ECON 289, Lecture 4: Bayesian Mechanism design ("The Revelation Principle holds!")

Instructor: Ben Brooks

Spring 2023

The mechanism design problem

- Thus far, our models have assumed as a primitive a particular game form, i.e.,
  - (i) sets of actions;
  - (ii) private information of agents;
  - (iii) preferences over outcomes (realized action profiles) that depend on private information
- Our question has been: what is "reasonable" (i.e., equilibrium) behavior in these environments

The mechanism design problem

- Thus far, our models have assumed as a primitive a particular game form, i.e.,
  (i) sets of actions;
  (ii) private information of agents;
  (iii) preferences over outcomes (realized action profiles) that depend on private information
- Our question has been: what is "reasonable" (i.e., equilibrium) behavior in these environments
- Mechanism design flips the question around
- We start an outcome we want to obtain, and try to engineer a game form such that this is the outcome
- Underlying premise is that the designer can control outcomes, but does not know agents' private information

The basic mechanism design problem

- Designer $i = 0$ can choose an outcome $x \in X$
- Payoff relevant state $\theta \in \Theta$
- Agents $i = 1, \ldots, N$ have private information about Θ represented by a Harsanyi type space $\mathcal{T} = (\{T_i\}, \pi)$
- Agent i's preferences are represented by $u_i : \Theta \times X \to \mathbb{R}$
- Designer can commit to a mechanism, which is just a game form consisting of
  - Actions A_i for each agent, $A = \prod_{i>0} A_i$
  - An outcome function $g : A \to \Delta(X)$
- Agents play an equilibrium of the game where strategies are $\sigma_i : T_i \to \Delta(A_i)$ and utility is

$$U_i(\sigma, t_i) = \sum_{\theta_i, t_{-i}, a, x} u_i(x, \theta) g(x|a) \sigma(a|t) \pi_i(\theta, t_{-i}|t_i)$$

Implementation theoretic approach

- Designer wants to implement a **social choice function** $f : T \to \Delta(X)$
- The game (A, g) and equilibrium σ implements f if for all t,

$$f(t) = \sum_a \sigma(a|t)g(a)$$

- Does there exist a game and equilibrium that implement f? (Partial/weak implementation)
- Does there exist a game for which **all** equilibria implement f? (Full/strong implementation)
- Notice that we don't take a stand on designer's preferences over (x, θ), but we are focusing on the implementation of a particular outcome
- Can also generalize to **social choice correspondences** $f : T \rightrightarrows \Delta(X)$

Bayesian mechanism design

- In Bayesian mechanism design, all of the agents have subjective expected utility preferences over outcomes
- Rather than specifying the social choice function, the designer too has preferences over outcomes, given by $u_0 : X \times \Theta \to \mathbb{R}$
- In addition, we typically assume that $\mathcal{T}$ satisfies the common prior assumption and that the designer is a "properly informed observer" in the sense of Harsanyi, i.e., the designer evaluates welfare according to the common prior $\pi \in \Delta(\Theta \times T)$, i.e.,

$$U_0(\sigma, t_i) = \sum_{\theta,t,a,x} u_0(x, \theta) g(x|a) \sigma(a|t) \pi(\theta, t)$$

- The objective is to identify mechanisms whose equilibrium outcomes are preferred by the designer
- NB we are adopting a "partial implementation" approach: The designer can pick both the mechanism and the equilibrium!

Participation constraints

- In many settings, it is natural to add further constraints on the mechanism designer, reflecting the option of agents to not participate in the mechanism
- In particular, we can model each agent as having an outside option $\underline{u}_i(\theta)$
- In that case, we may further want to restrict attention to mechanisms and equilibria that satisfy **interim participation constraints** (also called **interim individual rationality**)

$$U_i(\sigma, t_i) \geq \sum_{\theta, t_{-i}} \underline{u}_i(\theta)\pi_i(\theta, t_{-i}|t_i)$$

- This is the form of participation constraint that we will most often consider
- Depending on context, we may also want to consider other participation constraints, e.g., ex post: $u_i(x, \theta) \geq \underline{u}_i(\theta)$ for all outcomes that arise with positive probability

The revelation principle

- Finding a single equilibrium for a given type space and mechanism can be a monumental task
- How do we search through the vast space of all mechanisms and equilibria to find the best one for the designer?
- Fortunately, we can simplify the problem with using the celebrated **revelation principle**, which says that it is without loss to restrict attention to a relatively small class of mechanisms and equilibria
- In fact, the Bayesian implementation problem can be reduced to a linear program
- NB analogy with BCE/information design

Direct (revelation) mechanisms

- A **direct (revelation) mechanism** is one in which $A_i = T_i$, i.e., actions are identified with "reports" of a type
- The **truthful strategies** are such that $\sigma_i(t_i|t_i) = 1$, i.e., each agent reports their true signal to the mechanism
- The direct mechanism is **(interim) incentive compatible (IC)** if the truthful strategies are an equilibrium
- The direct mechanism is **(interim) individually rational (IR)** if the truthful strategies satisfy the (interim) participation constraints

Theorem

For any mechanism (A, g) and equilibrium σ, there is an IC direct mechanism for which the truthful equilibrium induces the same interim utilities for the agents and the same payoff for the designer. If (A, g) and σ satisfy the participation constraints, then the direct mechanism can also be taken to satisfy IR.

Proof of the revelation principle

- Just take $g'(t) = \sum_a g(a)\sigma(a|t)$, i.e., agents report their types and then the mechanism "plays" σ in (A, g) on the agents' behalf
- Clearly,

$$\sum_{\theta,t,a,x} u_0(x,\theta)g(x|a)\sigma(a|t)\pi(\theta,t) = \sum_{\theta,t,x} u_0(x,\theta)g'(x|t)\pi(\theta,t)$$

$$\sum_{\theta,t_{-i},a,x} u_i(x,\theta)g(x|a)\sigma(a|t)\pi_i(\theta,t_{-i}|t_i) = \sum_{\theta,t_{-i},x} u_0(x,\theta)g'(x|t)\pi_i(\theta,t_{-i}|t_i)$$

so IR will be satisfied if σ satisfies participation

Proof of the revelation principle

- Just take $g'(t) = \sum_a g(a)\sigma(a|t)$, i.e., agents report their types and then the mechanism "plays" σ in (A, g) on the agents' behalf
- Clearly,

$$\sum_{\theta,t,a,x} u_0(x,\theta)g(x|a)\sigma(a|t)\pi(\theta,t) = \sum_{\theta,t,x} u_0(x,\theta)g'(x|t)\pi(\theta,t)$$

$$\sum_{\theta,t_{-i},a,x} u_i(x,\theta)g(x|a)\sigma(a|t)\pi_i(\theta,t_{-i}|t_i) = \sum_{\theta,t_{-i},x} u_0(x,\theta)g'(x|t)\pi_i(\theta,t_{-i}|t_i)$$

  so IR will be satisfied if σ satisfies participation
- If IC is violated, then there is some i, t_i, t_i' such that

$$\sum_{\theta,t_{-i},x} u_i(x,\theta)(g'(x|t_i',t_{-i}) - g'(x|t_i,t_{-i}))\pi_i(\theta,t_{-i}|t_i) > 0$$

Proof of the revelation principle

- Just take $g'(t) = \sum_a g(a)\sigma(a|t)$, i.e., agents report their types and then the mechanism "plays" σ in (A, g) on the agents' behalf
- Clearly,

$$\sum_{\theta,t,a,x} u_0(x,\theta)g(x|a)\sigma(a|t)\pi(\theta,t) = \sum_{\theta,t,x} u_0(x,\theta)g'(x|t)\pi(\theta,t)$$

$$\sum_{\theta,t_{-i},a,x} u_i(x,\theta)g(x|a)\sigma(a|t)\pi_i(\theta,t_{-i}|t_i) = \sum_{\theta,t_{-i},x} u_0(x,\theta)g'(x|t)\pi_i(\theta,t_{-i}|t_i)$$

  so IR will be satisfied if σ satisfies participation
- If IC is violated, then there is some i, t_i, t_i' such that

$$\sum_{\theta,t_{-i},x} u_i(x,\theta)(g'(x|t_i',t_{-i}) - g'(x|t_i,t_{-i}))\pi_i(\theta,t_{-i}|t_i) > 0$$

