

Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments

Ruggiero Cavallo
DEAS, Harvard University
33 Oxford St.
Cambridge, MA 02138
cavallo@eecs.harvard.edu

ABSTRACT

Mechanisms for coordinating group decision-making among self-interested agents often employ a trusted center, capable of enforcing the prescribed outcome. Typically such mechanisms, including the ubiquitous Vickrey Clarke Groves (VCG), require significant transfer payments from agents to the center. While this is sought after in some settings, it is often an unwanted cost of implementation. We propose a modification of the VCG framework that—by using domain information regarding agent valuation spaces—is often able to achieve redistribution of much of the required transfer payments back among the agents, thus coming closer to budget-balance. The proposed mechanism is strategyproof, ex post individual rational, no-deficit, and leads to an efficient outcome; we prove that among all mechanisms with these qualities and an anonymity property it is *optimally balanced*, in that no mechanism ever yields greater payoff to the agents. We provide a general characterization of when strategyproof redistribution is possible, and demonstrate specifically that substantial redistribution can be achieved in allocation problems.

Categories and Subject Descriptors

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

General Terms

Algorithms, Economics, Theory

Keywords

Social choice, Mechanism design, VCG

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

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

1. INTRODUCTION

This paper is concerned with decision-making problems that involve a group of self-interested parties, when each party may be in competition with other members of the group. Instances of this problem abound in everyday life. Consider, to name just a few examples: a group of housemates that jointly own a car and must decide who should get to use it on a Friday night; a set of city neighborhoods competing for a grant to build a new playground; a group of astronomers competing for the rights to use a publicly-owned space telescope; or companies competing for government allocation of wireless spectrum.

How would one formulate a general framework for decision-making in domains such as these? One natural objective is that the decision process lead to the choice that maximizes the total value that is realized. Moreover, and to that end, it is highly desirable that the mechanism not encourage manipulation by dishonest agents—we want *truthfulness* to be in each agent's best-interest; this will allow us to reasonably predict the actual effects of whatever outcome is selected, and will simplify the task of the agents by taking strategic behavior out of the equation that maximizes expected reward.

The well-known Groves family of social choice mechanisms, by requiring specific payments to each agent, reaches outcomes with the desirable properties described above. However, implementing the payments specified in the most basic Groves mechanism requires a large external budget, making it infeasible in typical decision-making scenarios where there is no outside funding.

The Vickrey Clarke Groves (VCG) mechanism counteracts this potential budget imbalance by imposing “charges” on the agents, to be delivered to a “center” capable of enforcing the specified outcome. For typical decision-making problems,¹ VCG results in net transfers being delivered only in the direction of agents to the center. While there are some scenarios in which revenue to a center is sought after (e.g., some auctions), often it is preferable to maintain as much wealth as possible within the group of agents, and the transfer payment can be viewed as an undesirable “cost of implementation.” In other words, *budget-balance* is frequently a sought after mechanism property, and though it doesn't run a deficit, the VCG mechanism fails in this regard by producing a surplus of agent charges (hence called the “VCG surplus”).

¹Specifically, when there are no positive externalities.

The ideal social choice mechanism would have the desirable properties of VCG (truthfulness; social welfare maximization; non-negative payoff guarantee), and at the same time run neither a budget surplus nor deficit. As we will see, this *exact* budget-balance is usually not attainable. However, while it has often been stated that no improvement over VCG is possible, we demonstrate that this *is not the case* in a broad class of domains (e.g., allocation problems) where valuations have some basic structure.

We propose a modification of the VCG framework that incorporates redistribution of as much of the VCG surplus as possible back among the agents. We prove that—among all truthful, social welfare maximizing, and no-deficit mechanisms that meet basic anonymity and participation constraints—this “redistribution mechanism” is *optimally balanced*, in that no mechanism ever comes closer to budget-balance. We describe why in the case of completely unconstrained valuation functions no redistribution is possible, and characterize the domains in which it is.

We start by focusing our analysis on the special setting of “all-or-nothing” (AON) domains, those in which each outcome yields non-zero reward for just a single agent (every agent gets either all the reward or none). Here the redistribution mechanism has a particularly simple and elegant form, and the most redistribution is possible. In fact, in AON domains the mechanism is *completely* budget-balanced (i.e., redistributes all surplus) in the limit as the number of agents goes to infinity, for arbitrary valuations in an AON context. We demonstrate the applicability of the mechanism to general allocation problems, and describe empirical results that indicate the extent to which redistribution is possible for different degrees of dependency between agent valuations. Finally, we discuss implementation of the mechanism in the absence of a central enforcer.

2. BACKGROUND AND RELATED WORK

2.1 Notation

- Let I be the set of agents participating in a mechanism.
- Let $n = |I|$, the number of agents.
- Let O be the set of possible outcomes.
- Let $v_i(o)$ denote the *true* value to i of outcome $o \in O$.
- Let $\hat{v}_i(o)$ denote the *reported* value to agent i of outcome o .
- Let o^* denote the socially optimal outcome, according to reported valuations, i.e.:

$$o^* = \arg \max_{o \in O} \sum_{i \in I} \hat{v}_i(o)$$

- Let o_{-i}^* denote the outcome that maximizes payoff among agents other than i , according to reports, i.e.:

$$o_{-i}^* = \arg \max_{o \in O} \sum_{j \in I \setminus \{i\}} \hat{v}_j(o)$$

2.2 Mechanism design background

The field of mechanism design is concerned with reaching outcomes that meet certain system-level criteria, despite the assumption that individuals will pursue only their self-interest. A *mechanism* amounts to a specification of a decision rule, potentially accompanied by a set of requirements (typically monetary transfers) placed on participants. A

mechanism that maximizes social welfare is termed *socially optimal* or *efficient*. When no agent can ever benefit from lying, i.e., when being truthful is a dominant strategy, a mechanism is called *strategyproof*. A somewhat weaker property is *incentive-compatibility* (IC); in an IC mechanism truthfulness is a Nash equilibrium: no agent can benefit from being dishonest when all other agents are truthful.

