Auctioning or assigning an object: some remarkable VCG mechanisms

HERVÉ MOULIN
Rice University

1. AUCTIONING A GOOD

A second price Vickrey auction is a simple mechanism to transfer a valuable object (a good) between a seller and n potential buyers. It treats buyers fairly, elicits their truthful valuations for the good, and allocates the good efficiently. But the division of the surplus between the seller and the buyers is hardly compelling. Writing a^{*k} for the k-th highest valuation among buyers, and setting without loss of generality the seller's valuation at zero, of the total surplus a^{*1} the seller gets a^{*2} and the buyers $a^{*1} - a^{*2}$, thus either share can be arbitrarily large or small, depending on the distribution of valuations. Moreover at most one buyer gets any surplus at all.

The familiar Vickrey-Clarke-Groves (thereafter VCG, see [Green and Laffont 1979]) mechanisms preserve the incentives and efficiency properties of the Vickrey auction, but allow more flexibility in the division of the surplus. Suppose that explicit guidelines regulate the division of surplus: the seller should get λa^{*1}, the buyers $(1-\lambda)a^{*1}$. We construct a VCG mechanism achieving such division with a margin of error as small as possible. For any numbers λ_-, λ_+ such that

$$0 \leq \lambda_- \leq \lambda_+ \leq 1 \text{ and } \lambda_+ - \lambda_- = \frac{n}{2^{n-1}-1} \tag{1}$$

our mechanism guarantees to the seller, at *all* profiles of non-negative valuations $(a_j, j \in N)$, a revenue between $\lambda_- a^{*1}$ and $\lambda_+ a^{*1}$. For instance with 10 potential buyers, any given share λ can be approximated within 2% at all profiles.

The simplest way to describe this (or any other VCG) mechanism is by the net utility $U_i(a)$ of agent i at the profile $a = (a_j, j \in N) \in \mathbb{R}_+^N$, where no object and no cash transfer yields $U_i = 0$. Thus if i does not get the object, e.g., if $a_i < a^{*1}$, $U_i(a)$ is a cash transfer to i, whereas if i receives the object he pays $a^{*1} - U_i(a)$ for it. We use the notation $B_m^k = \sum_{l=k}^{m} \binom{m}{l}$, which decreases as k grows. Assume $n \geq 3$ and pick any λ_-, λ_+ as in (1). Then we define

$$U_i(a) = a^{*1} - \frac{n-(1-\lambda_-)}{n-1}\{\sum_{k=1}^{n-1} \frac{(-1)^{k-1}}{\binom{n-2}{k-1}} \frac{B_{n-1}^k}{B_{n-1}^1} a_{-i}^{*k}\} \text{ for any } a \in \mathbb{R}_+^N \tag{2}$$

where a^{*k}_{-i} is the k-th highest valuation among buyers other than i. We have for all a

$$(1-\lambda_+)a^{*1} \leq \sum_{i\in N} U_i(a) \leq (1-\lambda_-)a^{*1}$$

We also show that no tighter bounds on the respective shares of seller and buyers can be implemented by a strategyproof mechanism (including a non VCG one) treating equals equally.

An interesting special case is $\lambda_+ = 1$, i.e., we give most of the surplus to the seller. In this case the best lower bound on its share is $\lambda_- = 1 - \frac{n}{2^{n-1}}$ (slightly better than suggested by (1)), so the buyers get a non-negative share of surplus not larger than $\frac{n}{2^{n-1}}$.

Symmetrically, if we wish to give most of the surplus to the buyers, we set $\lambda_- = 0$. The corresponding mechanism (introduced in [Guo and Conitzer 2007a; Moulin 2007a])

$$U_i(a) = a^{*1} - \{\sum_{k=1}^{n-1} \frac{(-1)^{k-1}}{\binom{n-2}{k-1}} \frac{B^k_{n-1}}{B^1_{n-1}} a^{*k}_{-i}\} \tag{3}$$

is meaningful if the object is the common property of the agents, so the seller is replaced by a *residual claimant* burning the cash surplus generated by the mechanism. We speak in this case of an *assignment problem.* The cash transfer to the residual claimant never exceeds the share $\widehat{\rho} = \frac{n-1}{2^{n-1}-1}$ of the efficient surplus:

$$0 \leq a^{*1} - \sum_{i\in N} U_i(a) \leq \widehat{\rho} a^{*1} \text{ for all } a \in \mathbb{R}^N_+$$

The left-hand inequality above states that the mechanism is *self sufficient* (also known as *feasible*), it never generates a budget deficit.

