{i}} = \frac{1}{i}\sum_{j=i}^{n-1}\frac{i!(n-1-i)!}{j!(n-1-j)!} \ge \frac{1}{i+1}\sum_{j=i}^{n-2}\frac{i!(n-1-i)!}{j!(n-1-j)!} \ge$

$\frac{1}{i+1}\sum_{j=i}^{n-2}\frac{(i+1)!(n-1-i-1)!}{(j+1)!(n-1-j-1)!} = \frac{1}{i+1}\frac{\sum_{j=i+1}^{n-1}\binom{n-1}{j}}{\binom{n-1}{i+1}}$

This implies that the absolute value of c^*_i is decreasing as i increases (if the c^* contains more than one number). We further observe that the sign of c^*_i alternates, with the first element c^*_{m+1} positive. So $x^*_j = \sum_{i=m+1}^{j} c^*_i \ge 0$ for all j. Thus, we have shown that these $x_i = x^*_i$ together with $k = k^*$ form a feasible solution of the linear program. We proceed to show that it is in fact the unique optimal solution.

First we prove the following claim:

Claim 4. *If $\hat{k}, \hat{x}_i$, $i = m+1, m+2, \dots, n-1$ satisfy the following inequalities:*

$$\hat{k}m \le (n-m-1)\hat{x}_{m+1} \le m$$
$$\hat{k}m \le (m+i)\hat{x}_{m+i-1} + (n-m-i)\hat{x}_{m+i} \le m \text{ for } i = 2, \dots, n-m-1$$
$$\hat{k}m \le n\hat{x}_{n-1} \le m$$
$$\hat{k} \ge k^*$$

Then we must have that $\hat{x}_i = \hat{x}^*_i$ and $\hat{k} = k^*$.

Proof of claim. Consider the first inequality. We know that $(n-m-1)x^*_{m+1} = k^*m$, so $(n-m-1)\hat{x}_{m+1} \ge \hat{k}m \ge k^*m = (n-m-1)x^*_{m+1}$. It follows that $\hat{x}_{m+1} \ge x^*_{m+1}$ $(n-m-1 \ne 0)$.

Now, consider the next inequality for $i = 2$. We know that $(m+2)x^*_{m+1} + (n-m-2)x^*_{m+2} = m$. It follows that $(n-m-2)\hat{x}_{m+2} \le m-(m+2)\hat{x}_{m+1} \le m-(m+2)x^*_{m+1} = (n-m-2)x^*_{m+2}$, so $\hat{x}_{m+2} \le x^*_{m+2}$ $(i = 2 \le n-m-1 \Rightarrow n-m-2 \ne 0)$.

Now consider the next inequality for $i = 3$. We know that $(m+3)x^*_{m+2} + (n-m-3)x^*_{m+3} = m$. It follows that $(n-m-3)\hat{x}_{m+3} \ge \hat{k}m-(m+3)\hat{x}_{m+2} \ge k^*m-(m+3)x^*_{m+2} = (n-m-3)x^*_{m+3}$, so $\hat{x}_{m+3} \ge x^*_{m+3}$ $(i = 3 \le n-m-1 \Rightarrow n-m-3 \ne 0)$.

Proceeding like this all the way up to $i = n-m-1$, we get that $\hat{x}_{m+i} \ge x^*_{m+i}$ if i is odd and $\hat{x}_{m+i} \le x^*_{m+i}$ if i is even. Moreover, if one inequality is strict, then all subsequent inequalities are strict. Now, if we can prove $\hat{x}_{n-1} = x^*_{n-1}$, it would follow that the x^*_i are equal to the $\hat{x}_i$ (which also implies that $\hat{k} = k^*$). We consider two cases:

Case 1: $n-m$ is even. We have: $n-m$ even $\Rightarrow n-m-1$ odd $\Rightarrow \hat{x}_{n-1} \ge x^*_{n-1}$. We also have: $n-m$ even $\Rightarrow nx^*_{n-1} = m$. Combining these two, we get $m = nx^*_{n-1} \le n\hat{x}_{n-1} \le m \Rightarrow \hat{x}_{n-1} = x^*_{n-1}$.

Case 2: $n-m$ is odd. In this case, we have $\hat{x}_{n-1} \le x^*_{n-1}$, and $nx^*_{n-1} = k^*m$. Then, we have: $k^*m \le \hat{k}m \le n\hat{x}_{n-1} \le nx^*_{n-1} = k^*m \Rightarrow \hat{x}_{n-1} = x^*_{n-1}$.

This completes the proof of the claim. □

It follows that if $\hat{k}, \hat{x}_i$, $i = m+1, m+2, \dots, n-1$ is a feasible solution and $\hat{k} \ge k^*$, then since all the inequalities in Claim 4 are satisfied, we must have $\hat{x}_i = x^*_i$ and $\hat{k} = k^*$. Hence no other feasible solution is as good as the one described in the theorem. □

Knowing the analytical characterization of the worst-case optimal mechanism provides us with at least two major benefits. First, using these formulas is computationally more efficient than solving the linear program using a general-purpose solver. Second, we can derive the following corollary.

Corollary 1. *If the number of units m is fixed, then as the number of agents n increases, the worst-case percentage redistributed linearly converges to 1, with a rate of convergence $\frac{1}{2}$. (That is, $\lim_{n\to\infty}\frac{1-k^*_{n+1}}{1-k^*_n} = \frac{1}{2}$. That is, in the limit, the percentage that is not redistributed halves for every additional agent.)*

We note that this is consistent with the experimental data for the single-unit case, where the worst-case remaining percentage roughly halves each time we add another agent. The worst-case percentage that is redistributed under the Bailey-Cavallo mechanism also converges to 1 as the number of agents goes to infinity, but the convergence is much slower—it does not converge linearly (that is, letting k^C_n be the percentage redistributed by the Bailey-Cavallo mechanism in the worst case for n agents, $\lim_{n\to\infty}\frac{1-k^C_{n+1}}{1-k^C_n} = lim_{n\to\infty}\frac{n}{n+1} = 1$). We now present the proof of the corollary.

Proof. When the number of agents is n, the worst-case percentage redistributed is $k^*_n = 1 - \frac{\binom{n-1}{m}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}$. When the number of agents is $n+1$, the percentage becomes $k^*_{n+1} = 1 - \frac{\binom{n}{m}}{\sum_{j=m}^{n}\binom{n}{j}}$. For n sufficiently large, we will have $2^n - mn^{m-1} > 0$, and hence $\frac{1-k^*_{n+1}}{1-k^*_n} = \frac{\binom{n}{m}\sum_{j=m}^{n-1}\binom{n-1}{j}}{\binom{n-1}{m}\sum_{j=m}^{n}\binom{n}{j}} = \frac{n}{n-m}\frac{2^{n-1}-\sum_{j=0}^{m-1}\binom{n-1}{j}}{2^n-\sum_{j=0}^{m-1}\binom{n}{j}}$, and $\frac{n}{n-m}\frac{2^{n-1}-m(n-1)^{m-1}}{2^n} \le \frac{1-k^*_{n+1}}{1-k^*_n} \le \frac{n}{n-m}\frac{2^{n-1}}{2^n-mn^{m-1}}$ (because $\binom{n}{j} \le n^i$ if $j \le i$).

Since we have
$\lim_{n\to\infty}\frac{n}{n-m}\frac{2^{n-1}-m(n-1)^{m-1}}{2^n} = \frac{1}{2}$, and
$\lim_{n\to\infty}\frac{n}{n-m}\frac{2^{n-1}}{2^n-mn^{m-1}} = \frac{1}{2}$,
it follows that $\lim_{n\to\infty}\frac{1-k^*_{n+1}}{1-k^*_n} = \frac{1}{2}$. □

8. WORST-CASE OPTIMALITY OUTSIDE THE FAMILY

In this section, we prove that the worst-case optimal redistribution mechanism among linear VCG redistribution mechanisms is in fact optimal (in the worst case) among *all* redistribution mechanisms that are deterministic, anonymous, strategy-proof, efficient and satisfy the non-deficit constraint. Thus, restricting our attention to linear VCG redistribution mechanisms did not come at a loss.

To prove this theorem, we need the following lemma. This lemma is not new: it was informally stated by Cavallo [3]. For completeness, we present it here with a detailed proof.

Lemma 2. *A VCG redistribution mechanism is deterministic, anonymous and strategy-proof if and only if there exists a function $f : \mathbf{R^{n-1}} \to \mathbf{R}$, so that the redistribution payment z_i received by a_i satisfies*

$$z_i = f(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$$

for all i and all bid vectors.

Proof. First, let us prove the "only if" direction, that is, if a VCG redistribution mechanism is deterministic, anonymous and strategy-proof then there exists a deterministic function $f : \mathbf{R^{n-1}} \to \mathbf{R}$, which makes

$$z_i = f(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$$

for all i and all bid vectors.

