CS/Ec241 Notes : Winter 2008

John Ledyard

February 23, 2008

1 Social Decision Problems

- $I = \{i = 1 \dots N\}$: individuals
- $x \in X$: outcome space
- $u^i(x, \theta^i)$: i's utility where $\theta^i \in \Theta^i$ is i's type
- $e \in E$: environments, where $e = (I, X, \theta^1, \dots, \theta^N)$
- $P : E \rightarrow X$: performance standard, sometimes written as $P : \Theta \rightarrow X$.

  An example: $P(\theta) = \text{argmax}_{x \in X} \sum_i u^i(x, \theta^i)$
  An example: $P(\theta) = \{x | x \text{ is Pareto-optimal in X}\}$
  An example: $P(\theta) = \text{argmax}_X u^i(x, \theta^i)$. A dictatorial rule.

2 Mechanisms

A mechanism will be combined with an environment to create a game. We will use a behavioral assumption, how agents will play that game, to produce an output for each environment $e \in E$.

Begin with a set of environments E. A mechanism is

- $I = \{i = 1 \dots N\}$: individuals
- M^i : strategies (language, messages,...) for i
- $g : M^1 \times ... \times M^N \rightarrow X$: outcome function. $\gamma = (M, g)$ is called a game form where M can involve complex extensive forms.

Given e and the mechanism $\gamma = (M, g)$, each individual is now in a game. That game is described by:

- $I = \{i = 1 \ldots N\}$: individuals
- M^i : strategies for i
- $v^i(m, \theta^i) = u^i(g(m), \theta^i)$: i's payoff.

The behavioral model tells us how the individuals play this game. We let $\mu^i(\theta^i, \gamma) \to M^i$ be this behavior.

Combining mechanism and behavior we get

$$g(\mu^1(\theta^1, \gamma), \ldots, (\mu^1(\theta^N, \gamma)) = x$$
$$\text{or} \qquad g(\mu(e, \gamma)) = x$$

So given γ we get a performance mapping $g \circ \mu : \Theta \to X$. We can compare this to the performance standard. We will say that γ (weakly) implements P in E under the behavior μ if and only if $P(\theta) \cap g(\mu(\theta)) \neq 0 \quad \forall e \in E$.

3 Dominant Strategy Mechanisms

Definition 1 *(Dominant Strategy Equilibrium) A strategy* s^* *of the game* (I, S, v) *is a dominant strategy equilibrium iff*

$$v^i(s^*) \geq v^i(s^*/s^i), \quad \forall s^i \in S^i, \quad \forall i \tag{1}$$

Definition 2 *(Dominant Strategy Mechanism) We say a mechanism* γ *is a Dominant Strategy mechanism in E if for every* $e \in E$*, there is a strategy* $m^*(e)$ *such that* $m^*(e)$ *is a dominant strategy equilibrium for the game induced by the mechanism* γ *in e. For dominant strategy mechanisms, we let* $\mu^i(\theta^i, \gamma) = m^{*i}(\theta)$*.*

Note that not all mechanisms have dominant strategy equilibria. And some mechanisms may have many dominant strategy equilibria. For an example of the later, consider the constant mechanism, $g(m) = \hat{x}, \quad \forall m$. Then every m is a dominant strategy equilibrium.

3.1 Gibbard-Satterthwaite

Can we find dominant strategy mechanisms that implement arbitrary performance standards? The answer is no.

Theorem 1 *(Gibbard-Satterthwaite) Given a set of environments E such that* $|X| \geq 3$ *and such that* Θ *is rich enough, the only performance standards* $P : \Theta \to X$*, that can be implemented in dominant strategies are (a) dictatorial (that is, there is some i' such that* $P(\theta) = argmax\ u^i(x, \theta^i) \quad \forall \theta \in \Theta$*) or (b) constant (that is, there is some* $\hat{x} \in X \ni P(\theta) = \hat{x} \quad \forall \theta \in \Theta$*).*

3.2 Quasi-linearity and VCG mechanisms

If we restrict environments to those with quasi-linear utility functions, we can do better. $u^i(x, \theta^i)$ is quasi-linear if there is a commodity y^i such that $u(x, \theta^i) = w^i(x, \theta^i) + y^i$.

Suppose that our performance standard is $P(\theta) = \text{ argmax } \sum w^i(x, \theta^i)$. We can implement this in dominant strategies using VCG mechanisms.[1] VCG mechanisms are really a class of mechanisms. I will describe one.

Definition 3 *(VCG Mechanism) Let* $M^i = \Theta^i$*; that is, i is asked to report her utility function. The VCG outcome function is* $g(m) = (g_x(m), g_y(m))$ *where*

$$g_x(m) = argmax \sum_i w^i(x, m^i) \tag{2}$$

$$g_{y^i}(m) = -\left\{[max_{x\in X} \sum_{j\neq i} w^j(x, m^j)] - [\sum_{j\neq i} w^j(g_x(m), m^j)]\right\} \tag{3}$$

First note that if $b^i(\theta^i, VCG) = \theta^i$ then we will implement $P(\theta) = \text{argmax } \sum_i w^i(x, \theta^i)$ from (2). So the only question is whether θ is a dominant strategy equilibrium in θ. Player i will choose m^i to maximize $w^i(g_x(m), \theta^i) + \sum_{j\neq i} w^j(g_x(m), m^j)$. That is i wants the x that solves (2) when $m^i = \theta^i$. So θ is indeed a dominant strategy equilibrium of the game induced by the mechanism (2) and (3).