Importantly, participation in the mechanism (3) is voluntary ($U_i(a) \geq 0$ for all a), because the coefficients of a^{*k}_{-i} in equation (3) alternate in sign and decrease in absolute value.

By contrast, for any strictly positive share λ_-, participants in the mechanism (2) may end up with a net loss: the payment by the agent who gets the object is always non-negative, but can exceed a^{*1} slightly [1]; similarly an agent receiving no object generally receives some cash, but could end up losing some[2].

2. ASSIGNING A GOOD FAIRLY

In the assignment problem, the mechanism (3) is "fair" because it treats equals equally, and guarantees voluntary participation. The same is true of the Vickrey auction, but in the Vickrey auction the budget surplus (that we call a budget *loss*) may equal the entire efficient surplus a^{*1}, while for our mechanism (3) it never exceeds the share $\widehat{\rho}$ of a^{*1}.

A natural strenghtening of voluntary participation is often discussed in the literature (e.g., [Moulin 1992; Cramton et al. 1987]):

[1]It is as high as $(1 + \frac{\lambda_-}{n-1})a^{*1}$ at the profile where exactly two agents have the same positive valuation, and other valuations are zero.

[2]He will pay as much as $\frac{\lambda_-}{n-1}a^{*1}$ at the profile where only one agent has a positive valuation.

—*Unanimity lower bound*: $U_i(a) \geq \frac{a_i}{n}$, everyone has a claim on a fair chance to get the object.

Unfortunately, there exists no self sufficient strategyproof assignment mechanism meeting the Unanimity lower bound: see Lemma 1 in [Moulin 2007b][3]. A related but not directly comparable lower bound on individual net gains was recently introduced by [Porter et al. 2004] (see also [Atlamaz and Yengin 2006]):

—*q-fairness*: $U_i(a) \geq \frac{a^{*k}}{n}$, everyone has a claim on a fair share of the q-th highest valuation.

The 1-fairness property forces the equal division of the surplus a^{*1} (given self sufficiency), and is obviously out of reach for a self sufficient strategyproof mechanism (for instance, 1-fairness is stronger than the unanimity lower bound). It is easy to check that 2-fairness is equally impossible for such mechanisms. The following VCG mechanism (introduced by [Bailey 1997], see also [Cavallo 2006]) is 3-fair:

$$U_i^3(a) = a^{*1} - a_{-i}^{*1} + \frac{a_{-i}^{*2}}{n}$$

It is self sufficient and its relative budget loss is at most $\frac{2}{n}$:

$$\sum_{i\in N} U_i^3(a) = a^{*1} - \frac{2}{n}(a^{*2} - a^{*3}) \Rightarrow 0 \leq a^{*1} - \sum_{i\in N} U_i^3(a) \leq \frac{2}{n}a^{*1}$$

Letting q vary between 3 and n, we examine the tradeoff between the less and less generous individual guarantees under q-fairness, and our ability to minimize the relative budget loss. For each $q = 3, \cdots, n$, we seek the smallest number $\widehat{\rho}(q)$ for which we can find a q-fair strategyproof mechanism treating equals equally and such that

$$0 \leq a^{*1} - \sum_{i\in N} U_i(a) \leq \widehat{\rho}(q)a^{*1} \text{ for all } a \in \mathbb{R}_+^N$$

With the notation $B_m^{k,k'} = \sum_{l=k}^{k'} \binom{m}{l}$, and $q^* = q-1$ if q is odd,$= q-2$ if q is even, we find

$$\widehat{\rho}(q) = \frac{n-1}{B_{n-1}^{1,q^*}} \text{ for all } q = 3, \cdots, n \tag{4}$$

and the correponding optimal VCG mechanism is

$$U_i^q(a) = a^{*1} - \{\sum_{k=1}^{q^*} \frac{(-1)^{k-1}}{\binom{n-2}{k-1}} \frac{B_{n-1}^{k,q^*}}{B_{n-1}^{1,q^*}} a_{-i}^{*k}\} \text{ for all } a \in \mathbb{R}_+^N \tag{5}$$

Equations (4) and (5) imply that for any *odd* q, $\widehat{\rho}(q) = \widehat{\rho}(q+1)$ and $U^q = U^{q+1}$. Therefore q-fairness with q odd does not cost more in terms of the index $\widehat{\rho}$ than the weaker $(q+1)$-fairness. If we fix an odd q and let n grow, we see from (4) that

[3]This is true among *deterministic* mechanisms. If lotteries are allowed, random assignment with uniform probability on all participants is vacuously strategyproof and achieves the Unanimity lower bound.

