ECON 289, Lecture 5: Ex post implementation (Or, “the progress of game theory”!)

Instructor: Ben Brooks

Spring 2023

Common knowledge in mechanism design

- In the models of mechanism design discussed thus far, we have made some strong assumptions about what is common knowledge:
  - The rules of the game
  - A common prior over payoff relevant states
  - Higher order beliefs that are consistent with a common prior
  - Common knowledge of the strategies that are being used

Common knowledge in mechanism design

- In the models of mechanism design discussed thus far, we have made some strong assumptions about what is common knowledge:
  - The rules of the game
  - A common prior over payoff relevant states
  - Higher order beliefs that are consistent with a common prior
  - Common knowledge of the strategies that are being used
- On top of all of that, we have assumed extremely simple forms for information, e.g., private values, independence, symmetry, regularity
- Should we really expect economic agents to agree on all of these things, in a practical setting?

The Wilson critique

Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent it assumes other features to be common knowledge, such as one player's probability assessment about another's preferences or information.
I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality.
– Bob Wilson (1987)

Stronger implementation concepts

- In a Bayes Nash equilibrium, optimality of one's actions depends on beliefs about payoff relevant states and others' behavior
- Wilson's critique beseeches us to focus on mechanisms that achieves the designer's objective, regardless of the detailed structure of beliefs
- If we accept this premise, the next questions are:
  - How should we set up a "robust" mechanism design problem?
  - What is the set of feasible mechanisms?
  - What implementation concept should we use?

Review of basic setup

- Recall our general mechanism design setup:
- Finitely many agents $i = 1, \ldots, N$
- Finite set of payoff states Θ
- Finite set of outcomes X
- Expected utility preferences $u_i : X \times \Theta \to \mathbb{R}$
- A Harsanyi type space (T, π) to represent beliefs
- Designer chooses a mechanism (A, g)
- Agents play a Bayes Nash equilibrium

Review of the revelation princple

- Under these assumptions, the revelation principle applies: Any outcome that can be implemented by some mechanism/equilibrium can also be implemented by a direct revelation mechanism/truthful equilibrium
- Such a mechanism has to satisfy (interim) incentive compatibility: $\forall\ i,\ t_i,\ t_i'$,

$$\sum_{t_{-i},\theta,x} u_i(x,\theta)\left[g(x|t_i,t_{-i}) - g(x|t_i',t_{-i})\right]\pi_i(t_{-i},\theta|t_i) \geq 0$$

- Furthermore, we may also require (interim) individual rationality: $\forall\ i,\ t_i$,

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

From interim to ex post implementation

- Obviously, these conditions depend on both the sets of types T_i and the beliefs π_i
- Consistent with our desire for robustness, we may follow Dasgupta, Hammond, and Maskin (1979) and ask: When does a DRM satisfy IC/IR regardless of beliefs (but holding fixed the sets of types)?
- Clearly, there is a belief that puts probability one on a single (t_{-i}, θ) pair, so the only way for g to be an IC/IR DRM for all beliefs is if for all i, $(t_i, t_{-i}, \theta, t_i'$,

$$\sum_x u_i(x,\theta)\left[g(x|t_i,t_{-i}) - g(x|t_i',t_{-i})\right] \geq 0$$

$$\sum_x u_i(x,\theta) g(x|t_i,t_{-i}) \geq 0$$

- This are referred to as **ex post IC/IR**
- A DRM that satisfies EPIC/IR is said to be **ex post incentive compatible**

Different versions of ex post implementation

- In the DHM robustness-to-beliefs exercise I proposed above, we looked for robustness with respect to **all** beliefs that agents might have about (θ, t_{-i})
- We could also have looked at just robustness with respect to beliefs about t_{-i}, but held beliefs about θ fixed at $\pi_i(\theta|t_i, t_{-i})$ (i.e., just varying beliefs about epistemic types but not varying beliefs about the state given epistemic types)
- This would have given us an alternative version of EPIC/IR, which we may call **(epistemic) ex post IC/IR**:

$$\sum_{x,\theta} u_i(x,\theta)\left[g(x|t_i, t_{-i}) - g(x|t_i', t_{-i})\right]\pi_i(\theta|t_i, t_{-i}) \geq 0$$

