An Efficient Dynamic Mechanism

Susan Athey, Ilya Segal (2007)

May,13 2008

1 Purpose of the Paper

- Dynamic Setting with Private Info
- Construct an Efficient, BB, BIC mechanism
- Extension of AGV/Arrow (1979)
- Idea: charge an agent for the *change* in the opponent utilities induced by his report
- Conditions of Self-Enforcing Mechanism (in a Markov Setting)

2 A Motivating Example

Buyer, Seller trade $x_t \in [0,1]$ each period $t = 1, 2$
Seller's cost is $\frac{1}{2\theta_S}(x_t)^2$, parameter $\theta_S \in [0,1]$ is private info
Buyer's valuation is 1 in $t = 1$, and θ_B in $t = 2$
NB θ_B is learnt between periods 1 and 2

2.1 First best:

$$\max_{x_1,x_2} \left(x_1 - \frac{1}{2\theta_S}(x_1)^2 + \theta_B x_2 - \frac{1}{2\theta_S}(x_2)^2 \right)$$

$$\chi_1(\theta_S,\theta_B) = \theta_S; \qquad \chi_2(\theta_S,\theta_B) = \theta_S\theta_B$$

2.2 Expected externality transfers to induce truthful revelation:

$$\gamma_i(\theta_i) = E_{\widetilde{\theta}_{-i}} U_{-i}(\theta_i, \widetilde{\theta}_{-i})$$

$$\gamma_S(\theta_S) = E_{\widetilde{\theta}_B}\left[\chi_1(\theta_S) + \widetilde{\theta}_B \chi_2(\theta_S,\widetilde{\theta}_B)\right] = \theta_S \left(1 + E_{\widetilde{\theta}_B}\left[(\widetilde{\theta}_B)^2\right]\right)$$

$$\gamma_B(\theta_B) = E_{\widetilde{\theta}_S}\left[-\frac{1}{2\widetilde{\theta}_S}\left(\chi_1(\widetilde{\theta}_S)\right)^2 - \frac{1}{2\widetilde{\theta}_S}\left(\chi_2(\widetilde{\theta}_S,\theta_B)\right)^2\right] = -\frac{1}{2}(1+(\theta_B)^2)E_{\widetilde{\theta}_S}\left[\widetilde{\theta}_S\right]$$

To balance the budget, take $\psi_i(\theta_i,\theta_{-i}) = \gamma_i(\theta_i) - \gamma_{-i}(\theta_{-i})$

2.3 The Problem:

B cannot make his report before $t = 1$

By the moment he learns θ_B, he has already observed the reported θ_S from trade, $\chi_1(\theta_S) = \theta_S$

B's utility in the mechanism:

$$\underset{\widehat{\theta}_B}{argmax} U_B(\theta_B, \widehat{\theta}_B) = \underset{\widehat{\theta}_B}{argmax} \left[\theta_B \chi_2(\theta_S, \widehat{\theta}_B) + \gamma_B(\theta_B)\right]$$

$$= \underset{\widehat{\theta}_B}{argmax} \left[\theta_B \theta_S \widehat{\theta}_B - \tfrac{1}{2}(1 + (\theta_B)^2) E_{\widetilde{\theta}_S}\left[\widetilde{\theta}_S\right]\right] = \frac{E\left[\widetilde{\theta}_S\right]}{\theta_S}\theta_B$$

Knowing that B's strategy, S will also prefer to lie (a.s.).

2.4 Solution:

Let B's incentive transfer depend on both reports:

instead of $\psi_B(\theta_S, \theta_B) = E_{\widetilde{\theta}_S} U_S(\widetilde{\theta}_S, \theta_B) - E_{\widetilde{\theta}_B} U_B(\theta_S, \widetilde{\theta}_B)$

let $\psi_B(\theta_S, \theta_B) = U_S(\theta_S, \theta_B) + h(\theta_S) -$ will give him correct incentives

Seller's utility: $U_S(\theta_S, \theta_B) + \psi_S(\theta_S, \theta_B) = E_{\widetilde{\theta}_B}\left[U_S(\theta_S, \widetilde{\theta}_B) + U_B(\theta_S, \widetilde{\theta}_B)\right]$