$\widehat{\rho}(q)$ goes to zero as $\frac{1}{n^{q-2}}$; for instance

$$\widehat{\rho}(3) = \widehat{\rho}(4) = \frac{2}{n}; \widehat{\rho}(5) = \widehat{\rho}(6) = \frac{24}{n(n^2 - 5n + 18)}$$

Interestingly n-fairness is free if n is odd, as in this case U^n is precisely the mechanism (3) and $\widehat{\rho}(n) = \widehat{\rho}$. And if n is even, $\widehat{\rho}(n) = \frac{n-1}{2^{n-1}-2}$ is hardly larger than $\widehat{\rho} = \frac{n-1}{2^{n-1}-1}$ so the price of n-fairness is very small.

Writing $\{\frac{n}{2}\}$ for the integer $\frac{n}{2}$ or $\frac{n+1}{2}$, then $\{\frac{n}{2}\}$-fairness guarantees to everyone a fair share of the median valuation. Although this is a considerably stronger requirement than n-fairness, we note that it only requires to double the cap $\widehat{\rho}(q)$ on the relative budget loss. Indeed it is easy to check, with the help of Stirling formula, that $B_{n-1}^{1,n^*} \simeq 2B_{n-1}^{1,\{\frac{n}{2}\}^*}$ when n grows.

3. ASSIGNING A BAD FAIRLY

We now assume that one of the n agents must perform a costly task, for which they are equally responsible; individual costs c_i of doing the job are private information. This indivisible task is a common property "bad", and effciency requires to assign it to one of the agents with lowest cost. Cash transfers may be used to compensate this agent. For examples of this problem see [Porter et al. 2004] and the classic NIMBY problem ([Kunreuther 1996]).

We write c^{*k} for the k-th lowest cost, and V_i for agent i's net *disutility,* where performing no task and getting no cash yields $V_i = 0$. Thus in a budget balanced and efficient allocation we have $\sum_{i \in N} V_i = c^{*1}$, whereas if the task is assigned efficiently but transfers leave a surplus, $\sum_{i \in N} V_i \geq c^{*1}$.

We wish to compare the budget loss generated by a self sufficient and strategyproof mechanism to a meaningful notion of the "efficient surplus". The intuitive choice of the efficient cost c^{*1} proves misguided. Indeed the constant ρ caps the ratio of the budget loss to c^{*1} if and only if

$$0 \leq \sum_{i \in N} V_i(c) - c^{*1} \leq \rho c^{*1} \text{ for all } c \in \mathbb{R}_+^N \tag{6}$$

There is in fact only one strategyproof mechanism treating equals equally for which the ratio $\frac{\sum V_i - c^{*1}}{c^{*1}}$ remains non-negative and bounded above. This is the familiar *pivotal* mechanism ([Green and Laffont 1979; Moulin 1986]) defined by

$$V_i(c) = c^{*1} \text{ for all } c \in \mathbb{R}_+^N$$

Here the agent who is assigned the task is not compensated, while every other agent pays the minimal cost to the residual claimant. Thus $\rho = n-1$ and this huge waste of money occurs for all c, disqualifying the pivotal mechanism.

A better estimate of the surplus from assigning the task efficiently is $c^{*n} - c^{*1}$, namely the difference between the worst and the best possible assignment of the task[4]. Now the number ρ caps the relative budget loss of a given mechanism if and

[4] Another natural choice is $\overline{c} - c^{*1}$, where $\overline{c} = \frac{1}{n} \sum_N c_i$ is the average cost of the task. Here the benchmark is the random assignment of the taks, a perfectly incentive compatible and fair mechanism. The corresponding formulas are different, but the gist of the results is preserved.

only if

$$0 \leq \sum_{i \in N} V_i(c) - c^{*1} \leq \rho(c^{*n} - c^{*1}) \text{ for all } c \in \mathbb{R}_+^N \tag{7}$$