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

Dominant strategies

- Here is another variation; in the last version of ex post implementation, we allowed beliefs about others' types to vary, but we supposed that others would report truthfully
- But we could also look for robustness with respect to beliefs about others' types **and others' reports**
- In this case, we would get for all i, θ, θ'_{-i}, and θ'_i

$$\sum_{x,\theta} u_i(x,\theta)\left[g(x|t_i,t'_{-i}) - g(x|t'_i,t'_{-i})\right]\pi_i(\theta|t_i,t_{-i}) \geq 0$$

$$\sum_{x,\theta} u_i(x,\theta) g(x|t_i,t'_{-i})\pi_i(\theta|t_i,t_{-i}) \geq 0$$

- This condition is referred to as **dominant strategy IC/IR** (also called **strategyproof**)
- Clearly implies epistemic EPIC/IR, and in general it is stronger
- (Notice that it makes no difference for the first version of EPIC/IR, since then we are already conditioning on the realized θ, and others' types only matter through their reports)

Ex post implementation: The right solution?

- Part of my point is that we have to ask: ex post with respect to what? The "right" answer may depend on context
- Regardless, one way to proceed would be to use one of these stronger implementation concepts, instead of Bayes Nash
- This would deliver us robustness to beliefs (subject to knowing the right sets of types, and other features of beliefs)
- Many people seem to just run with this idea, and presume that ex post implementation resolves the Wilson critique
- The problem is: Why should we restrict ourselves to implementing the same DRM regardless of beliefs? Why not let the implemented outcome vary with beliefs, as long as it achieves our design goals?
- We will address this issue through two different approaches: implementation-theoretic with Bergemann and Morris (2005), and maxmin with Chung and Ely (2007)

Bergemann and Morris (2005)

- To narrow the distinction between interim and ex post, BM specialize to **known-payoff-type** (KPT) type spaces $\mathcal{T}$ that can be written in the form
  - $\Theta = \prod_i \Theta_i$
  - Finite sets of types T_i for each i
  - A mapping $\hat{\theta}_i : T_i \to \Theta_i$ for each i
  - A belief function $\hat{\pi}_i : T_i \to \Delta(T_{-i})$
- In effect, each player has a **payoff type** θ_i, which affects preferences, and an **epistemic type** t_i, which represents knowledge
- We impose an assumption on the belief hierarchies, that every player "knows" their payoff type
- Under this assumption, our two notions of ex post IC/IR collapse to the same thing, but because types are identified with the payoff state, there is still a gap between ex post and dominant strategies

Special cases of KPT type spaces

- **Private values**: u_i only depends on θ_i and not θ_{-i}
- In this case, ex post and dominant strategy coincide
- **Quasilinear auction model**:
  - $X = X_0 \times (\prod_i X_i)$
  - $X_0 = \Delta(\{0, 1, \ldots, N\})$, $X_i = \mathbb{R}$
  - $u_i(x, \theta) = v(y_0, \theta) - x_i$
    (x_0 is the allocation, x_i is the transfer)

Is it KPT?

- The IPV type space?
- The revenue minimizing information for the FPA?

Is it KPT?

- The IPV type space?
- The revenue minimizing information for the FPA?
- Depends on how we define θ! We could always define θ so that a particular type space is KPT
- But the set of type spaces we allowed in our robust predictions exercise were not all KPT type spaces for the same choice of θ

Implementation

- Rather than demanding that the same outcome is always implemented, the designer has some flexibility, represented by a **social choice correspondence** $F : \Theta \to 2^X \setminus \emptyset$

Implementation

- Rather than demanding that the same outcome is always implemented, the designer has some flexibility, represented by a **social choice correspondence** $F : \Theta \to 2^X \setminus \emptyset$
- Given $\mathcal{T} = (T, \hat{\theta}, \pi)$, a (direct) mechanism $f : T \to X$ is **interim incentive compatible (IIC)** if for all i and $t_i, t_i' \in T_i$,

$$\sum_{t_{-i}} u_i(f(t_i, t_{-i}), \hat{\theta}(t_i, t_{-i}))\pi(t_{-i}|t_i)$$
$$\geq \sum_{t_{-i}} u_i(f(t_i', t_{-i}), \hat{\theta}(t_i, t_{-i}))\pi(t_{-i}|t_i)$$