Note that if we add any function $h^i(m_{-i})$ to $g_{y^i}(m)$, where m_{-i} is the vector m with m^i removed, θ will still be a dominant strategy, so that there is a whole class of VCG mechanisms.

---

[1] VCG is Vickrey-Clarke-Groves.

Example 1 *(Allocating A Single Item) Consider the simplest allocation problem. Allocate an item to one of N people. Let* $x = (x^1, ..., x^N)$ *be the output decision where if i gets the item then* $x^i = 1$*. We will also let* x^i *be the probability that i gets the good. We require, for feasibility that* $x^i \in [0,1], \forall i$ *and that* $\sum x^i \leq 1$*. The utility of i is* $u^i = \theta^i x^i + y^i$*, so that* θ^i *is i's value (or willingness to pay) for the good. What is the VCG mechanism for this problem?*

The message space is $M^i = \Theta^i$*. The mechanism outcome rule is:* $g_x(m) =$ *argmax* $\sum m^i x^i$ *subject to* $\sum x^i \leq 0, x^i \in [0,1]$ *and* $g_{y^i}(m) = \{ \sum_{j \neq i} m^j g_x^j(m)\} - \{max \sum_{j \neq i} m^j x^j$ *subject to* $\sum_{j \neq i} x^j \leq 0\}$*. So the item is awarded to the highest* m^i *(where* m^i *can be thought of as i's bid for the item). If i does not win the item (i.e.,* $x^i = 0$*) then it is reasonably easy to see that* $g_{y^i} = 0$*. If i does win the item then* $\sum_{j \neq i} m^j g_x^j(m) = 0$ *and* $y^i = -$ *max* $[\sum_{j \neq i} m^j x^j$ *subject to* $\sum_{j \neq i} x^j \leq 0]$ *which is the second highest value of* m^k*. Thus the VCG mechanism is a sealed bid, second price auction. It awards the item to the highest valued bidder at a price equal to the second highest bid.*

It turns out that VCG mechanisms are essentially the only dominant strategy mechanisms that maximize the sum of utilities.

Theorem 2 *(Green-Laffont, Walker) If* γ *is a dominant strategy mechanism that implements* $p(\theta) = \sum_i w^i(x, \theta)$*, and* Θ *is rich enough,*[2] *then* γ *is essentially a VCG mechanism.*

What is meant by "essentially"? It means that the outcomes of any dominant strategy mechanism will be the same as a VCG mechanism. This follows from the Revelation Principle which uses the concept of a direct revelation mechanism.

Definition 4 *(Incentive Compatible Direct Mechanism) A mechanism* γ *is a Direct Revelation mechanism if* $M^i = \Theta^i \quad \forall i$*. That is, messages are announcements about a player's type. A direct mechanism is incentive compatible in E if* θ *is a dominant strategy equilibrium for the game induced by the mechanism. That is, truthful revelation is a dominant strategy equilibrium.*

---

[2] For example if Θ includes all concave functions and X is a convex set.

Theorem 3 *Revelation Principle (Gibbard, Myerson) If γ is any mechanism with dominant strategy equilibrium outcomes $g(m^*(\theta))$, then there is an incentive compatible direct revelation mechanism with exactly the same equilibrium outcomes.*

Proof: This is almost trivial. Let $g^*(\theta) = g(m^*(\theta)$.

It should be noted that VCG mechanisms are not efficient, unless one includes the mechanism manager in the calculations.

Definition 5 *(Efficient Mechanisms) A mechanism γ is Efficient for E under the behavior b iff for all $\theta \in \Theta$ there is no allocation $x' \in X$ such that*

$$u^i(x', \theta^i) \geq u^i(g(b(\theta), \theta^i) \quad \forall i \tag{4}$$

$$u^i(x', \theta^i) > u^i(g(b(\theta), \theta^i) \textit{ for some } i \tag{5}$$

It is easy to see that VCG mechanisms would be efficient if the budget balanced; that is, if $\sum_i {g_y}^i(\theta) = 0 \quad \forall \theta \in \Theta$. Unfortunately there are no VCG mechanisms which balance the budget.[3]

Theorem 4 *(Hurwicz-Walker) If Θ is rich enough, then for any VCG mechanism γ it is true that generically in $\Theta, \sum_i g_{y^i}(\theta) \neq 0$.*

This means that if we want to use dominant strategy mechanisms we will either have to burn money (i.e. throw away the excess y collected) or use a mechanism which does not maximize $\sum w^i$. We have to accept some inefficiency.

3.3 (Incentive) Efficient Dominant Strategy Mechanisms

One question might be: how well can we do? To answer this we need some measure of performance instead of a standard. One way to do this without

---

[3] If we used w^i in place of u^i, that is if we ignored the payments, then VCG mechanisms would be efficient. Or if we added the mechanism manager's utility $-\sum_i y^i$ to (4) and (5), then VCG mechanisms would be efficient.

fully committing to a social welfare function is to look for mechanisms that are not dominated by other mechanisms.[4]

Definition 6 *(Ex Post Efficient Mechanisms) Let G be a set of mechanisms on E. We will say that γ^* is Ex Post Efficient in G under the behavior b iff there is no other mechanism γ' such that*

$$u^i(g'(b(\theta,\gamma'),\theta^i) \geq u^i(g*(b(\theta,\gamma^*),\theta^i), \quad \forall i, \forall e \in E \tag{6}$$

$$u^i(g'(b(\theta,\gamma'),\theta^i) > u^i(g*(b(\theta,\gamma^*),\theta^i), \quad \textit{for some } i,\ \textit{for some } e \in E \tag{7}$$

If $x^*(\theta)$ were efficient in all $\theta \in \Theta$ and also Incentive Compatible on E, then $x^*()$ would be an ex post efficient mechanism. But as we have seen, there are no such x^* if Θ is rich enough.