BB: $\psi_B(\theta_S, \theta_B) = U_S(\theta_S, \theta_B) - E_{\widetilde{\theta}_B}\left[U_S(\theta_S, \widetilde{\theta}_B) + U_B(\theta_S, \widetilde{\theta}_B)\right]$

$\psi_S(\theta_S, \theta_B) = -U_S(\theta_S, \theta_B) + E_{\widetilde{\theta}_B}\left[U_S(\theta_S, \widetilde{\theta}_B) + U_B(\theta_S, \widetilde{\theta}_B)\right]$

The idea is to charge an agent for the *change* in the opponent utilities induced by his report.

3 Set-up

I agents
$t = 1, 2...$ countable/finite number of periods
Each period t each agent i privately observes $\theta_{i,t} \in \Theta_{i,t}$

individual type at t	$\Theta_{i,t}$	history of individual types up to t	$\Theta_i^t = \prod_{\tau=1}^{t} \Theta_{i,t}$
types at t	$\Theta_t = \prod_{i=1}^{I} \Theta_{i,t}$	history of types up to t	$\Theta^t = \prod_{i=1}^{I} \Theta_i^t$
the whole type space	$\Theta = \prod_{t=1}^{\infty} \Theta_t$		$\Theta = \Theta^\infty$

In each period t each agent i takes private decision $x_{i,t} \in X_{i,t}$ and there is a public decision $x_{0,t} \in X_{0,t}$.
Denote, equivalently, $X_t = \prod_{i=0}^{I} X_{i,t}$, $X^t = \prod_{\tau=1}^{t} X_\tau$, $X = X^\infty$.

Distribution of types over Θ_t is given by $\nu_t | x^{t-1}, \theta^{t-1}-$ conditioning on all past decisions and types.
A decision plan is a measurable function $\chi : \Theta \to X$,
where each $\chi_t(\theta)$ represents the decision made at t.
Any decision plan χ uniquely determines probability distribution $\mu[\chi]$ over the type space Θ.

Definition 1 We have independent types if for each t*,* x^{t-1}*, and* θ^{t-1}*, the probability measure* $\nu_t|x^{t-1}, \theta^{t-1}$ *over* Θ_t *can be written in the form:* $\nu_t|x^{t-1}, \theta^{t-1} = \prod_{i=1}^{I} \nu_{i,t}|x_o^{t-1}, x_i^{t-1}, \theta_i^{t-1}$ *for some probability measures* $\nu_{i,t}$ *over* $\Theta_{i,t}$ *that depend only on public decision history* x_o^{t-1} *and on agent* i*'s private history* $(x_i^{t-1}, \theta_i^{t-1})$.

NB $\mu[\chi]$ need not be independent across agents (influence through public decisions).

Agent i's Payoff:

$$\sum_{t=1}^{\infty} \delta^t \left(u_{i,t}(x^t, \theta^t) + y_{i,t} \right)$$

denote $U_i(x, \theta) = \sum_{t=1}^{\infty} \delta^t u_{i,t}(x^t, \theta^t)$.

Definition 2 We have private values if each agent i*'s utility* $U_i(x, \theta)$ *depends only on the public decisions* x_0 *and the agent's private history.*(x_i, θ_i).

NB The agent's expected utility conditional on time-t history (x^t, θ^t) is allowed to depend on other agents' information.

4 Mechanism

A mechanism is described by 2 observationally measurable functions:
decision plan $\chi : \Theta \to X$,
uniformly bounded transfer plan $\psi : \Theta \to (R^I)^\infty$

Public decisions $\chi_0(\theta)$ are implemented directly,
prescribed private decisions $\chi_i(\theta)$ for agents $i \geq 1$ can be disobeyed.

Transfer $\psi_{i,t}(\theta)$ is made to agent i in period t and depends only on the current history θ^t (observational measurability), the total discounted payments in the mechanism can be calculated as:

$$\Psi_i(\theta) = \sum_{t=1}^{\infty} \delta^t \psi_{i,t}(\theta^t)$$

Definition 3 A mechanism is budget balanced, if $\sum_i \psi_{i,t}(\theta^t) \equiv 0$ *for all* t.