- But then

$$\sum_{\theta,t_{-i},a,x} u_i(x,\theta)g(x|a)(\sigma(a|t_i',t_{-i}) - \sigma(a|t_i,t_{-i}))\pi_i(\theta,t_{-i}|t_i) > 0,$$

  i.e., agent i would strictly benefit when type t_i from playing $\sigma_i(t_i')$ rather $\sigma_i(t_i)$ in (A, g), which contradicts that σ is an equilibrium

Revelation principles

- There is a strong parallel between information design and mechanism design:
  - MD: Fix a type space (T, π). For any mechanism (A, g) and equilibrium σ, there is a direct "revelation" mechanism (T, g') and truthful equilibrium that induces the same outcome.
  - ID: Fix a game form (A, g). For any type space (T, π) and equilibrium σ, there is a "direct recommendation" type space (A, π') and obedient equilibrium that induce the same outcome.
- (For either problem, we can throw in participation constraints without conceptual difficulty, and we have done so for MD)
- So, we can either normalize actions to be types or types to be actions, depending on which one we treat as a primitive, and take strategies to be truthful/obedient
- Indeed, Myerson (1986) "Multistage games with communication" combines both RPs, in the context of a designer who can alternately elicit private information and privately recommend actions
- We will return to the parallels between information and mechanism design later in the course...

“The revelation principle holds!”

- People often say this casually, and then restrict attention to direct mechanisms/type spaces
- What are they actually assuming for the RP to hold? Two things:

"The revelation principle holds!"

- People often say this casually, and then restrict attention to direct mechanisms/type spaces
- What are they actually assuming for the RP to hold? Two things:
  1. The designer has available a rich space of mechanisms/type spaces that includes the direct ones

“The revelation principle holds!”

- People often say this casually, and then restrict attention to direct mechanisms/type spaces
- What are they actually assuming for the RP to hold? Two things:
    1. The designer has available a rich space of mechanisms/type spaces that includes the direct ones
    2. The designer can choose which equilibrium will be played

"The revelation principle holds!"

- People often say this casually, and then restrict attention to direct mechanisms/type spaces
- What are they actually assuming for the RP to hold? Two things:
  1. The designer has available a rich space of mechanisms/type spaces that includes the direct ones
  2. The designer can choose which equilibrium will be played
- Importantly, switching to the direct mechanism/information will generally change the **set** of equilibrium outcomes
- Simple example...

"The revelation principle holds!"

- People often say this casually, and then restrict attention to direct mechanisms/type spaces
- What are they actually assuming for the RP to hold? Two things:
    1. The designer has available a rich space of mechanisms/type spaces that includes the direct ones
    2. The designer can choose which equilibrium will be played
- Importantly, switching to the direct mechanism/information will generally change the **set** of equilibrium outcomes
- Simple example...
- So, the RP does not apply when the designer is restricted in their choices or we care about the set of equilibrium outcomes

Are we interested in direct mechanisms?

- All of this puts the type space and mechanism on equal footing
- But their conceptual basis is very different:
  - The mechanism often has a clear physical counterpart: When bidding in an auction, you actually do submit a bid, and there are clear rules about winning and payments
  - In some cases, like poker, there is a physical type, but more often the type is a highly abstract as-if description of how higher order uncertainty affects behavior, and to the extent that it exists, it is only in the mind of the agent
    Moreover, we often rely on untestable hypotheses about types, e.g., the CPA

Are we interested in direct mechanisms?

- All of this puts the type space and mechanism on equal footing
- But their conceptual basis is very different:
  - The mechanism often has a clear physical counterpart: When bidding in an auction, you actually do submit a bid, and there are clear rules about winning and payments
  - In some cases, like poker, there is a physical type, but more often the type is a highly abstract as-if description of how higher order uncertainty affects behavior, and to the extent that it exists, it is only in the mind of the agent

    Moreover, we often rely on untestable hypotheses about types, e.g., the CPA
- So, an inherent fragility of direct mechanisms is their dependence on a particular artificial language for describing players' knowledge

  (NB: This isn't resolved by eliciting belief hierarchies in the UTS!)

Are we interested in direct mechanisms?

- All of this puts the type space and mechanism on equal footing
- But their conceptual basis is very different:
  - The mechanism often has a clear physical counterpart: When bidding in an auction, you actually do submit a bid, and there are clear rules about winning and payments
  - In some cases, like poker, there is a physical type, but more often the type is a highly abstract as-if description of how higher order uncertainty affects behavior, and to the extent that it exists, it is only in the mind of the agent
    Moreover, we often rely on untestable hypotheses about types, e.g., the CPA
- So, an inherent fragility of direct mechanisms is their dependence on a particular artificial language for describing players' knowledge
  (NB: This isn't resolved by eliciting belief hierarchies in the UTS!)
- In many (most?) cases, direct mechanisms should be viewed as an analytical shortcut for identifying optimal **outcomes**
- After solving the MD problem, we should immediately pivot to the question of whether there are more natural "indirect" solutions

The mechanism design LP

- So, when maximizing the designer's payoff, we can WLOG restrict attention to IC/IR direct mechanisms
- But then the design problem (with participation) becomes

$$\max_{g(x|t)\geq 0} \sum_{\theta,t,x} u_0(x,\theta) g(x|t)\pi(\theta,t)$$

$$\text{s.t.} \quad \sum_{t_{-i},\theta,x} u_i(x,\theta)(g(x|t_i,t_{-i}) - g(x|t_i',t_{-i})) \underbrace{\pi(\theta,t_i,t_{-i})}_{\propto \pi_i(\theta,t_{-i}|t_i)} \geq 0 \ \forall i, t_i, t_i'$$

$$\sum_{t_{-i},\theta,x} (u_i(x,\theta) - \underline{u}_i(\theta)) g(x|t_i,t_{-i}) \underbrace{\pi(\theta,t_i,t_{-i})}_{\propto \pi_i(\theta,t_{-i}|t_i)} \geq 0 \ \forall i, t_i$$

$$\sum_x g(x|t) = 1 \ \forall t$$

- This is linear in g! So a finite dimensional LP problem

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

$$\min_{\alpha_i(t_i,t_i')\geq 0,\beta_i(t_i)\geq 0,\gamma(t)}$$

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

$$\min_{\alpha_i(t_i,t_i')\geq 0,\beta_i(t_i)\geq 0,\gamma(t)} \sum_t \gamma(t)$$

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

$$\min_{\alpha_i(t_i,t_i')\geq 0,\beta_i(t_i)\geq 0,\gamma(t)} \sum_t \gamma(t)$$

$$\forall t, x$$

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

$$\min_{\alpha_i(t_i,t_i')\geq 0, \beta_i(t_i)\geq 0, \gamma(t)} \sum_t \gamma(t)$$

$$\begin{aligned}
\text{s.t. } \gamma(t) \geq \sum_\theta \Bigg[& u_0(x,\theta)\pi(\theta,t) \\
& + \sum_{i,t_i'} u_i(x,\theta)\left[\pi(\theta,t)\alpha_i(t_i,t_i') - \pi(\theta,t_i',t_{-i})\alpha_i(t_i',t_i)\right] \\
& + \sum_i (u_i(x,\theta) - \underline{u}_i(\theta))\pi(\theta,t)\beta_i(t_i)\Bigg] \ \forall t, x
\end{aligned}$$

The dual problem

- As good analysts, and having just written down an LP, we will follow our instincts immediately write down its dual!

$$
\begin{aligned}
&\min_{\alpha_i(t_i,t_i')\geq 0,\beta_i(t_i)\geq 0,\gamma(t)} \sum_t \gamma(t) \\
\text{s.t. } \gamma(t) \geq \sum_\theta \Bigg[& u_0(x,\theta)\pi(\theta,t) \\
& + \sum_{i,t_i'} u_i(x,\theta)\left[\pi(\theta,t)\alpha_i(t_i,t_i') - \pi(\theta,t_i',t_{-i})\alpha_i(t_i',t_i)\right] \\
& + \sum_i (u_i(x,\theta) - \underline{u}_i(\theta))\pi(\theta,t)\beta_i(t_i)\Bigg] \ \forall t, x
\end{aligned}
$$

