.

1 The VCG mechanism

The GS theorem places a severe restriction on the set of SCR's that can be implemented in DSE when we don't restrict the set of preference profiles. However, in many applications we restrict attention to particular classes of utility functions/preference relations (e.g. quasi-linear preferences). For some restricted domains we can find non-dictatorial SCR's that are DSE implementable. One well known example is the efficient allocation of a public good when preferences are quasi-linear. The canonical mechanism used to implement this SCR is usually referred to as the VCG mechanism.

The basic problem: Set of individuals need to decide whether or not to take some costly action, and if so, how to divide the cost.

- $N = \{1, \ldots, n\}$
- $C = \{(x, m) : x \in \{0, 1\}, m \in \mathcal{R}^n\}$
- $\mathcal{P}$ is the set of quasi-linear preferences, represented by the function $\theta_i x - m_i$. Because an agent's utility function is uniquely determined by $\theta_i \in \mathcal{R}$, we set $\mathcal{P} = \mathcal{R}^n$.
- $\mathcal{G}$ is the class of normal game forms
- $f : \mathcal{R}^n \to C$ is a SCF satisfying $x = 1$ iff $\sum_{i \in N} \theta_i \geq \gamma$ where $\gamma \geq 0$ is the cost of the project.

Proposition. *f is DSE implementable in $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ iff for each $j \in N$ there exists a function h_j such that $m_j(\theta) = x(\theta)\left(\gamma - \sum_{i \in N \setminus \{j\}} \theta_i\right) + h_j(\theta_{-j})$.*

Proof. Assume first that for each j there exists a function h_j with above property. Consider the following mechanism.

- Each player submits a value $\hat{\theta}_i$ and the project is executed if the sum of declarations is at least γ.
- Each player pays a fee $h_j\left(\hat{\theta}_{-j}\right)$ regardless of whether the project is executed
- If the project *does* get executed he also pays the difference $\gamma - \sum_{i \in N \setminus \{j\}} \hat{\theta}_i$.

Case 1. $\sum_{i \in N \setminus \{j\}} \hat{\theta}_i < \gamma$ and $\theta_j + \sum_{i \in N \setminus \{j\}} \hat{\theta}_i \geq \gamma$

- If $\hat{\theta}_j$ satisfies $\hat{\theta}_j + \sum_{i \in N \setminus \{j\}} \hat{\theta}_i < \gamma$, then $u_i\left(\hat{\theta}\right) = -h_j\left(\hat{\theta}_{-j}\right)$
- If $\hat{\theta}_j$ satisfies $\hat{\theta}_j + \sum_{i \in N \setminus \{j\}} \hat{\theta}_i \geq \gamma$, then $u_i\left(\hat{\theta}\right) = \theta_i - \left(\gamma - \sum_{i \in N \setminus \{j\}} \hat{\theta}_i\right) - h_j\left(\hat{\theta}_{-j}\right)$, which is clearly *above* $-h_j\left(\hat{\theta}_{-j}\right)$.

Case 2. $\sum_{i \in N \setminus \{j\}} \hat{\theta}_i \geq \gamma$ and $\theta_i + \sum_{i \in N \setminus \{j\}} \hat{\theta}_i < \gamma$

- If $\hat{\theta}_j$ satisfies $\hat{\theta}_j + \sum_{i \in N \setminus \{j\}} \hat{\theta}_i < \gamma$, then $u_i\left(\hat{\theta}\right) = -h_j\left(\hat{\theta}_{-j}\right)$