- F is **interim implementable** on $\mathcal{T}$ if there exists an IIC $f : T \to X$ such that $f(t) \in F(\hat{\theta}(t))$ for all $t \in T$

Implementation

- Rather than demanding that the same outcome is always implemented, the designer has some flexibility, represented by a **social choice correspondence** $F : \Theta \to 2^X \setminus \emptyset$
- Given $\mathcal{T} = (T, \hat{\theta}, \pi)$, a (direct) mechanism $f : T \to X$ is **interim incentive compatible (IIC)** if for all i and $t_i, t_i' \in T_i$,

$$\sum_{t_{-i}} u_i(f(t_i, t_{-i}), \hat{\theta}(t_i, t_{-i}))\pi(t_{-i}|t_i)$$
$$\geq \sum_{t_{-i}} u_i(f(t_i', t_{-i}), \hat{\theta}(t_i, t_{-i}))\pi(t_{-i}|t_i)$$

- F is **interim implementable** on $\mathcal{T}$ if there exists an IIC $f : T \to X$ such that $f(t) \in F(\hat{\theta}(t))$ for all $t \in T$
- $f : \Theta \to X$ is **ex post incentive compatible (EPIC)** if for all i and θ, and θ_i',

$$u_i(f(\theta), \theta) \geq u_i(f(\theta_i', \theta_{-i}), \theta)$$

- F is **ex post implementable** if there is an $f : \Theta \to X$ that is EPIC and $f(\theta) \in F(\theta)$ for all $\theta \in \Theta$

Robustness of ex post implementation

Proposition

If F is ex post implementable, then F is implementable on all KPT type spaces.

Robustness of ex post implementation

Proposition

If F is ex post implementable, then F is implementable on all KPT type spaces.

- Proof: Suppose F is ex post implementable by, say, $f : \Theta \to X$
- Then F is implementable on $\mathcal{T}$ by the function $f' : T \to X$ defined by $f'(t) = f(\hat{\theta}(t))$ $\square$

BM's question

- Suppose we want to implement a SCC F, regardless of the details of higher-order beliefs
- Of course, we can do this if F is ex post implementable
- Are there F's that can be implemented on all KPT type spaces, exploiting flexibility in the outcome, even if they are not ex post implementable?

Separable environments

- We say that (Θ, X, u, F) is **separable** if
  - $X = X_0 \times \left(\prod_{i=1}^{N} X_i\right)$
  - $u_i(x, \theta) = \tilde{u}_i(x_0, x_i, \theta)$
  - There exists $f_0 : \Theta \to X_0$ and $F_i : \Theta \to 2^{X_i} \setminus \emptyset$ such that $F(\theta) = \{f_0(\theta)\} \times (\prod_i F_i(\theta))$
- X_0 is the public good component and X_i are private goods
- Substantive assumption: Options for private good for i does not depend on selection of private goods for other agents

Separable environments

- We say that (Θ, X, u, F) is **separable** if
  - $X = X_0 \times \left(\prod_{i=1}^{N} X_i\right)$
  - $u_i(x, \theta) = \tilde{u}_i(x_0, x_i, \theta)$
  - There exists $f_0 : \Theta \to X_0$ and $F_i : \Theta \to 2^{X_i} \setminus \emptyset$ such that $F(\theta) = \{f_0(\theta)\} \times (\prod_i F_i(\theta))$
- X_0 is the public good component and X_i are private goods
- Substantive assumption: Options for private good for i does not depend on selection of private goods for other agents
- In the quasilinear auction model, X_0 could represent the allocation, X_i is bidder i's transfer,

Separable environments

- We say that (Θ, X, u, F) is **separable** if
  - $X = X_0 \times \left(\prod_{i=1}^{N} X_i\right)$
  - $u_i(x, \theta) = \tilde{u}_i(x_0, x_i, \theta)$
  - There exists $f_0 : \Theta \to X_0$ and $F_i : \Theta \to 2^{X_i} \setminus \emptyset$ such that $F(\theta) = \{f_0(\theta)\} \times (\prod_i F_i(\theta))$
- X_0 is the public good component and X_i are private goods
- Substantive assumption: Options for private good for i does not depend on selection of private goods for other agents
- In the quasilinear auction model, X_0 could represent the allocation, X_i is bidder i's transfer,
- If there is a unique social welfare maximizing allocation given θ, then the problem of implementing a social welfare maximizing social choice function is separable