5 The Game

Each agent privately observes his private type $\theta_{i,t}$
Each agent makes a public report $\widehat{\theta}_{i,t}$
Each agent takes a private decision $x_{i,t}$. The mechanism implements the public decision $x_{0,t} = \chi(\widehat{\theta}^t)$ and the payments $y_{i,t} = \psi_{i,t}(\widehat{\theta}^t)$ to each agent.

6 Strategies

Agent i's reporting strategy $\beta_i : \Theta \times \Theta_i \times X_i \to \Theta_i$, where the report at time t is observationally measurable, in that $\beta_{i,t}(\widehat{\theta}, \theta_i, x_i)$ depends only on the history of reports $\widehat{\theta}^{t-1}$ and the agent's private history θ_i^t, x_i^{t-1}.
Agent i's decision strategy $\alpha_i : \Theta \times \Theta_i \times X_i \to X_i$, where the decision at time t is, again, observationally measurable.

A given strategy (β_i, α_i) induces a strategic plan $(\overline{\beta}_i, \overline{\alpha}_i)$, where $\overline{\beta}_i(\theta)$ denotes the agent i's reports and $\overline{\alpha}_i(\theta)$ his decisions from following the strategy given that his type is θ_i and his opponents report θ_{-i}.
The construction is recursive, $t = 1, 2.. :$

$$\overline{\beta}_{i,t}(\theta) = \beta_{i,t}\left(\left(\overline{\beta}_i^{t-1}(\theta_i^{t-1}), \theta_{-i}^{t-1}\right), \theta_i^t, \overline{\alpha}_i^{t-1}(\theta_i^{t-1})\right)$$
$$\overline{\alpha}_{i,t}(\theta) = \alpha_{i,t}\left(\left(\overline{\beta}_i^t(\theta_i^t), \theta_{-i}^t\right), \theta_i^t, \overline{\alpha}_i^{t-1}(\theta_i^{t-1})\right)$$

Observational measurability implies that $\overline{\beta}_{i,t}(\theta)$ depends only on $(\theta_i^t, \theta_{-i}^{t-1})$ and $\overline{\alpha}_{i,t}(\theta)$ only on θ^t.

Definition 4 For a given decision plan χ, agent i's strategy (β_i, α_i) is a truthful-obedient strategy, if the corresponding strategic plan $(\overline{\beta}_i, \overline{\alpha}_i)$ has $\overline{\beta}_i(\theta) = \theta_i$ and $\overline{\alpha}_i(\theta) = \chi_i(\theta)$ for all $\theta \in \Theta$.

NB Does not specify anything about reporting after a lie or disobedience. "If has been truthful and obedient so far, he'll continue to be truthful and obedient".

7 The Team Mechanism

Let χ^* be the efficient decision plan (assumed to exist) - it maximizes $E_{\widetilde{\theta}}^{\mu[\chi]}\left[\sum_i U_i\left(\chi\left(\widetilde{\theta}\right),\widetilde{\theta}\right)\right]$.

The Team Mechanism consists of the efficient decision plan χ^* and transfers:

$$\psi_{i,t}^M(\theta) = \sum_{j \neq i} u_{j,t}\left(\chi^*\left(\theta\right),\theta\right).$$

The resulting total discounted transfers are:

$$\Psi_i^M(\theta) = \sum_{j \neq i} U_j\left(\chi^*\left(\theta\right),\theta\right)$$

Proposition 1 Assume private values. Then for any mechanism with an efficient decision plan χ^* and total discounted transfers $\Psi_i^M(\theta)$ in the Team Mechanism, any truthful-obedient strategy profile is a BNE.