- First, note that if $x\left(\theta_{-j},\bar{\theta}\right)=x\left(\theta_{-j},\hat{\theta}\right)$, then $m_j\left(\theta_{-j},\bar{\theta}\right)=m_j\left(\theta_{-j},\hat{\theta}\right)$. To see why, note that if $m_j\left(\theta_{-j},\bar{\theta}\right)>m_j\left(\theta_{-j},\hat{\theta}\right)$, then a player of type $\bar{\theta}$ strictly prefer to lie and report a type of $\hat{\theta}$, in contradiction to our assumption that f is DSE implementable.
- Next, for any value of θ_j such that $x\left(\theta_{-j},\theta_j\right)=k\in\{0,1\}$, denote
$$m_j^k=m_j\left(\theta_{-j},\theta_j\right)$$
  - Since it is a dominant strategy for player j with preference parameter
$$\theta_j''=\gamma-\sum_{i\in N\backslash\{j\}}\theta_i$$
to report θ_j'', he must be no better off if instead he reported θ_j' when the other reported θ_{-j}. Hence,
$$\begin{aligned}\theta_j''-m_j^1 &\geq -m_j^0\\ \gamma-\sum_{i\in N\backslash\{j\}}\theta_i &\geq m_j^1-m_j^0\end{aligned}$$
  - For any $\epsilon>0$, it is also a dominant strategy for a player with preference parameter
$$\theta_j''=\gamma-\sum_{i\in N\backslash\{j\}}\theta_i-\epsilon$$
to report θ_j''. Hence, he cannot be better off if instead he reports θ_j' when the other reported θ_{-j}:
$$\begin{aligned}-m_j^0 &\geq \theta_j''-m_j^1\\ -m_j^0 &\geq \gamma-\sum_{i\in N\backslash\{j\}}\theta_i-\epsilon-m_j^1\\ m_j^1-m_j^0 &\geq \gamma-\sum_{i\in N\backslash\{j\}}\theta_i-\epsilon\end{aligned}$$
  - Since the last inequality holds for all $\epsilon>0$,
$$m_j^1-m_j^0\geq\gamma-\sum_{i\in N\backslash\{j\}}\theta_i$$
  - We conclude that
$$m_j^1-m_j^0=\gamma-\sum_{i\in N\backslash\{j\}}\theta_i$$

2 Assignment of individuals to positions

We proved that the VCG mechanism is the *only* strategy-proof mechanism that efficiently allocates a public good. This result can be generalized: the VCG mechanism is the only efficient mechanism that is DSE-implementable in an environment with quasi-linear preferences and strategic game forms. The basic idea of this mechanism can be stated in terms of a "generalized second-price" auction:
(1) each participants states his valuation for each of the possible outcomes,
(2) an outcome that maximizes the sum of valuations is picked,
(3) each player pays the externality that he imposes on the others: i.e., the difference between the highest sum of others' valuations attained in his absence, minus the highest sum of others' valuations attained in his presence.

To understand this better consider the problem of assigning a number of risk-neutral individuals to an equal number of positions (if these two are not really equal we can make them artificially equal by replicating individuals or positions). An efficient assignment satisfies that individuals have no incentive to trade the positions among themeselves. Consider the problem of devising a procedure that
(1) elicits honest preferences,
(2) results in an efficient allocation of individuals to positions, and
(3) charges prices that cause every individual to choose the best position

2.1 The assignment problem

We begin by first assuming that we have already succeeded in eliciting honest preferences, and are only faced with the problem of computing the efficient assignment. We later verify that this efficient assignment is compatible with truthful revelation.

1. Pure assignment
Let the set of individuals be I and the set of positions, J, so that we index an individual by i and a position by j. We let h_{ij} denote the value that individual i associated with being in position j (his "willingness to pay"). The "pure" assignment problem can be stated as a simple linear program:

$$\max \sum_{ij} x_{ij} h_{ij}$$

subject to

$$\begin{aligned}\sum_i x_{ij} &\leq 1 \\ \sum_j x_{ij} &\leq 1 \\ x_{ij} &\geq 0\end{aligned}$$

Since there is a finite number of possile assignments, this problem has at least one solution. Koopmans and Beckman (KB) in their 57 *Econometrica* paper noted that one of these solutions must be a corner one where $x_{ij} \in \{0, 1\}$ for all i and j. This solution can be obtained using the simplex method or any of its variants.

