

Undominated VCG Redistribution Mechanisms

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

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

ABSTRACT

Many important problems in multiagent systems can be seen as resource allocation problems. For such 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.

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, comparing and improving redistribution mechanisms

1. INTRODUCTION

Many important problems in multiagent systems 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,

Cite as: Undominated VCG Redistribution Mechanisms, Mingyu Guo and Vincent Conitzer, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

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 agents have no incentive to lie about their valuations. The best-known way of doing so is to use the *VCG mechanism* [22, 5, 11] 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 [20, 3, 12, 17, 13]. 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* $I = \{1, 2, \dots, n\}$, 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 $a : S \rightarrow \{0, 1, \dots, n\}$, where for any $s \in S$, $a(s)$ is the agent that obtains s (if no agent 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* $v_i : \Theta_i \times O \rightarrow \mathbb{R}$. (We assume that there always exists a $\theta_i \in \Theta_i$ such that $v_i(\theta_i, o) = 0$ for all $o \in O$; for example, in an auction, it is possible that an agent does not care for any of the items.) 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*

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

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

(subset) of the items. Often, the type space is assumed to be more restricted. For example, if each agent is only interested in a single bundle (that is, agents are *single-minded*), then a type θ_i consists of a pair (S'_i, x_i) , where S'_i is the bundle that the agent 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 j th number indicates the value for obtaining j units. A special case is a multi-unit auction with *unit demand*, in which each agent 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.³ Thus, a mechanism is given by an outcome selection function $f : \Theta_1 \times \dots \times \Theta_n \rightarrow O$, as well as n payment selection functions $\rho_i : \Theta_1 \times \dots \times \Theta_n \rightarrow \mathbb{R}$. As is common, we assume that preferences are *quasilinear*, that is, agent i 's utility is $u_i(\theta_i, (\hat{\theta}_1, \dots, \hat{\theta}_n)) = v_i(\theta_i, f(\hat{\theta}_1, \dots, \hat{\theta}_n)) - \rho_i(\hat{\theta}_1, \dots, \hat{\theta}_n)$. 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 $(\theta_1, \dots, \theta_n) \in \Theta_1 \times \dots \times \Theta_n$ and all $\hat{\theta}_i \in \Theta_i$, $u_i(\theta_i, (\theta_1, \dots, \theta_i, \dots, \theta_n)) \geq u_i(\theta_i, (\theta_1, \dots, \hat{\theta}_i, \dots, \theta_n))$.

Perhaps the most famous mechanism is the *Vickrey-Clarke-Groves* (VCG) mechanism [22, 5, 11]. 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 $\rho_j(\hat{\theta}_1, \dots, \hat{\theta}_n) = \sum_{i \neq j} v_i(\hat{\theta}_i, o_{-j}^*) - \sum_{i \neq j} v_i(\hat{\theta}_i, o^*)$. 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 $(\theta_1, \dots, \theta_n) \in \Theta_1 \times \dots \times \Theta_n$, for all i , $u_i(\theta_i, (\theta_1, \dots, \theta_i, \dots, \theta_n)) \geq 0$. That is, participating in the mechanism does not make anyone worse off.
- *non-deficit*: for all $(\theta_1, \dots, \theta_n) \in \Theta_1 \times \dots \times \Theta_n$, $\sum_{i=1}^n \rho_i(\theta_1, \dots, \theta_n) \geq 0$. 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. That is, for all $(\theta_1, \dots, \theta_n) \in \Theta_1 \times \dots \times \Theta_n$, if we swap θ_i and θ_j , then the results of the mechanism for agents i and j are swapped, and the results of the mechanism for agents other than i and j remain the same.

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

³We allow for payments to be negative, that is, agents may receive payments. For cases where payments must be nonnegative, several authors have proposed mechanisms that maximize the agents' combined utility after deducting the payments [14, 4].

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

2.2 Redistribution Mechanisms

For the VCG mechanism, sometimes, $\sum_{i=1}^n \rho_i(\theta_1, \dots, \theta_n) \neq 0$.