I believe it is an open question as to what dominant strategy mechanisms are ex post efficient. However, we can provide a statement of the problem in a form that might generate a solution. But first a theorem that might be helpful.

Theorem 5 *(Efficiency as Maximization) If $u^i(x,\theta^i)$ is concave for all i, and if X is convex, then A and B are equivalent.*

$$\begin{aligned}
&A)\ \textit{There is no } x' \in X \ni \\
&\qquad u^i(x',\theta^i) \geq u^i(x^*,\theta^i) \quad \forall i \\
&\qquad u^i(x',\theta^i) > u^i(x^*,\theta^i)\ \textit{some } i \\
&B) \exists \lambda = (\lambda^1,\ldots,\lambda^N) \geq 0 \ni \\
&\qquad x^* \in \textit{argmax}_{\,x\in X} \sum \lambda^i u^i(x,\theta^i)
\end{aligned}$$

We can use a version of this theorem along with the revelation principle to state

---

[4] One way to provide a performance measure would be to endow the mechanism designer with beliefs, say a distribution function $F(\theta)$, about what the true θ is and then ask them to chose an incentive compatible direct revelation mechanism that maximizes aggregate expected utility or $\int \sum (u^i(x(\theta),\theta^i) dF(\theta)$. We will look at something like this later. For now we will avoid beliefs.

Theorem 6 *γ is an Ex Post (Incentive) Efficient Dominant Strategy mechanism on E iff $\exists\lambda : \Theta \rightarrow R^N \ni x^*(\theta) = g(b(\theta), \theta)$ solves*

$$max_{x:\Theta\rightarrow X} \sum_{i,\theta\in\Theta} \lambda^i(\theta)u^i(x(\theta), \theta) \tag{8}$$

subject to

$$u^I(x(\theta), \theta^i) \geq u^i(x(\theta/\hat{\theta}^i), \theta^i) \quad \forall\hat{\theta}^i \in \Theta^i \tag{9}$$

(9) is sometimes summarized as "$x(\theta)$ is incentive compatible." Since the u are concave they describe a convex set. The trick in applying this result is to convert (9) into a form that is analytically tractable. There are a number of papers that characterize the set of (dominant strategy) Incentive Compatible mechanisms.[5] None that I know of is easily amenable to use in this maximization problem.

3.4 Computational Compatibility

In many cases, carrying out the computations required by VCG mechanisms is NP-hard. This is true for example for combinatorial auctions. So even if one were satisfied with dominant strategy mechanisms that just maximize $\sum w^i$ instead of $\sum u^i$ it may not be possible to use a direct revelation mechanism. So what happens if we restrict consideration to computable mechanisms? Not much is known.

We want to find mechanisms $\gamma = (g, M)$ that are incentive compatible on Θ (they have at least one dominant strategy equilibrium, that are computationally manageable, and that are good (where good could be ex post efficiency). The first thing we have to give up is the revelation principle. It is possible that no direct revelation mechanism is computationally manageable.

Given a class of environments, let G be the set of all mechanisms and $G^D \subset G$ be the set of all dominant strategy mechanisms. Let $G^C \subset G$ be the set of all computationally manageable mechanisms, suitably defined.[6] We are interested in $G^F = G^D \cap G^C$. Let $G^R \subset G$ be the set of all incentive compatible direct revelation mechanisms. Note that $G^R \subset G^D$. The revelation

---

[5] For example, a result of Roberts (1977) shows that if Θ is rich enough then a dominant strategy mechanism will maximize $F(x) + \sum \alpha^i u^i(x, \theta)$ for some α and $F()$.

[6] For example, it might be the set of all mechanisms such that the computations required by g are polynomial.

principle states that if $\gamma \in G^D$ then there is a $\hat{\gamma} \in G^R$ with the same performance. Unfortunately, we cannot show that if $\gamma \in G^D \cap G^C$, then there is a $\hat{\gamma} \in G^R \cap G^C$ with the same performance. We do know that if $\gamma \in G^D \cap G^C$, then there is a $\hat{\gamma} \in G^R$ with the same performance. Let G^B be the set of all of those. It is an open question as to what G^B looks like.

4 Bayesian Mechanisms

In this section, we relax somewhat the concept of incentive compatibility. We also allow the mechanism manager to use information about the players. We begin by considering a generalization of our class of environments.

4.1 Bayesian Social Decision Problems

In a Bayes Decision Problem, it is assumed that the decision maker has some knowledge about the environment which is embodied in beliefs. The structure is as follows:

- $I = \{i = 1 \dots N\}$: individuals
- $x \in X$: outcome space
- $u^i(x, \theta^i)$: i's utility where $\theta^i \in \Theta^i$ is i's type
- $f(\theta)$ a density function on Θ with distribution function, F.
- $e \in E$: environments, where $e = (I, X, \theta^1, \dots, \theta^N, F)$
- $P : E \to X$: performance standard, sometimes written as $P : \Theta \times \mathcal{F} \to X$, where $\mathcal{F}$ is the space of all F on Θ.

