

Handling Self-Interest in Groups, with Minimal Cost

Ruggiero Cavallo

Division of Engineering and Applied Sciences
Harvard University
33 Oxford St., Cambridge, MA 02138
cavallo@eecs.harvard.edu

Abstract

In group decision-making problems that involve self-interested agents with private information, reaching socially optimal outcomes requires aligning the goals of individuals with the welfare of the entire group. The well-known VCG mechanism achieves this by requiring specific payments from agents to a central coordinator. However, when the goal of coordination is to allow the group to jointly realize the greatest possible welfare, these payments amount to an unwanted cost of implementation, or waste. While it has often been stated that the payments VCG prescribes are *necessary* in order to implement the socially optimal outcome in dominant strategies without running a deficit, this is in fact not generally true. (Cavallo 2006) specified the mechanism that requires the minimal payments among all mechanisms that are socially optimal, never run a deficit, and are ex post individual rational with an anonymity property. The mechanism achieves significant savings over VCG in a broad range of practically relevant domains, including allocation problems, by using information about the structure of valuations in the domain. This paper gives a high-level overview of that result, and discusses some potential applications to AI.

Introduction

Consider an environment in which a group of independent agents have distinct values for realization of various events. In such settings, organized coordination of the group behavior can very often bring about outcomes that are significantly more desirable—under various metrics—than would be possible otherwise. There are two problems that any prospective coordinator must address: he must determine what coordination scheme is optimal, and must get the agents to participate in that scheme. The latter challenge essentially amounts to dealing with agent *self-interest*.

The field of *mechanism design* is concerned with this very problem: how can we bring about outcomes that meet desirable system-level criteria, despite the fact that agents are self-interested, driven only to maximize their own utility? Solutions generally involve specifying particular payments either to or from the agents, which can be a powerful tool in incenting desired behaviors. Of course, when no external budget is available net payments should flow only from the

agents to the coordinator, yet to ensure that agents want to participate, no agent should have to pay more than the value he derives from the outcome the mechanism chooses.

(Cavallo 2006) described the mechanism that—among all those that guarantee non-negative payoff, don't ever require an external budget, and always reach a system-welfare maximizing outcome in dominant strategy equilibrium¹—*minimizes* the payments that must be made by the agents to the coordinator. Since the point of coordinating the agents is often simply to yield greater utility for the group, from one point of view this mechanism achieves the desired outcome with the minimal “cost of implementation” or “waste”.

The goal of this paper is to give a high-level presentation of that mechanism, and to discuss some potential applications to AI. I first provide more motivation for trying to achieve minimal agent-to-coordinator transfers, as well as the necessary background.

Motivation and Background

Consider the following story: A group of independent farmers in one particular county have struggled to obtain high yield from their fields in recent years. By good fortune, the group receives a gift from an area philanthropist—a special fertilizer that will help them use their land more efficiently. The gift is given with the following condition: the fertilizer should be applied in a way that maximizes the benefit it produces, i.e., the social welfare derived from its use. However, there is only enough fertilizer for one farmer to cover his land, so the group faces the problem of deciding which farmer should get to use the fertilizer. Complicating matters, each farmer is self-interested, driven only to maximize his own benefit from the gift, and individual values for the fertilizer are private. The group seeks to apply a mechanism for allocating the fertilizer that meets the benefactor's condition, while keeping the maximum amount of wealth within the group.

This is just one of a practically limitless number of examples one can conjure up where a group of self-interested agents faces a decision problem, and either the group itself or an overseeing authority seeks to implement a solu-

¹In a *dominant strategy equilibrium*, each player's behavior maximizes his payoff, for any possible behavior the other agents may exhibit.

tion that maximizes total group welfare. Other instances include: a group of housemates deciding who gets to use a jointly-owned automobile on a given night; a group of siblings choosing who should get the last piece of an apple pie; and perhaps more relevantly, allocation of time on a publicly owned super-computer to researchers with distinct projects.