- As with BCE, we can interpret the α_i as a deviation in the normal form, and β_i is a deviation to the outside option

Optimal auctions

- Follows Myerson (1981), Bulow and Roberts (1989), Börgers (2015), Cai, Devanur, Weinberg (2019)
- N buyers, indexed by $i \in \{1, \ldots, N\}$
- Player $i = 0$ is the seller
- A single unit of a good for sale
- The buyers have **independent and private values** (IPV)
- $v_i \in V = \{0, \Delta, 2\Delta, \ldots, K\Delta\}$
- $f_i(v_i)$ is the PMF of v_i
  $f(v)$ is the joint PMF of $(v_1, \ldots, v_N)$

Optimal auctions

- Follows Myerson (1981), Bulow and Roberts (1989), Börgers (2015), Cai, Devanur, Weinberg (2019)
- N buyers, indexed by $i \in \{1, \ldots, N\}$
- Player $i = 0$ is the seller
- A single unit of a good for sale
- The buyers have **independent and private values** (IPV)
- $v_i \in V = \{0, \Delta, 2\Delta, \ldots, K\Delta\}$
- $f_i(v_i)$ is the PMF of v_i
  $f(v)$ is the joint PMF of $(v_1, \ldots, v_N)$
- The outcome consists of a **lottery** $q \in \Delta(\{0, 1, \ldots, N\})$ and a profile of **transfers** $t \in \mathbb{R}^n$

Optimal auctions

- Follows Myerson (1981), Bulow and Roberts (1989), Börgers (2015), Cai, Devanur, Weinberg (2019)
- N buyers, indexed by $i \in \{1, \dots, N\}$
- Player $i = 0$ is the seller
- A single unit of a good for sale
- The buyers have **independent and private values** (IPV)
- $v_i \in V = \{0, \Delta, 2\Delta, \dots, K\Delta\}$
- $f_i(v_i)$ is the PMF of v_i
  $f(v)$ is the joint PMF of $(v_1, \dots, v_N)$
- The outcome consists of a **lottery** $q \in \Delta(\{0, 1, \dots, N\})$ and a profile of **transfers** $t \in \mathbb{R}^n$
- Agent's have quasilinear preferences over probabilities of receiving the good and transfers (to the seller): for $i \geq 1$,

$$u_i(v_i, q, t) = v_i q_i - t_i$$

Optimal auctions

- Follows Myerson (1981), Bulow and Roberts (1989), Börgers (2015), Cai, Devanur, Weinberg (2019)
- N buyers, indexed by $i \in \{1, \dots, N\}$
- Player $i = 0$ is the seller
- A single unit of a good for sale
- The buyers have **independent and private values** (IPV)
- $v_i \in V = \{0, \Delta, 2\Delta, \dots, K\Delta\}$
- $f_i(v_i)$ is the PMF of v_i
  $f(v)$ is the joint PMF of $(v_1, \dots, v_N)$
- The outcome consists of a **lottery** $q \in \Delta(\{0, 1, \dots, N\})$ and a profile of **transfers** $t \in \mathbb{R}^n$
- Agent's have quasilinear preferences over probabilities of receiving the good and transfers (to the seller): for $i \geq 1$,

$$u_i(v_i, q, t) = v_i q_i - t_i$$

- $u_0(q, t) = \sum_i t_i$, i.e., seller wants to maximize revenue (and the good is worth nothing to the seller)

The optimal auctions LP

- NB we now have an unbounded variable (transfers)
- This wasn't allowed in our baseline model, so we should rewrite the LP:

$$\max_{q_i(v)\geq 0, t_i(v)} \sum_{v, i\geq 1} t_i(v) f(v)$$

$$\text{s.t.} \sum_{v_{-i}} \Big[(v_i q_i(v_i, v_{-i}) - t_i(v_i, v_{-i})) - (v_i q_i(v_i', v_{-i}) - t_i(v_i', v_{-i})) \Big] f(v_i, v_{-i}) \geq 0 \;\forall i \geq 1, v_i, v_i'$$

$$\sum_{v_{-i}} \left[v_i q_i(v_i, v_{-i}) - t_i(v_i, v_{-i}) \right] f(v_i, v_{-i}) \geq 0 \;\forall i \geq 1, v_i$$

$$\sum_i q_i(v) \leq 1$$

Lagrangian

- Rather than writing down the dual, let's consider the Lagrangian, with multipliers $\alpha_i(v_i, v_i') \geq 0$, $\beta_i(v_i) \geq 0$ on IC and IR
- If we add the constraints to the objective, what is the resulting term involving transfers? It will just be

$$\sum_{v,i\geq 1} t_i(v)\left[f(v) - \beta_i(v_i)f(v)\right.$$
$$\left. + \sum_{v_i'} (\alpha_i(v_i', v_i)f(v_i', v_{-i}) - \alpha_i(v_i, v_i')f(v))\right]$$

- Now, the transfers are a free variable, so unless the coefficient on $t_i(v)$ ends up being zero, the value of the Lagrangian will be infinite
- So, canceling $f_{-i}(v_{-i})$, (α_i, β_i) must satisfy **transfer neutrality**:

$$f_i(v_i) = \beta_i(v_i)f_i(v_i) + \sum_{v_i'} [\alpha_i(v_i, v_i')f_i(v_i) - \alpha_i(v_i', v_i)f_i(v_i')]$$

- (NB This is a general result: If a saddle point exists, then the optimal multipliers are such that the free variable drops out)

Virtual values

- Provided that we choose transfer-neutral (α, β), the resulting Lagrangian is

$$\max_q \sum_{i\geq 1} \sum_{v_{-i}} \left[v_i\beta_i(v_i)f_i(v_i) + \sum_{v_i'} \left(v_i\alpha_i(v_i, v_i')f_i(v_i) - v_i'\alpha_i(v_i', v_i)f_i(v_i') \right) \right] q_i(v)f_{-i}(v_{-i})$$

- But transfer neutrality can be rewritten

$$v_i f_i(v_i) + \sum_{v_i'} v_i\alpha_i(v_i', v_i)f_i(v_i') = v_i\beta_i(v_i)f_i(v_i) + v_i \sum_{v_i'} \alpha_i(v_i, v_i')f_i(v_i)$$

- Substituting this in, we obtain for the Lagrangian:

$$\max_q \sum_{i\geq 1} \sum_{v_{-i}} \underbrace{\left[v_i - \frac{1}{f_i(v_i)} \sum_{v_i'} (v_i' - v_i)\alpha_i(v_i', v_i)f_i(v_i') \right]}_{\equiv \phi_i^\alpha(v_i)} q_i(v)f(v)$$

- The object $\phi_i^\alpha(v_i)$ is referred to as the **virtual value**, and the Lagrangian is just choosing the allocation to maximize the virtual value of the winner

Maximizing the Lagrangian

- Now, it is easy to maximize this Lagrangian subject to the constraint $\sum_i q_i(v) \leq 1$:
- For each v, allocate to whichever bidder has the highest virtual value, as long as the highest virtual value is positive!
- The challenge is, we have to pick transfer-neutral α so that the resulting optimal allocation can actually be implemented
- Put slightly differently, we have to pick α so that there is an optimal (q, t) that satisfies complementary slackness, so that $\alpha_i(v_i, v_i')$ is positive only if v_i is indifferent reporting v_i', etc

Leveraging simulations

- In general, there are lots of dual solutions that make the free transfers drop out of the Lagrangian
- Which is the right one? Here is where simulations can give us some insight
- On your pset, you'll solve numerically for the optimal auction when there are two bidders and values are iid uniform on an evenly spaced grid $\{0, \Delta, 2\Delta, \ldots, 1\}$, where $1/\Delta$ is integral
- The simulations have several notable features:
  1. Optimal revenue is $\approx 5/12$
  2. The good is allocated to the bidder with the highest value, iff that value is at least $1/2$
  3. There is a lot of indeterminacy in the optimal transfers
  4. IR binds iff $v_i = 0$ and IC binds iff $v_i' = v_i - \Delta$ (local downward IC)

