Mechanism-design

Up to now we have been interested in guaranteeing that whatever private information is held by the agents, the corresponding desirable outcome - from the point of view of the "planner" - was realized. This approach makes two strong requirements:
(1) Every equilibrium outcome must be desirable
(2) The planner makes no trade-offs, his desirable outcome is either implemented or it is not.

We now relax both requirements. We endow the planner with a preference relation and search for the best outcome the planner can achieve by designing a mechanism in which agents coordinate on some equilibrium.

Optimal mechanisms (Monopolistic screening)

Consider a monopolist who offers access to a set of alternatives $[0, 1]$. Assume the monopolist's cost of providing a is given by ca. The monopolist charges an agent for each alternative that he chooses and his objective is to maximize profit.

An agent has quasi-linear preferences over pairs $(a, t) \in [0, 1] \times \mathbb{R}$, where t is a transfer they make to the monopolist. In particular, the utility of an agent from (a, t) is

$$\theta u(a) - t$$

where we normalize $u(0) = 0$. A pair $(a, t) \in [0, 1] \times \mathbb{R}$ will be referred to as a "contract". If the agent does not accept any contract from the monopolist, he chooses some outside option. We normalize the utility of the agent from this outside option to 0.

If the monopolist knew the value of θ, he would simply make a take-it-or-leave-it offer to the agent of choosing $a^* \in \arg\max_a [u(a) - ca]$ and charging $u(a^*)$ for it. We call this solution the *first-best* since the only constraint the monopolist faces is the "participation constraint" of the agent. It completely ignores the issue of incentive compatibility.

But what if the monopolist does not know the value of θ? Suppose he only knows that it is drawn from some c.d.f. F on $[0, 1]$ (an alternative interpretation is that there is a population of agents with different thetas). The monopolist can then offer agents a menu of contracts $((a_\theta, t_\theta))_{\theta \in [0,1]}$ that maximizes

$$\int_0^1 [t_\theta - ca_\theta)]dF(\theta)$$

subject to the constraints:

(IC_θ) An agent of type θ finds it optimal to choose the contract (a_θ, t_θ)

(IR_θ) If an agent of type θ chooses (a_θ, t_θ) his payoff is non-negative

A solution to this optimization problem is called the *second-best*. Since the planner now has a preference relation, we can talk about what is the best outcome that can be implemented. I.e., implementability is no longer a $0-1$ situation.

The *revelation-principle* allows us to find the menu $((a_\theta, t_\theta))_{\theta\in[0,1]}$ by restricting attention to a particular mechanism and solving for the best outcome that this mechanism can achieve. More specifically, instead of designing a menu of options we imagine a situation in which the monopolist asks each agent to report a value in $[0,1]$ interpreted as that agent's θ. For each report ϕ, the monopolist assigns each agent a pair $(a(\phi), t(\phi))$ such that:

($IC_{\theta,\phi}$) an agent of type θ has no incentive to report that his type is ϕ
(IR_θ) an honest agent is guaranteed a non-negative payoff,
(O) the monopolist cannot attain a higher expected profit by choosing a different schedule $(a(\cdot), t(\cdot))$ that satisfies the above constraints for all θ and ϕ.

After we solve the direct mechanism-design problem we are not done. For this solution to be economically meaningful we should find a plausible indirect mechanism that implements this solution. While there is a recipe for solving the direct mechanism-design proble, the task of finding an indirect mechanism has none.

We divide the problem of finding the optimal direct mechanism to two:
(1) the implementation problem - what is the set of IR and IC mechanisms?
(2) what is the best mechanism from the set identified in (1)?

The implementation problem

Incentive compatibility. Let $U(\phi, \theta)$ denote the payoff of an agent of type θ who pretends to be type ϕ. The IC and IR constraints are as follows: for all θ and ϕ,

$$U(\theta,\theta) \geq U(\phi,\theta) \quad \text{(IC)}$$
$$U(\theta,\theta) \geq 0 \quad \text{(IR)}$$

Although we are now focusing on the IC and IR constraint and not the issue of optimality, note that it cannot be the case that the IR constraint of all types has slack, since the monopolist could strictly increase his revenue without violating the IC constraint by raising the transfers of all types by the same small amount. Hence, there exists some type whose IR constraint binds. Because preferences are quasi-linear, the IC constraint implies that if $U(\theta,\theta) \geq 0$, then $U(\theta',\theta') > 0$ for all $\theta' > \theta$. Hence, if there exists some $\underline{\theta}$ with $U(\underline{\theta},\underline{\theta}) = 0$, then $U(\theta,\theta) > 0$ for all $\theta > \underline{\theta}$ and $U(\theta,\theta) = 0$ for all $\theta < \underline{\theta}$.

