

Generalized Vickrey Auctions

by

Jeffrey K. MacKie-Mason

Hal R. Varian

University of Michigan

July 1994

Abstract. We describe a generalization of the Vickrey auction. Our mechanism extends the auction to implement efficient allocations for problems with more than one good, multiple units for the goods, and externalities. The primary restriction on preferences is that they must be quasilinear.

Keywords. mechanism design, auctions

Address. Jeffrey K. MacKie-Mason and Hal R. Varian, Department of Economics, University of Michigan, Ann Arbor, MI 48109-1220. E-mail: jmm@umich.edu and Hal.Varian@umich.edu.

Generalized Vickrey Auctions

Jeffrey K. MacKie-Mason

Hal R. Varian

Suppose that you want to assign a good to one of two agents in a Pareto efficient way. Agent i values the good at v_i , so Pareto efficiency dictates that the good should go to the agent with the highest value. Vickrey (1961) suggests asking the agents to bid on the good and awarding the good to the highest bidder at the second highest price.

This mechanism has the following desirable property. It is a dominant strategy for each agent to truthfully reveal his true value; hence if each agent plays a dominant strategy the outcome will be Pareto efficient.

To see this, suppose that agent i bids b_i . The expected payoff to agent 1 is

$$\text{Prob}[b_1 > b_2][v_1 - b_2].$$

If $v_1 - b_2 > 0$, then agent 1 can make the probability of winning 1 by reporting $b_1 = v_1$. On the other hand if $v_1 < b_2$ agent 1 can make the probability of winning 0 by reporting $b_1 = v_1$. If $v_1 = b_2$ any bid is optimal. Hence, reporting the truth is always an optimal
 strategy, no matter what the other party reports.

The same argument works if there are n agents: award the good to the agent with the highest bid and charge him the second highest price. If there are k units of the good to sell, and each agent wants only a single unit, one can award the goods to the k highest bidders at the k^{th} highest price.

But what happens if agents want more than one unit of the good? What if there are externalities among the agents so that one agent cares about whether one of the other agents gets the good?

In this note we generalize the Vickrey auction to these situations. We show that the Vickrey auction is closely related to the Clarke-Groves mechanism (Clarke (1971), Groves (1973)), a fact that has been widely observed. However, perhaps it has not been sufficiently appreciated just how general the Vickrey auction is.

The closest analysis to this one is Demange, Gale, and Sotomayor (1986). See also Roth and Sotomayor (1990) for a nice treatment of incentive compatibility. However, it appears that these authors treat only the case where each agent wants a single unit of the good and there are no externalities.

1. Preliminaries

Suppose that there are $a = 1, \dots, A$ agents. Agent a can take some action x_a (e.g., choose a bundle of goods). The set of actions taken by all the agents is $x = (x_1, \dots, x_A)$. The set of all actions, except for x_a , is denoted by x_{-a} . A feasible set of actions satisfies the condition $F(x) = F(x_a, x_{-a}) = 0$.

Agent a 's utility function has the quasilinear form

$$u_a(x) + m_a,$$

This work was supported by the National Science Foundation Grants SES-8800114 and SES-9122240.

where m_a is the amount of “money” held by agent a . The initial endowment of money held by agent a is \bar{m}_a . Note that an agent’s utility depends on the whole set of actions taken; thus there can be arbitrary patterns of externalities among the agents.

A Pareto efficient allocation maximizes the sum of the utilities:

$$\begin{aligned} \max_{(x)} \sum_{a=1}^A [u_a(x) + m_a] \\ F(x) = 0 \\ \sum_{a=1}^A m_a = \sum_{a=1}^A \bar{m}_a. \end{aligned}$$

If we assume that the utility functions are strictly concave there will be a unique value of x^* that solves this problem. The Pareto efficient frontier is then described by $(u_a(x_a^*) + m_a)$ where $\sum_{a=1}^A m_a = \sum_{a=1}^A \bar{m}_a$.