The Groves class of mechanisms is *efficient* and *strategyproof* [Groves, 1973]. A Groves mechanism chooses the outcome that maximizes reported reward across all agents, and makes a transfer payment to each agent equal to the reward received by all other agents. In the most basic version, each agent’s payoff is equal to the reward received by the entire system; clearly such a mechanism would require a massive budget to implement. Fortunately, the Groves framework allows for the imposition of a charge on each agent, which does nothing to diminish the mechanism’s desirable incentive properties so long as each agent’s charge is completely beyond their control. Formally, a Groves mechanism chooses the outcome o^* that maximizes social welfare according to reported rewards, and defines the following transfer payment T_i to each agent i :

$$T_i = \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - C_{-i}$$

where C_{-i} can be *any quantity independent of agent i* .

In order to ensure that all agents want to participate in a mechanism, it is desirable that none ever be worse for having done so. A mechanism that provides this guarantee to every agent, for any possible set of agent valuation functions, is referred to as *ex post individual rational* (IR). A mechanism that yields non-negative payoff in expectation is termed *ex ante IR*. To obtain ex post IR in the Groves framework, we should ensure that for each i :

$$v_i(o^*) + \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) \geq C_{-i}$$

An important special case of the Groves class is the VCG mechanism [Vickrey, 1961; Clarke, 1971; Groves, 1973]. VCG defines C_{-i} above to be $\sum_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*)$, the reported reward that all other agents would have received if i were not a participant.² Thus VCG yields the following payoff, π_i , for each agent i :

$$\begin{aligned} \pi_i &= v_i(o^*) + T_i \\ &= v_i(o^*) + \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - \sum_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*) \end{aligned}$$

VCG is ex post IR, and is also *no-deficit*: the sum of payments made from agents to the center is never less than 0. But it is often greater—each agent’s net payment to the center is always between 0 and the reward it reported for the selected outcome, depending on the reported valuations of the other agents. In the case of AON domains, VCG is known simply as the Vickrey or second-price auction, and there is only one payment made—the winning agent must pay the second highest “price”.

Consider the following 4-agent AON example, where outcome o_i corresponds to the decision that yields value for agent i :

²Note that since all Groves mechanisms are strategyproof, $\hat{v}_j$ will equal v_j in dominant strategy equilibrium.

	v_1	v_2	v_3	v_4
o_1	10	0	0	0
o_2	0	8	0	0
o_3	0	0	5	0
o_4	0	0	0	4

Table 1: 4-agent AON decision problem.

VCG chooses outcome o_1 , as agent 1’s value (10) is highest, and the following payoffs result:

$$\pi_i = v_i(o^*) + \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - \sum_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*)$$

$$\pi_1 = 10 + 0 - 8 = 2$$

$$\pi_2 = 0 + 10 - 10 = 0$$

$$\pi_3 = 0 + 10 - 10 = 0$$

$$\pi_4 = 0 + 10 - 10 = 0$$

$$\text{total payment to center} = 8$$

In the case of an auction, this rather large payment of 8 may be desirable—it can go to the seller that is providing the item. But for a general social choice or decision problem, it may merely be an unwanted side-effect of choosing to implement VCG, the most famous strategyproof and efficient mechanism. In this work we explore the possibility of returning the surplus to the agents.

DEFINITION 1. (Redistribution mechanism) *An efficient social choice procedure that seeks to minimize net transfers from agents to an external body by return of VCG surplus to the agents.*

Our characterization of efforts to achieve minimal surplus in terms of redistribution is motivated by the centrality (and ubiquity) of VCG. In seeking to redistribute as much of the VCG surplus as possible (without violating truthfulness), we are effectively searching among all IR and no-deficit instances of the Groves class of mechanisms for the one that maximizes payoff to the agents. Green & Laffont [1977] showed that for general valuations, the only efficient and strategyproof mechanisms are the members of the Groves class. Holmstrom [1979] strengthened this result significantly, proving that Groves is also unique in any restricted domain in which each agent’s space of potential valuations is *smoothly connected*, i.e., when for any two valuations an agent could report, one can be differentiably deformed into the other. As Holmstrom notes, this result shows that “for all practical purposes” one must be content with Groves. Thus, our approach is comprehensive and will yield uniqueness properties for the mechanism we propose.

2.3 Related work

Work in the area of designing budget-balanced mechanisms is relatively scarce, presumably due to the primacy of the goal of social efficiency in combination with the strong negative result of the *Myerson-Satterthwaite impossibility theorem*: that no mechanism is capable of achieving individual rationality, efficiency, and budget balance at the same time for general valuation functions, even if we loosen our solution concept to Bayes-Nash equilibrium [Myerson and Satterthwaite, 1983]. This is an extension of the Hurwicz impossibility theorem [Hurwicz, 1975], which proves the result for dominant-strategy equilibria.