2. Supporting prices

Having obtained an efficient assignment we may ask what prices would support this in the sense that given those prices, each individual would find it optimal to purchase his assigned position. To answer this question we recall that in LP, the dual problem solves for the "shadow prices" of the constrained resources in the primal problem. In the primal problem these resources are the individuals and positions. Hence, the DP consists of finding a participation fee s_i for each individual and a price v_j for each position so as to

$$\min \sum_i s_i + \sum_j v_j$$

subject to

$$\begin{aligned}s_i + v_j &\geq h_{ij} \\ s_i, v_j &\geq 0\end{aligned}$$

(1) The DP says that the resources are to be assigned the *minimum* values compatible with the fact that they are valued by the recipients.
(2) The resources are assigned prices that reflect their values in the PP, i.e., in the *best possible use*.
(3) Since individuals choose positions and not vice versa, the participation fees are only used to extract the individuals' surplus and will not affect how individuals would choose positions given the prices v_j. Hence, we can focus on solutions to the DP that have $s_i = 0$ to all i.

3. Multiplicity of prices

Since the PP has a solution, the DP also has a solution and the objective function of both solutions attains the same value. Because the PP must have a degenerate solution, it will typically have multiple solutions and so the DP will have multiple solutions. Each of these solutions can serve as a "competitive price". To see this, consider the following 2×2 example:

	individual 1	individual 2
position 1	12	11
position 2	7	4

Setting the participation fees to zero, what positional prices support the efficient assignment?

$$\begin{aligned} 7 - v_2 &\geq 12 - v_1 \\ 11 - v_1 &\geq 4 - v_2 \end{aligned}$$

or

$$5 \leq v_1 - v_2 \leq 7$$

In addition, the individuals should not be worse off had they not been assigned, so that

$$\begin{aligned} v_1 &\leq 11 \\ v_2 &\leq 7 \end{aligned}$$

Any pair (v_1, v_2) that satisfies the above inequalities supports the efficient assignment. Hence, there is a continuum of "market-clearing" prices.

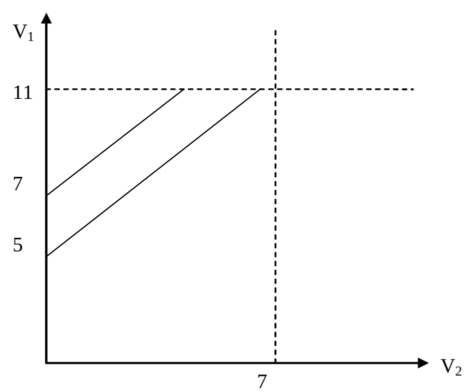

4. Incentive compatibility

Suppose we pick one of the supporting pair of prices. How do we know that given this pair, individuals would indeed want to reveal their valuations? (recall that in order to compute these prices we assumed individuals' valuations were known).

2.2 Incentive compatibility

We now ask, what set of prices in the region of "market-clearing prices" that are also incentive compatible? That is, imagine individuals are asked to submit their valuations and are told that based on these valuations an efficient assignment will be determined, and based on this assignment and the valuations positional prices will be determined according to a particular algorithm. We ask, what algorithm has the property that truthtelling will be a dominant strategy equilibrium? We shall see that there exists a unique set of prices with this property. Thus, by solving the incentive compatibility problem we also solve the multiplicity problem.

We know that any solution to the DP is a set of prices that support the efficient allocation. We also know that the individual fees do not affect the choices of individuals - they are like lump sum taxes. Ignoring these fees, the positional prices derived from the DP give us the "shadow prices" of either relaxing or tightening the constraint on the number of resources.
That is, v_j is either the amount by which the objective function ($\sum_{ij} x_{ij}h_{ij}$) would *rise* if we were to *add* a position j and reassign individuals optimally, or the amount by which the objective function would *fall* if we were to *remove* position j and reassign individuals optimally. Typically, the amount lost when a constraint is tightened may not be the same as that gained when the constraint is relaxed.

Suppose we took v_j to be the shadow price of *removing* position j. Then

$$v_j = V_I^J - V_I^{J-j}$$

where V_I^J is the value of assigning the individuals in I to positions in J, and V_I^{J-j} is the value when we remove j. While prices calculated this way will support the optimal assignment they do not guarantee truthtelling. This is because they depend on the values annouced by *all* individuals, and in particular, that individual who is ultimately assigned to j.