Efficient mechanisms: the Groves family

Reaching a socially optimal (also called *efficient*) outcome despite agent self-interest can be achieved by establishing specific monetary payment rules between agents and a coordinator. Let I be the set of agents, and O be the set of possible outcomes or decisions that are available. The Groves family of mechanisms (Groves 1973) achieves efficiency by: 1) querying each agent $i \in I$ for the value $\hat{v}_i(o)$ he places on each outcome $o \in O$; 2) implementing the socially highest valued outcome according to what the agents report, which I denote o^* ; and 3) making the following transfer payment T_i to each agent i , where C_{-i} is any quantity that is completely independent of what i reports:

$$T_i = \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - C_{-i}, \quad (1)$$

Each agent thus receives a payment equal to the total utility derived by all other agents combined, minus some charge that is out of his control. Since he can't affect the charge, he will act to maximize the first term plus his own value for the chosen outcome, which aligns his interest completely with that of the entire group and makes truth-telling a dominant strategy. In fact it is known (Green & Laffont 1977; Holmstrom 1979) that when agent valuation spaces are smoothly connected, *any* mechanism that achieves the efficient outcome in dominant strategy equilibrium is a Groves mechanism.² If this strong form of efficiency is what we're after, then, our flexibility is limited to how we define the agent-independent charge term, C_{-i} .

When C_{-i} is defined such that T_i is negative, agents are required to make payments, rather than receive them. While there are many examples of decision-problems where maximizing "revenue"—the sum of the payments from the agents to a decision-maker—is desired, in those mentioned above and many others, it would clearly be preferable to minimize such payments.

If we wish to guarantee that no agent will be worse off for having participated, i.e., achieve *ex post individual rationality* (IR), we must specify C_{-i} such that it is never greater than $\sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*)$ (the "Groves payment") plus $v_i(o^*)$ (the reward i naturally receives from o^* 's implementation). If we wish to maintain an *ex post no deficit* property, that places other constraints on each C_{-i} .

Improving on VCG

The result reviewed in this paper specifies the mechanism that minimizes the payments required of the agents while

²A *smoothly connected domain* is one in which any two valuations an agent could report have the property that one can be differentiably deformed into the other. This holds quite broadly (as Holmstrom (1979) notes, for "all practical purposes").

satisfying these constraints and maintaining an anonymity property, which I describe in the next section. I present the solution as a "redistribution mechanism" because of the relationship it bears to the well-known Vickrey-Clarke-Groves (VCG) mechanism. VCG specifies each C_{-i} above to be the net welfare that agents other than i would have achieved if i were not present. Under VCG:

$$T_i = \sum_{j \in I \setminus \{i\}} \hat{v}_j(o^*) - \sum_{j \in I \setminus \{i\}} \hat{v}_j(o_{-i}^*) \quad (2)$$

where o_{-i}^* is the outcome that maximizes welfare (according to reports) among agents other than i . It has been shown (Krishna & Perry 1998) that VCG *maximizes* revenue among mechanisms with the properties we desire. Thus, one can conveniently frame the goal of *minimizing* revenue as that of *redistributing* back among the agents the net payments made in VCG, which I will refer to as the "VCG surplus".

It has often been stated that no such redistribution is possible in dominant strategy equilibrium (see, e.g., Ephrati & Rosenschein 1991), but this is actually not true when there is certain known information about agent valuations prior to agent reports. In reality, it is quite rare (if ever) that nothing at all is known about agent valuations. To mention one broad and important domain where this doesn't hold, consider standard single-item allocation problems. Here we know that each agent values only the outcome in which he is allocated the good (in our farming example, each farmer only obtains value when he gets to use the fertilizer). In this domain and many others, significant redistribution of the VCG surplus is, in fact, possible.