An example: $P(e, F) = \text{argmax}_{x \in X} \int \sum u^i(x, \theta^i) dF(\theta)$.

Note: θ^i can be more than just utility function parameters. For example, it can include parameters about the beliefs of others, and about the beliefs of others about others beliefs, etc.

4.2 Bayesian Mechanisms

The standard approach assumes that F is common knowledge among players and mechanism manager. A Bayesian mechanism is then

- $I = \{i = 1 \dots N\}$: individuals
- M^i : strategies (language, messages,...) for i
- $g : M^1 \times ... \times M^N \times \mathcal{F} \rightarrow X$: outcome function. $\gamma = (M, g)$ is called a game form where M can involve complex extensive forms.

Note that in order to define a Bayes mechanism one needs to know E since one needs to know X and F.

Given e and the mechanism $\gamma = (M, g)$, each individual is now in a Bayes game. That game is described by:

- $I = \{i = 1 \dots N\}$: individuals
- $S^i = \{s^i : \Theta^i \rightarrow M^i\}$: strategies for i
- $v^i(s, \theta^i) = \int u^i(g(m(\theta), F), \theta^i) dF(\theta|\theta^i)$: i's payoff.

The behavioral model tells us how the individuals play this game. For Bayes mechanisms we will assume players choose a Bayes equilibrium.

Definition 7 *(Bayes Equilibrium) s^* is a Bayes equilibrium if*

$$s^{*i}(\theta^i, \gamma) = \ argmax_{m^i \in M^i} \int u^i(g(s^*(\theta)/m^i, F), \theta^i) dF(\theta|\theta^i), \forall \theta^i \in \Theta^i, \forall i. \tag{10}$$

So the performance of the mechanism γ is given by

$$x^*(\theta) = g(s^*(\theta), F) \tag{11}$$

There is a revelation principle for Bayesian mechanisms.

Theorem 7 *(Bayesian Revelation Principle) If γ is any Bayes mechanism with Bayes equilibrium outcomes $g(s^*(\theta, F), F)$, then there is a (Bayesian) incentive compatible direct revelation mechanism with the same equilibrium outcomes.*

Proof: It is straight forward to show that the mechanism given by $h(\theta, F) = g(s^*(\theta, F), F)$ does the job.QED

4.3 Quasi-Linearity

If we restrict attention to quasi-linear environments we can get some nice results. Remember, quasi-linear utility functions have the form $u^i(x, y^i, \theta^i) = w^i(x, \theta^i) + y^i$. The first hint about the gains from relaxing incentive compatibility comes from the following result:

Theorem 8 *(d'Aspremont & Gerard-Varet) If utility functions are quasi-linear then there is a Bayesian mechanism that selects efficient allocations for all* $\theta \in \Theta$.

Proof: We will provide a direct revelation mechanism. Let $h(\theta) = (h_x(\theta), h_y(\theta))$ be the VCG mechanism. Remember, $h_x(\theta) \in \operatorname{argmax}_{x \in X} \sum w^i(x, \theta)$. So if we can find a way to balance the budget and make $\sum h_{y^i}(\theta) = 0 \quad \forall \theta$, we will have an efficient mechanism. Let $T^i(\theta^i) = \int h_{y^i}(\theta) dF(\theta|\theta^i)$. Let $\eta^i(\theta_{-i}) = \frac{1}{N-1} \sum_{j \neq i} T^j(\theta^j)$. Finally, let $\hat{h}^i(\theta) = h_{y^i}(\theta) - \eta^i(\theta_{-i})$. It is easy to see that $\sum_i \hat{h}^i(\theta) = 0 \quad \forall \theta$. Consider $U^i(s^i, \theta^i) = \int w^i(h_x(\theta/s^i), \theta^i) - \hat{h}^i_{y^i}(\theta^i) dF(\theta|\theta^i) = \int w^i(h_x(\theta/s^i), \theta^i) - h^i_{y^i}(\theta) dF(\theta) + \int \eta^i(\theta_{-i}) dF(\theta|\theta^i)$. Since θ is a dominant strategy equilibrium for all $\theta \in \Theta$ for (h_x, h_y), then $s(\theta) = \theta$ is a Bayes equilibrium for the mechanism given by $g = (h_x, \hat{h})$. So $g = (h_x, \hat{h})$ is Bayesian incentive compatible and balances the budget. Thus, it chooses an efficient allocation in each θ. QED

The Bayesian mechanism is a VCG mechanism. The magic happens because we can let $\hat{h}^i$ depend on F. But, the good news needs to be tempered. There are at least two "problems". First, the AGV mechanism may leave participants worse off than if they didn't participate at all.[7] That is, it is possible that $\int w^i(g(\theta), \theta^i) - y^i(\theta) dF(\theta) < 0$. So what happens if we try to avoid this by asking that the Bayes mechanisms satisfy (Bayes) Voluntary Participation.

Definition 8 *(Voluntary Participation) A Bayes mechanism satisfies Voluntary Participation iff, when* s^* *is the Bayes equilibrium,*

$$\int w^i(g(s^*(\theta)), \theta^i) - y^i(s^*(\theta)) dF(\theta|\theta^i) \geq 0, \quad \forall \theta^i, \forall i. \tag{12}$$

---

[7] This is also true of the VCG mechanism.