Suppose instead we took v_j to be the shadow price of *adding* position j. Then

$$v_j = V_I^{J+j} - V_I^J$$

This also appears to depend on the values annoucned by the person who ends up assigned to j. But in fact it does not. To see this, note that if we add another postion of type j and reassign individuals, it must be the case that there exists an efficient assignment in which the individual who was previously assigned to j, say individual

i, is still be assigned to this position. Suppose this was not true. Then we can strictly increase the objective function by moving i to a position that was previosy occupied and move a different individual to position j (either the new one or the old one).

Suppose *only one* of the two j positions are occupied in the new assignment. Then this assignment was also possible before we added another j. Since it was not optimal before, it cannot be optimal now.

Suppose *both* positions of type j are now occupied. Compare the original assignment and the new one. Change the original assignment by making all the new shuffles except for moving an individual to the newly added j position. This rearrangement was possible before, but yet was not made. Hence, it cannot be optimal now.

original assignment

positions	1	2	3	4
individuals	1	2	3	4

new assignment

positions	1	1	2	3	4
individuals	4	2		1	3

suhuffle available before

positions	1	1	2	3	4
individuals		4	2	1	3

CONCLUSION: *when a new position j is added, the new optimal assignment must assign i to j*

Hence,

$$V_I^{J+j} = h_{ij} + V_{I-i}^{J}$$

and since

$$V_I^{J} = h_{ij} + V_{I-i}^{J-j}$$

we have that

$$v_j = V_I^{J+j} - V_I^{J} = V_{I-i}^{J} - V_{I-i}^{J-j}$$

and this independent of i's annouced values. I.e., if positional prices are set in this way, each individual knows that the price he'll pay for his assigned position will not be affected by the value he announces. His announced value will only affect the probability in which he will be assigned to his favored position (just like in a second price auction). Hence, individuals have no incentive to misrepresent their preferences.

Is there ever a strict incentive to tell the truth? Suppose i misrepresented his preferences and was assigned as a result to position k and not j. Then i's payoff from this deviation is

$$h_{ik} - (V_{I-i}^{J} - V_{I-i}^{J-k})$$

While his payoff from truthtelling is

$$h_{ij} - (V_{I-i}^{J} - V_{I-i}^{J-j})$$

The latter is at least as high as the former iff

$$h_{ik} + V_{I-i}^{J-k} \leq h_{ij} + V_{I-i}^{J-j} = V_{I}^{J}$$

But this must be true since the LHS is the value of the objective function in the constrained maximization problem where we restrict i to be assigned to position k.

Since the VCG pricing scheme is the unique strategy-proof pricing scheme that guarantees efficiency, these must be the VCG prices. Hence, by imposing strategy-proofness we selected a unique set of prices out of the entire set that "clears the market".

2.3 Efficient computation of the IC prices

It is straightforward but tedious to show the IC positional prices we found, together with a set of participation fees, constitute a solution to the DP. The question is, how can we compute these prices for a given assignment problem.

A tedious method would be to solve $N+1$ assignment problems. To compute v_j for each j we need to know V_{I-i}^{J-j} and V_{I-i}^{J}. Since $V_{I-i}^{J-j} = V_{I}^{J} - h_{ij}$ we can compute this by solving the (I,J) assignment problem, compute the solution V_{J}^{J}, and for each j subtract h_{ij} (where i is the individual assigned to j). To compute V_{I-i}^{J} we need to solve the $(I-i,J)$ assignment problem for each i.

A more efficient and illuminating way is to identify the unique properties of the IC solution and impose this as an additional constraint on the DP. This way we directly identify the IC dual solution. Think of a second price auction. There the price of the good is the second highest bid. This is the *lowest* price at which the winner prefers to purchase the good, while the loser prefers not to purchase it. SImilarly, v_j is the value of position j as obtained from the *individuals not assigned to it.* Hence, it is the *lowest* price of position j at which each individual is best off buying his assigned position. This suggests (but can be proven formally) that the VCG prices have the property that they maximize the sum of *net* consumer surpluses. In other words, they have the lowest sum out of all the positional prices that solve the DP.

