Classic Mechanism Design (III)

David C. Parkes

Division of Engineering and Applied Science,

Harvard University

CS 286r–Spring 2002

Vickrey-Clarke-Groves Mechanism

[Vickrey61, Clarke71, Groves73] (VCG or "Pivotal" mechanism.)

Def. [VCG mechanism] Implement efficient outcome, $k^* = \max_k \sum_j v_j(k, \hat{\theta}_j)$, and compute *transfers*

$$t_i(\hat{\theta}) = \sum_{j \neq i} v_j(k^{-i}, \hat{\theta}_j) - \sum_{j \neq i} v_j(k^*, \hat{\theta}_j)$$

where $k^{-i} = \max_k \sum_{j \neq i} v_j(k, \hat{\theta}_j)$.

Thm. The VCG mechanism is strategyproof, efficient, and *interim* IR.

Alternative description:

$$p_{\text{vick},i}(\theta) = v_i(k^*, \theta_i) - [V(\mathcal{I}) - V(\mathcal{I} \setminus i)]$$

where $V(\mathcal{K})$ is the value of the efficient allocation in the subproblem restricted to agents $\mathcal{K} \subseteq \mathcal{I}$.

[Note 1:] given strategies $\hat{\theta}_{-i}$, each agent's adjusted payment, $v_i(k^*, \hat{\theta}_i) - [V(\mathcal{I}) - V(\mathcal{I} \setminus i)]$, sets

$$v_i(k^*, \hat{\theta}_i) - [V(\mathcal{I}) - V(\mathcal{I} \setminus i)] + \sum_{j \neq i} v_j(k^*, \hat{\theta}_j)$$

$$= V(\mathcal{I} \setminus i)$$

i.e., this is the least value agent i could have bid for outcome k^*.

[Note 2:] Each agent's equilibrium utility is:

$$\begin{aligned} \pi_{\mathrm{vick},i} &= v_i(k^*, \theta_i) - [v_i(k^*, \theta_i) - V(\mathcal{I}) + V(\mathcal{I} \setminus i)] \\ &= V(\mathcal{I}) - V(\mathcal{I} \setminus i) \end{aligned}$$

i.e., equal to its marginal contribution to the welfare of the system.

Example: Shortest Path.

[Nisan 99]

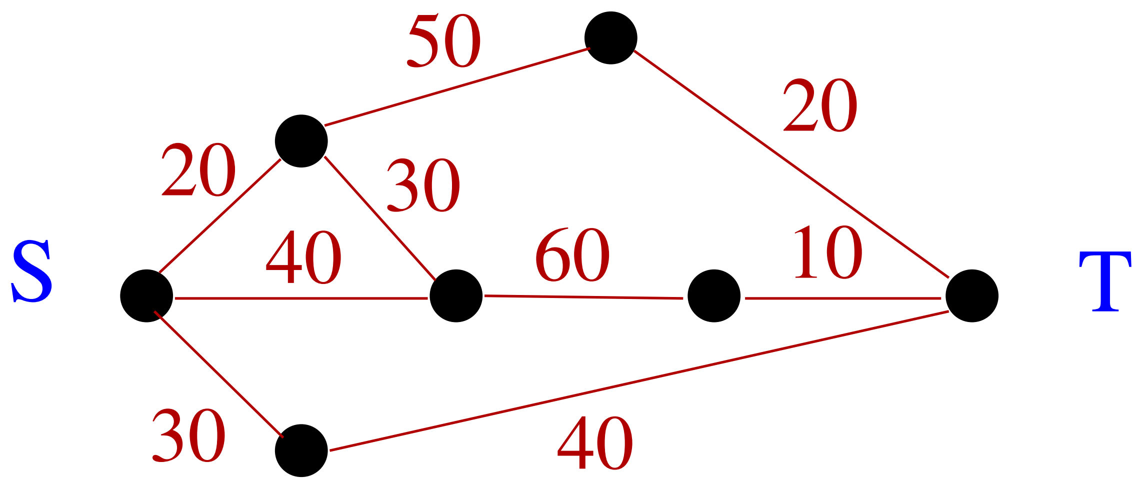

Biconnected graph, $G = (N, E)$, cost $c_l \geq 0$ per edge $l \in E$, edges srategic. Assume *large* value V to send message.

Goal: route packets along the lowest-cost path from S to T.

VCG Payment edge e:

$$p_{\text{vick},l} = -c_l - \left[(V - d_G) - (V - d_{G/l})\right]$$
$$= -c_l - (d_{G/l} - d_G)$$

Bayesian-Nash Implementation

Drop dominant-strategy implementation, try to achieve budget-balance.