Local downward IC

- Let us use this last hint to simplify the Lagrangian
- Write $\tilde{\alpha}_i(v_i) \equiv \alpha_i(v_i, v_i - \Delta)$ for $v_i > 0$ and $\tilde{\alpha}_i(0) = 0$
- Then transfer neutrality reduces to

$$0 = f_i(v_i) - \tilde{\alpha}_i(v_i) f_i(v_i) + \tilde{\alpha}_i(v_i + \Delta) f_i(v_i + \Delta)$$
$$\iff \tilde{\alpha}_i(v_i) = \frac{1}{f_i(v_i)} \sum_{v_i' \geq v_i} f_i(v_i')$$

- This is an **inverse hazard rate**
- If we substitute these multipliers into the virtual value, we get

$$\phi_i^\alpha(v_i) = \tilde{\phi}_i(v_i) \equiv v_i - \Delta \frac{\sum_{v_i' > v_i} f_i(v_i')}{f_i(v_i)}$$

- NB: Even if these are not the optimal multipliers, they still give us an upper bound on the optimal value!
- NB: ϕ_i^α for this particular choice of α is what's been traditionally called the "virtual value", following Myerson (1981)

Back to the uniform example

- Passing to the continuum limit with a CDF G_i and density g_i, we have that

$$v_i - \frac{\sum_{v_i' > v_i} f_i(v_i')}{f_i(v_i)/\Delta} \approx v_i - \frac{1 - G_i(v_i)}{g_i(v_i)}$$

Back to the uniform example

- Passing to the continuum limit with a CDF G_i and density g_i, we have that

$$v_i - \frac{\sum_{v_i' > v_i} f_i(v_i')}{f_i(v_i)/\Delta} \approx v_i - \frac{1 - G_i(v_i)}{g_i(v_i)}$$

- In the uniform case, this reduces to

$$v_i - \frac{1 - v_i}{1} = 2v_i - 1,$$

- So, the virtual value is the highest for bidders with the highest value, and the Lagrangian is maximized by allocating to the high value bidder as long as $2v_i - 1 \geq 0 \iff v_i \geq 1/2$, exactly as in the simulation!

Back to the uniform example

- Passing to the continuum limit with a CDF G_i and density g_i, we have that

$$v_i - \frac{\sum_{v_i' > v_i} f_i(v_i')}{f_i(v_i)/\Delta} \approx v_i - \frac{1 - G_i(v_i)}{g_i(v_i)}$$

- In the uniform case, this reduces to

$$v_i - \frac{1 - v_i}{1} = 2v_i - 1,$$

- So, the virtual value is the highest for bidders with the highest value, and the Lagrangian is maximized by allocating to the high value bidder as long as $2v_i - 1 \geq 0 \iff v_i \geq 1/2$, exactly as in the simulation!
- Moreover, the optimal value with two bidders will be simply

$$\int_{v \in [0,1]^2} \max\{0, 2v_1 - 1, 2v_2 - 1\} dv$$
$$= \int_{x=1/2}^{1} (2x - 1) d(x^2) = 5/12$$

But can the bound be attained?

But can the bound be attained?

- Yes! Just run a second price auction with reserve price of $1/2$

But can the bound be attained?

- Yes! Just run a second price auction with reserve price of $1/2$
- In fact, this is the solution in the symmetric case as long as the virtual value is non-decreasing in v_i, as it is in the uniform example
  This is referred to as the **regular case**
- But if the virtual value is not non-decreasing, these may not be the optimal multipliers, we may have to be more clever

But can the bound be attained?

- Yes! Just run a second price auction with reserve price of $1/2$
- In fact, this is the solution in the symmetric case as long as the virtual value is non-decreasing in v_i, as it is in the uniform example
  This is referred to as the **regular case**
- But if the virtual value is not non-decreasing, these may not be the optimal multipliers, we may have to be more clever
- Indeed, on your pset, you will also solve an example with where the optimal multipliers are more complicated, and we can have both up and down constraints binding
- That's what we turn to next...

Single agent problem and monotonicity

- Let's make things simpler for a bit and look at the single bidder problem
- For now we drop the i subscript, and write v, v', $\alpha(v, v')$, etc.
- In this context, there is a well-known characterization of optimal mechanisms, that goes like this:
- For two values $v > v'$, IC implies that

$$\begin{aligned} vq(v) - t(v) &\geq vq(v') - t(v') \\ v'q(v') - t(v') &\geq v'q(v) - t(v) \\ \implies (v - v')(q(v) - q(v')) &\geq 0 \\ \implies q(v) &\geq q(v') \end{aligned}$$

- So, any incentive compatible allocation must be **monotonic** (i.e., non-decreasing)

Monotonicity and posted prices

- In fact, monotonicity is not only necessary for an allocation to be implementable, but it is sufficient as well:
- A posted price mechanism is to simply offer the good at a price p, so that type v buys iff $v \geq p$
  - This induces an allocation $q_p(v) = \mathbb{I}_{v \geq p}$ and transfer $t_p(v) = q_p(v)p$, which is obviously IC
  - NB $q_p(v)$ is a **step function** that is zero below p and one above p
- Next, offer the buyer a **random posted price**, where p is drawn independently of the value with probability $h(p)$
  - This induces an allocation

$$q_h(v) = \sum_p h(p)\mathbb{I}_{p \leq v} = \sum_{p \leq v} h(p)$$

    and the associated transfer is $t_h(v) = \sum_{p \leq v} ph(p)$
  - But picking $h(p) = q(v) - q(v - \Delta)$ (non-negative as long as q is monotonic), we get precisely the allocation $q(v)$!

Revenue from posted prices

- But more than that, any non-decreasing allocation can be implemented in a way that attains our upper bound on revenue
- Let $R(p)$ be the revenue from posted price p
- Then for $v \in V$

$$\begin{aligned} R(v) - R(v+\Delta) &= v \sum_{v' \geq v} f(v') - (v+\Delta) \sum_{v' \geq v+\Delta} f(v') \\ &= vf(v) - \Delta \sum_{v' \geq v+\Delta} f(v') \\ &= \tilde{\phi}(v) f(v) \end{aligned}$$

- So, $\tilde{\phi}(v)$ is the marginal revenue from selling to buyers of value v, per "probability unit" of buyer
- This implies that

$$R(v) = \sum_{v' \geq v} \tilde{\phi}(v') f(v'),$$

which is exactly our upper bound from the local relaxation!

Revenue from random posted prices

- And by exactly the same logic, revenue from the lottery over posted prices h that implements the allocation q will be

$$R(q) = \sum_{p} h(p) \sum_{v' \geq p} \tilde{\phi}(v')$$
$$= \sum_{v'} \tilde{\phi}(v') f(v') \sum_{p \leq v'} h(p)$$
$$= \sum_{v'} \tilde{\phi}(v') q(v') f(v')$$

- The bottom line is that any monotonic allocation can be implemented in a way that attains the upper bound from local multipliers

Optimal allocations

- So, one way to proceed would be to just optimize the bound $R(q)$ over monotonic allocations
- In fact, since the objective $R(q)$ is linear in q, we know it is WLOG to look at extreme points of the set of non-decreasing allocation rules, which are...

Optimal allocations

- So, one way to proceed would be to just optimize the bound $R(q)$ over monotonic allocations
- In fact, since the objective $R(q)$ is linear in q, we know it is WLOG to look at extreme points of the set of non-decreasing allocation rules, which are... the step functions $q_p(v)$!