Previous approaches

There have been several previous approaches seeking to achieve improved budget-balance over VCG, all of which give up one or more of the following criteria: efficiency, *ex post* IR, *ex post* no deficit. The d'AGVA mechanism (d'Aspremont & Gerard-Varet 2002) is perfectly budget-balanced but not *ex post* IR; (Faltings 2004) also achieves perfect budget-balanced, but in a mechanism that is efficient for only a subgroup of agents; the threshold rule of (Parkes, Kalagnanam, & Eso 2001) achieves perfect budget-balance in an exchange setting at the expense of truthfulness (and thus efficiency). Finally, the mechanism of (Bailey 1997) also redistributes VCG surplus, but there are instances for which it runs a deficit.

The Result

The key insight of the mechanism proposed in (Cavallo 2006) is this: if it can be determined that a certain amount of VCG surplus will definitely exist, independent of what a particular agent reports to the coordinator, then we can redistribute surplus back to that agent without creating an incentive to misreport valuation information. It will be useful to explicitly define this agent-independent guarantee on surplus. First, observe that the total VCG surplus is simply the sum of the payments made from agents to the coordinator:

$$\text{VCG surplus} = \sum_{i \in I} \sum_{j \in I \setminus \{i\}} [\hat{v}_j(o_{-i}^*) - \hat{v}_j(o^*)] \quad (3)$$

Definition 1. (Surplus-guarantee S_i) The minimum amount of VCG surplus that could be realized, taken over all possible instantiations of agent i 's valuation function v_i :

$$S_i = \min_{\hat{v}_i} \left\{ \sum_{i \in I} \sum_{j \in I \setminus \{i\}} \left[\hat{v}_j(o'_{-i}) - \hat{v}_j(o') \right] \right\} \quad (4)$$

where o' and o'_{-i} are analogs of o^* and o^*_{-i} computed with the hypothetical \hat{v}_i in place of the actually reported value.

The proposed redistribution mechanism pays out a portion of the VCG surplus to each agent that is directly proportional to the agent's surplus-guarantee level.

Definition 2. (Redistribution Mechanism \mathbf{RM}_g)

1. Each agent i communicates to the coordinator his valuation function \hat{v}_i .
2. The coordinator implements the outcome that maximizes social welfare according to reports.
3. Each agent is assigned the following transfer payment, where n is the number of agents:

$$T_i = \sum_{j \in I \setminus \{i\}} \left[\hat{v}_j(o^*) - \hat{v}_j(o^*_{-i}) \right] + \frac{S_i}{n} \quad (5)$$

\mathbf{RM}_g is *strategyproof*: reporting \hat{v}_i equal to the true valuation v_i is always a dominant strategy, for every agent i . Then T_i above will always be non-positive, between $-v_i(o^*)$ and 0, so the mechanism is *ex post IR* yet never runs a deficit. \mathbf{RM}_g is also *surplus-anonymous*: any two agents i and j with $S_i = S_j$ receive the same amount of “redistribution” of the VCG surplus. Note that any traditionally *anonymous* mechanism (i.e., one that does not discriminate at all based on agent identity) is also *surplus-anonymous*.

The mechanism has a very simple and intuitive form when agent valuations are completely exclusive, as in single-item allocation problems, where only one agent has non-zero value for any given outcome. In such settings, values are “bids” for the item, and \mathbf{RM}_g reduces to the following:

Definition 3. (\mathbf{RM}_g in single-item allocation settings)

1. Each agents reports its value (bid) for the item.
2. The coordinator allocates the item to the agent with the highest bid.
3. The winning agent pays the coordinator an amount equal to the second highest bid.
4. The coordinator pays the winner and the second highest bidder an amount equal to the third highest bid divided by the number of agents, n , and pays all other agents the second highest bid divided by n .