Bilateral trading problem: single seller, single buyer. One good. Values drawn from $v_1 \in [0, 1], v_2 \in [0, 1]$.

Thm. [Myerson-Satterthwaite 83] In the bilateral trading problem, no mech. can implement an efficient, interim IR, and *ex post* (weak) budget-balanced SCF, even in Bayes-Nash eq.

Note: this is a negative result for a very simple problem, therefore quite a “strong” negative result!

The Centrality of VCG

[Krishna & Perry 98]

Thm. Among all efficient and interim IR mechanisms, the VCG maximizes the expected transfers from agents.

Note: this is interesting, shows that the best mechanism amongst all (Bayes-Nash), etc. is dominant strategy.

Thm. Given preferences, Θ, there exists a (weak) BB and efficient mechanism, with interim IR, if and only if the VCG has positive expected surplus.

. . . leads to quite direct proofs of Myerson-Satterthwaite, other negative results.

Expected Externality Mechanism

[Arrow79,d'Aspremont&Gerard-Varet79] **Retain** Bayes-Nash, and **relax** interim IR to ex ante IR, and try to achieve BB.

The d'AGVA mechanism (or *expected-Groves* mechanism), uses the same allocation as the Groves, but computes an transfer term averaged across all possible types of agents. [P.55, Parkes-Diss]

Thm. The d'AGVA mechanism is efficient, *ex post* budget-balanced, but only *ex ante* IR.

Demonstrates: (wrt Eff. mech. des.):

(a) *ex ante* IR really does make mechanism design "easier" than *interim* IR (compare Myerson-Satterthwaite with d'AGVA)

(b) Bayes-Nash implementation really does make mechanism design "easier" than dominant-strategy equilibrium (compare Green-Laffont impossibility with d'AGVA).

Summary

Name	Pref	Solution	**Possible**
Median	no transfers single-peaked	dominant	Parto opt.
Groves	quasi-linear	dominant	Eff
dAGVA	quasi-linear	Bayesian-Nash	Eff,BB,*ex ante* IR
Clarke	quasi-linear	dominant	Eff & IR

Name	Preferences	Solution concept	**Impossible**	Environment
GibSat	general	dominant	Non-dictatorial (incl. Pareto Optimal)	general
HGL	quasi-linear	dominant	Eff& BB	simple-exchange
MyerSat	quasi-linear	Bayesian-Nash	Eff& weak BB & IR	simple-exchange

Eff: *ex post* efficiency; *BB*: *ex post* strong budget-balance; *IR*: *interim* IR.

Cost-Sharing Problems

Choice set $\mathcal{K}$, N buyers, 1 seller. Transfers $t_1, \ldots, t_N$ and t_s. Values $v_i(k) \geq 0$ for buyers, and cost $c_s(k) \geq 0$ for seller. Quasi-linear utility functions, $u_i(k, t_i) = v_i(k) - t_i$, and $u_s(k, t_i) = -c_s(k) - t_i$.

Example: Multi-cast cost sharing.