Proof. Define a new decision plan and transfers to incorporate a deviation to (β_i, α_i) :

$$\widehat{\chi}(\theta) = \left(\chi_{-i}\left(\overline{\beta}_i(\theta), \theta_{-i}\right), \overline{\alpha}_i(\theta)\right)$$
$$\widehat{\psi}(\theta) = \psi\left(\overline{\beta}_i(\theta), \theta_{-i}\right) \qquad and \qquad \widehat{\Psi}_i(\theta) = \sum_{t=1}^{\infty} \delta^t \widehat{\psi}_{i,t}(\theta)$$

A truthful-obedient strategy in the new mechanism $\left(\widehat{\chi}, \widehat{\psi}\right)$ will induce the same decisions and transfers as deviation (β_i, α_i) in the original mechanism (χ, ψ). The respective payoff is $E_{\widetilde{\theta}}^{\mu[\widehat{\chi}]}\left[U_i\left(\widehat{\chi}(\widetilde{\theta}), \widetilde{\theta}\right) + \widehat{\Psi}_i(\widetilde{\theta})\right]$.

For the Team Mechanism the new transfers are

$$\widehat{\Psi}_i(\theta) = \sum_{j \neq i} U_j\left(\chi_{-i}^*\left(\overline{\beta}_i(\theta), \theta_{-i}\right), \chi_i^*\left(\overline{\beta}_i(\theta), \theta_{-i}\right), \overline{\beta}_i(\theta), \theta_{-i}\right)$$
$$= \sum_{j \neq i} U_j\left(\widehat{\chi}_{-i}(\theta), \chi_i^*\left(\overline{\beta}_i(\theta), \theta_{-i}\right), \overline{\beta}_i(\theta), \theta_{-i}\right) = \sum_{j \neq i} U_j\left(\widehat{\chi}(\theta), \theta\right)$$

Thus agent i's expected payoff from deviation is:

$$E_{\widetilde{\theta}}^{\mu[\widehat{\chi}]}\left[U_i\left(\widehat{\chi}(\widetilde{\theta}), \widetilde{\theta}\right) + \sum_{j \neq i} U_j\left(\widehat{\chi}\left(\widetilde{\theta}\right), \widetilde{\theta}\right)\right].$$

Recall that χ^* maximizes this expression, and $\widehat{\chi} = \chi^*$ if (β_i, α_i) is a truthful-obedient strategy. Thus there is no profitable deviation from truthful-obedience. ■

8 Balancing the Transfers

So far, transfers of the Team Mechanism are not balanced.
Recall $\psi_{i,t}^{M}(\theta) = \sum_{j \neq i} u_{j,t}\left(\chi^{*}\left(\theta\right),\theta\right) \qquad \Psi_{i}^{M}(\theta) = \sum_{j \neq i} U_{j}\left(\chi^{*}\left(\theta\right),\theta\right).$

Construct balanced transfers for a mechanism (χ, ψ):

$$\psi_{i,t}^{B}(\theta^{t}) = \gamma_{i,t}\left(\theta_{i}^{t}, \theta_{-i}^{t-1}\right) - \tfrac{1}{I-1}\sum_{j \neq i} \gamma_{j,t}\left(\theta_{j}^{t}, \theta_{-j}^{t-1}\right)$$, where

$$\gamma_{i,t}\left(\theta_{i}^{t}, \theta_{-i}^{t-1}\right) = \delta^{-t}\left(E_{\widetilde{\theta}}^{\mu[\chi]|\theta_{i}^{t},\theta_{-i}^{t-1}}\left[\Psi_{i}(\widetilde{\theta})\right] - E_{\widetilde{\theta}}^{\mu[\chi]|\theta^{t-1}}\left[\Psi_{i}(\widetilde{\theta})\right]\right)$$

Proposition 2 Consider any mechanism (χ, ψ) in which there is a truthful-obedient BNE. If the agents' types are independent, then there is also a truthful-obedient BNE in the balanced mechanism $\left(\chi, \psi^{B}\right)$.

Corollary 1 With independent types and private values, there is a truthtelling BNE in the Balanced Team Mechanism.

9 Markov games with Private States

A Folk-theorem like result for Markov games: efficiency can be sustained in equilibrium when the discount factor is close enough to 1.

9.1 The Game

Each agent i privately observes his private type $\theta_{i,t}$
Each agent i makes a public report $\widehat{\theta}_{i,t}$
Each agent i chooses a publicly observable action $x_{0,i,t} \in X_{0,i,t}$, a privately observed action $x_{i,t} \in A_{i,t}$, and a publicly observable payment $z_{i,j,t} \geq 0$ to each agent *j*.