One proposal for achieving budget-balance is the d’AGVA mechanism [d’Aspremont and Gerard-Varet, 2002]. d’AGVA selects the efficient outcome according to reported types, but determines transfer payments based on a model of agent valuations that the center maintains. This mechanism is interesting primarily because it leads to exact budget-balance; however, it is implementable only in Bayes-Nash (rather than dominant-strategy) equilibrium, and if the center’s model is produced via iterative execution, serious problems regarding incentive compatibility can arise. Moreover, d’AGVA is only ex ante IR—instances in which an agent is worse off for participating are possible.

Parkes et al. [2001] describe a payment rule that approximates VCG and achieves perfect budget-balance in exchange settings where VCG runs a deficit, though truthfulness (and thus efficiency) is sacrificed. Faltings [2004] formulates the problem more closely to the way we do in this work, though his approach also attains exact budget-balance at the expense of efficiency. His mechanism chooses the outcome that is socially efficient among a subset of agents, and distributes the VCG surplus among agents that are not part of that subset. In a similar vein, Feigenbaum et al. [2001] analyze the Shapley-value mechanism for sharing multicast transmission costs, which comes closest to the efficient outcome among all budget-balanced mechanisms for that domain, though their results suggest its implementation is computationally intractable. Most similar to our proposal is that of Bailey [1997]. He specifies an efficient mechanism that refunds surplus to participants, but he forfeits the no-deficit guarantee in order to achieve a mechanism that is budget-balanced in expectation.

Our approach is significantly different than those described above—we characterize the extent to which budget-balance can be approximated (and sometimes reached) in dominant strategies, without sacrificing individual-rationality, efficiency, or no-deficit guarantees. Rather than giving up one or another of these properties, we are able to achieve budget-balance superiority over VCG by using domain information that constrains the space of possible valuations.

3. REDISTRIBUTION IN AON DOMAINS

It is instructive to start with the special case of AON domains, for which our solution is especially simple and elegant, as it provides good intuition for the general case. Formally, we define an AON domain as one in which the number of outcomes is equal to the number of agents, and each agent i ’s valuation function v_i is known to be constrained in the following way: $v_i(o_j) = 0, \forall j \neq i$. We identify o_i with the unique outcome that may yield positive value to agent i , and use the following short-hand notation:

- Let a_c denote the agent with the c^{th} highest reported value. (So in a welfare-maximizing mechanism a_1 is the “winning” agent—the one that receives positive value from the selected outcome.)
- Let $\mathbb{V}_{a_c}$ denote the *true* value to a_c of the outcome he favors (i.e., $v_{a_c}(o_{a_c})$), and $\hat{\mathbb{V}}_{a_c}$ the *reported* value (i.e., $\hat{v}_{a_c}(o_{a_c})$).

We propose the following redistribution mechanism for AON domains that maintains much of the VCG surplus among the agents:

DEFINITION 2. (Redistribution mechanism RM)

1. Each agent i communicates to the center $\mathbb{V}_i$, the value he places on the outcome under which he obtains non-zero reward.
2. The center implements the outcome that maximizes social welfare.
3. The winner (a_1) pays each agent (including himself) a VCG surplus redistribution of Z_i , and pays the center $\hat{\mathbb{V}}_{a_2} - \sum_{i \in I} Z_i$, where

$$Z_i = \begin{cases} \frac{\hat{\mathbb{V}}_{a_3}}{\hat{\mathbb{V}}_{a_2}} & \text{for } i = a_1, a_2 \\ \frac{\hat{\mathbb{V}}_{a_2}}{\hat{\mathbb{V}}_{a_2}} & \text{for } i = a_3, \dots, a_n \end{cases}$$

THEOREM 1. *RM is ex post individual-rational in AON domains.*

PROOF. The payoff to each agent for the implemented outcome in **RM** is:

$$\begin{aligned} \pi_{a_1} &= \mathbb{V}_{a_1} - \hat{\mathbb{V}}_{a_2} + Z_{a_1} \\ \pi_{j \neq a_1} &= Z_j \end{aligned}$$

The payoff to any “non-winner” j is always the non-negative Z_j payment term. When the mechanism is followed, $\mathbb{V}_{a_1} = \hat{\mathbb{V}}_{a_1}$; then since $\hat{\mathbb{V}}_{a_1} \geq \hat{\mathbb{V}}_{a_2}$ by definition, and since $Z_{a_i} \geq 0$, a_1 ’s payoff is also always non-negative. $\square$

THEOREM 2. *RM is strategyproof in AON domains.*

PROOF. **RM** differs from the VCG mechanism only in the Z payment term—each agent receives the standard VCG payoff plus a non-negative Z payment. Since we know VCG is strategyproof, it suffices to show that no agent i can increase its Z_i payment term by misreporting its reward. a_k below denotes the agent with the k^{th} highest bid under truthful reporting.

Agent a_1 receives redistribution payment (from himself) $\hat{\mathbb{V}}_{a_3}/n$. Over-reporting $\hat{\mathbb{V}}_{a_1}$ changes nothing. Under-reporting could put a_1 in the second or third position (beyond the third it’s obvious that nothing could change). In the second position, he would still receive $\hat{\mathbb{V}}_{a_3}/n$. In the third he would receive $\hat{\mathbb{V}}_{a_3}/n$ as well, since the second position would then be held by the actual a_3 .