BM's main result

Proposition

If (Θ, X, u, F) *is separable and* F *is implementable on all KPT type spaces, then* F *is ex post implementable.*

BM's main result

Proposition

If (Θ, X, u, F) *is separable and* F *is implementable on all KPT type spaces, then* F *is ex post implementable.*

- For some i and θ_{-i}, look at the type space where $T_i = \Theta_i$ and $T_j = \{\theta_j\}$

BM's main result

Proposition

If (Θ, X, u, F) is separable and F is implementable on all KPT type spaces, then F is ex post implementable.

- For some i and θ_{-i}, look at the type space where $T_i = \Theta_i$ and $T_j = \{\theta_j\}$
- Separability and the fact that F is interim implementable $\implies$ there exist $g_{i,\theta_{-i}} : \Theta \to X_i$ such that

$$\tilde{u}_i(f_0(\theta), g_{i,\theta_{-i}}(\theta), \theta) \geq \tilde{u}_i(f_0(\theta_i, \theta_{-i}), g_{i,\theta_{-i}}(\theta'_i, \theta_{-i}), \theta)$$

BM's main result

Proposition

If (Θ, X, u, F) *is separable and* F *is implementable on all KPT type spaces, then* F *is ex post implementable.*

- For some i and θ_{-i}, look at the type space where $T_i = \Theta_i$ and $T_j = \{\theta_j\}$
- Separability and the fact that F is interim implementable $\implies$ there exist $g_{i,\theta_{-i}} : \Theta \to X_i$ such that

$$\tilde{u}_i(f_0(\theta), g_{i,\theta_{-i}}(\theta), \theta) \geq \tilde{u}_i(f_0(\theta_i, \theta_{-i}), g_{i,\theta_{-i}}(\theta_i', \theta_{-i}), \theta)$$

- From separability, the function $f' : \Theta \to Y$ where $f_0'(\theta) = f_0(\theta)$ and $f_i'(\theta) = g_{i,\theta_{-i}}(\theta)$ is feasible, and it is EPIC, so F is ex post implementable $\square$

Example without separability

- BM show by example that separability cannot be dropped:
- $N = \{1, 2\}$, $\Theta_i = \{\theta_i, \theta_i'\}$, $X = \{a, b, c\}$
- Payoffs:

a	θ_2	θ_2'	b	θ_2	θ_2'	c	θ_2	θ_2'
θ_1	$(1, 0)$	$(-1, 2)$	θ_1	$(-1, 2)$	$(1, 0)$	θ_1	$(0, 0)$	$(0, 0)$
θ_1'	$(0, 0)$	$(0, 0)$	θ_1'	$(0, 0)$	$(0, 0)$	θ_1'	$(1, 1)$	$(1, 1)$

Example without separability

- BM show by example that separability cannot be dropped:
- $N = \{1, 2\}$, $\Theta_i = \{\theta_i, \theta_i'\}$, $X = \{a, b, c\}$
- Payoffs:

$$\begin{array}{ccc|ccc|ccc} a & \theta_2 & \theta_2' & b & \theta_2 & \theta_2' & c & \theta_2 & \theta_2' \\ \theta_1 & (1,0) & (-1,2) & \theta_1 & (-1,2) & (1,0) & \theta_1 & (0,0) & (0,0) \\ \theta_1' & (0,0) & (0,0) & \theta_1' & (0,0) & (0,0) & \theta_1' & (1,1) & (1,1) \end{array}$$

- Social choice correspondence

$$F(\theta_1, \theta_2) = F(\theta_1, \theta_2') = \{a, b\}$$
$$F(\theta_1', \theta_2) = F(\theta_1', \theta_2') = \{c\}$$

- This SCC is always interim implementable (let player 1 choose the outcome), but it is not implementable ex post

What goes wrong with separability?