2. Incentive Compatible Implementation

Consider the following Generalized Vickrey Auction (GVA).

1. Each agent a reports a utility function $r_a(\cdot)$.
2. The planner computes

$$x^* = \arg \max_x \sum_a r_a(x)$$

subject to

$$F(x) = 0$$

and assigns action x_a^* to agent $a = 1, \dots, A$. Then compute

$$W_{-a}(x^*) = \sum_{b \neq a} r_b(x^*)$$

which is the total valuation of all agents other than a according to their reported utility functions.

3. Agent a receives payoff

$$u_a(x^*) - [G_a(r_{-a}) - W_{-a}(x_a^*)],$$

where $G_a(r_{-a})$ is an arbitrary function of the other agents’ reported utilities. We claim that it is a dominant strategy for each agent to report his or her true utility function. To see this, note first that a necessary condition for the maximization in item (2) is that x^* maximizes $r_a(x) + W_{-a}(x_a)$. Agent a ’s true payoff is given in item (3). It follows that agent a will maximize his payoff by setting $r_a(\cdot) = u_a(\cdot)$. The set of actions taken will then be the actions that maximize the sum of the true utility functions.

There is no guarantee that the government’s budget will balance; in fact, in general it will not. However, the function $G_a(r_{-a})$ can be chosen to help with this problem. One interesting choice is given by

$$G_a(r_{-a}) = \max_x \sum_{b \neq a} r_b(0, x_{-a})$$

subject to

$$F(0, x_{-a}) = 0.$$

This is the “second-price” analogue to the original Vickrey auction. Each agent is charged the total social surplus that would be possible if that agent did not participate in the auction at all. The result, then, is that the net payoff received by agent a is the net increment in total surplus that his participation creates.

If the only externalities are negative, not positive, then the auction raises non-negative revenue. With negative externalities, $G_i > W_{-i}$, for all agents i , because the other agents are better off with a not participating.¹ That means that each agent pays a non-negative amount for her allocation. If there are costs of production, these can be simply recovered by adding them to G_a without changing the problem, since the production costs do not depend on the bid values.

As a simple example, consider the original Vickrey problem: one unit of one good, winner-takes-all; assume there are only two bidders, a and b . In that case, the person with the highest reported utility gets $u_a + [0 - u_b]$. u_b is the surplus that the second bidder would get if a *did not participate*, because b would then get the good. This is the standard Vickrey rule: bidder a gets the good, pays bidder b 's bid (which truthfully revealed b 's valuation), and keeps the net, $u_a - u_b$. The loser, of course, gets $0 + [u_a - u_a] = 0$.

3. Examples

For these examples, we implement the GVA using the specific $G_a(r_{-a})$ function suggested above.

GVA with 0-1 good

Suppose that there are two agents and a single good so that x_a is zero or one. The utility of agent a is $v_a x_a$, and suppose that agent 2 has the highest value of the good. Using a GVA, agent 2's payoff is

$$v_x x_2 - [G_2(r_{-2}) - W_{-2}(x^*)].$$

Let $b_1 = r_1(\{1, 0\})$, that is, the value from getting the good as reported by agent 1. Then, we can evaluate G_2 and W_{-2} to get

$$v_2 x_2 - [b_1 - b_1(1 - x_2)] = [v_2 - b_1]x_2.$$

If $b_2 > b_1$ the auctioneer will set $x_2 = 1$ and if $b_2 < b_1$ the auctioneer will set $x_2 = 0$. Truthfully setting $b_2 = v_2$ ensures that agent 2 achieves the optimal utility in either case.

¹ Though not usually phrased this way, the *second price* in the traditional Vickrey auction is actually the correct congestion price: when a gets the good, that causes congestion and crowds out b 's consumption. The social cost of that congestion is the utility foregone by b .