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 [16, 10, 18]. In light of this impossibility result, several authors have obtained budget balance by sacrificing some of the other desirable properties [2, 9, 19, 8]. 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 [3].⁴ Thus, a redistribution mechanism is defined by a function $r_i : \Theta_1 \times \dots \times \Theta_{i-1} \times \Theta_{i+1} \times \dots \times \Theta_n \rightarrow \mathbb{R}$ for each agent i . That is, letting θ_{-i} be the vector of types submitted by agents other than i , $r_i(\theta_{-i})$ 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 agents.

For example, Cavallo's mechanism [3] is given by $r_i(\theta_{-i}) = \frac{1}{n} \min_{\theta_i \in \Theta_i} VCG(\theta_i, \theta_{-i})$, where $VCG(\theta_i, \theta_{-i})$ is the total VCG payment collected for those reports.⁵ In the special case of a single-item auction, under Cavallo's mechanism, an agent's redistribution payment is $1/n$ times the second-highest bid among *other* agents' bids. That is, the agent with the highest bid wins and pays the second-highest bid, as in a second-price sealed-bid (Vickrey) auction; then, the agents with the highest and the second-highest bids each receive a redistribution payment of v_3/n , where v_3 is the third-highest bid; and the remaining agents each receive a redistribution payment of v_2/n , where v_2 is the second-highest bid. Hence, the total redistributed is $2v_3/n + (n-2)v_2/n \leq nv_2/n = v_2$, so that there is never a deficit. It should also be noted that no agent can affect his redistribution payment (since it is the second-highest *other* bid, divided by n), hence the incentives are the same

⁴Mechanisms which differ from the VCG mechanism only by an additional term in the payment function that does not depend on the agent's own bid are known as *Groves mechanisms*. Hence, all the mechanisms in this paper are Groves mechanisms. In sufficiently general settings, Groves mechanisms are the only incentive compatible mechanisms that satisfy efficiency [10, 15].

⁵We 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).

as in the Vickrey auction, which is incentive compatible. In this (single-item) case, Cavallo’s mechanism coincides with the mechanisms proposed by Bailey [2] and Porter *et al.* [20]. Cavallo’s mechanism and Bailey’s mechanism are in fact the same in any setting under which VCG satisfies *revenue monotonicity*,⁶ which includes multi-unit auctions with unit demand. For multi-unit auctions with unit demand, we have previously characterized a redistribution mechanism that maximizes the worst-case redistribution percentage [12]. The same mechanism was independently proposed by Moulin [17], who pursues a different worst-case objective: whereas the objective in our paper 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. 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. Our earlier work [12] focused on maximizing the percentage of VCG payments redistributed in the worst case. However, that paper 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 a different criterion is needed.⁷

We will require the following claim for our examples.

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

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

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

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 $\theta_1, \dots, \theta_{i-1}, \theta'_i, \theta_{i+1}, \dots, \theta_n$.

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

⁶Bailey’s mechanism redistributes to each agent $1/n$ of the total VCG payment that would result if this agent were removed from the auction. Cavallo’s mechanism considers how small an agent could make the total VCG payment by changing her type, and redistributes $1/n$ of that to the agent. If the total VCG payment is monotonically nondecreasing in the agents, then the type that would minimize the total VCG payment is the one that has a valuation of 0 for everything, which is equivalent to not participating in the auction. It is well known that in general, the VCG mechanism does *not* satisfy this revenue monotonicity criterion [1, 7, 23, 24, 25] (this is in fact true for a much wider class of mechanisms [21]). However, in more restricted settings, revenue monotonicity often holds.

⁷In other work, we study settings where a prior distribution over agents’ preferences is available, and try to maximize the *expected* redistribution [13]. However, in this paper, we continue the prior-free approach.

is satisfied. For any i and $\theta_1, \dots, \theta_n$, Equation 2 implies that $r_i(\theta_{-i}) \leq VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})$ for any $\theta'_i \in \Theta_i$. If we let $\theta'_i = \theta_i$, we obtain $r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta_{-j}) \leq VCG(\theta_i, \theta_{-i})$. 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 $\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n$, 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. \square

Example 1. Consider a combinatorial auction with two items $\{a, b\}$ and three agents $\{1, 2, 3\}$. Agent 1 bids 10 on the bundle $\{a, b\}$; agent 2 bids ϵ on $\{a\}$; agent 3 bids $10 - 2\epsilon$ on $\{b\}$. For sufficiently small ϵ , agent 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 $r_1(\{a, b\}, \epsilon, \{b\}, 10 - 2\epsilon) = 0$ (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 agents who bid $(\{a\}, 0)$, then the bounds on the first three agents’ redistribution payments remain the same, and each additional agent can have a redistribution payment of at most 2ϵ (if any one of them bids more than 10 on $\{a\}$, 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 agents. This is in contrast to the case of multi-unit auctions with unit demand, where additional agents improve the worst-case redistribution percentage [12].