Let $x_t = ((x_{0,1,t}, .., x_{0,I,t}), x_{1,t}, .., x_{I,t})$,
$y_{i,t} = \sum_{j=1}^{I} (z_{j,i,t} - z_{i,j,t})$
then the payoff of agent i is, as before: $\sum_{t=1}^{\infty} \delta^t \left(u_{i,t}(x^t, \theta^t) + y_{i,t}\right)$

9.2 Assumptions

- (a) $\Theta_t = \overline{\Theta}$ and $X_t = \overline{X}$ for all t, and all sets are finite.
  (b) Per-period utilities $u_{i,t}(x^t, \theta^t)$ depend only on period t variables and do not depend on t, so $u_i(x_t, \theta_t)$.

(c) Similarly, the probability measures $\nu_t|x^{t-1},\theta^{t-1} = \nu|x_{t-1},\theta_{t-1}$.

(d) Independent types, private values.

(e) $\exists\widehat{x}_t \in X_t$ such that $u_i(x_{i,t},\widehat{x}_{-i,t},\theta_t) = 0$ for all i, all $x_{i,t} \in \overline{X}$, and all $\theta_{i,t} \in \overline{\Theta}$.

Definition 5 Decision policy χ is a Markov policy if each $\chi_t(\theta^t)$ depends only on the current state and does not depend on t directly, and thus can be written as $\chi(\theta_t)$. A decision policy χ^* is a Blackwell policy, if it is a Markov policy and there exists $\overline{\delta} < 1$ such that for all $\delta \in \left[\overline{\delta},1\right]$ it achieves a weakly greater present expected total surplus starting at any state than any other observationally measurable policy.

Puterman (1994) proves existence and constructs a Blackwell policy.

In the Markov game the Team transfers can be written as:

$$\psi_{i,t}(\theta_t) = \sum_{j\neq i} u_j\left(\chi\left(\theta_t\right),\theta_t\right)$$

and the Balanced transfers can be written as:

$$\psi^B_{i,t}\left(\theta_{i,t},\theta_{t-1}\right) = \gamma_i\left(\theta_{i,t},\theta_{t-1}\right) - \tfrac{1}{I-1}\sum_{j\neq i}\gamma_j\left(\theta_{j,t},\theta_{t-1}\right)$$

$$\gamma^B_i\left(\theta_{i,t},\theta_{t-1}\right) = \sum_{\tau=t}^{\infty}\delta^{\tau-t}\sum_{j\neq i}\left[E^{\mu^i_t[\chi]|\theta_{i,\tau},\theta_{\tau-1}}_{\widetilde{\theta}}u_j\left(\chi\left(\widetilde{\theta}_\tau\right),\widetilde{\theta}_\tau\right) - E^{\mu_t[\chi]|\theta_{\tau-1}}_{\widetilde{\theta}}u_j\left(\chi\left(\widetilde{\theta}_\tau\right),\widetilde{\theta}_\tau\right)\right]$$

Definition 6 Given a Markov process, a set of states $S \subseteq \overline{\Theta}$ is called ergodic, if, starting in S, all system remains remains there with probability 1, and this is not true for any proper subset of S.

Proposition 3 Consider dynamic game satisfying assumptions 1-5. Let χ^* be a Blackwell decision policy for the game, and suppose that the Markov process induced by χ^* has a single ergodic set, and that its invariant distribution yields a positive expected total surplus. Then there exists $\delta^* < 1$ such that, for all $\delta \in (\delta^*, 1)$, the game has an efficient PBE that sustains decision plan χ^*.

10 Extensions

- Relax Markovian properties, look at limited dependence
- Allocations other than efficient ones
- Costly computation and communication
- Correlated values (efficiency with BB may be achievable more easily)

11 Summary

- Dynamic Setting with Private Info
- Construct an Efficient, BB, BIC mechanism
- Extension of AGV/Arrow (1979)
- Idea: charge an agent for the *change* in the opponent utilities induced by his report
- Conditions of Self-Enforcing Mechanism (in a Markov Setting)