Agent a_2 receives redistribution payment $\hat{\mathbb{V}}_{a_3}/n$. Over-reporting could move him to the first position, in which case his payment would be the same. Under-reporting could put him in the third position or beyond, but then the second position would be held by the actual a_3 , so a_2 ’s payoff would be the same.

Agent a_3 receives redistribution payment $\hat{\mathbb{V}}_{a_2}/n$. Under-reporting changes nothing. Over-reporting could put him in the first or second position. In both cases he receives the same Z payment, since the third position would then be held by the actual a_2 . The same holds for all $a_{j>3}$. $\square$

Theorems 1 and 2 together show that **RM** for AON domains always achieves the efficient outcome in dominant strategy equilibrium.

THEOREM 3. *As the number of participating agents n goes to ∞ , the amount of extracted wealth that cannot be redistributed among the agents under **RM** goes to 0. That is, **RM** is asymptotically budget-balanced for AON domains.*

PROOF. The amount of wealth redistributed among the agents is:

$$\sum_{i \in I} Z_i = \frac{n-2}{n} \cdot \hat{\mathbb{V}}_{a_2} + \frac{2}{n} \cdot \hat{\mathbb{V}}_{a_3}$$

The amount not redistributed is the payment to the center:

$$\begin{aligned} Z_c &= \hat{\mathbb{V}}_{a_2} - \sum_{i \in I} Z_i \\ &= \hat{\mathbb{V}}_{a_2} - \frac{n-2}{n} \cdot \hat{\mathbb{V}}_{a_2} - \frac{2}{n} \cdot \hat{\mathbb{V}}_{a_3} \\ &= \frac{2}{n} \cdot (\hat{\mathbb{V}}_{a_2} - \hat{\mathbb{V}}_{a_3}) \end{aligned}$$

We assume existence of some finite bound on valuations, and thus on $\hat{\mathbb{V}}_{a_2}$. $\square$

As n increases, we may expect Z_c to be pushed down by a convergence of $\mathbb{V}_{a_2}$ and $\mathbb{V}_{a_3}$, but regardless of this **RM** achieves perfect budget balance in the limit. *VCG will always lose $\hat{\mathbb{V}}_{a_2}$, no matter the number of agents.*

Consider once more the example illustrated in Table 1. The following payoffs are obtained under **RM**:

$$\begin{aligned} \pi_1 &= 10 - 8 + \frac{5}{4} = \frac{13}{4} \\ \pi_2 &= \frac{5}{4}, \quad \pi_3 = \frac{8}{4}, \quad \pi_4 = \frac{8}{4} \\ Z_c &= 8 - \left(\frac{5}{4} + \frac{5}{4} + \frac{8}{4} + \frac{8}{4} \right) = \frac{3}{2} \end{aligned}$$

Even in this example with just 4 agents, the vast majority (81%) of the VCG surplus has been redistributed. If $\mathbb{V}_{a_3}$ were 8 rather than 5, 100% would be redistributed.

3.1 Optimality of RM

We will now characterize the relationship between **RM** and all other possible redistribution mechanisms through a series of incremental observations. We seek the optimal way of redistributing the VCG payment back among agents in AON domains. More specifically, we wish to achieve the following:

DEFINITION 3. (Optimally balanced mechanism) *A mechanism is optimally balanced for domain D if—among all efficient, no-deficit, and ex post IR mechanisms that can be implemented in dominant strategy equilibrium—it maintains the most wealth within the set of agents, for every instance $d \in D$.*

A domain D specifies a space of possible valuations for every agent, and an instance $d \in D$ may be any realization of actual valuations consistent with D .

LEMMA 1. *In any strategyproof redistribution mechanism for smoothly connected valuation spaces, for every agent i , the amount of surplus redistributed to i , Z_i , must be independent of the valuation function i reports.*

PROOF. The Lemma follows from uniqueness of the Groves class. For intuition, consider an agent i with arbitrary valuation $v_i = \{v_i(o_1), \dots, v_i(o_n)\}$ yielding payment Z_i . If there is some outcome o_k and $\epsilon \in \mathbb{R}$ such that reporting $\hat{v}_i(o_k) = v_i(o_k) + \epsilon$ yields $Z_i + \gamma$ for some $\gamma > 0$ without changing the selected outcome, agent i will misreport to maximize payment. $\square$

To achieve efficiency in dominant strategy equilibrium, a redistribution mechanism must be an instantiation of the Groves class, so the redistribution can be considered part of the necessarily agent-independent Groves “charge” function.

Given Lemma 1, it will be useful now to consider each agent’s redistribution payment as a function of the reported rewards of all other agents in the system. So we have, $\forall i \in I$:

$$Z_i = f_i(\hat{v}_0, \hat{v}_1, \dots, \hat{v}_{i-1}, \hat{v}_{i+1}, \dots, \hat{v}_n)$$

for some f_i . The task of constructing a redistribution mechanism can then be viewed as the specification of n functions, each of which is defined over a distinct set of $n - 1$ of the valuation reports. It’s important to note that the above characterization is comprehensive, as each f_i is defined over all variables in the system save i ’s report, which it can’t depend on by Lemma 1.

But it is often preferred (or even required) that we do not discriminate amongst agents based on personal identity; we often want to implement mechanisms that are *anonymous*:

DEFINITION 4. (Anonymous mechanism) *A mechanism that chooses an outcome and defines transfer payments (including redistribution) according to a single deterministic function that is invariant to domain information that does not apply identically to every agent.*

In other words, anonymity requires that transfer payments be computed independent of “personal identity”. To construct an anonymous redistribution mechanism, then, in fact we define only a single function f that is used universally, but applied to a distinct set of $n - 1$ parameters for each agent. We know that an *optimally balanced* anonymous redistribution mechanism exists, since to define it amounts to specifying the value of f , for every possible set of $n - 1$ values, that is greatest while still ensuring that no deficit could result for any possible instantiation of the excluded n^{th} value.