The above suggests a direct 2-step approach of computing the VCG prices. We first solve the PP and obtain the value V_I^J. Then we look for the positional prices with the lowest sum, subject to the constraint that together with an appropriate sum of individual fees we obtain the value V_I^J :

$$\max \sum_{ij} x_{ij} h_{ij}$$

subject to

$$\begin{aligned} \sum_i x_{ij} &\leq 1 \\ \sum_j x_{ij} &\leq 1 \\ x_{ij} &\geq 0 \end{aligned}$$

and then using the value of the objective function at the solution to solve

$$\min \sum_j v_j$$

subject to

$$\begin{aligned} s_i + v_j &\geq h_{ij} \\ s_i, v_j &\geq 0 \\ \sum_i s_i + \sum_j v_j &= V_I^J \end{aligned}$$

Note that by minimizing $\sum_j v_j$, while holding $\sum_i s_i + \sum_j v_j$ fixed, we maximize the net individual surplus. Hence, the VCG prices are the best market clearing prices for individuals. This raises the following question: can we solve a 2-sided matching problem in which positions have also say in the matter?

3 School choice

A. The School choice problem:

- There are a number of students, each of whom should be assigned a seat at one of a number of schools.
- Each school has a maximum capacity, but there is no shortage of the total seats.
- Each student has strict preferences over all schools
- Each school has a strict priority ordering of all students - these do not represent the school's preferences, but rather are imposed by state or local laws ("ties" are broken by lotteries):
  - a student who has a sibling already attending the school
  - a student who lives a walking distance from school

B. Matching.

- The outcome of a school choice problem is an assignment of students to schools such that (i) each student is assigned one school, and (ii) no school is assigned to more students than its capacity. Such an outcome will be called a *matching*.
- A matching is *Pareto efficient* if there exists no other matching, which assigns each student to a weakly better school and at least one student to a strictly better school [*note that schools are objects to be consumed, their priorities are not taken into account in terms of welfare*]
- A matching is *stable* if there exists no unmatched student-school pair (i, s) such that student i prefers school s to his assigned school, and also has a higher priority than some other student who is assigned a seat at school s.

C. A Student assignment mechanism. This is a systematic procedure that selects a matching for each school choice problem.

- a *direct* mechanism requires students to submit their rankings of schools, and based on this info, and the schools' priorities, a matching is produced.

- a direct mechanism is *strategy-proof* if it is weakly dominating for each student to submit his true preferences.
- an *efficient* mechanism selects a Pareto efficient matching for every school choice problem
- a *stable* mechanism selects a stable matching for every school choice problem
- a *partially-stable* mechanism has the following property: for every school s and for every pair of students i and j, where i has a higher priority than j at s, student i has a "better opportunity" to get into school s, other things being equal.

D. The Gale-Shapley optimal stable mechanism

Step 1:

- Each student proposes to her first choice.
- Each school tentatively assigns its seats to its proposers, one at a time following their priority order.
- Any remaining proposers are rejected.

⋮

Step k:

- Each student who was rejected in the previous step proposes to her next choice.
- Each school considers the students it has been holding together with its new proposers and tentatively assigns its seats to these students one at a time following their priority order.
- Any remaining proposers are rejected.

⋮

End: Procedure terminates when no student proposal is rejected and each student is assigned his final tenetative assignment.

PROPOSITION 1. *The Gale-Shapley mechanism has the following properties:*
(1) It Pareto dominates any other stable mechanism

(2) It is strategy-proof

However, this mechanism is not efficient:

s_1	s_2	s_3		i_1	i_2	i_3
i_1	i_2	i_2		s_2	s_1	s_1
i_3	i_1	i_1		s_1	s_2	s_2
i_2	i_3	i_3		s_3	s_3	s_3

the unique stable matching is:

s_1	s_2	s_3
i_1	i_2	i_3

but this is pareto dominated by:

s_1	s_2	s_3
i_2	i_1	i_3

How is the stable matching derived? notice that both i_2 and i_3 have s_1 as their top school, but i_3 has a priority over i_2 in that school. This means that any matching in which i_2 is placed in s_1 will be unstable. Thus, in any stable matching i_2 can be assigned either to his 2nd best school, s_2, or to his worst school, s_3. But any matching in which i_2 is assigned to s_3 must be unstable since i_2 prefers s_2 to s_3 and i_2 has the highest priority in s_2. Therefore, in any stable matching, i_2 must be assigned to s_2. By a similar argument, i_1 must be assigned to his 2nd best school, s_1.