Optimal allocations

- So, one way to proceed would be to just optimize the bound $R(q)$ over monotonic allocations
- In fact, since the objective $R(q)$ is linear in q, we know it is WLOG to look at extreme points of the set of non-decreasing allocation rules, which are... the step functions $q_p(v)$!
- Indeed, we already proved this, when we showed that any non-decreasing q can be implemented by an allocation where $q(v)$ is the probability that $p \leq v$
- So, extreme allocations are all induced by posted prices, and there is an optimal mechanism that is a posted price

Optimal multipliers

- You can approach the many-bidder problem using monotonicity, but we will take a slightly different tack, deriving the optimal multipliers
- Recall that the dual LP is

$$\min_{\alpha} \sum_{v} \max\{0, \phi^{\alpha}(v)\} f(v)$$

$$\text{s.t. } f(v) = \beta(v) f(v) + \sum_{v'} \left[\alpha(v, v') f(v) - \alpha(v', v) f(v')\right]$$

$$\phi^{\alpha}(v) = v - \frac{1}{f(v)} \sum_{v'} (v' - v) \alpha(v', v) f(v')$$

Local virtual values and monotonicity

- In general, the local downward solution to transfer neutrality need not induce a monotonic allocation
- For example, suppose the values are $V = \{1, 2, 3\}$, with likelihoods $3/4$, ϵ, and $1/4 - \epsilon$
- The local virtual values are

$$\tilde{\phi}(1) = 1 - \frac{1/4}{3/4} = 2/3$$

$$\tilde{\phi}(2) = 2 - \frac{1/4 - \epsilon}{\epsilon}$$

$$\tilde{\phi}(3) = 3$$

- Clearly, when ϵ is small, we have $\tilde{\phi}(1)$ and $\tilde{\phi}(3)$ both positive and $\tilde{\phi}(2)$ negative, so the Lagrangian will be maximized by selling just to low and high, which is not IC

Modifying the local solution

- In a sense, the problem with the local virtual value is that it understates the cost of not selling to some types
- In particular, not selling to a type v means that you can't sell to types higher than v
- So what we have to do is price in the forgone revenue from selling to those higher types
- What is that forgone revenue?

Revenue as a function of price

Switch to quantity units

- It will help to change the domain from prices to quantities
- In particular, we could think of a posted price p as being the mechanism that sells to the $q = \sum_{v \geq p} f(v)$ buyers with the highest values
- Moreover, random posted prices can be viewed as simply random quantities
- With this change of units, the picture becomes...

Revenue as a function of quantity

Randomized posted prices

- As we know, **any** mechanism can be implemented with a randomized posted price
- So, suppose you wanted to sell a fixed quantity q in expectation?
- What is the best randomized posted price to sell that quantity, and what is the associated expected revenue?
- The answer to this question will then tell us what is the "right" marginal revenue from selling to additional buyers

Revenue from randomized posted prices

Towards the right virtual values

- We can see that the maximum revenue from a randomized posted price is just the concave hull of the revenue function
- That is, the smallest concave function that is everywhere above the revenue
- This is attained by randomizing over two posted prices
- The "correct" virtual value should be the probability weighted average of the local virtual values between the two prices that we are randomizing between
- This is a heuristic, but it will be vindicated by our derivation of the optimal multipliers for the mechanism design LP
- But first, we have to more precisely define the virtual value that we are shooting for

The "ironed" solution

- We will define a new **ironed virtual value** $\overline{\phi}$ that correctly prices the marginal revenue of selling to a given type
- For each v, v', $v' < v$, define the "average virtual revenue" between v and v' as

$$A(v', v) = \frac{R(v' - \Delta) - R(v)}{\sum_{v' \leq v'' \leq v} f(v'')} = \frac{\sum_{v' \leq v'' \leq v} \tilde{\phi}(v'') f(v'')}{\sum_{v' \leq v'' \leq v} f(v'')}$$

- Define $\overline{\phi}$ recursively as follows
  1. Start with the highest **unironed type** being $v = \max\{V\}$
  2. If the highest unironed type is v, let $v' \in \arg\max\{A(v', v)\}$, then we define $\bar{\phi}(v'') = A(v', v)$ for $v' \leq v'' \leq v$, and we redefine the highest unironed type to be v' and go back to 1
- In each iteration of step 2, we refer to $[v', v]$ as an **ironed interval** (may be degenerate if $v' = v$)

Properties of $\overline{\phi}$

- NB that $\overline{\phi}$ must be non-decreasing:
  - If not, say $\bar{\phi}(v) < \bar{\phi}(v + \Delta)$
  - Let v' and v'' be the types such that $A(v', v) = \bar{\phi}(v)$ and $A(v + \Delta, v'') = \bar{\phi}(v + \Delta)$
  - Then we could have made average marginal revenue higher by "merging" the two intervals, to make a big ironed interval $[v', v'']$
- As a result, if $\bar{\phi}$ can be implemented, then the Lagrangian will be maximized by simply setting a posted price at $\min\{v|\bar{\phi}(v) \geq 0\}$, which can obviously implemented

Shifting virtual value up

- So, the big question is whether there exist multipliers that induce $\bar{\phi}$...

Shifting virtual value up

- So, the big question is whether there exist multipliers that induce $\bar{\phi}$...Yes, there are!
- Recall the transfer-neutrality constraint:

$$f(v) = \beta(v)f(v) + \sum_{v'} \left[\alpha(v,v')f(v) - \alpha(v',v)f(v')\right]$$

- How can we modify α to preserve transfer neutrality? For $v < v'$, if we increase $\alpha(v,v')$ $\epsilon/f(v)$, then we have to increase $\alpha(v',v)$ by $\epsilon/f(v')$

Shifting virtual value up

- So, the big question is whether there exist multipliers that induce $\bar{\phi}$...Yes, there are!
- Recall the transfer-neutrality constraint:

$$f(v) = \beta(v)f(v) + \sum_{v'} \left[\alpha(v,v')f(v) - \alpha(v',v)f(v')\right]$$

- How can we modify α to preserve transfer neutrality? For $v < v'$, if we increase $\alpha(v,v')$ $\epsilon/f(v)$, then we have to increase $\alpha(v',v)$ by $\epsilon/f(v')$
- Note the change in the virtual value:

$$\Delta\phi^{\alpha}(v) = -(v'-v)\frac{\epsilon}{f(v)}$$

$$\Delta\phi^{\alpha}(v') = -(v-v')\frac{\epsilon}{f(v')} = (v'-v)\frac{\epsilon}{f(v')}$$

so we shift virtual value from v up to v'

Law of conservation of virtual value

- Note that the change in the expected virtual value is

$$\Delta\phi^{\alpha}(v)f(v) + \Delta\phi^{\alpha}(v')f(v')$$
$$= -(v' - v)\frac{\epsilon}{f(v)}f(v) + (v' - v)\frac{\epsilon}{f(v')}f(v') = 0!$$

- In other words, **expected virtual value is conserved**:
- We can shift virtual value up but we cannot change the total amount of virtual value
- Thus, what we are trying to do is redistribute the virtual value so as to smooth out on ironed intervals

Construct the multipliers to implement $\overline{\phi}$

- Suppose an ironed interval consists of $v_1 < \ldots < v_K$
- Note that for any k, we must have $A(v_1, v_k) \geq A(v_1, v_K)$ (otherwise, we would have $A(v_{k+1}, v_K) > A(v_1, v_K)$, and we contradict v_1 maximizing $A(v_1, v_K)$)
- Take as an inductive hypothesis that we have modified the values so that $A(v_1, v_K) = \phi^{\alpha}(v_1) = \cdots = \phi^{\alpha}(v_{k-1}) \leq A(v_1, v_k)$
- Base step: $k = 0$ with the local multipliers and $\phi^{\alpha} \equiv \tilde{\phi}$
- Inductive step:
  - It must be that $\phi^{\alpha}(v_k) \geq A(v_1, v_K)$ (if not then $A(v_1, v_k) \leq A(v_1, v_K)$)
  - Then we can increase $\alpha(v_k, v_{k+1})$ and $\alpha(v_{k+1}, v_k)$ to shift virtual value up from v_k to v_{k+1} so that $\phi^{\alpha'}(v_k) = A(v_1, v_K)$
  - But virtual value is conserved, so $A(v_1, v_{k+1})$ is unchanged, and we still have $A(v_1, v_{k+1}) \geq A(v_1, v_K)$, and the inductive hypothesis is now satisfied for $k + 1$
- After K steps, we will have shifted virtual value up so as to attain $\phi^{\alpha} = \bar{\phi}$ on $[v, v']$
- Denote the associated multipliers by $(\overline{\alpha}, \overline{\beta})$

The “level free” solution

- This completes the characterization of a solution to the single-bidder problem
- We have binding local downward constraints, and within ironed interval, local upward constraints bind as well
- NB The solution is not unique! Why?