Theorem 9 *(Myerson-Satterthwaite) For very simple environments there are no Bayesian incentive compatible mechanisms that are efficient and satisfy voluntary participation.*

Second, suppose the mechanism manager doesn't know the common prior and g cannot depend on F. Then we have the following result.

Theorem 10 *(Ledyard) Consider direct revelation mechanisms. For a mechanism* $\gamma = (g, \Theta)$*, h* θ *is the Bayes equilibrium for all* $\theta \in \Theta$*, and all* $F \in \mathcal{F}$ *iff* θ *is a dominant strategy for all* $\theta \in \Theta$*.*

With the revelation principle, this means that, unless g can depend on F, the class of Bayesian incentive compatible mechanisms is exactly the same as the class of dominant strategy mechanisms. This observation has been morphed (by Bergemann and Morris) into a more sophisticated argument using generalized priors and a concept of "Robust Mechanisms"[8] to say that a robust mechanism is Bayesian Incentive compatible for all priors on Θ iff the mechanism is ex post incentive compatible on all Θ.

Definition 9 *(Ex post Incentive Compatibility) A direct revelation mechanism is ex post incentive compatible on* Θ *iff* $s^*(\theta) = \theta$ *is an ex post Nash equilbrium for all* $\theta \in \Theta$*.* $s^*(\theta)$ *is an ex post Nash equilibrium if*

$$u^i(s^*(\theta), \theta^i) \geq u^i(s^*(\theta/m^i), \theta^i), \quad \forall m^i \in \Theta^i, \forall i. \tag{13}$$

Note that for the utility functions that depend only on θ^i, a mechanism is ex post incentive compatible iff that mechanism is dominant strategy incentive compatible. If we broaden the class of utility functions and allow them to have interdependent values, so that $u^i = u^i(x, \theta)$, then ex post Nash is stronger than dominant strategy. There are many environments, including quasi-linear ones, in which ex post Nash mechanisms do not exist.

---

[8] Robust mechanisms can depend only on payoff relevant parts of θ.

4.4 (Incentive) Efficient Bayesian Mechanisms

We can ask the same second-best question for Bayes mechanisms that we did for Dominant strategy mechanisms. But we have a new problem here. What information do we use in making our judgement? There are, at least, 3 possibilities: (i) before agents know their types, (ii) after agents know their own types but before they know all types, and (iii) after all agents know all types. These have been dubbed, respectively, by Holmstrom and Myerson, as ex ante, interim, and ex post. "Ex post" is what we used for dominant strategy mechanisms, the idea being that the mechanisms should be immune to manipulation even if all information is revealed and known. For Bayesian mechanisms, the interim stage seems more appropriate.

Definition 10 *(Interim Incentive Efficient Mechanisms) Let G be a set of mechanisms on E. We will say that γ^* is Interim Efficient in G under behavior, b, iff there is no other mechanism γ' such that*

$$\int u^i(g'(b(\theta, \gamma', F), \theta^i) dF(\theta|\theta^i) \geq \int u^i(g(b(\theta, \gamma, F), \theta^i) dF(\theta|\theta^i), \forall i, \forall \theta^i \in \Theta \tag{14}$$

$$\int u^i(g'(b(\theta, \gamma', F), \theta^i) dF(\theta|\theta^i) > \int u^i(g(b(\theta, \gamma, F), \theta^i) dF(\theta|\theta^i), \textit{some } i, \textit{some } \theta^i \in \Theta \tag{15}$$

"Some" in the definition means "for a set of positive measure". A mechanism is Interim efficient if there is no other mechanism that will make some type of some agent better off and no type of any i worse off. It should be clear that if a mechanism is Ex Post Efficient then it is Interim Efficient. It can be shown that a mechanism is Interim Efficient iff it is common knowledge that no one can be made better off. If a mechanism were not Interim Efficient, it could be replaced, prior to the revelation of information, with one that everyone knows would make everyone at least as well off and some better off.

We can use a version of Theorem 5 along with the Revelation Principle (see Theorem 7) to state

Theorem 11 *γ is an Ex Post (Incentive) Efficient Dominant Strategy mechanism on E iff $\exists \lambda : \Theta^i \rightarrow R^N \ni x^*(\theta) = g(b(\theta), \theta)$ solves*

$$max_{x:\Theta \rightarrow X} \int \sum_{i,\theta \in \Theta} \lambda^i(\theta^i) u^i(x(\theta), \theta) dF(\theta) \tag{16}$$

subject to

$$\int u^I(x(\theta), \theta^i) dF(\theta|\theta^i) dF(\theta|\theta^i) \geq \int u^i(x(\theta/\hat{\theta}^i), \theta^i) dF(\theta|\theta^i) \quad \forall \hat{\theta}^i \in \Theta^i \tag{17}$$

Note that the λ^i depend only on θ^i and not on all of θ as they did in the ex post version. The constraint (17) is sometimes summarized as "$x(\theta)$ is Bayesian (or Interim) incentive compatible." If $x^*(\theta)$ were efficient in all $\theta \in \Theta$ and also Incentive Compatible, then $x^*()$ would be an interim efficient mechanism. So the AGV mechanism is interim efficient. (Let $\lambda^i = 1, \forall \theta^i$.)