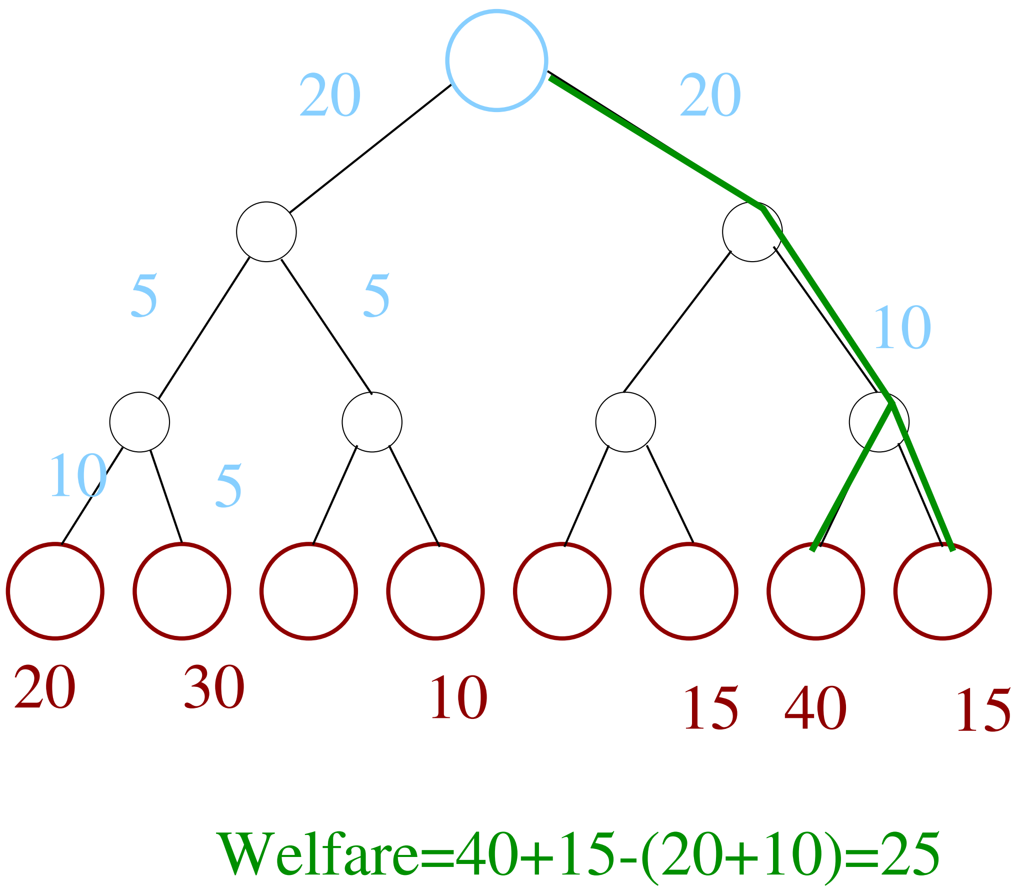

Desirable Properties

[Assume the seller is truthful.] Compute outcome $k^*(\theta)$ and transfers $t_i(\theta), t_s(\theta)$.

Use *revelation principle*, focus on IC mechanisms.

- **EFF**. Select $\max_k \sum_i v_i(k^*(\theta), \theta_i) - c_s(k^*(\theta))$, for all $\theta \in \Theta$.
- **BB**. Transfers $\sum_i t_i(\theta) + t_s(\theta) = 0$, for all $\theta \in \Theta$.
- **No-profit**. Transfers $-c_s(k^*(\theta)) - t_s(\theta) = 0$, for all $\theta \in \Theta$.
- **Buyer-SP**. Satisfy: $v_i(k_i^*(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \geq v_i(k_i^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i(\hat{\theta}_i, \theta_{-i})$ for all $\hat{\theta}_i \neq \theta_i$, θ_i, and θ_{-i}.
- **IR**. Satisfy: $v_i(k_i^*(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \geq 0$, for all $\theta \in \Theta$.

Dominant Strategy BB & EFF Impossibility

[Green&Laffont 79]

Thm. SP, EFF, and BB with No-Profit are impossible.

Proof. By contradiction. Suppose $\mathcal{M}$ is SP, EFF, BB, and No-Profit. Consider a problem in which $c_s(k) = 0$, for all $k \in \mathcal{K}$. Then, we must have $\sum_i t_i = c_s(k^*) = 0$ by BB and No-Profit, *and* $k^* = \arg\max_k \sum_i v_i(k)$ by Eff; this violates the Green-Laffont imposs. theorem.

[SP and EFF:] VCG mechanism: Given receiver set $R \subseteq \mathcal{I}$, let $c_s^*(R)$ denote the minimal cost tree. Select R to maximize $W(v) = \sum_{i \in R} v_i - c_s^*(R)$, and charge each user $i \in R$, $p_{\text{vick},i} = v_i - [W(v) - W(v_{-i})]$.

Group Strategyproofness

[this comes for “free” when we worry about BB]

- **Buyer-GSP**. For all coalitions $S \subseteq \mathcal{I}, \hat{\theta}_S \neq \theta_S, \theta_S, \theta_{-S}$, either $u_i(\hat{\theta}_S, \theta_{-S}, \theta_i) \leq u_i(\theta_S, \theta_{-S}, \theta_i), \forall i \in S$, or $\exists i \in S$ s.t. $u_i(\hat{\theta}_S, \theta_{-S}, \theta_i) < u_i(\theta_S, \theta_{-S}, \theta_i)$.

No coalition of agents can manipulate the outcome of the mechanism without making one of the agents in the coalition worse off.

Simplifying: A Binary Choice Model

[Moulin&Shenker 99]

Suppose $\mathcal{I}$ agents, either receive the service or not (binary choice). Let $R \subseteq \mathcal{I}$ denote the *receiver set*. Define $C(R)$ as the cost of providing service to R agents.