The "level free" solution

- This completes the characterization of a solution to the single-bidder problem
- We have binding local downward constraints, and within ironed interval, local upward constraints bind as well
- NB The solution is not unique! Why?
  1. There's more than one choice of multipliers that induce $\bar{\phi}$
  2. Moreover, we don't have to induce $\bar{\phi}$... Any ϕ^{α} that is single-crossing at the optimal posted price would do

The "level free" solution

- This completes the characterization of a solution to the single-bidder problem
- We have binding local downward constraints, and within ironed interval, local upward constraints bind as well
- NB The solution is not unique! Why?
  1. There's more than one choice of multipliers that induce $\bar{\phi}$
  2. Moreover, we don't have to induce $\bar{\phi}$... Any ϕ^{α} that is single-crossing at the optimal posted price would do
- But what is nice about the ironed solution is that it is "level free," in the following sense:
- Suppose the seller has a cost c, positive or negative
- This just shifts values down by c, but doesn't change differences in values, so the induced virtual value is just $\bar{\phi}(v) - c$, which is still non-decreasing

The "level free" solution

- This completes the characterization of a solution to the single-bidder problem
- We have binding local downward constraints, and within ironed interval, local upward constraints bind as well
- NB The solution is not unique! Why?
  1. There's more than one choice of multipliers that induce $\bar{\phi}$
  2. Moreover, we don't have to induce $\bar{\phi}$... Any ϕ^α that is single-crossing at the optimal posted price would do
- But what is nice about the ironed solution is that it is "level free," in the following sense:
- Suppose the seller has a cost c, positive or negative
- This just shifts values down by c, but doesn't change differences in values, so the induced virtual value is just $\bar{\phi}(v) - c$, which is still non-decreasing
- Clearly, the maximizer of the Lagrangian will be to set a posted price at $\min\{v | \bar{\phi}(v) - c \geq 0\}$, which is implementable

Random posted prices and complementary slackness

- Does this solution satisfy complementary slackness?

Random posted prices and complementary slackness

- Does this solution satisfy complementary slackness? Yes!
- NB random posted prices make all local downward IC constraints hold as equality:
  1. If your value is strictly above or strictly below the price, then changing your report doesn't change whether or not you get the good or what you pay
  2. If your value is equal to the price, then you get zero surplus from reporting truthfully, and also zero surplus from a local downward deviation!

Random posted prices and complementary slackness

- Does this solution satisfy complementary slackness? Yes!
- NB random posted prices make all local downward IC constraints hold as equality:
  1. If your value is strictly above or strictly below the price, then changing your report doesn't change whether or not you get the good or what you pay
  2. If your value is equal to the price, then you get zero surplus from reporting truthfully, and also zero surplus from a local downward deviation!
- The posted price $p = \min\{v | \bar{\phi}(v) - c \geq 0\}$ also makes local upward constraints bind on ironed intervals
  In particular, p must be at the bottom of an ironed interval, so local upward deviations within ironed intervals don't change the allocation or the payment
- So, this solution does in fact satisfy complementary slackness with the "ironed" IC multipliers

Back to auctions

- Let's apply what we've learned to solve the N bidder problem
- Note that because values are private, we can rewrite IC more compactly as follows: Define

$$Q_i(v_i) = \sum_{v_{-i}} q_i(v_i, v_{-i}) f_{-i}(v_{-i})$$

$$T_i(v_i) = \sum_{v_{-i}} t_i(v_i, v_{-i}) f_{-i}(v_{-i})$$

- These are the interim expected allocation and transfer, respectively
- Then type v_i's payoff from reporting v_i' is just

$$v_i Q_i(v_i') - T_i(v_i')$$

and IC just says that this expression is maximized at $v_i' = v_i$

Monotonicity and random posted prices

- By exactly the same argument as in the single bidder case, monotonicity of $Q_i(v_i)$ is equivalent to its implementability by incentive compatible transfers
- Moreover, if Q_i is non-decreasing, then q be implemented with transfers that attain the upper bound from local downward virtual values
- How?

Monotonicity and random posted prices

- By exactly the same argument as in the single bidder case, monotonicity of $Q_i(v_i)$ is equivalent to its implementability by incentive compatible transfers
- Moreover, if Q_i is non-decreasing, then q be implemented with transfers that attain the upper bound from local downward virtual values
- How? Well, draw a price p_i independent of v_{-i}, distributed according to the CDF Q_i
- Given a report of v_i, bidder i wins with probability $q_i(v_i, v_{-i})$, and pays p_i if $v_i \geq p_i$... This implements exactly the same interim allocation and payment as the random posted price!

The complete solution

- Let's use complementary slackness to verify a saddle point
- Let $\overline{\phi}_i$ be bidder i's ironed virtual value
- Define

$$W(v) = \arg\max_i \overline{\phi}_i(v_i)$$

$$q_i(v) = \begin{cases} \frac{1}{|W(v)|} & \text{if } i \in W(v) \\ 0 & \text{otherwise} \end{cases}$$

Theorem

The aforementioned allocation is implementable with transfers that achieve the local upper bound on revenue. Therefore, this is an optimal allocation.

Proof of the optimal auction

- Clearly, for the choice of multipliers $(\overline{\alpha}, \overline{\beta})$, we will induce the ironed virtual values
- Moreover, the allocation we have proposed maximizes the resulting Lagrangian:

$$\sum_{v, i \geq 1} \overline{\phi}_i(v) q_i(v) f(v)$$

- So, it just remains to check that this allocation is implementable, which it is because $\overline{\phi}_i(v_i)$ is non-decreasing, so that $q_i(v_i, v_{-i})$ is non-decreasing in v_i, and therefore $Q_i(v_i)$ is non-decreasing as well
- Moreover, we have already verified that Q_i can be implemented with posted prices that attain the local upper bound $\square$

Finite vs convex domain

- Throughout our analysis, we have focused on the model with discrete values, so that we could apply techniques from linear programming
- In the classic treatment of Myerson (1981), the domain of values is an interval $[\underline{v}, \overline{v}]$, and the probability distribution has a density
- The solution more or less corresponds to what I have shown you, but there is an important difference
- In the discrete model, there may be more than one transfer rule that implements a given allocation
- For example, any price in $[v - \Delta, v]$ would implement the allocation q_v (but the price of $p = v$ is the one that maximizes revenue and makes local downward constraints bind, and therefore attains the local IC upper bound on revenue)
- But when we go to a convex domain, transfers are uniquely pinned down, as I now explain

The envelope formula

- Note that type v_i's surplus in equilibrium is

$$U_i(v_i) = v_i Q_i(v_i) - T_i(v_i) = \max_w v_i Q_i(w) - T_i(w)$$

- Can use monotonicity of Q_i to show that U_i is continuous and a.e. differentiable, and the envelope formula holds, i.e.,

$$\frac{d}{dv_i} U_i(v_i) = Q_i(v_i)$$

The envelope formula

- Note that type v_i's surplus in equilibrium is

$$U_i(v_i) = v_i Q_i(v_i) - T_i(v_i) = \max_w v_i Q_i(w) - T_i(w)$$

- Can use monotonicity of Q_i to show that U_i is continuous and a.e. differentiable, and the envelope formula holds, i.e.,

$$\frac{d}{dv_i} U_i(v_i) = Q_i(v_i)$$

- Thus,

$$U_i(v_i) = U_i(0) + \int_{x=0}^{v_i} Q_i(x) dx$$

Solving out transfers

- The envelope formula pins down transfers:

$$T_i(v_i) = v_i Q_i(v_i) - U_i(v_i) = v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx - U_i(0)$$

- Since U_i is increasing, IR is equivalent to $U_i(0) \geq 0$, and obviously R is maximized by setting $U_i(0) = 0$
- The seller's revenue is therefore

$$R = \sum_{i=1}^{n} \int_{v_i \in [0,\overline{v}]} \left(v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx \right) f_i(v_i)dv_i$$

Solving out transfers

- The envelope formula pins down transfers:

$$T_i(v_i) = v_i Q_i(v_i) - U_i(v_i) = v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx - U_i(0)$$

- Since U_i is increasing, IR is equivalent to $U_i(0) \geq 0$, and obviously R is maximized by setting $U_i(0) = 0$
- The seller's revenue is therefore

$$R = \sum_{i=1}^{n} \int_{v_i \in [0,\overline{v}]} \left(v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx \right) f_i(v_i)dv_i$$