4.4.1 The Optimal Auction Problem (Myerson)

Example 2 *(Auctioning a Single item) Consider the problem of allocating one good to one of N people. As before assume that $u^i = \theta^i x^i + y^i$. Suppose we include the auctioneer as one of the players by letting them have a utility function of $u^0 = -\sum_{=1}^N y^i$. If we let $\lambda^i = 0 \forall i = 1, \ldots, N$ and $\lambda^0 = 1$, then we have the optimal auction problem:*

$$max \int \sum_{i=1}^N y^i dF(\theta)$$

subject to $[x(\theta), y(\theta)]$ is incentive compatible

and satisfies voluntary participation

To make the problem easy we will assume that beliefs are that the θ^i are independently and identically distributed. That is, $dF(\theta) = dH(\theta^1)...dH(\theta^N)$.

Since the u are linear, the incentive compatibility constraints describe a convex set. The trick in applying this result is to convert (17) into a form that is analytically tractable. First note that Bayesian Incentive Compatibility of $[x(\theta), y(\theta)]$ is equivalent to

$$\theta^i \in \text{ argmax}_{m^i \in \Theta^i} \Big\{ \int \theta^i x(\theta/m^i) dF(\theta|\theta^i) - \int y^i(\theta/m^i), \theta^i) dF(\theta|\theta^i) \Big\} \tag{18}$$

Let $Q^i(m^i) = \int x(\theta/m^i)dF(\theta|\theta^i)$ and $Y^i(m^i) = \int y^i(\theta/m^i)dF(\theta|\theta^i)$. Looking at first and second order conditions[9] for the maximum, we get

$$\theta^i(dQ^i(\theta^i)/dm^i) + dY^i(\theta^i)/dm^i = 0 \tag{19}$$

$$dQ^i(\theta^i)/d\theta^i + \theta^i(d^2Q^i(\theta^i)/d^2m^i) + d^2Y^i(\theta^i)/d^2m^i \leq 0 \tag{20}$$

I will often write this as

$$\theta\dot{Q} + \dot{Y} = 0 \tag{21}$$

$$\theta\ddot{Q} + \ddot{Y} \leq 0 \tag{22}$$

Several wonderful things follow from these. First since $u = \theta Q + Y$, $\dot{u} = Q + \theta\dot{Q} + \dot{Y} = Q$. So $u^i(\theta^i) = u^i(\theta_0) + \int_{\theta_0}^{\theta} Q(s)ds = \theta Q + Y$. This means that[10] $Y^i(\Theta^i) = u^i(\theta_0) + \int_{\theta_0}^{\theta} Q(s)ds - \theta Q$. So we can write the objective function of the optimal auction problem as

$$-\int\sum Y^i(\theta^i)dH = -[\sum u^i(\theta_0)] + [\sum\int \theta^i Q^i(\theta^i)dH \tag{23}$$

$$-\sum[\int\int_{\theta_0}^{\theta} Q(s)dsdH] \tag{24}$$

Second, because $u^i(\theta^i) = u^i(\theta_0) + \int_{\theta_0}^{\theta} Q(s)ds$ we know that to satisfy voluntary participation, we want $u(\theta_0) \geq 0$. Therefore to maximize expected revenue we want $u(\theta_0) = 0$.

Third, reconsider the second order conditions (SOC, 22). If we differentiate the FOC (21) with respect to θ, we get $\dot{Q} + \theta\ddot{Q} + \ddot{Y} = 0$. So the SOC are equivalent to $\dot{Q} \geq 0$.

Now integrate $\int\int_{\theta_0}^{\theta} Q(s)dsdH$ by parts to get $\int(Q/h)dH \int(HQ/hdH)$ so that the objective to the optimal auction problem is $\sum_i \int \{\theta - \frac{1-H(\theta)}{h(\theta} big\}Q(\theta dH$. We can expand Q back to what it was and so the optimal auction problem

---

[9] These are necessary and sufficient for a local maximum in this case. One usually needs to worry that they might miss the global maximum, but in this case we do not. We also need to assume some differentiability of the mechanisms.

[10] Kim Border has reminded me that one can use subgradients to get this result without assuming the continuity necessary to use derivatives. See, e.g., Rochet.

is: choose $x(\theta)$ to

$$\max \sum_i \int \{\theta^i - \frac{1-H(\theta^i)}{h(\theta^i)}\} x(\theta) dF(\theta) \tag{25}$$

$$\text{subject to } \sum x^i(\theta) \leq 1, \forall\theta \tag{26}$$

$$\text{and } \dot{Q} \geq 0. \tag{27}$$

If we ignore the SOC (27) for now, the optimal auction mechanism awards the item to the individual with the highest value of[11] $\theta^i - \frac{1-H(\theta^i)}{h(\theta^i}$, as long as that value is greater than 0. It would be nice if this also satisfied the SOC. To this end, notice that if $d\{\theta^i - \frac{1-H(\theta^i)}{h(\theta^i}\}/d\theta^i \geq 0$, then the probability of winning the item conditional on θ^i is increasing in θ, so $\dot{Q} \geq 0$. So we assume[12] that $\frac{1-H(\theta^i)}{h(\theta^i}$, the inverse hazard rate, is decreasing in θ. Under this assumption our solution satisfies the SOC and, so, is a solution to the optimal auction problem. Further, since θ^i are iid, the solution is equivalent to awarding the item to the highest bidder as long as their bid is greater than or equal to $\bar{R}$ where $\bar{R}$ is the solution to $R = \frac{1-H(R)}{h(R)}$. You can think of $\bar{R}$ as a reservation price.