Let $p_i^{(k)}$ denote the k^{th} largest parameter to the function that defines Z_i . $p_i^{(2)}$ is then the second highest bid among all agents other than i , and is also the greatest lower-bound (*guarantee*) on the VCG surplus that will exist, independent of the valuation i reports.

LEMMA 2. *The optimally balanced anonymous redistribution mechanism for AON domains must define redistribution payments Z_i , $\forall i \in I$, that are monotonically increasing in $p_i^{(2)}$.*

PROOF. Suppose the Lemma didn’t hold. Then there would be reported reward profiles A and $B \in \mathbb{R}^n$ with $\hat{v}_{a_2}(A) \leq \hat{v}_{a_2}(B)$, yet with $Z(A) > Z(B)$. In such a case the mechanism could not be optimal, as implementing $Z(A)$ in case B would keep more wealth among the agents. $\square$

THEOREM 4. RM *is the optimally balanced anonymous mechanism for AON domains.*

PROOF. Any anonymous, efficient, no-deficit, and ex post IR redistribution mechanism defines payments $Z_i \leq \frac{p_i^{(2)}}{n}$, $\forall i \in I$. Consider the parameters of f_i that would lead to Z_i being $> \frac{p_i^{(2)}}{n}$. Simply note that there exists a problem instance such that all $p_j^{(2)}$, $\forall j \in I$ are identical. This, when viewed in light of Lemma 2, shows that if Z_i were $> \frac{p_i^{(2)}}{n}$,

the mechanism would not be guaranteed IR and no-deficit, as $\sum_j Z_j$ could then be $> \hat{v}_{a_2}$.

Now observe that another way of describing **RM** is as follows: redistribute a portion of $\hat{v}_{a_2}$ to each agent i equal to $\frac{1}{n}$ times the value of the second highest reported reward of agents other than i . $\square$

4. REDISTRIBUTION IN THE GENERAL CASE

We have thus far only considered redistribution of VCG payments in AON domains. In this section we explore the general conditions under which strategyproof redistribution is possible, and characterize how to redistribute surplus optimally.

We first show that in the case of domains in which value functions are completely unrestricted, redistributing any of the transfer payments while maintaining the desirable properties of VCG is impossible. One way of understanding this is to consider the problem of finding a lower-bound on the total amount of VCG surplus that is guaranteed to exist, independent of any given agent’s reported valuation.

LEMMA 3. *In any efficient, strategyproof, ex post IR and no-deficit mechanism, an upper-bound on the redistribution payment to any agent i is the minimum VCG surplus that could be realized, taken over all possible instantiations of $\hat{v}_i$.*

PROOF. Among all efficient, strategyproof, and ex post IR mechanisms, VCG yields the greatest surplus [Krishna and Perry, 1998]. Thus a mechanism with these properties that redistributes a quantity greater than the VCG surplus will run a deficit, since charging any agent more would violate ex post IR.

Let $\hat{v}_i$ be the instantiation of $\hat{v}_i$ that results in minimum VCG surplus, S_i . Suppose for contradiction that i receives redistribution payment $Z_i > S_i$. By Lemma 1, Z_i must be the same regardless of the valuation i reports; thus when i reports $\hat{v}_i$, he will have been redistributed a quantity greater than the total VCG surplus, and the system will have run a deficit. $\square$

Lemma 3 points to a way of generally characterizing when redistribution is and is not possible. Whether or not a VCG surplus will result is intimately connected to the *similarity* of agents’ valuations. In the extreme, all agents may favor the same outcome, in which case there will be no surplus. In some domains, e.g., AON, we know a priori that agents’ valuations are distinct; such distinctions are in fact what enables redistribution. We now formalize these intuitions.

DEFINITION 5. (Potential for Universal Relevance Nullification [PURN]) *An agent i has the potential for universal relevance nullification, given the valuation reports of the other agents, if i could report a valuation $\hat{v}_i$ that would yield $o^* = o_{-j}^*$, $\forall j \in I$ (including i).*

The **PURN** property indicates that an agent could potentially report a value that renders null the influence on the outcome of any single agent’s report alone, even his own. **PURN** is often dependent on the values other agents have reported, and is then a property of one agent’s valuation domain combined with the other agents’ actual reports.

THEOREM 5. *No efficient, strategyproof, ex-post IR, and no-deficit mechanism redistributes any surplus to an agent in an instance in which PURN holds for that agent.*

PROOF. The total VCG surplus is:

$$\sum_{j \in I} \left(\sum_{k \in I \setminus \{j\}} \hat{v}_k(o_{-j}^*) - \sum_{k \in I \setminus \{j\}} \hat{v}_k(o^*) \right)$$

Here we sum over all agents $j \in I$ the VCG surplus j creates: the difference between j 's VCG "charge" delivered to the center and the Groves payment he receives. By definition, any agent i with *PURN* may report a $\hat{v}_i$ that would lead to this quantity being zero. The theorem follows directly from this and Lemma 3. $\square$

LEMMA 4. *Any agent with an unconstrained valuation always has PURN, regardless of the other agents' valuations.*