- We can't take $X_0 = \{a, b, c\}$, since then we wouldn't be allowed flexibility in what to implement at, e.g., (θ_1, θ_2)
- But if X_1 or X_2 is $\{a, b, c\}$, then we have flexibility, but we violate the assumption that players' preferences only depend on the common outcome and their individual outcome
- Bottom line: BM's separability condition is pretty restrictive, and may even be viewed as a negative result

Auctions with ex post implementation

- Let's now take a different tack, and study the implications of ex post implementation in the context of auctions
- Suppose there are n bidders
- Values $v_i \in V$, where values are evenly spaced in increments of Δ
- We continue to assume private values, so each bidder's v_i is their known payoff type
- The environment is separable, so an outcome is interim implementable in all KPT type spaces if and only if is ex post implementable
- EPIC takes the following form:

$$v_i q_i(v) - t_i(v) \geq v_i q_i(v_i', v_{-i}) - t_i(v_i', v_{-i}) \ \forall i, v$$

- It is standard to also ask that bidders be willing to participate for all KPT type spaces, so we get an ex post participation constraint, i.e.,

$$v_i q_i(v) - t_i(v) \geq 0 \ \forall i, v$$

Profit maximization with ex post implementation

- What are the ex post mechanisms that maximize expected profit?
- Let $f \in \Delta(V^n)$ be the seller's prior
- The revenue maximization program is:

$$\max_{(q,t)} \sum_{v \in V^n} \sum_{i=1}^{N} t_i(v) f(v)$$

$$\text{s.t. } q_i(v) \geq 0 \ \forall i, v, \sum_{i=1}^{N} q_i(v) \leq 1 \ \forall v; \tag{P}$$

$$v_i q_i(v) - t_i(v) \geq 0 \ \forall i, v;$$

$$v_i q_i(v) - t_i(v) \geq v_i q_i(v_i', v_{-i}) - t_i(v_i', v_{-i}) \ \forall i, v;$$

- NB: Very different if we use interim participation!

Applying what we learned from the single agent solution

- We have basically already characterized the solution, with our analysis of the single agent case
- If we fix v_{-i}, then we have a (sub-)distribution $f(v_i, v_{-i})$ of bidder i's value
- We have multipliers $\alpha_i(v, v_i')$ and $\beta_i(v)$
- Transfer neutrality:

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

- Generalized virtual value:

$$\phi_i^\alpha(v) = v_i - \frac{\sum_{v_i'} \alpha(v_i', v_{-i}, v_i)(v_i' - v_i) f(v_i', v_{-i})}{f(v)}$$

- Lagrangian:

$$\max_{q:V \to [0,1]^N} \sum_{v,i} q_i(v) \phi_i^\alpha(v) f(v)$$

The optimal multipliers

- A mechanism is IC for bidder i given v_{-i} iff the allocation $q_i(v_i, v_{-i})$ is non-decreasing in v_i, and the transfers are those that we would implement with the randomized posted price whose CDF is $q_i(\cdot, v_{-i})$

The optimal multipliers

- A mechanism is IC for bidder i given v_{-i} iff the allocation $q_i(v_i, v_{-i})$ is non-decreasing in v_i, and the transfers are those that we would implement with the randomized posted price whose CDF is $q_i(\cdot, v_{-i})$
- Moreover, treating this as bidder i's value distribution, we have Lagrange multipliers $(\overline{\alpha}_i(v_i, v_{-i}), \overline{\beta}_i(v_i, v_{-i}))$ that induce a conditional ironed virtual value $\overline{\phi}_i(v_i|v_{-i})$, which is non-decreasing
- Local downward IC binds everywhere, local upward IC binds on ironed intervals, and IR binds only for the lowest type
- This is the "level free" solution to the dual LP

Analyzing the Lagrangian

- By construction, these multipliers satisfy transfer neutrality and they result in the non-decreasing ironed virtual value
- With this choice of multipliers, the Lagrangian is

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

- An optimal allocation is just to allocate the good to the bidder with the highest virtual value, as long as it is positive:

$$W(v) = \{i|\overline{\phi}_i(v) = \max_j \overline{\phi}_j(v), \overline{\phi}_i(v) \geq 0\}$$

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

- How does this differ from the solution with independence?