GVA with multiple units

Suppose that there are 3 units of the good to be sold. Agent 1 is willing to pay (10,8,6) for 1,2 or 3 units of the good; agent 2's willingness to pay is (9,7,5). Pareto efficiency requires that the 2 units of the good go to agent 1, and that 1 unit goes to agent 2. The competitive equilibrium involves a price p where $8 \geq p > 7$.

The GVA assigns the Pareto efficient amounts to the agents. Agent 1 pays $12 = 7 + 5$. This is how much welfare would be generated by the two goods that agent 1 holds if agent 1 were not present. Similarly, agent 2 pays 6 for his one unit of the good. Hence each agent is charged the opportunity cost of his holdings.

GVA with buyer and seller

There are two agents who value a good at v_1 and v_2 . Agent 2 has an initial endowment of 1 unit of the good; agent 1 has 0 units of the good. Let x denote how much of the good he transfers to agent 1. Under the dominant strategy of truthful revelation, the agents' payoffs are:

$$\begin{aligned}v_1 x + v_2(1 - x) - v_2 &= (v_1 - v_2)x \\v_2(1 - x) + v_1 x - v_1 &= (v_1 - v_2)x + (v_2 - v_1).\end{aligned}$$

It is easy to verify that if $v_1 > v_2$, agent 1 receives the good and has utility $v_1 - v_2$. If $v_2 > v_1$, agent 2 receives the good and gets utility $v_2 - v_1$. In both cases the agent that doesn't have the good receives 0 utility.

Essentially each person states a value; the person who states the highest value gets the good at the price stated by the *other* person. The person who states the lower value gets 0.

GVA with congestion externalities

Suppose that there are 2 agents and agent a has a concave differentiable utility function $u_a(x) = U_a(x_a) - C_a(x_b)$. Suppose further that agent a satiates at \bar{x}_a . The dominant strategy equilibrium is then (\bar{x}_a, \bar{x}_b) but Pareto efficiency would require a choice (x_a^*, x_b^*) that satisfies the condition

$$u'_a(x_a^*) = c'_b(x_b^*) \quad \text{for } a = 1, 2.$$

Under the GVA agent a would report $R_a(\cdot) - D_a(\cdot)$. The auctioneer would choose x^* to maximize the sum of the reported utilities:

$$R_a(x_a) - D_a(x_b) + R_b(x_b) - D_b(x_a).$$

The auctioneer then calculates $W_{-a}(x^*) = R_b(x_b^*) - D_b(x_a^*)$, and likewise for $W_{-b}(x^*)$. Agent a 's payoff is

$$U_a(x_a^*) - C_a(x_b^*) + R_b(x_b^*) - D_b(x_a^*) - G_a(r_{-a})$$

so clearly a wants to report

$$R_a(\cdot) - D_a(\cdot) \equiv U_a(\cdot) - C_a(\cdot).$$

Now, if we let the arbitrary term $G_a(r_{-a}) = \max_x \sum_{b \neq a} r_b(0, x_{-a})$, which in this case is $G_a(r_{-a}) = R_2(\bar{x}_2) - D_2(0)$, a 's payoff is clearly less than his direct utility, $U_a(x_a^*) - C_a(x_b^*)$. Indeed, a 's payoff is reduced below his direct utility by precisely the incremental surplus that b loses as a result of the congestion externality imposed by a 's participation. Thus, each agent is charged for the congestion cost that it imposes on the other agent.

References

- Clarke, E. (1971) "Multipart pricing of public goods," *Public Choice*, **11**, 17–33.
- Demange, G., Gale, D., and Sotomayor, M. (1986) "Multi-item auctions," *Journal of Political Economy*, **94**, 863–872.
- Groves, T. (1973) "Incentives in teams," *Econometrica*, **41**, 617–631.
- Roth, A. E., and Sotomayor, M. A. O. (1990) *Two-sided Matching*. Cambridge, England: Cambridge University Press.
- Vickrey, W. (1961) "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, **16**, 8-37.