Each agent's redistribution payment is $\frac{1}{n}$ times the second highest bid among the $(n - 1)$ other agents. In single-item allocation settings, \mathbf{RM}_g is *perfectly budget-balanced* in the limit. That is, as the number of agents approaches infinity the payments go to 0, regardless of the agent valuations. In VCG waste is always equal to the second highest

bid, regardless of the number of agents. To more clearly see the relationship between the basic Groves mechanism (BG) with C_{-i} defined to be 0, VCG, and \mathbf{RM}_g , consider the outcomes they yield in the following 4-agent, single-item allocation example. Let $\pi(M)$ be the payoff (value of the outcome, plus transfer) received under mechanism M , and let \mathbb{C} denote the coordinator.

agent	bid	$\pi(\mathbf{BG})$	$\pi(\mathbf{VCG})$	$\pi(\mathbf{RM}_g)$
1	15	15	3	5
2	12	15	0	2
3	8	15	0	3
4	4	15	0	3
\mathbb{C}	-	-45	12	2

Table 1: Comparison of mechanisms in an allocation problem.

The basic Groves mechanism runs an enormous deficit, while VCG eliminates the deficit and in fact yields a large revenue. \mathbf{RM}_g keeps significantly more wealth among the agents, while still avoiding a deficit. It turns out that no mechanism that always maintains the budgetary and incentive properties we seek keeps more wealth among the agents than \mathbf{RM}_g in *any* problem instance (where an instance is defined by a specific set of agent valuations).

Theorem 1. (Cavallo 2006) *For smoothly connected valuation spaces, among all efficient, ex post IR, no deficit, and surplus-anonymous mechanisms that can be implemented in dominant strategy equilibrium, \mathbf{RM}_g maintains the most wealth within the group of agents in every problem instance.*

Potential Applications to AI

The result described here is obviously not one that pertains specifically to artificial intelligence; it is rather a general observation about how one can achieve a system of decision-making that reaches socially-optimal outcomes while incurring minimal cost. But there are many problems within AI to which the proposed mechanism may be fruitfully applied.

Designing mechanisms to effectively deal with self-interest in multi-agent systems has been a compelling and growing topic of study in recent AI literature (see Dash, Jennings, & Parkes 2003); moreover, often these mechanisms are designed primarily to yield payoff to the participating agents. The application of \mathbf{RM}_g would seemingly be a more appropriate solution than VCG or other mechanisms in these environments. Perhaps the most natural applications for the mechanism are problems in which scarce resources must be allocated. Here I will go through one such example in detail, then touch briefly on a second application area.

Application: Job scheduling

Consider the problem of allocating processor time on a server to computational tasks, each of which is “owned” by a selfish human client who has private knowledge of his own value for having the task completed. Say there are 3 jobs requesting processor time in a single day, each of which requires half a day to complete; so there are 3 clients, each owning 1 job, competing for 2 time-slots. Moreover, assume clients obtain different value depending on whether the job is completed in the first or second time-slot, and that no value

is obtained if the job is not completed (the day is a “hard deadline”). Consider valuations of the following form:

Agent	$v(\text{time}_1)$	$v(\text{time}_2)$
1	10	5
2	7	4
3	6	2

Table 2: Valuations of 3 agents over 2 time-slots.

It is helpful to consider expansion of the above representation into a table of values over every possible outcome. I denote the outcome in which agent i 's job is processed in the first slot and j 's in the second as (i, j) .

outcome	v_1	v_2	v_3
(1, 2)	10	4	0
(1, 3)	10	0	2
(2, 1)	5	7	0
(2, 3)	0	7	2
(3, 1)	5	0	6
(3, 2)	0	4	6

Table 3: Enumeration of agent values, with outcome o^* in bold.

If the server is a public resource established with the mandate of increasing benefit to the social good, it will be desirable to maintain as much wealth as possible among the job-owners. We can compare the results that VCG and RM_g yield when applied to this example. The outcome o^* that maximizes social welfare is (1, 2), which yields total value 14. To compute the transfer payments, we need also to recognize the outcomes that *would* have been implemented in the absence of each of the agents. We have the following: $o_{-1}^* = (3, 2)$; $o_{-2}^* = (1, 3)$; $o_{-3}^* = (1, 2)$.