Example 2. Consider a multi-unit auction with two units and three agents $\{1, 2, 3\}$. Agent 1 bids $(0, 10)$ (0 for getting one unit and 10 for getting two units). Agent 2 bids (ϵ, ϵ) . Agent 3 bids $(10 - 2\epsilon, 10 - 2\epsilon)$. For sufficiently small ϵ , agent 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 agents (which can be shown by adding agents 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 $\theta_1, \dots, \theta_n$, $r'_i(\theta_{-i}) \geq r_i(\theta_{-i})$.
- for some i , for some $\theta_1, \dots, \theta_n$, $r'_i(\theta_{-i}) > r_i(\theta_{-i})$.

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 $\theta_1, \dots, \theta_n$

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

$$r_i(\theta_{-i}) = \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})\} \quad (4)$$

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 $r'_i(\theta_{-i}) > r_i(\theta_{-i}) = \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r'_j(\theta'_{-j})\}$. 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, $r'_i(\theta_{-i}) = a + r_i(\theta_{-i}) = \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r'_j(\theta'_{-j})\}$. Thus, by Claim 1, \mathbf{r}' is feasible. This contradicts that \mathbf{r} is undominated. Hence, Equation 4 must hold. \square

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.⁸ Therefore, all of our results apply to this modified definition as well.

⁸This 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.

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 $\theta'_i \in \Theta_i$, $VCG(\theta'_i, \theta_{-i}) = r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta'_{-j}) > 0$. Thus the redistribution percentage for (θ'_i, θ_{-i}) is 100%. \square

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 $n \geq 3$ agents. Agent i bids $\theta_i \in [0, \infty)$. Let $p(j, \theta)$ be the j th 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 \mathbf{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 $r_i(\theta_{-i}) + \sum_{j \neq i} r_j(\theta_{-j}) = \frac{1}{n}p(2, \theta_{-i}) + \frac{1}{n} \sum_{j \neq i} p(2, \theta'_{-j})$, it follows that for all θ'_i , $r_i(\theta_{-i}) \leq VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})$. 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 $n \geq 5$ agents. Agent i bids θ_i . Let \mathbf{r} be the following anonymous redistribution mechanism: $r(\theta_{-i}) = \frac{1}{n-2}p(2, \theta_{-i}) - \frac{2}{(n-2)(n-3)}p(3, \theta_{-i}) + \frac{6}{n(n-2)(n-3)}p(4, \theta_{-i})$. 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 the 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 earlier paper [12]/Moulin [17] (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 $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ 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 k th-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 k th 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}^π satisfies the following properties:

CLAIM 3. For all i , for all θ_{-i} , $r_i^\pi(\theta_{-i}) \geq r_i(\theta_{-i})$.

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 $r_i(\theta_{-i}) \leq \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r_j(\theta'_{-j})\}$. Hence $r_i^\pi(\theta_{-i}) \geq r_i(\theta_{-i})$.

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

$$\min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j) > \pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j) < \pi(i)} r_j^\pi(\theta'_{-j})\} \geq 0.$$