- Using Fubini, this can be rewritten as

$$R = \sum_{i=1}^{n} \int_{v_i=0}^{\overline{v}} \left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right) Q_i(v_i)f_i(v_i)dv_i$$

Solving out transfers

- The envelope formula pins down transfers:

$$T_i(v_i) = v_i Q_i(v_i) - U_i(v_i) = v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx - U_i(0)$$

- Since U_i is increasing, IR is equivalent to $U_i(0) \geq 0$, and obviously R is maximized by setting $U_i(0) = 0$
- The seller's revenue is therefore

$$R = \sum_{i=1}^{n} \int_{v_i \in [0,\overline{v}]} \left(v_i Q_i(v_i) - \int_{x=0}^{v_i} Q_i(x)dx \right) f_i(v_i)dv_i$$

- Using Fubini, this can be rewritten as

$$\begin{aligned} R &= \sum_{i=1}^{n} \int_{v_i=0}^{\overline{v}} \left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right) Q_i(v_i) f_i(v_i) dv_i \\ &= \sum_{i=1}^{n} \int_{v \in [0,\overline{v}]^n} \underbrace{\left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right)}_{=\tilde{\phi}_i(v_i)} q_i(v) f(v) dv \end{aligned}$$

Further observations about convex domain

- The uniqueness of the transfers is the main advantage of working in the continuum model
- Indeed, the multiplicity of tranfer rules in the discrete example from two slides ago is because we didn't specify the allocation for types in $[v-\delta, v]$
- In fact, convex domain is in a sense without loss; even if the value v has zero probability, we could always specify an allocation and transfer for type v that satisfies IC: Just set $(q(v), t(v)) \in \arg\max_{v' \in V} vq(v') - t(v')$, i.e., what v would choose if they could mimic any type in V!

Revenue equivalence

Theorem

With a convex domain, revenue generated by an auction is completely determined by the induced allocation and the utilities of the players' lowest types.

- Thus, if two mechanisms induce the same allocation and give the lowest types the same payoff, then they generate the same revenue
- This resolved a long-standing mystery in auction theory going back to the 60's, that first- and second-price auctions generate the same revenue in symmetric IPV environments

Optimal auctions versus competition

- We might ask, how much does the seller benefit from running the optimal auction vs something simpler, e.g., a second-price auction
- A famous paper of Bulow and Klemperer (1996) shows that in the regular and symmetric case, the revenue from running the exact optimal auction is no greater than the revenue from a second-price auction with one additional buyer, assuming that buyer's value is drawn from the same distribution
- Thus, there is a bound on how much revenue can be gained from fine-tuning the reserve prices

Proof

- Note that the second-price auction maximizes revenue subject to the constraint that the good is always allocated, since it maximizes the virtual value of the buyer who receives the good (although this virtual value may be negative)

Proof

- Note that the second-price auction maximizes revenue subject to the constraint that the good is always allocated, since it maximizes the virtual value of the buyer who receives the good (although this virtual value may be negative)
- Consider the following mechanism when there are $n+1$ buyers: pick a subset of n and run the Myerson optimal auction, and if the good would be unallocated, give it to the $n+1$th buyer for free
- This mechanism obviously generates the same revenue as the optimal auction with n buyers
- It also always allocates the good, so it generates less revenue than the SPA with $n+1$ buyers $\square$

Proof

- Note that the second-price auction maximizes revenue subject to the constraint that the good is always allocated, since it maximizes the virtual value of the buyer who receives the good (although this virtual value may be negative)
- Consider the following mechanism when there are $n+1$ buyers: pick a subset of n and run the Myerson optimal auction, and if the good would be unallocated, give it to the $n+1$th buyer for free
- This mechanism obviously generates the same revenue as the optimal auction with n buyers
- It also always allocates the good, so it generates less revenue than the SPA with $n+1$ buyers $\square$
- (This result fascinating, though it is a rather artificial comparison... when is this the tradeoff made by a designer? And why should the marginal buyer have a value drawn from the same distribution as other participants?)

The importance of independence

- Myerson's result firmly relies on independent values... where?

The importance of independence

- Myerson's result firmly relies on independent values... where?
- The incentive constraint! It was essential that the interim probability of receiving the good $Q_i(w)$ and transfer $T_i(w)$ only depend on buyer i's report w, not on his type

The importance of independence

- Myerson's result firmly relies on independent values... where?
- The incentive constraint! It was essential that the interim probability of receiving the good $Q_i(w)$ and transfer $T_i(w)$ only depend on buyer i's report w, not on his type
- If types were correlated, then different buyers might have different beliefs about v_{-i}, and thereby face different interim expected allocations and transfers
- Such correlation might allow the seller to extract more revenue than with Myersonian auctions...

A model with correlated types

- Follows Crémer and McLean (1988)
- Each bidder has finite set of types S_i, $S = \prod_{i=1}^{n} S_i$
- There is a valuation function $v_i : S_i \to \mathbb{R}$
- Common prior $\pi \in \Delta(S)$, which induces conditional distributions $\pi(s_{-i}|s_i)$

Mechanisms

- The revelation principle continues to hold (and only relies on the fact that the designer can choose any mechanism and can pick the equilibrium), so it is WLOG to restrict attention to direct mechanisms, i.e.,

$$q : S \to [0,1]^n, \sum_i q_i(s) \leq 1, \quad t : S \to \mathbb{R}^n$$

- $q_i(s)$ is the probability that buyer i receives the good and $t_i(s)$ is buyer i's net transfer to the seller

Mechanisms

- The revelation principle continues to hold (and only relies on the fact that the designer can choose any mechanism and can pick the equilibrium), so it is WLOG to restrict attention to direct mechanisms, i.e.,

$$q : S \to [0,1]^n, \sum_i q_i(s) \leq 1, \quad t : S \to \mathbb{R}^n$$

- $q_i(s)$ is the probability that buyer i receives the good and $t_i(s)$ is buyer i's net transfer to the seller
- The mechanism is **incentive compatible** (IC) if for all i, s_i, and s_{-i},

$$\sum_{s_{-i}} \pi(s_{-i}|s_i)\left(v_i(s_i)q_i(s_i, s_{-i}) - t_i(s_i, s_{-i})\right)$$

$$\geq \sum_{s_{-i}} \pi(s_{-i}|s_i)\left(v_i(s_i)q_i(s_i', s_{-i}) - t_i(s_i', s_{-i})\right)$$

Mechanisms

- The revelation principle continues to hold (and only relies on the fact that the designer can choose any mechanism and can pick the equilibrium), so it is WLOG to restrict attention to direct mechanisms, i.e.,

$$q : S \to [0,1]^n, \sum_i q_i(s) \leq 1, \quad t : S \to \mathbb{R}^n$$

- $q_i(s)$ is the probability that buyer i receives the good and $t_i(s)$ is buyer i's net transfer to the seller
- The mechanism is **incentive compatible** (IC) if for all i, s_i, and s_{-i},

$$\sum_{s_{-i}} \pi(s_{-i}|s_i)\left(v_i(s_i)q_i(s_i, s_{-i}) - t_i(s_i, s_{-i})\right)$$
$$\geq \sum_{s_{-i}} \pi(s_{-i}|s_i)\left(v_i(s_i)q_i(s_i', s_{-i}) - t_i(s_i', s_{-i})\right)$$

- The mechanism is **individually rational** (IR) if for all i and s_i,

$$\sum_{s_{-i}} \pi(s_{-i}|s_i)\left(v_i(s_i)q_i(s_i, s_{-i}) - t_i(s_i, s_{-i})\right) \geq 0$$

Towards full surplus extraction

- Let $\overline{TS}$ denote the efficient surplus

$$\overline{TS} = \sum_{s \in S} \pi(s) \max_{i=1,\ldots,n} v_i(s_i)$$

- We will show that, given enough linear independence in interim beliefs/correlation in values, there exist IC and IR mechanisms such that revenue is equal to $\overline{TS}$
- The basic strategy is as follows:
  - Allocate the good efficiently to maximize welfare
  - Extract agents' rents by exploiting differences in beliefs

The main result

Theorem (Crémer and McLean)

Suppose that for all i and s_i, there do not exist $\rho(s_i') \geq 0$ for $s_i' \in S_i \setminus \{s_i\}$ such that ($\pi(s_{-i}|s_i) = \sum_{s_i' \in S_i \setminus \{s_i\}} \rho(s_i') \pi(s_{-i}|s_i')$ for all $s_{-i} \in S_{-i}$). Then, there exists an IC and IR mechanism such that $R = \overline{TS}$.