If a VCG redistribution mechanism is deterministic and anonymous, then for any bid vector $\hat{v}_1 \geq \hat{v}_2 \geq \ldots \geq \hat{v}_n$, the mechanism outputs a unique redistribution payment list: $z_1, z_2, \ldots, z_n$. Let $G : \mathbf{R^n} \to \mathbf{R^n}$ be the function that maps $\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n$ to $z_1, z_2, \ldots, z_n$ for all bid vectors. Let $H(i, x_1, x_2, \ldots, x_n)$ be the ith element of $G(x_1, x_2, \ldots, x_n)$, so that $z_i = H(i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n)$ for all bid vectors and all $1 \leq i \leq n$. Because the mechanism is anonymous, two agents should receive the same redistribution payment if their bids are the same. So, if $\hat{v}_i = \hat{v}_j$, $H(i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n) = H(j, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n)$. Hence, if we let $j = min\{t | \hat{v}_t = \hat{v}_i\}$, then $H(i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n) = H(j, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n)$.

Let us define $K : \mathbf{R^n} \to \mathbf{N} \times \mathbf{R^n}$ as follows: $K(y, x_1, x_2, \ldots, x_{n-1}) = [j, w_1, w_2, \ldots, w_n]$, where $w_1, w_2, \ldots, w_n$ are $y, x_1, x_2, \ldots, x_{n-1}$ sorted in descending order, and $j = min\{t | w_t = y\}$. ($\{t | w_t = y\} \neq \emptyset$ because $y \in \{w_1, w_2, \ldots, w_n\}$). Also let us define $F : \mathbf{R^n} \to \mathbf{R}$ by
$F(\hat{v}_i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$
$= H \circ K(\hat{v}_i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$
$= H(min\{t | \hat{v}_t = \hat{v}_i\}, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n)$
$= H(i, \hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n) = z_i$.
That is, F is the redistribution payment to an agent that bids $\hat{v}_i$ when the other bids are $\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n$.

Since our mechanism is required to be strategy-proof, and the space of valuations is unrestricted, z_i should be independent of $\hat{v}_i$ by Lemma 1 in Cavallo [3]. Hence, we can simply ignore the first variable input to F; let $f(x_1, x_2, \ldots, x_{n-1}) = F(0, x_1, x_2, \ldots, x_{n-1})$. So, we have for all bid vectors and i, $z_i = f(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$. This completes the proof for the "only if" direction.

For the "if" direction, if the redistribution payment received by a_i satisfies $z_i = f(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$ for all bid vectors and i, then this is clearly a deterministic and anonymous mechanism. To prove strategy-proofness, we observe that because an agent's redistribution payment is not affected by her own bid, her incentives are the same as in the VCG mechanism, which is strategy-proof. □

Now we are ready to introduce the next theorem:

Theorem 2. *For any m and n with $n \geq m+2$, the worst-case optimal mechanism among the family of linear VCG redistribution mechanisms is worst-case optimal among all mechanisms that are deterministic, anonymous, strategy-proof, efficient and satisfy the non-deficit constraint.*

While we needed individual rationality earlier in the paper, this theorem does not mention it, that is, we can not find a mechanism with better worst-case performance even if we sacrifice individual rationality. (The worst-case optimal linear VCG redistribution mechanism is of course individually rational.)

Proof. Suppose there is a redistribution mechanism (when the number of units is m and the number of agents is n) that satisfies all of the above properties and has a better worst-case performance than the worst-case optimal linear VCG redistribution mechanism, that is, its worst-case redistribution percentage $\hat{k}$ is strictly greater than k^*.

By Lemma 2, for this mechanism, there is a function $f : \mathbf{R^{n-1}} \to \mathbf{R}$ so that $z_i = f(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_{i-1}, \hat{v}_{i+1}, \ldots, \hat{v}_n)$ for all i and all bid vectors. We first prove that f has the following properties.

Claim 5. *$f(1, 1, \ldots, 1, 0, 0, \ldots, 0) = 0$ if the number of 1s is less than or equal to m.*

Proof of claim. We assumed that for this mechanism, the worst-case redistribution percentage satisfies $\hat{k} > k^* \geq 0$. If the total VCG payment is x, the total redistribution payment should be in $[\hat{k}x, x]$ (non-deficit criterion). Consider the case where all agents bid 0, so that the total VCG payment is also 0. Hence, the total redistribution payment should be in $[\hat{k} \cdot 0, 0]$—that is, it should be 0. Hence every agent's redistribution payment $f(0, 0, \ldots, 0)$ must be 0.

Now, let $t_i = f(1, 1, \ldots, 1, 0, 0, \ldots, 0)$ where the number of 1s equals i. We proved $t_0 = 0$. If $t_{n-1} = 0$, consider the bid vector where everyone bids 1. The total VCG payment is m and the total redistribution payment is $nf(1, 1, \ldots, 1) = nt_{n-1} = 0$. This corresponds to 0% redistribution, which is contrary to our assumption that $\hat{k} > k^* \geq 0$. Now, consider $j = min\{i | t_i \neq 0\}$ (which is well-defined because $t_{n-1} \neq 0$). If $j > m$, the property is satisfied. If $j \leq m$, consider the bid vector where $\hat{v}_i = 1$ for $i \leq j$ and $\hat{v}_i = 0$ for all other i. Under this bid vector, the first j agents each get redistribution payment $t_{j-1} = 0$, and the remaining $n - j$ agents each get t_j. Thus, the total redistribution payment is $(n - j)t_j$. Because the total VCG payment for this bid vector is 0, we must have $(n - j)t_j = 0$. So $t_j = 0$ ($j \leq m < n$). But this is contrary to the definition of j. Hence $f(1, 1, \ldots, 1, 0, 0, \ldots, 0) = 0$ if the number of 1s is less than or equal to m. □

Claim 6. *f satisfies the following inequalities:*

$$\hat{k}m \leq (n - m - 1)t_{m+1} \leq m$$

$$\hat{k}m \leq (m + i)t_{m+i-1} + (n - m - i)t_{m+i} \leq m \text{ for } i = 2, 3, \ldots, n - m - 1$$

$$\hat{k}m \leq nt_{n-1} \leq m$$

Here t_i is defined as in the proof of Claim 5.

Proof of claim. For $j = m+1, \dots, n$, consider the bid vectors where $\hat{v}_i = 1$ for $i \leq j$ and $\hat{v}_i = 0$ for all other i. These bid vectors together with the non-deficit constraint and worst-case constraint produce the above set of inequalities: for example, when $j = m+1$, we consider the bid vector $\hat{v}_i = 1$ for $i \leq m+1$ and $\hat{v}_i = 0$ for all other i. The first $m+1$ agents each receive a redistribution payment of $t_m = 0$, and all other agents each receive t_{m+1}. Thus, the total VCG redistribution is $(n-m-1)t_{m+1}$. The non-deficit constraint gives $(n-m-1)t_{m+1} \leq m$ (because the total VCG payment is m). The worst-case constraint gives $(n-m-1)t_{m+1} \geq \hat{k}m$. Combining these two, we get the first inequality. The other inequalities can be obtained in the same way. □

We now observe that the inequalities in Claim 6, together with $\hat{k} \geq k^*$, are the same as those in Claim 4 (where the t_i are replaced by the $\hat{x}_i$). Thus, we can conclude that $\hat{k} = k^*$, which is contrary to our assumption $\hat{k} > k^*$. Hence no mechanism satisfying all the listed properties has a redistribution percentage greater than k^* in the worst case. □

So far we have only talked about the case where $n \geq m+2$. For the purpose of completeness, we provide the following claim for the $n = m+1$ case.

Claim 7. *For any m and n with $n = m+1$, the original VCG mechanism (that is, redistributing nothing) is (uniquely) worst-case optimal among all redistribution mechanisms that are deterministic, anonymous, strategy-proof, efficient and satisfy the non-deficit constraint.*

We recall that when $n = m+1$, Claim 1 tells us that the only mechanism inside the family of linear redistribution mechanisms is the original VCG mechanism, so that this mechanism is automatically worst-case optimal inside this family. However, to prove the above claim, we need to show that it is worst-case optimal among *all* redistribution mechanisms that have the desired properties.

Proof. Suppose a redistribution mechanism exists that satisfies all of the above properties and has a worst-case performance as good as the original VCG mechanism, that is, its worst-case redistribution percentage is greater than or equal to 0. This implies that the total redistribution payment of this mechanism is always nonnegative.

By Lemma 2, for this mechanism, there is a function $f : \mathbf{R^{n-1}} \rightarrow \mathbf{R}$ so that $z_i = f(\hat{v}_1, \hat{v}_2, \dots, \hat{v}_{i-1}, \hat{v}_{i+1}, \dots, \hat{v}_n)$ for all i and all bid vectors. We will prove that $f(x_1, x_2, \dots, x_{n-1}) = 0$ for all $x_1 \geq x_2 \geq \dots \geq x_{n-1} \geq 0$.