Consider $p = \pi^{-1}(\pi(i) - 1)$ (the agent immediately before i in terms of priority). For any θ_i, θ_{-i} , we have $VCG(\theta_i, \theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j) > \pi(i)} r_j(\theta_{-j}) - \sum_{\pi(j) < \pi(i)} r_j^\pi(\theta_{-j}) = VCG(\theta_i, \theta_{-i}) - \sum_{\pi(j) > \pi(p)} r_j(\theta_{-j}) - \sum_{\pi(j) < \pi(p)} r_j^\pi(\theta_{-j}) - r_p^\pi(\theta_{-p}) \geq \min_{\theta'_p \in \Theta_p} \{VCG(\theta'_p, \theta_{-p}) - \sum_{\pi(j) > \pi(p)} r_j(\theta'_{-j}) - \sum_{\pi(j) < \pi(p)} r_j^\pi(\theta'_{-j})\} - r_p^\pi(\theta_{-p}) = 0$. (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 $\min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - r_i(\theta_{-i}) - \sum_{\pi(j) > \pi(i)} r_j(\theta'_{-j}) - \sum_{\pi(j) < \pi(i)} r_j^\pi(\theta'_{-j})\} \geq 0$. It follows that $r_i^\pi(\theta_{-i}) \geq r_i(\theta_{-i})$ for all i and θ_{-i} . \square

CLAIM 4. \mathbf{r}^π is an undominated redistribution mechanism.

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

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

Example 5. Consider a single-item auction with four agents 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 k th 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 $\frac{17}{7}$.)

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 $r_4^\pi(8, 10, 13) = \min_{\theta'_4 \in [0, \infty)} \{VCG(8, 10, 13, \theta'_4) - r_1^\pi(10, 13, \theta'_4) - r_2^\pi(8, 13, \theta'_4) - r_3^\pi(8, 10, \theta'_4)\}$. 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, 0.315 remains undistributed (compared to 3 for the WCO mechanism). The following table shows the results for all priority orders for this example.

Bids	Increase	Remaining
5,13,10,8	6/7,9/7,4/7,0	2/7
5,13,8,10	6/7,9/7,0,1/7	5/7
5,10,13,8	6/7,9/7,4/7,0	2/7
5,10,8,13	6/7,9/7,0,1/7	5/7
5,8,10,13	6/7,15/7,0,0	0
5,8,13,10	6/7,15/7,0,0	0
13,5,10,8	9/7,6/7,6/7,0	0
13,5,8,10	9/7,6/7,0,0	6/7
13,10,5,8	9/7,6/7,6/7,0	0
13,10,8,5	9/7,6/7,0,6/7	0
13,8,10,5	9/7,0,0,6/7	6/7
13,8,5,10	9/7,0,6/7,0	6/7
10,13,5,8	9/7,6/7,6/7,0	0
10,13,8,5	9/7,6/7,0,6/7	0
10,5,13,8	9/7,6/7,6/7,0	0
10,5,8,13	9/7,6/7,0,0	6/7
10,8,5,13	9/7,0,6/7,0	6/7
10,8,13,5	9/7,0,0,3/7	9/7
8,13,10,5	15/7,6/7,0,0	0
8,13,5,10	15/7,6/7,0,0	0
8,10,13,5	15/7,6/7,0,0	0
8,10,5,13	15/7,6/7,0,0	0
8,5,10,13	15/7,6/7,0,0	0
8,5,13,10	15/7,6/7,0,0	0
Average (1)	1.39,0.89,0.26,0.14	0.315
Average (2)	0.71,0.64,0.64,0.70	

Increase in redistribution payments relative to WCO, and total VCG payments that are not redistributed, for different priority orders. Note that increases are ordered according to the priority order. The "average" item gives the average increase to the agent ordered in the k th place (first), as well as the average increase to agent i (second).

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 $\{\mathbf{r}^{(1)}, \dots, \mathbf{r}^{(t)}\}$ 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 $\theta_1, \dots, \theta_n$, for any k , $\sum_{i=1}^n r_i^{(k)}(\theta_{-i}) \leq VCG(\theta_1, \dots, \theta_n)$, hence $\sum_{i=1}^n r_i(\theta_{-i}) =$

$\sum_{k=1}^t \alpha_k \sum_{i=1}^n r_i^{(k)}(\theta_{-i}) \leq VCG(\theta_1, \dots, \theta_n)$. This implies the non-deficit property. \square

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.⁹ 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 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 $\theta_1, \dots, \theta_n$, $r^{k+1}(\theta_{-i}) = \frac{n-1}{n} r^k(\theta_{-i}) + \frac{1}{n} \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\}$.