Let $q(\theta) \equiv u(a(\theta))$ and let $V(\theta) \equiv U(\theta, \theta)$. Then the collection of $IC_{\theta,\phi}$ constraints may be written as

$$V(\theta) = \max_{\phi \in [0.1]} \{\theta q(\phi) - t(\phi)\}$$

By IR_θ, for all $\phi \in [0, 1]$,

$$q(\phi) \geq 0$$

It follows that $V(\theta)$ is a maximum of a family of affine functions in θ (i.e., linear functions with positive slopes), and hence it is *convex*.

Note that the $IC_{\theta,\phi}$ constraints can also be written as follows: for all θ and ϕ,

$$\begin{aligned} V(\theta) &\geq \theta q(\phi) - t(\phi) \\ &= [\phi q(\phi) - t(\phi)] + [\theta q(\phi) - \phi q(\phi)] \\ &= V(\phi) + q(\phi)(\theta - \phi) \end{aligned}$$

This implies that for all ϕ, $q(\phi)$ is the slope of a line that supports the function $V(\cdot)$ at the point ϕ. A convex function is absolutely continuous and thus it is differentiable almost everywhere in the interior of its domain. This has two important implications:

(1) At every point θ that V is differentiable,

$$\frac{dV(\theta)}{d\theta} = q(\theta)$$

Since $V(\cdot)$ is convex, this implies that $q(\theta)$ *is non-decreasing.*

(2) Every absolutely continuous function is the definite integral of its derivative, so that

$$V(\theta) = V(0) + \int_0^\theta q(x)dx \tag{1}$$

The integral representation of the IC constraint helps us prove two instrumental results:

Lemma 1 *A mechanism is IC iff (1) and $q(\cdot)$ is non-decreasing*

Proof. We already proved the "only of" part of the lemma. To prove the "if" part, note that a non-decreasing $q(\cdot)$ implies that

$$\int_\phi^\theta q(x)dx \geq q(\phi)(\theta - \phi)$$

implying $IC_{\theta,\phi}$. ■

Lemma 2 *If a direct mechanism is IC, then for θ,*

$$t(\theta) = t(0) + \theta q(\theta) - \int_0^\theta q(x)dx$$

Thus, the expected payment in any two IC mechanisms that implement the same schedule $a(\theta)$ is equivalent up to a constant.

Proof.

$$\begin{aligned} V(\theta) &= \theta q(\theta) - t(\theta) \\ V(\theta) &= V(0) + \int_0^\theta q(x)dx \\ V(0) &= -t(0) \end{aligned}$$

■

Individual rationality. Since $IC_{0,\theta}$ implies that $V(0) = -t(0)$, IR_0 and $IC_{0,\theta}$ together imply that $t(0) \leq 0$.

The optimization problem

In light of Lemma 1 the monopolist's optimization problem may be written as:

$$\max_{(a(\theta),t(\theta))} \int_0^1 [t(\theta) - ca(\theta)]dF(\theta)$$

subject to:

(1) $t(\theta) = t(0) + \theta u(a(\theta)) - \int_0^\theta u(a(x))dx$ for all θ

(2) $u(a(\theta))$ is non decreasing for all θ

(3) $t(0) \leq 0$

The standard way to solve this is to substitute (1) into the objective function, solve the *unconstrained* optimization problem (which is often referred to as the *relaxed* problem), and then verify that (2) and (3) hold. The relaxed problem is then to choose a schedule $a(\theta)$ that maximizes

$$\begin{aligned} &= \int_0^1 [t(0) + \theta u(a(\theta)) - \int_0^\theta u(a(x))dx - ca(\theta)]dF(\theta)\} \\ &= \int_0^1 [\theta u(a(\theta)) - ca(\theta)]dF(\theta_i) - \int_0^1 \int_0^\theta [u(a(x))dx]dF(\theta) \end{aligned}$$

To simplify this problem the "usual trick" is to apply integration-by-parts to the double integral. Recall that

$$\int_0^1 gh' = [gh]_0^1 - \int_0^1 g'h$$

Here, let $h' = dF(\theta)$ and $g = \int_0^\theta u(a(x))dx$ so that

$$\begin{aligned}\int_0^1 \int_0^\theta [u(a(x))dx]dF(\theta) &= \left[\int_0^\theta [u(a(x))dx]F(\theta)\right]_0^1 - \int_0^1 u(a(\theta))F(\theta)d\theta \\ &= \int_0^1 u(a(\theta))[1-F(\theta)]d\theta\end{aligned}$$

Hence, the relaxed problem reduces to choose a menu $a(\theta)$ to maximize

$$\pi = \int_0^1 [(\theta - \frac{1-F(\theta)}{f(\theta)})u(a(\theta)) - ca(\theta)]f(\theta)d\theta$$

To simplify the solution it is customary to restrict attention to a cdf F with an increasing *hazard rate*. I.e., that

$$\frac{f(\theta)}{1-F(\theta)}$$

is increasing in θ, which then implies that

$$\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$$

increases with θ. This term is usually referred to as the "*virtual valuation*" of a buyer with value θ.