Finally, for the optimal mechanism, $Q(\theta) = Pr\{\theta^i \geq \theta^j, \forall j \neq i | \theta^i \geq \bar{R}\} = H^{N-1}(\theta^i)$ if $\theta^i \geq \bar{R}$ and $Q(\theta) = 0$ otherwise. Using this, the payment of i can be calculated to be $Y^i(\theta^i) = -\int_{\bar{R}}^{\theta^i} sdQ^i(s) = -\int_{\bar{R}}^{\theta^i} s(N-1)H^{N-2}(s)h(s)ds$. But this is just the expected value of the second highest bid if it is less than θ^i and bigger than $\bar{R}$. So a second price auction with reserve price of $\bar{R}$ is an optimal auction.

We can also get a remarkable additional result from this analysis.

Theorem 12 *(Revenue Equivalence Theorem If two auctions have the same allocation rule $x(\theta)$ then they will yield the same expected revenue to the auctioneer.*

The proof follows from the fact that $Q^i(\theta^i)$ will be the same for all i for both auctions. So $Y^i(\theta^i) = -\int_{\bar{R}}^{\theta^i} sdQ^i(s)$ will be the same for both auctions. So any auction that awards the item to the highest bidder will yield identical expected revenue to the seller. A simple example of this[13] is the fact that

---

[11] $\theta^i - \frac{1-H(\theta^i)}{h(\theta^i}$ is called i's virtual valuation by Myerson.

[12] Myerson calls this the "regular case".

[13] This example comes from Vickrey in the same article as his second price auction.

the first price and second price auction are revenue equivalent. Any attempt to get around this, e.g. by using fancy payment rules, will be undone by the bidder's choice of bid.

4.5 Relation Between Bayesian and Dominant Strategy Mechanisms

There is a neat result of Mookherjee and Reichelstein that allows us to often convert a Bayes Incentive Compatible mechanism into a Dominant Strategy mechanism when the type space is one dimensional. If the outcome rule of the Bayes mechanism (i.e., $x(\theta)$) has a certain monotonicity property then we can often unbundle $Y = -\theta Q + \int_{\theta_0}^{\theta} Q(s)ds$ and let $y^i(\theta) = -\theta^i x^i(\theta) + \int_{\bar{R}}^{\theta} x^i(\theta_{-i}, s)ds$.

Theorem 13 *(Mookherje - Reichelstein) If* $u^i = w^i(x, \theta^i) + y^i$ *where* w^i *satisfies a single-crossing property,*[14] *then any Bayesian Incentive Compatible mechanism can be implemented as a Dominant Strategy mechanism if and only if* $\partial w^i(x(\theta_{-i}, s), \theta^i)/\partial \theta^i$ *is increasing in s.*

For the optimal auction, $w^i = \theta^i x^i$ and $\partial w^i(x, \theta^i)/\partial\theta^i = x^i$. So w^i trivially satisfies the single crossing property. Also $x^i(\theta_{-i}, s)$ is increasing in s since the item goes to the bidder with the highest θ^i.

Let $z^i(\theta_{-i}) = \max\{\bar{R}, \theta^j, j \neq i\}$. Then for the optimal auction $x^i(\theta) = 1$ if $\theta^i \geq z^i(\theta_{-i})$ and $x^i(\theta) = 0$ otherwise. So $-\theta^i x^i(\theta) + \int_{\bar{R}}^{\theta} x^i(\theta_{-i}, s)ds = -\theta^i + (\theta^i - z^i(\theta_{-i}) = -z^i(\theta_{-i}$ if i wins and is 0 otherwise. This means if i wins they pay the highest of $\bar{R}$ or the second highest bid. This is a dominant strategy auction. And the only use of the prior is by the auctioneer in setting the reserve price, $\bar{R}$.

5 Nash Mechanisms

We return to the structure we used at the beginning without beliefs. We are now interested in mechanisms that are not dominant strategy mechanisms. One reason for this might be computational compatibility. Using a direct

---

[14] w^i satisfies a single crossing property iff for all x, x' if there is a $\hat{\theta}^i \ni \partial w^i(x, \hat{\theta}^i)/\partial\theta^i \geq \partial w^i(x', \hat{\theta}^i)/\partial\theta^i$ then $\partial w^i(x, \theta^i)/\partial\theta^i \geq \partial w^i(x', \theta^i)/\partial\theta^i \forall \theta^i$.

mechanism or any dominant strategy mechanism requires players to submit messages that are at least as large as θ. We might want to work with mechanisms whose message spaces are smalller. So we need a behavioral model of what agents will do in this situation without dominant strategies. One option is Nash Behavior.[15]

Remember that given e and the mechanism $\gamma = (M, g)$, each individual is in a game. That game is described by:

- $I = \{i = 1 \ldots N\}$: individuals
- $S^i = M^i$: strategies for i
- $v^i(s, \theta^i) = \ u^i(g(m), \theta^i)$ i's payoff.

The behavioral model tells us how the individuals play this game. For Nash mechanisms we will assume players choose a Nash equilibrium.

Definition 11 *(Nash Equilibrium) s^* is a Nash equilibrium for θ if*

$$s^{*i} \in \ argmax_{m^i \in M^i} v^i(s^*/s^i, \theta^i), \forall s^i \in S^i, \forall i. \tag{28}$$

This probably makes the most sense as the stationary point in a repeated or dynamic situation in which agents grope their way along.