PROOF. Let i be an agent with an unconstrained valuation. Consider the outcome o_{-i}^* that maximizes utility among the $n - 1$ other agents. Let i 's valuation be as follows: $\hat{v}_i(o_{-i}^*) = \max_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*)$, and $\hat{v}_i(o) = 0$, $\forall o \neq o_{-i}^*$. Considering this full set of n valuations, the overall efficient outcome o^* will be identical to o_{-i}^* , $\forall j \in I$; this is precisely the condition for *PURN*. $\square$

PROPOSITION 1. *The VCG mechanism with no redistribution is optimally balanced for unconstrained valuations.*

PROOF. When valuations are unconstrained, by Lemma 4 each agent has *PURN*. By Theorem 5, then, no agent can receive any redistribution payment. $\square$

While this is a strong negative result, its applicability is hardly universal; in practice no valuation function is *entirely* unconstrained. Specifically, in determining the payment Z_i for any agent i , v_i need not be treated as a complete unknown. The AON case, for instance, is not caught by Proposition 1 because we need only consider valuations that are non-zero for just a single outcome.

In any given domain there may be several ways in which valuations are known to be restricted. Even the mere existence (and knowledge) of distinct finite bounds on each agent's valuation for any outcome can be enough to enable redistribution. *PURN* can be seen as one extreme in a spectrum of possible relationships between agent valuation spaces—rather than having the capacity to *completely* nullify the relevance of other agents in determining the outcome, an agent may have the potential to diminish others' impact only partially, or perhaps not at all. This capacity is of interest to us here because of the bearing it has on the amount of VCG surplus that will result. *Our ability to redistribute VCG surplus to an agent is directly tied to the extent to which we can know surplus will exist independent of that agent's reported valuation.*

DEFINITION 6. (**Surplus-guarantee S_i**) *The lower-bound on VCG surplus computed over all possible instantiations of $\hat{v}_i$, i.e., the greatest surplus guarantee that is independent of agent i .*

$$\begin{aligned} S_i &= \min_{\hat{v}_i} \left[\sum_{j \in I} \left(\max_{o' \in O} \left\{ \sum_{k \in I \setminus \{j\}} \hat{v}_k(o') \right\} - \sum_{k \in I \setminus \{j\}} \hat{v}_k(o^\dagger) \right) \right] \quad (1) \\ &= \min_{\hat{v}_i} \left[\sum_{j \in I} \left(\max_{o' \in O} \sum_{k \in I \setminus \{j\}} \hat{v}_k(o') \right) - (n-1) \left(\sum_{j \in I} \hat{v}_j(o^\dagger) \right) \right] \end{aligned}$$

subject to the constraint that the $o^\dagger$ and $\hat{v}_i$ in the above minimization must satisfy the following:

$$o^\dagger = \arg \max_{o \in O} \sum_{j \in I} \hat{v}_j(o) \quad (2)$$

Breaking down equation (1), within brackets is an expression representing the net payment—summed over all agents—to the center under VCG; this is the surplus we'll have to redistribute. For each agent i , we take the minimum of this expression over all possible valuations i could report.

Now note that a payment based on S_i will not always be *completely* anonymous, since we're using information about valuation constraints that are potentially unique to each agent. Specifically, it will not be anonymous in domains that are not *symmetric*.³ But interestingly, it turns out that a generalization of mechanism **RM** is optimal not only across the space of completely anonymous mechanisms, but also among mechanisms that meet just the following weaker requirement:

DEFINITION 7. (**Surplus-anonymous redistribution mechanism**) *A mechanism that chooses an outcome and maps agent-specific VCG surplus lower-bounds (S_i) to redistribution payments (Z_i) according to a single deterministic function that is invariant to domain information that does not apply identically to every agent.*

In a *surplus-anonymous* redistribution mechanism, then, for any two agents i and j : if $S_i = S_j$, then $Z_i = Z_j$.

THEOREM 6. *For all domains with smoothly connected valuation spaces, the redistribution mechanism **RM_g** that modifies the VCG mechanism by paying each agent i a quantity equal to $\frac{S_i}{n}$ is strategyproof, and is optimally balanced across all surplus-anonymous mechanisms.*

PROOF SKETCH. Strategyproofness follows immediately from strategyproofness of VCG plus the fact that S_i is independent of i 's reported valuation. The argument for optimality is omitted here for lack of space. $\square$

Any *completely-anonymous* mechanism is also *surplus-anonymous*, and the conditions for *complete-anonymity* and *surplus-anonymity* coincide in symmetric domains (e.g., AON). **RM_g** is thus *completely-anonymous* when applied to symmetric domains, and *surplus-anonymous* otherwise. It is easy to verify that the AON mechanism **RM** is a special case of **RM_g**— S_i in an AON domain is equal to the second highest of the bids placed by agents other than i .

It can be shown that no mechanism, including **RM_g**, is *optimally-balanced* when anonymity considerations are completely dropped. For instance, a mechanism that redistributes S_i to a randomly selected agent i , and none to other agents, will in some cases redistribute more surplus than **RM_g**; however, such a mechanism is clearly not *optimally-balanced*.

4.1 Redistribution in Allocation Problems

We've been discussing redistribution in the context of general social choice or decision-making problems. An important subclass of social choice problems consists of those in which a decision must be made on how to allocate goods amongst a group of competing agents. The most basic allocation problem is that in which a single item is available; this limited case (usually) falls within the AON class.⁴ Generally, an allocation problem consists of a number of sellers

³A *symmetric domain* is one in which the space of possible valuation reports is identical for each agent.