First, consider the bid vector where $\hat{v}_i = 0$ for all i. Here, each agent receives a redistribution payment $f(0, 0, \dots, 0)$. The total redistribution payment is then $nf(0, 0, \dots, 0)$, which should be both greater than or equal to 0 (by the above observation) as well less than or equal to 0 (using the non-deficit criterion and the fact that the total VCG payment is 0). It follows that $f(0, 0, \dots, 0) = 0$. Now, let us consider the bid vector where $\hat{v}_1 = x_1 \geq 0$ and $\hat{v}_i = 0$ for all other i. For this bid vector, the agent with the highest bid receives a redistribution payment of $f(0, 0, \dots, 0) = 0$, and the other $n-1$ agents each receive $f(x_1, 0, \dots, 0)$. By the same reasoning as above, the total redistribution payment should be both greater than or equal to 0 and less than or equal to 0, hence $f(x_1, 0, \dots, 0) = 0$ for all $x_1 \geq 0$.

Proceeding by induction, let us assume $f(x_1, x_2, \dots, x_k, 0, \dots, 0) = 0$ for all $x_1 \geq x_2 \geq \dots \geq x_k \geq 0$, for some $k < n-1$. Consider the bid vector where $\hat{v}_i = x_i$ for $i \leq k+1$, and $\hat{v}_i = 0$ for all other i, where the x_i are arbitrary numbers satisfying $x_1 \geq x_2 \geq \dots \geq x_k \geq x_{k+1} \geq 0$. For the agents with the highest $k+1$ bids, their redistribution payment is specified by f acting on an input with only k non-zero variables. Hence they all receive 0 by induction assumption. The other $n-k-1$ agents each receive $f(x_1, x_2, \dots, x_k, x_{k+1}, 0, \dots, 0)$. The total redistribution payment is then $(n-k-1)f(x_1, x_2, \dots, x_k, x_{k+1}, 0, \dots, 0)$, which should be both greater than or equal to 0, and less than or equal to the total VCG payment. Now, in this bid vector, the lowest bid is 0 because $k+1 < n$. But since $n = m+1$, the total VCG payment is $m\hat{v}_n = 0$. So we have $f(x_1, x_2, \dots, x_k, x_{k+1}, 0, \dots, 0) = 0$ for all $x_1 \geq x_2 \geq \dots \geq x_k \geq x_{k+1} \geq 0$. By induction, this statement holds for all $k < n-1$; when $k+1 = n-1$, we have $f(x_1, x_2, \dots, x_{n-2}, x_{n-1}) = 0$ for all $x_1 \geq x_2 \geq \dots \geq x_{n-2} \geq x_{n-1} \geq 0$. Hence, in this mechanism, the redistribution payment is always 0; that is, the mechanism is just the original VCG mechanism. □

Incidentally, we obtain the following corollary:

Corollary 2. *No VCG redistribution mechanism satisfies all of the following: determinism, anonymity, strategy-proofness, efficiency, and (strong) budget balance. This holds for any $n \geq m+1$.*

Proof. *For the case $n \geq m+2$:* If such a mechanism exists, its worst-case performance would be better than that of the worst-case optimal linear VCG redistribution mechanism, which by Theorem 1 obtains a redistribution percentage strictly less than 1. But Theorem 2 shows that it is impossible to outperform this mechanism in the worst case.

For the case $n = m+1$: If such a mechanism exists, it would perform as well as the original VCG mechanism in the worst case, which implies that it is identical to the VCG mechanism by Claim 7. But the VCG mechanism is not (strongly) budget balanced. □

9. CONCLUSIONS

For allocation problems with one or more items, the well-known Vickrey-Clarke-Groves (VCG) mechanism is efficient, strategy-proof, individually rational, and does not incur a deficit. However, the VCG mechanism is not (strongly) budget balanced: generally, the agents' payments will sum to more than 0. If there is an auctioneer who is selling the items, this may be desirable, because the surplus payment corresponds to revenue for the auctioneer. However, if the items do not have an owner and the agents are merely interested in allocating the items efficiently among themselves, any surplus payment is undesirable, because it will have to flow out of the system of agents. In 2006, Cavallo [3] proposed a mechanism that redistributes some of the VCG payment back to the agents, while maintaining efficiency, strategy-proofness, individual rationality, and the non-deficit property. In this paper, we extended this result in a restricted setting. We studied allocation settings where there are multiple indistinguishable units of a single good, and agents have unit demand. (For this specific setting, Cavallo's mechanism coincides with a mechanism proposed by Bailey in 1997 [2].) Here we proposed a family of mechanisms that redistribute some of the VCG payment back to the agents. All mechanisms in the family are efficient, strategy-proof, individually rational, and never incur a deficit. The family includes the Bailey-Cavallo mechanism as a special case. We then provided an optimization model for finding the optimal mechanism—that is, the mechanism that maximizes redistribution in the worst case—inside the family, and showed how to cast this model as a linear program. We gave both numerical and analytical solutions of this linear program, and the (unique) resulting mechanism shows significant improvement over the Bailey-Cavallo mechanism (in the worst case). Finally, we proved that the obtained mechanism is optimal among *all* anonymous deterministic mechanisms that satisfy the above properties.

One important direction for future research is to try to extend these results beyond multi-unit auctions with unit demand. However, it turns out that in sufficiently general settings, the worst-case optimal redistribution percentage is 0. In such settings, the worst-case criterion provides no guidance in determining a good redistribution mechanism (even redistributing nothing achieves the optimal worst-case percentage), so it becomes necessary to pursue other criteria. Alternatively, one can try to identify other special settings in which positive redistribution in the worst case is possible.

Another direction for future research is to consider whether this mechanism has applications to collusion. For example, in a typical collusive scheme, there is a *bidding ring* consisting of a number of colluders, who submit only a single bid [10, 17]. If this bid wins, the colluders must allocate the item amongst themselves, perhaps using payments—but of course they do not want payments to flow out of the ring.

This work is part of a growing literature on designing mechanisms that obtain good results in the worst case. Traditionally, economists have mostly focused either on designing mechanisms that always obtain certain properties (such as the VCG mechanism), or on designing mechanisms that are optimal with respect to some prior distribution over the agents' preferences (such as the Myerson auction [20] and the Maskin-Riley auction [18] for maximizing expected revenue). Some more recent papers have focused on designing mechanisms for profit maximization using worst-case competitive analysis (*e.g.* [9, 1, 15, 8]). There has also been growing interest in the design of *online* mechanisms [7] where the agents arrive over time and decisions must be taken before all the agents have arrived. Such work often also takes a worst-case competitive analysis approach [14, 13]. It does not appear that there are direct connections between our work and these other works that focus on designing mechanisms that perform well in the worst case. Nevertheless, it seems likely that future research will continue to investigate mechanism design for the worst case, and hopefully a coherent framework will emerge.

10. REFERENCES

[1] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of auctions. *STOC*, 619–625, 2005.

[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. *AAMAS*, 882–889, 2006.

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

[5] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. *AMEC*, 30–43, 2005.

[6] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. *JCSS*, 63:21–41, 2001.

[7] E. Friedman and D. Parkes. Pricing WiFi at Starbucks – Issues in online mechanism design. *EC*, 240–241, 2003.

[8] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

[9] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. *SODA*, 735–744, 2001.

[10] D. A. Graham and R. C. Marshall. Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy*, 95(6):1217–1239, 1987.

[11] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

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

[13] M. T. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. C. Parkes. Online auctions with re-usable goods. *EC*, 165–174, 2005.

[14] M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. *EC*, 71–80, 2004.

[15] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. *EC*, 175–182, 2005.

[16] 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.

[17] K. Leyton-Brown, Y. Shoham, and M. Tennenholtz. Bidding clubs in first-price auctions. *AAAI*, 373–378, 2002.

[18] E. Maskin and J. Riley. Optimal multi-unit auctions. In F. Hahn, editor, *The Economics of Missing Markets, Information, and Games*, chapter 14, 312–335. Clarendon Press, Oxford, 1989.

[19] H. Moulin. Efficient and strategy-proof assignment with a cheap residual claimant. Working paper, March 2007.

[20] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[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. *IJCAI*, 1161–1168, 2001.

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

Undominated VCG Redistribution Mechanisms

Mingyu Guo
Dept. of Computer Science
Duke University
Durham, NC, USA

Vincent Conitzer
Dept. of Computer Science
Duke University
Durham, NC, USA

Abstract

For allocation problems, the well-known Vickrey-Clarke-Groves (VCG) mechanism is efficient, incentive compatible, individually rational, and does not incur a deficit. However, the VCG mechanism is not (strongly) budget balanced: generally, the agents' payments will sum to more than 0. Very recently, several mechanisms have been proposed that *redistribute* a significant percentage of the VCG payments back to the agents while maintaining the other properties. This increases the agents' utilities.