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 $r^1(\theta_{-i}) = \frac{1}{n} \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i})\}$,

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. \square

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 $r^{k+1}(\theta_{-i}) \geq r^k(\theta_{-i})$. \square

⁹Computational 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 [6], 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.

CLAIM 8. As $k \rightarrow \infty$, 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 $k \rightarrow \infty$, $d_k = \min_{\theta'_i \in \Theta_i} \{VCG(\theta'_i, \theta_{-i}) - \sum_{j \neq i} r^k(\theta'_{-j})\} - r^k(\theta_{-i}) \rightarrow 0$. So in the limit, Equation 4 is satisfied. Thus, r^k converges (pointwise) to an undominated redistribution mechanism. \square

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 . \square

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. For the purpose of completeness, in the combinatorial auction setting, we also apply the nonanonymous (priority-based) technique to the trivial redistribution mechanism that redistributes nothing, and compare the resulting mechanism’s performance with that of Cavallo’s mechanism. (We omit the result of applying the anonymity-preserving technique to the trivial redistribution mechanism because, as we mentioned, after one iteration, we just obtain Cavallo’s mechanism. We also omit the result of applying the nonanonymous technique to the trivial redistribution mechanism in multi-unit auctions with unit demand, because the resulting mechanism always has the same expected redistribution amount as Cavallo’s mechanism: $m(m+1)/n$ times the $m+2$ th highest bid, plus $m(n-m-1)/n$ times the $m+1$ th highest bid.)

Improving the WCO mechanism. The WCO mechanism applies only to multi-unit auctions with unit demand (*i.e.* in which each agent 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 agents), we generated 100 random instances with each agent’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 agents. This technique redistributes more than the anonymity-preserving technique.

n	m	Nonanon. 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 agents and 2 items, we generated 100 random instances. For each agent, 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 agent 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 1 iteration unknown	Nonanon. 2 updates known	Anonymous 1 iteration known
5	81%	84%	61%	75%
6	76%	82%	64%	69%
7	73%	81%	54%	68%
8	78%	83%	59%	66%

For the same set of 100 random instances, the table below shows the ratio between the average amount that is not redistributed by the mechanism which results from applying the nonanonymous technique to the trivial redistribution mechanism, and the average amount that is not redistributed by Cavallo’s mechanism.

n	Nonanon. 3 updates, unknown	Nonanon. 3 updates, known
5	88%	68%
6	91%	67%
7	95%	51%
8	96%	81%

6. CONCLUSIONS

For resource 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 agents, the improved mechanisms redistributed (on average) between 16% and 46% of what Cavallo's mechanism failed to redistribute.

Future research on the dominance concept proposed in this paper can take a number of directions. For one, it is possible to apply the techniques in this paper to other mechanisms, including mechanisms that allocate inefficiently. It may also be worthwhile to try to find other techniques for improving a given mechanism; it is possible that such techniques will scale to larger auctions than the ones presented in this paper. Another direction is to try to derive analytical characterizations of undominated mechanisms, perhaps in more restricted settings. Finally, one can try to identify circumstances under which there is a unique undominated mechanism.

7. REFERENCES

- [1] L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1. MIT Press, 2006.
- [2] M. J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.
- [3] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.
- [4] S. Chakravarty and T. Kaplan. Manna from heaven or forty years in the desert: Optimal allocation without transfer payments, October 2006. Working Paper.
- [5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [6] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, Edmonton, Canada, 2002.
- [7] V. Conitzer and T. Sandholm. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 521–528, Hakodate, Japan, 2006.
- [8] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce (AMEC), LNAI, 3435*, pages 30–43, 2005.
- [9] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.
- [10] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [11] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [12] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 30–39, San Diego, CA, USA, 2007.
- [13] M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Estoril, Portugal, 2008.
- [14] J. Hartline and T. Roughgarden. Money burning and implementation, January 2007. Working Paper.
- [15] B. Holmström. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.
- [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] H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment, March 2007. Working Paper.
- [18] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [19] D. Parkes, J. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1161–1168, Seattle, WA, 2001.
- [20] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118:209–228, 2004.
- [21] B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.
- [22] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [23] M. Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 733–742, Acapulco, Mexico, 2003.
- [24] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.
- [25] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.