How can we interpret these virtual valuations? Suppose the monopolist makes a take-it-or-leave-it offer to the agent at a price of p. The probability that the agent will accept this offer is $1 - F(p)$. We can think of the probability of purchase as the "quantity" demanded by the agent. We may therefore write his "implied demand curve" as

$$q(p) = 1 - F(p)$$

The "inverse demand curve" is, therefore:

$$p(q) = F^{-1}(1-q)$$

where q is the "quantity purchased" or, equivalently, the probability of making a purchase. The resulting "revenue function" for the monopolist is, therefore

$$R(q) = q \cdot p(q) = q \cdot F^{-1}(1-q)$$

Hence,

$$\frac{d}{dq}[q \cdot p(q)] = F^{-1}(1-q) - \frac{q}{F'[F^{-1}(1-q)]}$$

But since $F^{-1}(1-q) = p$, we have that

$$MR(p) = p - \frac{1-F(p)}{f(p)} = \psi(p)$$

Hence, the "virtual valuation" of an agent may be interpreted as a "marginal revenue". The "monopoly price", p^M, is obtained by solving $MR(p^M) = 0$. To interpret p^M consider a risk-neutral agent with $u(a) = a$. In this case, p^M is the minimal agent type (the type with the smallest willingness-to-pay for $a = 1$) who agrees to purchase from the monopolist at the optimum.

Assuming an agent's virtual valuation increases with his type, the usual procedure for solving the optimization problem is as follows:

1. Let $\underline{\theta}$ satisfy $\psi(\underline{\theta}) = 0$. Since the hazard rate is increasing, $\psi(\theta) < 0$ for all $\theta < \underline{\theta}$. For any $\theta \leq \underline{\theta}$ set $a(\theta) = t(\theta) = 0$. I.e., exclude all types below $\underline{\theta}$.

2. For all $\theta > \underline{\theta}$ the maximization of π with respect to the schedule $a(\cdot)$ requires that the term under the integral be maximized with respect to $a(\theta)$ for all θ (note that since the integral IC representation is already included, there is no link between types, so we can maximize the integral point-by-point). Thus, for each $\theta > \underline{\theta}$,

$$\theta u'(a(\theta)) = c + \frac{1-F(\theta)}{f(\theta)} u'(a(\theta))$$

or

$$\psi(\theta)\, u'(a(\theta)) = c$$

Since first-best efficiency requires $\theta u'(a(\theta)) = c$, there is *under-consumption* for all types $\theta < 1$. This is the well-known, "no distortion at the top" result.

3. To compute $t(\theta)$ for all $\theta > \underline{\theta}$, we use the equality $t(\theta) = t(0) + \theta u(a(\theta)) - \int_{\underline{\theta}}^{\theta} u(a(x))dx$.

4. It remains to verify that $u(a(\theta))$ is non-decreasing in θ. From step 1 we know that $u(a(\theta)) = 0$ for all $\theta \leq \underline{\theta}$. Since $\psi(\theta)\, u'(a(\theta)) = c$ for all $\theta > \underline{\theta}$, it follows that for all these types, $u(a(\theta)) > 0$ and $u'(a(\theta)) > 0$ (where the latter inequality follows from our assumption on $\psi(\theta)$).

Example: Optimal auction with risk-neutral buyers and uniform dist.

- $B = \{(a_i, a_{-i}) \in [0,1]^n : \sum_i a_i = 1\}$
- $F(\theta) = \theta$
- $\psi(\theta) = 2\theta - 1 \Longrightarrow \underline{\theta} = \frac{1}{2} \Longrightarrow I = \{i \in N : \theta_i > \frac{1}{2}\}$
- *Relaxed problem*:

$$\max \sum_{i \in I} \theta_i a_i \text{ subject to } \sum_{i \in I} a_i = 1$$

- *Solution*:

$$a_i = \begin{cases} 1 & if \quad \theta_i > \{\frac{1}{2}, \max_{j \neq i} \theta_j\} \\ 0 & if \quad \quad other \end{cases}$$

$$V(\theta_i) = \begin{cases} 0 & if \quad \theta_i \leq \frac{1}{2} \\ \theta^{n-1} & if \quad \theta_i > \frac{1}{2} \end{cases}$$

$$\begin{aligned} T(\theta_i) &= \theta_i^n - \int_{\frac{1}{2}}^{\theta_i} x^{n-1} dx \\ &= \theta_i^n - [\frac{x^n}{n}]_{\frac{1}{2}}^{\theta_i} \\ &= \theta_i^n - \frac{\theta_i^n}{n} + \frac{1}{n}(\frac{1}{2})^n \\ &= \frac{n-1}{n}\theta_i^n + \frac{1}{n}(\frac{1}{2})^n \end{aligned}$$

Notice that $T(\theta_i)$ is exactly the expected payment of a bidder with value $\theta_i > \frac{1}{2}$ in a second-price auction with a reserve price of $\frac{1}{2}$:

$$T(\theta_i) = (\frac{1}{2})^n + \int_{\frac{1}{2}}^{\theta_i} y(n-1)y^{n-2} dy$$