One redistribution mechanism *dominates* another if it always redistributes at least as much to each agent (and sometimes more). In this paper, we provide a characterization of undominated redistribution mechanisms. We also propose several techniques that take a dominated redistribution mechanism as input, and produce as output another redistribution mechanism that dominates the original. One technique immediately produces an undominated redistribution mechanism that is not necessarily anonymous. Another technique preserves anonymity, and repeated application results in an undominated redistribution mechanism in the limit. We show experimentally that these techniques improve the known redistribution mechanisms.

1 Introduction

Many important problems in artificial intelligence can be seen as resource allocation problems, in which we want to allocate resources (or *items*) to the agents that value them the most. However, agents' valuations are private knowledge, and self-interested agents will lie about their valuations if this is to their benefit. One solution is to *auction* off the items. By carefully deciding how much the winning agents pay, it is possible to create an auction in which bidders have no incentive to lie about their valuations. The best-known way of doing so is to use the *VCG mechanism* [13, 3, 8] for determining the payments. This mechanism has various desirable properties. One disadvantage of this approach is that the payments that the agents make flow out of the system, and this reduces the agents' utilities. To minimize this disadvantage, very recently, several mechanisms have been proposed that *redistribute* a significant percentage of the VCG payments back to the agents while maintaining the other properties [2, 9]. In this paper, we continue this line of research. We introduce several general techniques that can be applied to *any* redistribution mechanism to obtain a new mechanism. The resulting mechanism redistributes at least as much, and typically more, for any prior distribution over agents' valuations.

2 Mechanism Design Background

In this section, we briefly review basic elements of mechanism design, as well as redistribution mechanisms.

2.1 Mechanism Design Basics

A typical setting in mechanism design is given by the following. There is a set of *agents* $I = \{1, 2, \ldots, n\}$, and a set of possible *outcomes* O. For example, in a *combinatorial auction*, a set of items S is (simultaneously) for sale, and the set of outcomes is the set of all possible *allocations* of the items to the agents (also known as *bidders*). An allocation is given by a function $a : S \to \{0, 1, \ldots, n\}$, where for any $s \in S$, $a(s)$ is the bidder that obtains s (if no bidder obtains s, then $a(s) = 0$[1]). Each agent has privately held preferences over the outcomes. As is common in mechanism design, these preferences are represented as follows. For each agent i, there is a set of possible *types* Θ_i. Some θ_i is the agent's actual type; this is information that is private to i. There is also a (commonly known) *valuation function* $v_i : \Theta_i \times O \to \mathbb{R}$. For example, in a single-item auction, $\theta_i \in \mathbb{R}$ is agent i's valuation for the item, and $v_i(\theta_i, o) = \theta_i$ if o allocates the item to i (and it is 0 otherwise).[2] In a combinatorial auction, in general, θ_i consists of $2^{|S|} - 1$ real numbers, where each number represents the valuation for receiving a certain nonempty *bundle* (subset) of the items. Often, the type space is assumed to be more restricted. For example, if each bidder is only interested

[1] The assumption that items can remain unallocated is known as the *free disposal* assumption.

[2] This is assuming *no externalities*: if an agent does not receive the item, the agent does not care which other agent receives it.

in a single bundle (that is, bidders are *single-minded*), then a type θ_i consists of a pair (S'_i, x_i), where S'_i is the bundle that the bidder is interested in, and $v_i(\theta_i, o) = x_i$ if the bundle that o allocates to i includes S'_i (and it is 0 otherwise). Another special case is a *multi-unit* auction, in which m indistinguishable items are for sale (equivalently, there are multiple units of the same item for sale). Here, a type consists of m real numbers, where the jth number indicates the value for obtaining j units. A special case is a multi-unit auction with *unit demand*, in which each bidder wants to obtain only one unit—that is, all m numbers are always the same, so a type effectively consists of a single number.

In a *(direct-revelation) mechanism*, each agent reports a type $\hat{\theta}_i \in \Theta_i$ (not necessarily equal to θ_i), and based on this, an outcome is chosen, as well as a payment to be made by each agent. Thus, a mechanism is given by an outcome selection function $f : \Theta_1 \times \ldots \times \Theta_n \rightarrow O$, as well as n payment selection functions $\pi_i : \Theta_1 \times \ldots \times \Theta_n \rightarrow \mathbb{R}$. As is common, we assume that preferences are *quasi-linear*, that is, agent i's utility is $u_i(\theta_i, (\hat{\theta_1}, \ldots, \hat{\theta_n})) = v_i(\theta_i, f(\hat{\theta_1}, \ldots, \hat{\theta_n})) - \pi_i(\hat{\theta_1}, \ldots, \hat{\theta_n})$. A mechanism is *(dominant-strategies) incentive compatible* if it is a dominant strategy for each agent to reveal his true type, that is, for all $(\theta_1, \ldots, \theta_n) \in \Theta_1 \times \ldots \times \Theta_n$ and all $\hat{\theta}_i \in \Theta_i$, $u_i(\theta_i, (\theta_1, \ldots, \theta_i, \ldots, \theta_n)) \geq u_i(\theta_i, (\theta_1, \ldots, \hat{\theta}_i, \ldots, \theta_n))$.

Perhaps the most famous mechanism is the *Vickrey-Clarke-Groves (VCG)* mechanism [13, 3, 8]. This mechanism chooses an outcome o^* that maximizes the sum of agents' reported valuations, that is, $o^* \in \arg\max_o \sum_{i=1}^{n} v_i(\hat{\theta}_i, o)$. That is, the mechanism is *efficient*. Then, to determine agent j's payment, it computes an outcome o^*_{-j} that would have been optimal if agent j had not been present, that is, $o^*_{-j} \in \arg\max_o \sum_{i \neq j} v_i(\hat{\theta}_i, o)$. Finally, it determines agent j's payment as $\pi_j(\hat{\theta_1}, \ldots, \hat{\theta_n}) = \sum_{i \neq j} v_i(\hat{\theta}_i, o^*_{-j}) - \sum_{i \neq j} v_i(\hat{\theta}_i, o^*)$. This mechanism is well-known to be incentive compatible. It has several other nice properties. Under certain minimal assumptions (which are satisfied in (combinatorial) auctions with free disposal), it also satisfies:

- *individual rationality*: for all $(\theta_1, \ldots, \theta_n) \in \Theta_1 \times \ldots \times \Theta_n$, for all i, $u_i(\theta_i, (\theta_1, \ldots, \theta_i, \ldots, \theta_n)) \geq 0$. That is, participating in the mechanism does not make anyone worse off.
- *non-deficit*: for all $(\theta_1, \ldots, \theta_n) \in \Theta_1 \times \ldots \times \Theta_n$, $\sum_{i=1}^{n} \pi_i(\theta_1, \ldots, \theta_n) \geq 0$. That is, the mechanism does not need to be subsidized by external funds, because the total payments to agents never exceed the total payments from agents.
- *anonymity*: the mechanism treats all agents the same.

We will assume throughout that we are in a setting where the VCG mechanism obtains all of the above properties.

For single-item auctions, the VCG mechanism coincides with the *second-price sealed-bid auction*, that is, the highest bidder wins the item and pays the bid of the second highest bidder.

2.2 Redistribution Mechanisms

For the VCG mechanism, sometimes, $\sum_{i=1}^{n} \pi_i(\theta_1, \ldots, \theta_n) \neq 0$. That is, the VCG mechanism is not *(strongly) budget balanced*. In general, no mechanism that is budget balanced also satisfies all of efficiency, incentive compatibility, and individual rationality [10, 7, 11]. In light of this impossibility result, several authors have obtained budget balance by sacrificing some of the other desirable properties [1, 6, 12, 5]. Another approach that is perhaps preferable is to use a mechanism that is "more" budget balanced than the VCG mechanism, and maintains all the other desirable properties. One way of trying to design such a mechanism is to redistribute some of the VCG payment back to the agents in a way that will not affect the agents' incentives (so that incentive compatibility is maintained), and that will maintain the other properties. This idea has resulted in a few recent papers on *(VCG) redistribution mechanisms*. Such a mechanism works as follows. First, the agents report their types, and the VCG mechanism is run (so that the efficient outcome is chosen). Second, some of the VCG payments collected in the first step are redistributed back to the agents, in a way that maintains incentive compatibility, individual rationality, and non-deficit. To maintain incentive compatibility, an agent's redistribution payment should not depend on his own reported type [2]. Thus, a redistribution mechanism is defined by a function $r_i : \Theta_1 \times \ldots \times \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_n \rightarrow \mathbb{R}$ for each agent i. That is, letting $\hat{\theta}_{-i}$ be the vector of types submitted by agents other than i, $r_i(\hat{\theta}_{-i})$ indicates the amount redistributed to i. For an anonymous redistribution mechanism, $r_i = r$ for all i.

Let us say that a redistribution mechanism is *feasible* if it satisfies individual rationality and non-deficit. (Efficiency and incentive compatibility follow immediately from the above definition of a redistribution mechanism.) The trivial redistribution mechanism that redistributes nothing is always feasible. In some settings, this is the only feasible redistribution mechanism—for example, in a single-item auction with two bidders.