⁴The exception is when there are externalities, e.g., when agent i is happy if his friend agent j is allocated the item.

who bring goods to the system, and a number of interested buyers, each with preferences over the goods that are potentially combinatorial in nature.

Sometimes in such domains allocation mechanisms are selected to maximize revenue to the seller(s), but there are many significant cases in which it is preferable to keep as much wealth as possible in the hands of the agents. For example, consider the allocation of job time on a publicly owned super-computer among a community of researchers. Resources like this are often established with the mandate of maximizing benefit to the public good (efficiency); if a VCG-based allocation were implemented a surplus would result, fair redistribution of which would go further toward satisfying this mandate.

In the vast majority of allocation problems, the following are generally accepted to hold: agents receive no value if they aren't allocated anything (normalization); agents' values monotonically increase as they receive more goods (free disposal); and agents have no preferences over allocations to other agents (no externalities). These intrinsic elements of the allocation domain map directly to constraints on valuation functions that can allow for significant strategyproof redistribution of VCG surplus. Formally, let G be the set of goods to be allocated, and for any bundle of goods $b \subset G$ let $v_i(b)$ be agent i 's value for obtaining b . Each agent i 's valuation conforms to the following:

$$\begin{aligned} v_i(\emptyset) &= 0 \\ v_i(b) &\leq v_i(b \cup \{g\}), \quad \forall b \subset G, g \in G \end{aligned}$$

Consider the following valuations of three agents in an allocation problem with two goods, A and B , where $\{X_1, X_2, X_3\}$ represents the outcome in which bundles X_1 , X_2 , and X_3 are allocated to agents 1, 2, and 3 respectively.

	v_1	v_2	v_3
$\{AB, \emptyset, \emptyset\}$	12	0	0
$\{\emptyset, AB, \emptyset\}$	0	10	0
$\{\emptyset, \emptyset, AB\}$	0	0	11
$\{A, B, \emptyset\}$	4	6	0
$\{A, \emptyset, B\}$	4	0	5
$\{B, A, \emptyset\}$	5	7	0
$\{B, \emptyset, A\}$	5	0	7
$\{\emptyset, A, B\}$	0	7	5
$\{\emptyset, B, A\}$	0	6	7

Table 2: 2-good, 3-agent allocation problem.

The efficient outcome is to allocate A to agent 3 and B to agent 2. Applying $\mathbf{RM}_g$, agent payoffs are as follows:

$$\begin{aligned} \pi_i &= v_i(o^*) + \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - \sum_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*) + Z_i \\ \pi_1 &= 0 + 13 - 13 + \frac{8}{3} = \frac{8}{3} \\ \pi_2 &= 6 + 7 - 12 + \frac{11}{3} = \frac{14}{3} \\ \pi_3 &= 7 + 6 - 12 + \frac{10}{3} = \frac{13}{3} \end{aligned}$$

Using VCG with no redistribution, total social utility is 2 and payment to the center is 11. 9.67 of this (88% of surplus) can be redistributed to the agents under $\mathbf{RM}_g$, yielding a social utility of 11.67, a nearly 5-fold improvement.

5. SIMULATIONS

5.1 Computing redistribution payments

Determining redistribution payments under $\mathbf{RM}_g$ amounts to computing surplus guarantee S_i for each agent i , and then merely dividing by n . In a broad range of domains, a simple algorithm for computing S_i exists. For example, in AON and in fact any typical⁵ allocation domain, S_i can be computed as follows: consider v_i to be as high as possible for outcome o^* , and 0 for all other outcomes; then compute the VCG surplus with this v_i and all other agents' reported valuations. This simple algorithm does not hold in general (i.e., not for *all* possible sets of value constraints); however, we can *always* compute S_i through a mixed-integer programming (MIP) specification of equation (1). We briefly outline the formulation here.

The objective in the MIP for agent i 's payment is to minimize VCG surplus, and the primary variables are $v_i(o)$ for each $o \in O$. Representing the program constraints is relatively straightforward, but some care must be taken in handling equation (2) and the inner maximization in (1). In both cases there are non-linearities. In (2), for instance, we must specify $c_o \cdot v_i(o)$ for each outcome o , where c_o is a boolean variable representing whether or not o is chosen as the outcome that maximizes social welfare. We can get around this issue by representing $c_o \cdot v_i(o)$ with a new variable $v'_i(o)$, and including the following constraints:

$$\begin{aligned} v'_i(o) &\leq v_i(o) \\ v'_i(o) &\leq c_o \cdot M \end{aligned}$$

where M is a value larger than the maximum possible value an agent could have for any outcome. $v'_i(o)$ will then be $v_i(o)$ if $o^* = o$, and 0 otherwise, as desired.

While solving a mixed-integer program has exponential worst-case running time, in practice we were able to quickly find solutions to very large problems. Determining redistribution payments in a 100 agent, 100 outcome problem took 24 seconds for each agent.⁶ Note that the MIPs for calculating agent payments (one for each agent) are independent of each other, and thus all can be solved in parallel.