There is not a revelation principle for Nash mechanisms. This is partly because the behavior is not described by $b^i(\theta^i, \gamma)$ but instead by $b^i(\theta, \gamma) = s*^i(\theta)$ where $s^*(\theta)$ is an Nash equilibrium in θ. And it can be shown that if we were to ask everyone for their type and then played the appropriate Nash Equilibrium for those types, (that is, we used the mechanism $h(\theta) = g(s^*(\theta))$) then telling the truth would not always be a Nash equilibrium of this new game.[16] Think of a Stackleberg player vs Cournot players. In fact,

Theorem 14 *If truth is a Nash equilibrium of a direct revelation game $h(\theta), \forall\theta \in \Theta$ then h is a dominant strategy mechanism for Θ.*

It is possible to design mechanisms whose Nash Equilibria are efficient. One way to do this is to think of employing a VCG mechanism where the

---

[15] Nash was not the first to model this. Cournot among others was much earlier.

[16] Hurwicz called the equilibrium of the new game a manipulable Nash equilibrium.

message space is a subset of the true Θ. One thing that works is quadratic, quasi-linear functions with the same linear terms. Then the message space, M^i, has the same dimension as X.[17]

But there is another easier way. Suppose $u^i = u^i(x, y^i, \theta^i)$ and feasibility is defined by $F(x) = \sum y^i$. You can think of F(x) as the cost of x in terms of the transferable resource. Then an efficient allocation (x^*, y^*) solves, for some $\lambda^i > 0$,

$$\max \sum \lambda^i u^i(x, y^i) \tag{29}$$

$$\text{subject to } F(x) - \sum y^i = 0 \tag{30}$$

The first order conditions for this are:

$$\sum \lambda^i \partial u^i(x^*, y^{*i})/\partial x - \delta \partial F/\partial x = 0 \tag{31}$$

$$\lambda^i \partial u^i(x^*, y^{*i})/\partial y^i + \delta = 0 \tag{32}$$

Solving (32) for λ^i and putting it into (31) yields

$$\sum \frac{u^i_x}{u^i_{y^i}} = -F_x \tag{33}$$

$$\sum y^i = F(x) \tag{34}$$

(33) and (34) characterize efficient allocations. Note that the λ's have disappeared. That is because these equations describe a manifold in the (X,Y) space and any point on that manifold is an efficient allocation. There is not a unique efficient x when quasi-linearity is not present.

We now turn to constructing a Nash mechanism. The Nash equilibrium under a mechanism $g = (x(m), y(m))$ satisfies

$$(\partial u^i/\partial x)(\partial x/\partial m^i) + (\partial u^i/\partial y^i)(\partial y^i/\partial m^i) = 0 \tag{35}$$

Suppose we let $x(m) = \sum m^i$. Then $\partial x/\partial m^i = 1$ and we can rewrite this as

$$(\partial u^i/\partial x) + (\partial u^i/\partial y^i)(\partial y^i/\partial m^i) = 0 \tag{36}$$

$$\text{or } (\partial y^i/\partial m^i) = -\frac{\partial u^i/\partial x}{\partial u^i/\partial y^i} \tag{37}$$

---

[17] This is the approach taken by Groves-Ledyard. It has since been rediscovered in several engineering applications.

Comparing (36) and (33), we can see that to get efficient Nash equilibria, we need the mechanism to satisfy

$$\sum \partial y^i(m)/\partial m^i = F_x(\sum m^i) \tag{38}$$

$$\sum y^i(m) = F(\sum m^i) \tag{39}$$

Let's consider $y^i = -\frac{1}{N}F(\sum m^k) - t^i(m)$. To get efficiency we need $t^i(m)$ to satisfy

$$\sum y^i_{m^i} = 0 \tag{40}$$

$$\sum y^i = 0 \tag{41}$$

Suppose we try $t^i = \alpha^i + \beta^i m$ where $\beta^i = (\beta^i_1, ..., \beta^i_N)$. Then we need

$$\sum \beta^i_i = 0 \tag{42}$$

$$\sum \alpha^i + (\sum \beta^i)m = 0 \tag{43}$$

It is fairly easy to see that if we did this, then Nash equilibria would generally not exist. So let's try a quadratic $t^i = \alpha^i + \beta^i m + m'\Omega^i m$. Now it is required that

$$\sum \beta^i + 2\Omega^i m = 0 \tag{44}$$

$$\sum \alpha^i + \sum \beta^i m + \sum m'\Omega^i m = 0 \tag{45}$$

I won't go through the math here but it can be shown that the following leads to (among many possibilities):

$$t^i(m) = \xi\big[\frac{N-1}{N}(\mu^i - m^i)^2 - \sigma^i\big] \tag{46}$$

where $\xi > 0$ is an arbitrary constant, $\mu^i = \frac{1}{N-1}\sum_{j\neq i} m^j$, and $\sigma^i = \frac{1}{N-2}\sum_{j\neq i}(m^j - \mu^i)^2$.

Note that $\partial t^i/\partial m^i = 2\xi\frac{N-1}{N}(\mu^i - m^i)$ so $\sum \partial t^i/\partial m^i = 0$. It can also be shown that $\sum t^i = 0$ so the Nash equilibria allocations of this mechanism will be efficient.[18]

---

[18] The mechanism given by (34) and $x(m) = \sum m^i$ is a Groves-Ledyard mechanism.