For example, Cavallo's mechanism [2] is given by $r_i(\theta_{-i}) = (1/n) \min_{\theta_i \in \Theta_i} VCG(\theta_i, \theta_{-i})$, where $VCG(\theta_i, \theta_{-i})$ is the total VCG payment collected for those reports.[3] In the special case of a single-item auction, an agent's redistribution payment is $1/n$ times the second-highest bid among *other* agents' bids. In this

[3] We use θ_i rather than $\hat{\theta}_i$ when there is no need to emphasize the difference between reported and true types (since the mechanism is incentive compatible).

case (or, more generally, multi-unit auctions with unit demand), Cavallo's mechanism coincides with a mechanism proposed by Bailey [1]. For multi-unit auctions with unit demand, we characterized a redistribution mechanism that maximizes the worst-case redistribution percentage [9]. We do not present the (complex) general form of this worst-case optimal (WCO) redistribution mechanism here.

3 Undominated Redistribution Mechanisms

How should we select a redistribution mechanism? In general, we prefer to redistribute as much as possible. However, two redistribution mechanisms may be incomparable in the sense that one redistributes more for one vector of reported types, and the other redistributes more for another vector. In earlier work [9], we focused on maximizing the percentage of VCG payments redistributed in the worst case. However, we only studied multi-unit auctions with unit demand. It turns out that in more general settings, the worst-case redistribution percentage is often 0 (we will see examples shortly). This does not mean that nothing can ever be redistributed, but it does mean that we need to change our criterion.

We will require the following claim for our examples.

Claim 1 *A redistribution mechanism* $\mathbf{r} = (r_1, \ldots, r_n)$ *is* feasible *if and only if for all* i *and all* $\theta_1, \ldots, \theta_n$

$$r_i(\theta_{-i}) \geq 0 \quad (1)$$

$$r_i(\theta_{-i}) \leq \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})\} \quad (2)$$

Here, θ'_{-j} *are the reported types of the agents other than* j *when* θ_i *is replaced by* θ'_i. $VCG(\theta'_i, \theta_{-i})$ *is the total VCG payments for the type vector* $\theta_1, \ldots, \theta_{i-1}, \theta'_i, \theta_{i+1}, \ldots, \theta_n$.

Proof: We first prove the "if" direction. Because the VCG mechanism is individually rational, and by Equation 1 the redistribution can only increase agents' utilities, individual rationality is satisfied. For any $\theta_1, \ldots, \theta_n$, Equation 2 implies that $r_1(\theta_{-1}) \leq VCG(\theta'_1, \theta_{-1}) - \sum_{j \neq 1} r_j(\theta'_{-j})$ for any $\theta'_1 \in \Theta_1$. If we let $\theta'_1 = \theta_1$, we obtain $r_1(\theta_{-1}) + \sum_{j \neq 1} r_j(\theta_{-j}) \leq VCG(\theta_1, \theta_{-1})$. Thus, the non-deficit property holds.

We now prove the "only if" direction. For any i and θ_{-i}, there exists some θ_i such that i will not derive any utility from the allocation. Thus, if $r_i(\theta_{-i}) < 0$, i would have negative utility, contradicting individual rationality. Thus Equation 1 must hold. By the non-deficit property, for any i, any $\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n$, and any θ'_i, we must have $r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta'_{-j}) \leq VCG(\theta'_i, \theta_{-i})$, or equivalently $r_i(\theta_{-i}) \leq VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})$. Since θ'_i is arbitrary, Equation 2 follows. ∎

Example 1. Consider a combinatorial auction with two items $\{a, b\}$ and three bidders $\{1, 2, 3\}$. Bidder 1 bids 10 on the bundle $\{a, b\}$; bidder 2 bids ϵ on $\{a\}$; bidder 3 bids $10 - 2\epsilon$ on $\{b\}$. For sufficiently small ϵ, bidder 1 wins both items and pays $10 - \epsilon$. For any feasible redistribution mechanism $\mathbf{r}$, Equation 1 and Equation 2 together imply $r_i(\theta_{-i}) \leq \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i})\}$. For $\theta'_1 = (\{a, b\}, 0)$ (*i.e.* if 1 had bid 0 on $\{a, b\}$ instead), $VCG(\theta'_1, \theta_{-1}) = 0$, hence it must be that $r_1((\{a\}, \epsilon), (\{b\}, 10 - 2\epsilon)) = 0$ (*i.e.* nothing is redistributed to 1). For $\theta'_2 = (\{a\}, 11)$, $VCG(\theta'_2, \theta_{-2}) = 2\epsilon$, so $r_2((\{a, b\}, 10), (\{b\}, 10 - 2\epsilon)) \leq 2\epsilon$. Finally, for $\theta'_3 = (\{b\}, 0)$, $VCG(\theta'_3, \theta_{-3}) = \epsilon$, so $r_3((\{a, b\}, 10), (\{a\}, \epsilon)) \leq \epsilon$. Hence, the percentage redistributed is at most $\frac{3\epsilon}{10-\epsilon}$, which approaches 0 as ϵ approaches 0. Thus, *every* redistribution mechanism has a worst-case redistribution percentage of 0 in this setting.

If we add any number of additional bidders who bid $(\{a\}, 0)$, then the bounds on the first three bidders' redistribution payments remain the same, and each additional bidder can have a redistribution payment of at most 2ϵ (if any one of them bids more than 10, then the resulting total VCG payment is 2ϵ). By letting $\epsilon \rightarrow 0$, it can be seen that the worst-case percentage redistributed remains 0 for any number of bidders. This is in contrast to the case of multi-unit auctions with unit demand, where additional bidders improve the worst-case redistribution percentage [9].

Example 2. Consider a multi-unit auction with two units and three bidders $\{1, 2, 3\}$. Bidder 1 bids $(0, 10)$ (0 for getting one unit and 10 for getting two units). Bidder 2 bids (ϵ, ϵ). Bidder 3 bids $(10 - 2\epsilon, 10 - 2\epsilon)$. For sufficiently small ϵ, bidder 1 wins both units and pays $10 - \epsilon$. As in the previous example, for any feasible redistribution mechanism $\mathbf{r}$,$r_i(\theta_{-i}) \leq \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i})\}$. For $\theta'_1 = (0, 0)$, $VCG(\theta'_1, \theta_{-1}) = 0$, so $r_1((\epsilon, \epsilon), (10-2\epsilon, 10-2\epsilon)) = 0$. For $\theta'_2 = (11, 11)$, $VCG(\theta'_2, \theta_{-2}) = 2\epsilon$, so $r_2((0, 10), (10 - 2\epsilon, 10 - 2\epsilon)) \leq 2\epsilon$. For $\theta'_3 = (0, 0)$, $VCG(\theta'_3, \theta_{-3}) = \epsilon$, so $r_3((0, 10), (\epsilon, \epsilon)) \leq \epsilon$. Hence, the percentage redistributed is at most $\frac{3\epsilon}{10-\epsilon}$, which approaches 0 as ϵ approaches 0. It follows that *every* redistribution mechanism has a worst-case redistribution percentage of 0 in this setting. As in the previous example, this remains true for any number of bidders (which can be shown by adding bidders that bid $(0, 0)$).

The previous examples show that the worst-case criterion is not a helpful guide in designing redistribution mechanisms for more complex auction settings. Instead, we will pursue a new objective: we will design redistribution mechanisms that are *undominated*. A redistribution mechanism is undominated if there does not exist another redistribution mechanism that always redistributes at least as much to each agent, and, in at least one case, strictly more. The following definition makes this precise.

Definition 1 *A redistribution mechanism* $\mathbf{r}$ *is* undominated *if it is feasible, and there does not exist a feasible redistribution mechanism* $\mathbf{r}'$ *that* dominates *it, that is,*

- *for all* i*, for all* $\theta_1, \ldots, \theta_n$, $r_i'(\theta_{-i}) \geq r_i(\theta_{-i})$.
- *for some* i*, for some* $\theta_1, \ldots, \theta_n$, $r_i'(\theta_{-i}) > r_i(\theta_{-i})$.

For example, the trivial redistribution mechanism that redistributes nothing is dominated by both WCO and Cavallo's mechanism; neither of WCO and Cavallo's mechanism dominates the other; and in general, WCO and Cavallo's mechanism are not undominated (as we will see later). The following theorem provides an alternative characterization.

Theorem 1 *A redistribution mechanism* $\mathbf{r}$ *is undominated if and only if for all* i *and all* $\theta_1, \ldots, \theta_n$

$$r_i(\theta_{-i}) \geq 0 \tag{3}$$

$$r_i(\theta_{-i}) = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\} \tag{4}$$

Here, θ_{-j}' *are the reported types of the agents other than* j *when* θ_i *is replaced by* θ_i'.