5.2 Numerical results

In order to understand how much redistribution can be achieved for valuations under different levels of mutual constraint, we performed an empirical analysis on large sets of randomly generated problem instances (sets of valuations), each with the same number of outcomes as agents. We generated valuations according to the following process, where e is an "exclusivity" (between agent valuations) parameter representing the extent to which a domain has "all-or-nothing properties." We chose a maximum value for each agent i 's valuation function uniformly at random between 0 and 100. We then chose each $v_i(o_i)$ uniformly at random between 0 and $\max val_i$, and $v_i(o_{j \neq i})$ uniformly at random between 0 and $(1 - e) \cdot \max val_i$. So when $e = 1$, we have a completely AON domain; when $e = 0$ there are no exclusivity constraints between agent valuations. The graph

⁵One without valuation space constraints beyond the standard normalization, free disposal, and no externalities assumptions.

⁶Solutions were obtained using the commercial solver CPLEX, run on a 1.6 GHz Pentium 4 PC.

below plots the percentage of VCG surplus that could be redistributed as a function of the parameter e , for problems with various numbers of agents. For each number of agents, 100 samples were computed for each value of e between 0 and 1 in increments of 0.05.




Figure 1: Percent of surplus redistributed as a function of mutual “exclusivity” between agent valuations.

Notably, redistribution remains nearly constant for values of e between 0 and 0.5, and then increases roughly linearly with e from 0.5 to 1. As expected, the possibility for redistribution grows with the number of agents.

6. A CENTER-LESS IMPLEMENTATION

Consider the problem of building a decision-making mechanism that doesn’t require a center. This is intuitively desirable, as it’s easy to imagine scenarios in which a group of agents must determine an allocation or decide on an action to take without the benefit of a trusted central authority.⁷

THEOREM 7. $\mathbf{RM}_g$ can be implemented in Nash equilibrium without the facilitation of a center when the following conditions hold:

- i) no agent acting alone has the power to obstruct realization of a specified outcome or transfer payments, while $n - 1$ agents acting together do.
- ii) agents can simultaneously “broadcast” valuations to all other agents, and can compute and perform transfers.
- iii) agents have the capacity to publicly destroy money.

PROOF SKETCH. The first condition brings abiding by the mechanism into Nash equilibrium—unilaterally deviating cannot be beneficial. The second and third conditions ensure that agents will be able to execute what the mechanism prescribes. With (ii), all agents can recognize the efficient outcome, and deliver the appropriate redistribution payment to each “receiving” agent. (iii) allows the “paying” agents (just a_1 in AON domains) to demonstrably receive the appropriate payoff. The quantity Z_c that cannot stay within the group of agents would normally be transferred to a center, but its destruction is a satisfactory substitute. $\square$

In AON domains, at least, where communicating a valuation amounts merely to a public announcement of a single (value, outcome) pair, satisfaction of these conditions is

⁷See [Shneidman and Parkes, 2004; Petcu et al., 2006] for further discussion of distributed implementations that seek to minimize the role of a center.

quite plausible. Then using $\mathbf{RM}$, it is possible in equilibrium for a large group of self-interested agents to independently reach the socially optimal outcome, and jointly reap nearly all fruits of the chosen action via redistribution.

7. CONCLUSION

In this paper we argued for the desirability of redistributing VCG payments back among participants in a mechanism, and showed it is feasible in a broad range of settings, including allocation problems. We presented the *optimally balanced* redistribution mechanism, that which always redistributes the maximal amount that any strategyproof, efficient, IR, and no-deficit mechanism could. In “all-or-nothing” domains, a class which encompasses all typical single-item allocation problems, the mechanism is perfectly budget-balanced in the limit.

The VCG mechanism maximizes transfers *from* the group of agents, which is sometimes appealing because transfers amount to revenue for a center. But in many decision problems the transfers are, in a sense, actually waste; in this paper we specified the mechanism that minimizes that waste.

Acknowledgments

I’d like to thank David C. Parkes for inspiration, as well as very helpful suggestions and comments on a draft.

8. REFERENCES

- [Bailey, 1997] Martin J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91:107–126, 1997.
- [Clarke, 1971] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.
- [d’Aspremont and Gerard-Varet, 2002] Claude d’Aspremont and Louis-Andre Gerard-Varet. Incentives and incomplete information. In *Game Theory in the Tradition of Bob Wilson*, 2002.
- [Faltings, 2004] Boi Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-mediated E-commerce AMEC*, 2004.
- [Feigenbaum et al., 2001] Joan Feigenbaum, Christos H. Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1):21–41, 2001.
- [Green and Laffont, 1977] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [Groves, 1973] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [Holmstrom, 1979] Bengt Holmstrom. Groves’ scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.
- [Hurwicz, 1975] Leonid Hurwicz. On the existence of allocation systems whose manipulative nash equilibria are pareto optimal. (presented at the 3rd World Congress of the Econometric Society), 1975.
- [Krishna and Perry, 1998] Vijay Krishna and Motty Perry. Efficient mechanism design. Game theory and information, Economics Working Paper Archive at WUSTL, December 1998.
- [Myerson and Satterthwaite, 1983] Roger Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [Parkes et al., 2001] David C. Parkes, J. R. Kalagnanam, and M. Eso. Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proc. 17th Int. Joint Conf. on Artificial Intelligence*, pages 1161–1168, 2001.
- [Petcu et al., 2006] Adrian Petcu, Boi Faltings, and David C. Parkes. MDPOP: Faithful distributed implementation of efficient social choice problems. In *Proc. 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems*, 2006.
- [Shneidman and Parkes, 2004] Jeffrey Shneidman and David C. Parkes. Specification faithfulness in networks with rational nodes. In *Proc. 23rd ACM Symp. on Principles of Distributed Computing*, 2004.
- [Vickrey, 1961] William Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.