To execute RM_g we additionally need to determine the agent-independent surplus-guarantee S_i for each agent i . In this domain the VCG surplus is minimized when an agent reports value 0 for each time-slot, so it is simple to compute $S_1 = 3$, $S_2 = 4$, and $S_3 = 3$. We can now determine the payoff π_i for each agent i under VCG and RM_g (both mechanisms will reach the efficient outcome o^*). Under VCG $\pi_i = v_i(o^*) + \sum_{j \neq i} v_j(o^*) - \sum_{j \neq i} v_j(o_{-i}^*)$. Under RM_g , $\pi_i = v_i(o^*) + \sum_{j \neq i} v_j(o^*) - \sum_{j \neq i} v_j(o_{-i}^*) + \frac{S_i}{n}$. We have the following:

	VCG	RM_g
π_1	$10 + 4 - 10 = 4$	$10 + 4 - 10 + \frac{3}{3} = 5$
π_2	$4 + 10 - 12 = 2$	$4 + 10 - 12 + \frac{4}{3} = 3.33$
π_3	$0 + 14 - 14 = 0$	$0 + 14 - 14 + \frac{3}{3} = 1$
π_C	8	4.67

Table 4: Payoffs in the job scheduling example.

A significantly greater portion of the wealth derived from the server use is maintained among the agents under RM_g .

Application: Distributed constraint optimization

Distributed constraint optimization (Modi *et al.* 2005) represents another potential application area for RM_g . Consider a setting in which there is a set of agents that seeks to maximize reward by achieving an optimal joint state, where

each agent has control over some subset of the problem variables, which are constrained by the values of other agents' variables. Here agents' costs/rewards may be private, and achieving the optimal solution could require their communication amongst the group. There have been VCG-based proposals for handling incentive issues in this environment (see, e.g., Faltings & Macho-Gonzalez 2003); the “waste” incurred by VCG would be minimized by RM_g , which would be significant in settings where valuations have some degree of mutual exclusivity.

Conclusion

Coordinating decision-making in groups of independent agents requires disarming self-interest, by creating an environment in which each individual is best off when the group is best off. It is desirable that the mechanism used to achieve this solution not run a deficit, and at the same time guarantee that individuals won't be worse off from participating; it is also often desirable that the mechanism not discriminate among agents based on their identity. The mechanism RM_g specified in (Cavallo 2006) minimizes implementation cost among all solutions that reach system-welfare maximizing outcomes with these properties. RM_g can be considered a superior solution to VCG for domains in which the objective is to maximize payoff to the agents.

References

- Bailey, M. J. 1997. The demand revealing process: To distribute the surplus. *Public Choice* 91:107–126.
- Cavallo, R. 2006. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proc. 5th Int. Conf. on Auton. Agents and Multi-Agent Systems (AAMAS'06)*.
- Dash, R. K.; Jennings, N. R.; and Parkes, D. C. 2003. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems* 40–47.
- d'Aspremont, C., and Gerard-Varet, L.-A. 2002. Incentives and incomplete information. In *Game Theory in the Tradition of Bob Wilson*.
- Ephrati, E., and Rosenschein, J. S. 1991. The clarke tax as a consensus mechanism among automated agents. In *Proc. 9th Annual Conf. on Artificial Intelligence (AAAI-91)*, 173–178.
- Faltings, B., and Macho-Gonzalez, S. 2003. Incentive compatible open constraint optimization. In *AAMAS'03*, 986–987.
- Faltings, B. 2004. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-mediated E-commerce VI*.
- Green, J., and Laffont, J.-J. 1977. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45:427–438.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Holmstrom, B. 1979. Groves' scheme on restricted domains. *Econometrica* 47(5):1137–1144.
- Krishna, V., and Perry, M. 1998. Efficient mechanism design. Game theory and information, Economics Working Paper Archive at WUSTL.
- Modi, P.; Shen, W.; Tambe, M.; and Yokoo, M. 2005. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence Journal* 161:149–180.
- Parkes, D. C.; Kalagnanam, J. R.; and Eso, M. 2001. Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proc. 17th Int. Joint Conf. on Artif. Intel.*, 1161–1168.