It should be noted that the only difference between Equation 2 and Equation 4 is that "$\leq$" is replaced by "$=$".

Proof: We prove the "if" direction first. Any redistribution mechanism $\mathbf{r}$ that satisfies Equation 3 and Equation 4 is feasible by Claim 1. Now suppose that $\mathbf{r}$ is dominated, that is, there exists a feasible redistribution mechanism $\mathbf{r}'$ such that for all i and θ_{-i}, we have $r_i'(\theta_{-i}) \geq r_i(\theta_{-i})$, and for some i and θ_{-i}, we have $r_i'(\theta_{-i}) > r_i(\theta_{-i})$. For the i and θ_{-i} that make this inequality strict, we have $r_i'(\theta_{-i}) > r_i(\theta_{-i}) = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\} \geq \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j'(\theta_{-j}')\}$. But this contradicts the feasibility of $\mathbf{r}'$. It follows that $\mathbf{r}$ is undominated.

Now we prove the "only if" direction. An undominated mechanism is feasible by definition, so by Claim 1, Equation 3 must hold. Suppose Equation 4 is not satisfied. Then, there exists some i and θ_{-i} such that $r_i(\theta_{-i}) < \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\}$. Let $a = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\} - r_i(\theta_{-i})$ (so that $a > 0$), and let $\mathbf{r}'$ be the same as $\mathbf{r}$, except that for the aforementioned i and θ_{-i}, $r_i'(\theta_{-i}) = r_i(\theta_{-i}) + a$. To show that this does not break the non-deficit constraint, consider any type vector (θ_i, θ_{-i}) where i and θ_{-i} are the same as before (that is, any type vector that is affected). Then, $r_i'(\theta_{-i}) = a + r_i(\theta_{-i}) = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\} = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j'(\theta_{-j}')\}$. Thus, by Claim 1, $\mathbf{r}'$ is feasible. This contradicts that $\mathbf{r}$ is undominated. Hence, Equation 4 must hold. ■

As an aside, suppose we were only interested in anonymous mechanisms, and we would therefore only consider a mechanism dominated if it were dominated by an *anonymous* mechanism. Then, the characterization in Theorem 1 remains identical.[4] Therefore, all of our results apply to this modified definition as well.

One interesting property of nontrivial undominated redistribution mechanisms is that there is always *some* case where they redistribute 100% of the VCG payments. (A redistribution mechanism is *trivial* if it never redistributes anything.) So (non-trivial) undominated VCG redistribution mechanisms are also optimal in the sense of best-case redistribution percentage.

Claim 2 *If a nontrivial redistribution mechanism* $\mathbf{r}$ *is undominated, then there exists a case where it redistributes* 100% *of the (nonzero) total VCG payments.*

Proof: If $\mathbf{r}$ is not trivial, then for some i and θ_{-i}, we have $r_i(\theta_{-i}) > 0$. By Theorem 1, $r_i(\theta_{-i}) = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')\}$, so for some $\theta_i' \in \Theta_i$, $VCG(\theta_i', \theta_{-i}) = r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta_{-j}') \geq 0$. Thus the redistribution percentage for (θ_i', θ_{-i}) is 100%. ■

An undominated redistribution mechanism always exists; in general, it is not unique. We now give two examples of undominated redistribution mechanisms.

Example 3. Consider a single-item auction with $n \geq 3$ bidders. Bidder i bids $\theta_i \in \mathbb{R}$. Let $p(j, \theta)$ be the jth highest element of θ. If $\mathbf{r}$ is Cavallo's mechanism, then $r(\theta_{-i}) = \frac{1}{n} p(2, \theta_{-i})$ (Cavallo's mechanism is anonymous, so we omit the subscript of r.) To show r is undominated, it suffices to show Equation 3 and Equation 4 are satisfied. For Equation 3, this is clear. For Equation 4, we first observe that for all θ_i', $VCG(\theta_i', \theta_{-i}) = p(2, (\theta_i', \theta_{-i})) \geq p(2, \theta_{-i})$ and for all $j \neq i$, $VCG(\theta_i', \theta_{-i}) = p(2, (\theta_i', \theta_{-i})) \geq p(2, \theta_{-j}')$. Because $r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta_{-j}') = p(2, \theta_{-i})/n + \sum_{j \neq i} p(2, \theta_{-j}')/n$, it follows that $r_i(\theta_{-i}) \leq VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j(\theta_{-j}')$ for all θ_i'. Moreover, if $\theta_i' = p(2, \theta_{-i})$, then all of the above inequalities become equalities. Hence Equation 4 holds. It follows that Cavallo's mechanism is undominated in this setting. (We will show that it is not undominated in more general settings.)

Example 4. Consider again a single-item auction with $n \geq 5$ bidders. Bidder i bids θ_i. Let $\mathbf{r}$ be the following anonymous redistribution mechanism: $r(\theta_{-i}) = \frac{1}{n-2} p(2, \theta_{-i}) -$

[4] This can be proved by modifying the proof of Theorem 1, adding a/n to each agent's redistribution function instead of adding a to one agent's redistribution function.

$\frac{2}{(n-2)(n-3)}p(3,\theta_{-i}) + \frac{6}{n(n-2)(n-3)}p(4,\theta_{-i})$. Equation 3 and Equation 4 can be shown to hold (the equality in Equation 4 is achieved by setting $\theta'_i = p(4,\theta_{-i})$).

Because in general, there are multiple undominated redistribution mechanisms, it is not clear which one is best. If a prior distribution over agents' types is available, then we would prefer the one that redistributes the most in expectation; however, in this paper, we do not wish to assume that such a prior is available. Nevertheless, for any (feasible) redistribution mechanism that we might consider using, if it is dominated, then there exists another (feasible) redistribution mechanism that always redistributes at least as much to each agent, and more in some cases. Thus, in expectation, the latter mechanism redistributes at least as much for any prior distribution, and strictly more if the prior assigns positive probability to the set of type vectors on which the latter mechanism redistributes more. Hence, we would certainly prefer the latter mechanism—and if that mechanism is not undominated, we would prefer to find one that dominates it, *etc.* But how do we find such an improved mechanism? This is what we study in the rest of the paper.

4 Methods for Constructing Undominated Redistribution Mechanisms

In this section, we propose several techniques that, given a redistribution mechanism that is feasible and dominated, find a feasible redistribution mechanism that dominates it. (If the initial mechanism is already undominated, then the techniques will return the same mechanism.) One technique immediately produces an undominated mechanism that is not anonymous; the other techniques preserve anonymity, and after repeated application converge to an undominated mechanism. We emphasize that we can start with *any* feasible redistribution mechanism, including Cavallo's mechanism, the WCO mechanism from our previous work (which, even though is optimal in the worst case, is generally not undominated), or even the trivial redistribution mechanism that redistributes nothing. These techniques can also be useful in settings where we do have a prior distribution. For example, after designing a redistribution mechanism based on a prior distribution, we can further improve it and make it undominated, which will never decrease the redistribution payment to any agent.

4.1 A Priority-Based Technique

Given a feasible redistribution mechanism $\mathbf{r}$ and a priority order over agents π, we can improve $\mathbf{r}$ into an undominated redistribution mechanism that is not anonymous. The technique works as follows.

1) Let $\pi : \{1,\ldots,n\} \to \{1,\ldots,n\}$ be a permutation representing the priority order. That is, $\pi(i)$ is agent i's priority value (the lower the value, the higher the priority). $\pi^{-1}(k)$ is the agent with the kth-highest priority.

2) Let $i = \pi^{-1}(1)$, and update i's redistribution function to $r_i^\pi(\theta_{-i}) = \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - \sum_{\pi(j)>1} r_j(\theta'_{-j})\}$. That is, we redistribute as much as possible to this agent without breaking feasibility.

3) We will now consider the remaining agents in turn, according to the order π. In the kth step, we update the redistribution function of agent $i = \pi^{-1}(k)$ to $r_i^\pi(\theta_{-i}) = \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - \sum_{\pi(j)>k} r_j(\theta'_{-j}) - \sum_{\pi(j)<k} r_j^\pi(\theta'_{-j})\}$ That is, we redistribute as much as possible to this agent without breaking feasibility, taking the previous $k-1$ updates into account.

Thus, for every agent i, $r_i^\pi(\theta_{-i}) = \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta'_{-j})\}$. The new redistribution mechanism $\mathbf{r}^\pi$ satisfies the following properties:

Claim 3 *For all i, for all θ_{-i}, $r_i^\pi(\theta_{-i}) \ge r_i(\theta_{-i})$.*

Proof: First consider $i = \pi^{-1}(1)$, the agent with the highest priority. For any θ_{-i}, we have $r_i^\pi(\theta_{-i}) = \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - \sum_{j\neq i} r_j(\theta'_{-j})\}$. Since the original redistribution mechanism $\mathbf{r}$ is feasible, by Equation 2, we have $r_i(\theta_{-i}) \le \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - \sum_{j\neq i} r_j(\theta'_{-j})\}$. Hence $r_i^\pi(\theta_{-i}) \ge r_i(\theta_{-i})$.