Eff: $R(\theta) = \arg\max_R \sum_{i \in R} v_i - C(R), \quad \forall \theta$

BB,No-Profit: $\sum_{i \in R(\theta)} t_i(\theta) = C(R(\theta)), \quad \forall \theta$.

GSP, IP.

Def. Mechanism $\mathcal{M} = (\Theta, R, t_i)$ satisfies the *core property* if and only if

$$\sum_{i \in Q} t_i(\theta) \leq C(Q), \quad \forall Q \subseteq \mathcal{I}$$

i.e., there is no incentive for a subset of agents to break from the grand coalition.

Cost-Sharing Methods

Let $\xi(Q, i) \geq 0$ define the payment made by agent i whenever $Q \subseteq \mathcal{I}$ that receiver service. Let $C(Q) \geq 0$ denote the cost of providing service to Q.

Def. Cost-sharing method, $\xi(Q, i)$, is a well-formed cost sharing method if and only if

$$i \neq Q \Rightarrow \xi(Q, i) = 0$$

$$\sum_{i \in Q} \xi(Q, i) = C(Q)$$

Use $\xi(Q, i)$ to define transfers $t_i = \xi(R^*, i)$, will satisfy BB,No-Profit,and IR.

Cross-monotonic Cost Sharing

Def. Cost sharing method, $\xi(Q,i)$ is cross-monotonic if and only if

$$\xi(Q,i) \geq \xi(R,i), \quad \forall Q \subseteq R$$

Note: this is also *weak* cross-monotonic, satisfying:

$$\sum_{i \in Q} \xi(Q,i) \geq \sum_{i \in Q} \xi(R,i), \quad \forall Q \subseteq R$$

Prop. A mechanism that implements transfers, $t_i = \xi(R^*, i)$, for some weak cross-mononotic, $\xi(\cdot)$, implements a core allocation for each subset of users.

Coalitional StrategyProof Cost-Sharing Mechanisms

[Moulin& Shenker, 99]

Given cross-monotonic cost-sharing method, $\xi(Q, i)$, mechanism $\mathcal{M}(\xi)$ computes the receiver set R^* and transfers $t_i = \xi(R^*, i)$ as follows:

Def. Mechanism $\mathcal{M}(\xi)$:

Agents report values, $\hat{v}$; initialize $R^* \leftarrow \mathcal{I}$.

Select an agent $i \in R^*$ at random, if $\hat{v}_i < \xi(R^*, i)$ then drop i from R^*.

Continue until $\hat{v}_i \geq \xi(R^*, i)$ for all $i \in R^*$.

Implement R^* and transfers $\xi(R^*, i)$.

Thm. Given cross-monotonic, $\xi(Q, i)$, then $\mathcal{M}(\xi)$ is BB and GSP.

Shapley Value

[Moulin & Shenker, 99]

Def. Submodular: $C(i \cup T) - C(T) \leq C(i \cup S) - C(S)$, for all $S \subseteq T \subseteq I$, and $i \notin S$.

Thm. If the cost function $C(S)$ is submodular, then all GSP and BB mechanisms can be characterized by $\mathcal{M}(\xi)$, for some cross-monotonic cost-sharing method $\xi(Q, i)$.

Thm. If costs are submodular, then the *Shapley Value*, $\xi_{\text{Shapley}}(Q, i)$ defines a cross-monotonic cost-sharing method, and $\mathcal{M}(\xi_{\text{Shapley}})$ defines the GSP and BB mechanism that *maximizes the worst-case eff. loss*.

Example: Shapley Mechanism

Jain & Vazirani: assume a general biconnected network, propose a centralized *approximation* mechanism; not submodular and Shapely does not apply.

Additional Implementation Concepts

- **Repeated implementation**: can begin to implement more [Kalai 97]
  – if the planner learns and is more patient than the agents, and agents in a multi-round game, then can achieve dom. strategy implementation (in limit if center has no time discounting)
  – reduce to a one-shot revelation game
- **Large societies**:
  – can get approx. EFF and approx. balance in large double auctions [McAfee92, Satterthwaite&Williams89, Rustichini et al.95]

What is Missing?

- No computational constraints
- Focus on efficiency (social-welfare), little considerations of alternative objectives (e.g. fairness, max-min, make-span, etc.)
- Little discussion of special preference structure in resource allocation (beyond quasilinear preferences, some concavity assumptions)
- No use of randomization in the mechanism itself
- Revelation principle is the central paradigm, and there is no attention to **indirect** mechanisms