How is the efficient matching derived? recall that in evaluating welfare, only the students' preferences are taken into account. Hence, by moving from the stable matching to the matching directly below it, both i_1 and i_2 strictly improve their positions, while the welfare of i_3 remains the same.

E. Top trading cycle mechanism

An alternative mechanism that guarantees efficiency is based on a milder interpretation of schools' priorities: If student i_1 has higher priority than student i_2 for school

s, then i_1 should have a *better opportunity* to get into school s. In other words, i_1 is not guaranteed a seat in s before i_2. The top trading cycle mechanism operates by starting with students who have the highest priority in a set of schools and allows these students to trade seats among themselves.

Step 0:

- Assign a counter for each school which keeps track of how many seats are still available at the school.
- Set the counters equal to the capacities of the schools.
- Ask students to report their ranking of schools.

Step 1:

- Each student points to his top ranked school according to his reported rankings.
- Each school points to the student with the highest priority
- Since the number of students and schools is finite, there is at least one cycle:

  $$s_1 \to i_1 \to s_2 \to \ldots \to i_k \to s_1$$

  A cycle *begins with a school* and *ends with a student* who points to that school.
- Each school and each student can be part of at most one cycle
- Every student in a cycle is assigned a seat in the school that he points to and is removed.
- The counter of each school in a cycle is reduced by one, and if it is reduced to zero, the school is removed.
- Counters of all schools not in a cycle are unchanged.

Step k:

- Each remaining studnet points to his favorite school among the remaining schools

- Each remaining school points to the student with the highest priority among the remaining students
- Since there is at least one cycle, each student in a cycle is assigned a seat in the school he points to and the counter of each school in a cycle is reduced by one.
- All students in a cycle and schools whose counter reached zero are removed.

Since the number of schools and students is finite, the algorithm terminates when all students are assigned seats. Note also that the number of steps does not exceed the number of students.

Example: 4 stdudents and 4 schools with 1 seat each:

s_1	s_2	s_3	s_4
i_1	i_2	i_3	i_4
i_4	i_3	i_4	i_1
i_3	i_2	i_1	i_2
i_2	i_1	i_2	i_3

i_1	i_2	i_3	i_4
s_2	s_1	s_3	s_1
s_3	s_2	s_2	s_2
s_4	s_3	s_3	s_3
s_1	s_4	s_4	s_4

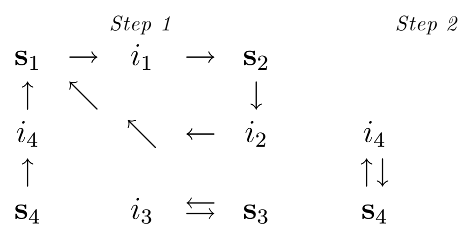

The resulting matching is:

i_1	i_2	i_3	i_4
s_2	s_1	s_3	s_4

Proposition 2. *The Top-Trading Cycle mechanism is both efficient and strategy-proof.*

To see why this mechanism is efficient note that any student who is assigned a seat in the first step is assigned a seat in his top ranked school, hence no student who exists in Step 1 can be made better off. Any student who is assigned a seat in Step 2 is assigned a seat in his top ranked school among those that remain after Step 1.

Because preferences are strict, none of these students can be made better off without hurting the students who were assigned seats in Step 1. Proceeding this way we conclude that no student can be made better off without hurting another.

The reason the TTC mechanism is efficient is analogous to the reason why serial dictatorship is efficient. In fact, when all schools have the same exact priorities, the TTC mechanism reduces to serial disctatoriship, where the order of studnets is determined by the priority ordering.

PROPOSITION 3. *The Top-Trading Cycle mechanism is not stable.*

To see this, consider the above example. i_4 prefers s_1 to his assigned school, and he also has a higher priority than i_2 who ends up being assigned to s_1.