For any $i \neq \pi^{-1}(1)$, $r_i^\pi(\theta_{-i}) = r_i(\theta_{-i}) + \min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta'_{-j})\}$. We must show $\min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta'_{-j})\} \ge 0$.

Consider $p = \pi^{-1}(\pi(i)-1)$ (the agent immediately before i in terms of priority). For any θ_i, θ_{-i}, we have $VCG(\theta_i,\theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta_{-j}) = VCG(\theta_i,\theta_{-i}) - \sum_{\pi(j)>\pi(p)} r_j(\theta_{-j}) - \sum_{\pi(j)<\pi(p)} r_j^\pi(\theta_{-j}) - r_p^\pi(\theta_{-p}) \ge \min_{\theta'_p\in\Theta_p}\{VCG(\theta'_p,\theta_{-p}) - \sum_{\pi(j)>\pi(p)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(p)} r_j^\pi(\theta'_{-j})\} - r_p^\pi(\theta_{-p}) = 0$. (For the above inequality only, θ'_{-j} is the set of types reported by the agents other than j when θ_p is replaced by θ'_p.) Because θ_i is arbitrary, it follows that $\min_{\theta'_i\in\Theta_i}\{VCG(\theta'_i,\theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta'_{-j})\} \ge 0$. It follows that $r_i^\pi(\theta_{-i}) \ge r_i(\theta_{-i})$ for all i and θ_{-i}. ■

Claim 4 *$\mathbf{r}^\pi$ is an undominated redistribution mechanism.*

Proof: By Claim 3, for all i and θ_{-i}, $r_i^\pi(\theta_{-i}) \ge r_i(\theta_{-i}) \ge 0$. So, $\mathbf{r}^\pi$ is individually rational.

Let $i = \pi^{-1}(n)$. For all $\theta_1, \ldots, \theta_n$, the total VCG payment that is not redistributed by $\mathbf{r}^\pi$ is $VCG(\theta_1, \ldots, \theta_n) - \sum_{j=1,\ldots,n} r_j^\pi(\theta_{-j}) \geq \min_{\theta_i' \in \Theta_i}\{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r_j^\pi(\theta'_{-j})\} - r_i^\pi(\theta_{-i}) = 0$. Hence $\mathbf{r}^\pi$ never incurs a deficit. So, $\mathbf{r}^\pi$ is feasible.

Using Claim 3, we have $r_i^\pi(\theta_{-i}) = \min_{\theta_i' \in \Theta_i}\{VCG(\theta_i', \theta_{-i}) - \sum_{\pi(j)>\pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j)<\pi(i)} r_j^\pi(\theta'_{-j})\} \geq \min_{\theta_i' \in \Theta_i}\{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq n} r_j^\pi(\theta'_{-j})\}$. Because $\mathbf{r}^\pi$ is feasible, the opposite inequality must also be satisfied (Equation 2)—hence we must have equality, that is, Equation 4 must hold. Because Equation 3 is also satisfied by Claim 1, it follows that $\mathbf{r}^\pi$ is undominated. ■

Example 5. Consider a single-item auction with four bidders $1, 2, 3, 4$. In this setting, the redistribution under the WCO mechanism to agent i is $r(\theta_{-i}) = (2/7)p(2, \theta_{-i}) - (1/7)p(3, \theta_{-i})$ (where $p(k, \theta_{-i})$ is the kth highest bid among bids other than i's). Consider a specific set of bids $(8, 10, 13, 5)$ and let $\pi(i) = i$ for all i. (That is, agent 1 bids 8 for the item and has the highest priority, *etc.*) If we apply the above technique, the resulting redistribution payment to agent 1 is $r_1^\pi(10, 13, 5) = \min_{\theta_1' \in [0,\infty)}\{VCG(\theta_1', 10, 13, 5) - r(\theta_1', 13, 5) - r(\theta_1', 10, 5) - r(\theta_1', 10, 13)\}$ (where r is the WCO mechanism). It turns out that the expression is minimized at $\theta_1' = 0$, so that $r_1^\pi(10, 13, 5) = \frac{30}{7}$. This is twice the amount 1 would have received under WCO: $r(10, 13, 5) = (2/7) \cdot 10 - (1/7) \cdot 5 = \frac{15}{7}$.

For agent 2, $r_2^\pi(8, 13, 5) = \min_{\theta_2' \in [0,\infty)}\{VCG(8, \theta_2', 13, 5) - r_1^\pi(\theta_2', 13, 5) - r(8, \theta_2', 5) - r(8, \theta_2', 13)\}$. This expression is minimized at $\theta_2' = 8$, so that $r_2^\pi(8, 13, 5) = \frac{17}{7}$. (Under WCO, 2 receives only $\frac{11}{7}$.)

For agent 3, $r_3^\pi(8, 10, 5) = \min_{\theta_3' \in [0,\infty)}\{VCG(8, 10, \theta_3', 5) - r_1^\pi(10, \theta_3', 5) - r_2^\pi(8, \theta_3', 5) - r(8, 10, \theta_3')\}$. This expression is minimized at $\theta_3' = 8$, so that $r_3^\pi(8, 10, 5) = \frac{11}{7}$. (Under WCO, 3 receives $\frac{11}{7}$ as well.)

For agent 4 $r_4^\pi(8, 10, 13) = \min_{\theta_4' \in [0,\infty)}\{VCG(8, 10, 13, \theta_4') - r^\pi(10, 13, \theta_4') - r^\pi(8, 13, \theta_4') - r^\pi(8, 10, \theta_4')\}$. This expression is minimized at $\theta_4' = 5$, so that $r_4^\pi(8, 10, 13) = \frac{12}{7}$. (Under WCO, 4 receives $\frac{12}{7}$ as well.)

We note that for this priority order, the total amount redistributed is $\frac{30+17+11+12}{7} = 10$, that is, all of the VCG payments are redistributed. This is not true for all priority orders; averaging over all priority orders, 0.304 remains undistributed (compared to 3 for the WCO mechanism).

Generally, most of the increase in redistribution payment goes to high-priority agents. Hence, a reasonable approximation can be obtained by only updating the redistribution payment functions of the first few agents. This still results in a feasible mechanism that dominates the original (or is the same), but it is no longer guaranteed to be undominated.

4.2 Iterative Techniques that Preserve Anonymity

The technique from the previous subsection will, in general, not produce an anonymous redistribution mechanism, even if the original mechanism $\mathbf{r}$ is anonymous. This is because agents higher in the priority order tend to receive higher redistribution payments. In this subsection, we will introduce techniques that preserve anonymity.

One way to obtain an anonymous mechanism is to consider r^π for *all* permutations π, and take the average. That is, let $\bar{\mathbf{r}}$ be defined by $\bar{r}_i = \frac{1}{n!}\sum_{\pi \in S_n}(r_i^\pi)$, where S_n is the set of all permutations of n elements. Given that the setting and the initial mechanism are anonymous, this results in an anonymous mechanism. It is also feasible:

Claim 5 *Any convex combination of a set* $\{\mathbf{r^{(1)}}, \ldots, \mathbf{r^{(t)}}\}$ *of feasible redistribution mechanisms is itself feasible.*

Proof: Let $\sum_{k=1}^{t} \alpha_k = 1$ with each $\alpha_k \geq 0$; we must show that $\mathbf{r} = \sum_{k=1}^{t} \alpha_k \mathbf{r^{(k)}}$ is feasible. For any i and θ_{-i}, for any k, we have $r_i^{(k)}(\theta_{-i}) \geq 0$, hence $r_i(\theta_{-i}) = \sum_{k=1}^{t} \alpha_k r_i^{(k)}(\theta_{-i}) \geq 0$. This implies individual rationality. Also, for any $\theta_1, \ldots, \theta_n$, for any k, $\sum_{i=1}^{n} r_i^{(k)}(\theta_{-i}) \leq VCG(\theta_1, \ldots, \theta_n)$, hence $\sum_{i=1}^{n} r_i(\theta_{-i}) = \sum_{k=1}^{t} \alpha_k \sum_{i=1}^{n} r_i^{(k)}(\theta_{-i}) \leq VCG(\theta_1, \ldots, \theta_n)$. This implies the non-deficit property. ■

Because $\bar{\mathbf{r}}$ is anonymous, all $\bar{r}_i$ are the same, so we will simply use $\bar{r}$. Even though $\bar{r}$ is an average of a set of undominated redistribution mechanisms, in general, it itself is not undominated. In principle, we can take the resulting mechanism and apply the technique again. Unfortunately, this approach is not practical—in fact, it may not be feasible to perform even one iteration of this technique if n is large, since we have to take an average over $n!$ mechanisms.[5] However, as we mentioned, it is also possible to apply the priority-based technique only to the first h agents. This still results in a feasible (but not necessarily undominated) mechanism, and tends to obtain most of the increase in redistribution payments. Taking the average over all such

[5] Computational limitations often prevent us from using certain mechanisms. As an extreme example, it is possible to have a computer search over the space of all possible (incentive compatible) mechanisms for the setting at hand and find the best one [4], but this does not scale to very large instances. By contrast, here, we have an analytical characterization of the mechanism, but computing its outcomes is still hard.

mechanisms is feasible for sufficiently small h (there will be $P_h^n = n!/(n-h)!$ such mechanisms), and will result in an anonymous mechanism. We will consider the extreme case where $h = 1$ (*i.e.* we only change one agent's redistribution function), so that we have to take an average over only n mechanisms. This we can do iteratively.