The main result

Theorem (Crémer and McLean)

Suppose that for all i and s_i, there do not exist $\rho(s_i') \geq 0$ for $s_i' \in S_i \setminus \{s_i\}$ such that ($\pi(s_{-i}|s_i) = \sum_{s_i' \in S_i \setminus \{s_i\}} \rho(s_i')\pi(s_{-i}|s_i')$ for all $s_{-i} \in S_{-i}$). Then, there exists an IC and IR mechanism such that $R = \overline{TS}$.

- Proof: The allocation is defined by

$$W(s) = \{i | v_i(s) = \max_{j=1,\ldots,n} v_j(s_j)\}$$

$$q_i(s) = \frac{1}{|W(s)|}\mathbb{I}_{i \in W(s)}$$

  i.e., $q_i(s)$ randomizes the allocation among the bidders with high values

- Now, we will construct transfers such that the IC constraints are satisfied and IR is satisfied as an **equality** for all i

Proof, continued

- The hypothesis of the theorem implies that $\pi(s_{-i}|s_i)$ (viewed as an element of $\mathbb{R}^{S_{-i}}$ is not in the closed convex cone generated by

$$\{\pi(s_{-i}|s_i')|s_i' \in S_i \setminus \{s_i\}\}$$

Proof, continued

- The hypothesis of the theorem implies that $\pi(s_{-i}|s_i)$ (viewed as an element of $\mathbb{R}^{S_{-i}}$ is not in the closed convex cone generated by

$$\{\pi(s_{-i}|s_i')|s_i' \in S_i \setminus \{s_i\}\}$$

- By Farkas, there exists a separating hyperplane $g_i(s_i) \in \mathbb{R}^{S_{-i}}$ such that

$$\sum_{s_{-i}\in S_{-i}} g_i(s_{-i}|s_i)\pi(s_{-i}|s_i) = 0$$

$$\sum_{s_{-i}\in S_{-i}} g_i(s_{-i}|s_i)\pi(s_{-i}|s_i') > 0 \ \forall s_i' \neq s_i$$

Proof, continued

- The hypothesis of the theorem implies that $\pi(s_{-i}|s_i)$ (viewed as an element of $\mathbb{R}^{S_{-i}}$ is not in the closed convex cone generated by

$$\{\pi(s_{-i}|s'_i)|s'_i \in S_i \setminus \{s_i\}\}$$

- By Farkas, there exists a separating hyperplane $g_i(s_i) \in \mathbb{R}^{S_{-i}}$ such that

$$\sum_{s_{-i}\in S_{-i}} g_i(s_{-i}|s_i)\pi(s_{-i}|s_i) = 0$$

$$\sum_{s_{-i}\in S_{-i}} g_i(s_{-i}|s_i)\pi(s_{-i}|s'_i) > 0 \ \forall s'_i \neq s_i$$

- We then set the transfers to be

$$t_i(s) = q_i(s)v_i(s_i) - \gamma g_i(s_{-i}|s_i)$$

for some large γ TBD

Proof, continued continued

- Now, observe that the transfer is

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) q_i(s_i, s_{-i}) v_i(s_i)$$

if the player tells the truth, so that equilibrium surplus is zero

Proof, continued continued

- Now, observe that the transfer is

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) q_i(s_i, s_{-i}) v_i(s_i)$$

  if the player tells the truth, so that equilibrium surplus is zero
- The transfer from misreporting s_i' is

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) q_i(s_i', s_{-i}) v_i(s_i) - \gamma \sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) g_i(s_{-i}|s_i')$$

Proof, continued continued

- Now, observe that the transfer is

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) q_i(s_i, s_{-i}) v_i(s_i)$$

if the player tells the truth, so that equilibrium surplus is zero

- The transfer from misreporting s_i' is

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) q_i(s_i', s_{-i}) v_i(s_i) - \gamma \sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) g_i(s_{-i}|s_i')$$

- Thus, if we pick γ so that

$$\gamma \geq \max_{s_i, s_i'} \frac{v_i(s_i)}{\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|s_i) g_i(s_{-i}|s_i')},$$

then the utility from every misreport will be non-positive □

Genericity of full surplus extraction

- We could view π as being a vector in $\mathbb{R}^{|S|}$
- It is topologically generic that the convex independence assumption will be satisfied (indeed, generically the vectors $\pi(\cdot|s_i)$ are linearly independent)
- In that sense, full surplus extraction is "almost always" possible
- Also, there are other results demonstrating that full surplus extraction is "generically" possible in common value settings, e.g., McAfee, McMillan, and Reny (1989) and McAfee and Reny (1992)
- These results are a bit disturbing and paradoxical: They sugest that we should see full surplus extraction everywhere, but in fact it is nowhere!

Genericity of full surplus extraction

- We could view π as being a vector in $\mathbb{R}^{|S|}$
- It is topologically generic that the convex independence assumption will be satisfied (indeed, generically the vectors $\pi(\cdot|s_i)$ are linearly independent)
- In that sense, full surplus extraction is "almost always" possible
- Also, there are other results demonstrating that full surplus extraction is "generically" possible in common value settings, e.g., McAfee, McMillan, and Reny (1989) and McAfee and Reny (1992)
- These results are a bit disturbing and paradoxical: They sugest that we should see full surplus extraction everywhere, but in fact it is nowhere!
- There is a significant literature debating whether or not this is the right notion of genericity, and whether or not full surplus extraction is truly generic, e.g., Neeman (2004), Heifetz and Neeman (2006), Barelli (2009), Chen and Xiong (2011), Chen and Xiong (2013)

Genericity of full surplus extraction

- We could view π as being a vector in $\mathbb{R}^{|S|}$
- It is topologically generic that the convex independence assumption will be satisfied (indeed, generically the vectors $\pi(\cdot|s_i)$ are linearly independent)
- In that sense, full surplus extraction is "almost always" possible
- Also, there are other results demonstrating that full surplus extraction is "generically" possible in common value settings, e.g., McAfee, McMillan, and Reny (1989) and McAfee and Reny (1992)
- These results are a bit disturbing and paradoxical: They sugest that we should see full surplus extraction everywhere, but in fact it is nowhere!
- There is a significant literature debating whether or not this is the right notion of genericity, and whether or not full surplus extraction is truly generic, e.g., Neeman (2004), Heifetz and Neeman (2006), Barelli (2009), Chen and Xiong (2011), Chen and Xiong (2013)
- This debate is a tad esoteric, but the real point is to highlight a deficiency in the standard model...

Critiques of full surplus extraction

- In my view, there are (at least) two problems with the model that resolve the full-surplus extraction paradox

Critiques of full surplus extraction

- In my view, there are (at least) two problems with the model that resolve the full-surplus extraction paradox
- First, if the matrices $\{\pi(s_{-i}|s_i)\}_{s_i \in S_i}$ are close to singular, then the γ's may have to be enormous to deter deviations
- In other words, very large transfers would be required after certain signal realizations
- This is problematic if there is limited liability or risk aversion

Critiques of full surplus extraction

- In my view, there are (at least) two problems with the model that resolve the full-surplus extraction paradox
- First, if the matrices $\{\pi(s_{-i}|s_i)\}_{s_i \in S_i}$ are close to singular, then the γ's may have to be enormous to deter deviations
- In other words, very large transfers would be required after certain signal realizations
- This is problematic if there is limited liability or risk aversion
- Still, as long as there is correlation, we would expect to see sellers exploit some correlation of this form...

Towards robust mechanism design

- But perhaps more importantly, calibrating these “side bets” requires the seller to have very precise knowledge of beliefs
- If the probability law is misspecified, the seller may go from breaking even to losing millions on average
- (The same comment applies the buyers)
- Addressing this issue, either from the perspective of the designer or the agents, has been the focus of a large part of the literature on robust mechanism design, which we turn to next