The pivotal mechanism does not satisfy (7) for any bounded number ρ. On the other hand the tightest cap $\widetilde{\rho}$ that a strategyproof and self sufficient mechanism treating equals equally can achieve is

$$\widetilde{\rho} = \frac{n-1}{2^{n-1}-1} \text{ if } n \text{ is odd; } = \frac{n-1}{2^{n-1}-2} \text{ if } n \text{ is even}$$

and a mechanism implementing $\widetilde{\rho}$ is

$$V_i(c) = c^{*1} - \{\sum_{k=1}^{n^*} \frac{(-1)^{k-1}}{\binom{n-2}{k-1}} \frac{B_{n-1}^k}{B_{n-1}^1} c_{-i}^{*k}\} \text{ where } n^* = n-1 \text{ if } n \text{ odd, } = n-2 \text{ if } n \text{ even}$$

Note the close analogy with the tightest cap $\widehat{\rho}$ and the optimal mechanism (3) in the good assignment problem: the formulas are in fact identical for n odd. The parallel between the two models extends to the discussion of individual guarantees.

The natural *Stand Alone upper bound* $V_i(c) \leq c_i$ cannot be met by a self sufficient strategyproof mechanism unless its relative budget loss in (7) is unbounded; and the much stronger *Unanimity upper bound* $V_i(c) \leq \frac{c_i}{n}$ cannot be true for any self sufficient strategyproof mechanism.

Now q-fairness places the upper bound $V_i(c) \leq \frac{c^{*k}}{n}$ on individual disutility. It is not comparable to the Stand Alone upper bound, yet it is typically much tighter. We find that equation (5) where c replaces a defines a q-fair mechanism achieving the cap $\widetilde{\rho}(q) = \widehat{\rho}(q)$ ((4)). Moreover this is the smallest feasible cap on the relative budget loss of any q-fair strategyproof mechanism treating equals equally.

Therefore the tradeoffs between q-fairness and the worst relative budget loss are identical in both models.

4. CONCLUDING COMMENTS

1. For proofs and more detailed discussion of the results described here, the reader is referred to [Moulin 2007a; 2007b].

2. Our results in section 1 beg for a generalization to the auction or assignment of multiple, possibly heterogenous, objects.

For the case of multiple *identical* objects, the generalization of the VCG mechanisms discussed here is straightforward: see [Guo and Conitzer 2007a; 2007b; Moulin 2007a]. However it can be shown that VCG mechanisms no longer minimize the (worst case) relative efficiency loss, therefore there is no reason to believe they achieve the optimal tradeoff between efficiency and q-fairness.

Finally the case of multiple heterogenous objects remains entirely open: a plausible conjecture is that we can improve somewhat the efficiency performance of the canonical pivotal mechanism.

REFERENCES

M. Atlamaz and D. Yengin, *Fair Groves mechanisms*, mimeo, Rochester University, 2006.

M.J. Bailey, *The demand revealing process: to distribute the surplus,* Public Choice, 91:107-126, 1997.

R. Cavallo, *Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments,* International Conference on Autonomous Agents and Multi-agents Systems, (AA-MAS) Hakodate, Japan, 2006.

P. Cramton, R. Gibbons and P. Klemperer, *Dissolving a partnership efficiently,* Econometrica, 55, 3, 615-632, 1987.

M. Guo and V. Conitzer, *Worst case optimal redistribution of VCG payments,* Proceedings of the 8th ACM Conference on Electronic Commerce (EC-07), pp. 30-39, San Diego, CA, USA.

M. Guo and V. Conitzer, *Worst case optimal redistribution of VCG payments* (extended version)

J. Green, and J.J. Laffont , *Incentives in public decision making,* Amsterdam: North-Holland, 1979.

H. Kunreuther, *The Role of Compensation in Siting Hazardous Facilities,* Journal of Policy Analysis and Management, September 1996

H. Moulin, *Characterizations of the Pivotal Mechanism*, Journal of Public Economics, 31, 53–78, 1986.

H. Moulin, *An Application of the Shapley Value to Fair Division with Money*, Econometrica, 60, 6, 1331–1349, 1992.

H. Moulin, *Efficient, strategy-proof and almost budget-balanced assignment,* www.ruf.rice.edu/˜econ/faculty/Moulin/spassgnt.pdf

H. Moulin, *Auctioning or assigning an object: some remarkable VCG mechanisms,* www.ruf.rice.edu/˜econ/faculty/Moulin/auctass1obj.pdf

R. Porter, Y. Shoham and M. Tennenholtz, *Fair imposition,* Journal of Economic Theory, 118, 209-228, 2004.