Given a feasible and anonymous redistribution mechanism r, let $r^0 = r$, and let r^k be the mechanism that results after k iterations of the above technique (with $h = 1$). Then, for all i and $\theta_1, \ldots, \theta_n$, $r^{k+1}(\theta_{-i}) = \frac{n-1}{n} r^k(\theta_{-i}) + \frac{1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\}$.

This technique can be interpreted as a generalization of the basic idea underlying Cavallo's mechanism. We can rewrite $r^{k+1}(\theta_{-i}) = r^k(\theta_{-i}) + \frac{1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j}) - r^k(\theta_{-i})\}$. If the starting mechanism $r = r^0$ is the trivial redistribution mechanism that redistributes nothing, then we have $r^1(\theta_{-i}) = \frac{1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i})\}$, which is exactly Cavallo's mechanism.

Claim 6 *If r^k is feasible, r^{k+1} is feasible.*

Proof: r^{k+1} is an average of feasible mechanisms, so Claim 5 applies. ■

Claim 7 *For any i and θ_{-i}, $r^k(\theta_{-i})$ is nondecreasing in k.*

Proof: $r^{k+1}(\theta_{-i}) = r^k(\theta_{-i}) + \frac{1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j}) - r^k(\theta_{-i})\}$. Because r^k is feasible by Claim 6, $\min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j}) - r^k(\theta_{-i})\} \geq 0$. Hence $r^{k+1}(\theta_{-i}) \geq r^k(\theta_{-i})$. ■

Claim 8 *As $k \to \infty$, r^k converges (pointwise) to an undominated redistribution mechanism.*

Proof: By Claim 7, the $r^k(\theta_{-i})$ are nondecreasing in k, and since every r^k is feasible by Claim 6, they must be bounded; hence they must converge (pointwise). For any i and θ_{-i}, let $d_k = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - r^k(\theta_{-i})$. Using Claim 7, we derive the following inequality: $d_{k+1} = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^{k+1}(\theta'_{-j})\} - r^{k+1}(\theta_{-i}) \leq \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - r^{k+1}(\theta_{-i}) = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - \frac{n-1}{n} r^k(\theta_{-i}) - \frac{1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} = \frac{n-1}{n} \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - \frac{n-1}{n} r^k(\theta_{-i}) = \frac{n-1}{n} d_k$. As $k \to \infty$, $d_k = \min_{\theta_i' \in \Theta_i} \{VCG(\theta_i', \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - r^k(\theta_{-i}) \to 0$. So in the limit, Equation 4 is satisfied. Thus, r^k converges (pointwise) to an undominated redistribution mechanism. ■

Corollary 1 *If $r^{k+1} = r^k$, then r^k is undominated.*

Claim 9 *If r^k is not undominated, then r^{k+1} dominates r^k.*

Proof: r^{k+1} always redistributes at least as much as r^k to each agent by Claim 7. Moreover, $r^{k+1} \neq r^k$ (otherwise Corollary 1 would imply r^k is undominated). Hence there must be a case where r^{k+1} redistributes more than r^k. ■

5 Numerical Results

In this section, we present the results of some experiments in which we use the techniques from the previous sections to improve both the WCO mechanism and Cavallo's mechanism. (We do not present any results starting from the trivial redistribution mechanism that redistributes nothing, because, as we mentioned, after one iteration of the anonymity-preserving technique, we would obtain Cavallo's mechanism.)

Improving the WCO mechanism. The WCO mechanism applies only to multi-unit auctions with unit demand (*i.e.* in which each bidder only wants a single unit); in this setting, this mechanism maximizes the percentage that is redistributed in the worst case. This, however, does not mean that it is undominated, because it could be dominated by another mechanism that does equally well in the worst case, and better in other cases. Indeed, we can improve the WCO mechanism using the techniques from this paper (resulting in another, better, worst-case optimal mechanism).

For various m (number of units) and n (number of bidders), we generated 100 random instances with each bidder's valuation drawn uniformly from $[0, 1]$. The table below shows the ratio between the average amount that is not redistributed by the new mechanism (which results from applying one of our techniques to the WCO mechanism), and the average amount that is not redistributed by the (original) WCO mechanism. That is, it is the percentage of the amount that WCO fails to redistribute that the new mechanism also fails to redistribute. Lower numbers are better—100% indicates no improvement over WCO, 0% indicates that everything is redistributed. For the nonanonymous (priority-based) technique, to save computation time, we only update the redistribution payments for the first three bidders. This technique redistributes more than the anonymity-preserving technique.

n	m	Nonanoymous (3 updates)	Anonymous (1 iteration)	Anonymous (2 iterations)
4	1	42%	66%	52%
5	1	49%	69%	55%
6	1	32%	55%	39%
5	2	44%	68%	54%
6	3	45%	68%	54%

Improving Cavallo's mechanism. We recall that Cavallo's mechanism is undominated in the single-item auction setting (in fact, this remains true for multi-unit auctions with unit demand). However, as the experiment below shows, it is not undominated in general.

For a combinatorial auction with n single-minded bidders and 2 items, we generated 100 random instances. For each bidder, we randomly chose a nonempty bundle of items, and randomly chose a per-item value from $[0, 1]$ (which is multiplied by two if the bidder desires the bundle of two items). The percentages have the same meaning as before. We distinguish between the known single-minded case (where the auctioneer knows which bundle the agent wants) and the unknown case. Again, the nonanonymous technique redistributes more; also, more is redistributed in the known case.

n	Nonanon. (2 updates) (unknown)	Anonymous (2 iterations) (unknown)	Nonanon. (2 updates) (known)	Anonymous (2 iterations) (known)
5	83%	85%	61%	76%
6	76%	82%	57%	70%
7	72%	81%	52%	69%
8	78%	83%	57%	67%

6 Conclusion

For allocation problems, the well-known VCG mechanism is efficient, incentive compatible, individually rational, and does not incur a deficit. However, the VCG mechanism is not (strongly) budget balanced: generally, the agents' payments will sum to more than 0. Very recently, several mechanisms have been proposed that *redistribute* a significant percentage of the VCG payments back to the agents while maintaining the other properties. This increases the agents' utilities. In this paper, we provided a characterization of undominated redistribution mechanisms. We also proposed several techniques that take a dominated redistribution mechanism as input, and produce as output another redistribution mechanism that dominates the original. The dominating redistribution mechanism always redistributes at least as much, and in some cases more. Hence, for any prior distribution over agents' types, the dominating mechanism redistributes at least as much as the original in expectation; if the prior assigns positive probability to the set of type vectors where the dominating mechanism redistributes more, then the dominating mechanism redistributes strictly more in expectation.

One of the techniques that we proposed takes as input a priority order over the agents. It first redistributes as much as possible to the highest-priority agent, then it redistributes as much of the remainder as possible to the second-highest priority agent, *etc.* At the end of this process, the mechanism is guaranteed to be undominated—but it is generally not anonymous. Another technique that we proposed does preserve anonymity, and can be seen as taking the average over all priority orders of the first step of the priority-based technique. It can also be seen as a generalization of the basic idea underlying Cavallo's mechanism, and Cavallo's mechanism results after one iteration of the technique when starting with the mechanism that redistributes nothing. Repeated application of this technique produces an undominated mechanism in the limit.

Finally, we showed experimentally that these techniques improve both the WCO mechanism and Cavallo's mechanism. In our experiment on multi-unit auctions with unit demand, the improved mechanisms redistributed (on average) between 31% and 68% of what WCO failed to redistribute. In our experiment on combinatorial auctions with single-minded bidders, the improved mechanisms redistributed (on average) between 15% and 48% of what Cavallo's mechanism failed to redistribute.

References

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

[2] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. *AAMAS*, 882–889, 2006.

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

[4] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. *UAI*, 103–110, 2002.

[5] Boi Faltings. A budget-balanced, incentive-compatible scheme for social choice. *AMEC*, 30–43, 2005.

[6] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of mulicast transmissions. *JCSS*, 63:21–41, 2001.

[7] J Green and J-J Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[8] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[9] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. *EC*, 2007.

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

[11] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[12] David Parkes, Jayant Kalagnanam, and Marta Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. *IJCAI*, 1161–1168, 2001.

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