Analyzing the Lagrangian

- By construction, these multipliers satisfy transfer neutrality and they result in the non-decreasing ironed virtual value
- With this choice of multipliers, the Lagrangian is

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

- An optimal allocation is just to allocate the good to the bidder with the highest virtual value, as long as it is positive:

$$W(v) = \{i|\overline{\phi}_i(v) = \max_j \overline{\phi}_j(v), \overline{\phi}_i(v) \geq 0\}$$

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

- How does this differ from the solution with independence? $\overline{\phi}_i$ depends on the whole value profile v!
- Nonetheless, this allocation is non-decreasing, and can be implemented with personalized randomized posted prices p_i, whose distribution depends on v_{-i}

Regular case

- Recall the **local virtual value**, $\alpha_i(v, v_i - \Delta) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$, and is zero otherwise, and associated local virtual value

$$\tilde{\phi}_i(v) - \Delta \frac{\sum_{v_i' > v_i} f(v_i', v_{-i})}{f(v)} \tag{1}$$

- f is **regular** if $\tilde{\phi}_i(v)$ is non-decreasing in v_i for all i, v_{-i}
- If f is regular, then $\tilde{\phi}_i = \overline{\phi}_i$, and the local multipliers are optimal

Maxmin foundations

- As discussed in the first lecture, when there is correlation, the seller can generally do strictly better with interim implementation than with ex post, e.g., Crémer and McLean
- One foundation for ex post mechanisms is we want the outcome to be implemented on all KPT type spaces
- Of course, if the real goal is revenue maximization, we might ask: why care whether the same outcome is always implemented, as long as the mechanism performs well in terms of revenue?
- Regardless of the type space, the seller has the option of running the optimal ex post mechanism, and obtain a payoff of Π^* (as long as they can select the equilibrium)
- Natural question: Would any mechanism generate uniformly higher revenue, regardless of the type space?

Chung and Ely (2007)

Theorem (Chung and Ely (2007))

If f is the seller's prior and it is regular, then for any mechanism, there is a KPT type space such that revenue is no greater than Π^*

- In fact, there is a "worst case" type space $\mathcal{T}^*$ such that maximum revenue across all Bayesian mechanisms is Π^*
- Thus, an optimal ex post mechanism $\mathcal{M}^*$ and $\mathcal{T}^*$ are a "saddle point", in the sense that $\mathcal{M}^*$ maximizes revenue on $\mathcal{T}^*$, and $\mathcal{T}^*$ minimizes revenue on $\mathcal{M}^*$
- We will subsequently return to this notion of a saddle point

A worst-case belief structure

- In $\mathcal{T}^*$, each bidder's signal is just their valuation, so it is described by beliefs $\pi_i(v_{-i}|v_i)$
- The corresponding revenue maximization problem is:

$$
\begin{aligned}
&\max_{(q,t)} \sum_{v \in V^n} \sum_{i=1}^{N} t_i(v) f(v) \\
\text{s.t. } \; & q_i(v) \geq 0 \;\forall i, v, \sum_{i=1}^{N} q_i(v) \leq 1 \;\forall v; \\
& \sum_{v_{-i}} \pi_i(v_{-i}|v_i)(v_i q_i(v_i, v_{-i}) - t_i(v_i, v_{-i}) \geq 0 \;\forall i, v_i; \qquad \text{(P}'\text{)} \\
& \sum_{v_{-i}} \pi_i(v_{-i}|v_i)(v_i q_i(v) - t_i(v)) \\
& \qquad \geq \sum_{v_{-i}} \pi_i(v_{-i}|v_i)(v_i q_i(v_i', v_{-i}) - t_i(v_i', v_{-i})) \;\forall i, v;
\end{aligned}
$$

Deriving π^*

- Chung and Ely construct π^* so (P) and (P′) have the same value
- In fact, π^* can be derived from the optimal multipliers for (P)
- Recall that under the regularity hypothesis, only local downward IC and IR for the lowest type are binding
- The optimal multiplier for (i, v) is $\alpha_i(v) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$

Deriving π^*

- Chung and Ely construct π^* so (P) and (P′) have the same value
- In fact, π^* can be derived from the optimal multipliers for (P)
- Recall that under the regularity hypothesis, only local downward IC and IR for the lowest type are binding
- The optimal multiplier for (i, v) is $\alpha_i(v) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$
- Basic fact about linear programs: The value remains the same if a subset of the binding constraints are replaced by a weighted sum of those constraints, with weights that are proportional to the optimal multipliers

Deriving π^*

- Chung and Ely construct π^* so (P) and (P′) have the same value
- In fact, π^* can be derived from the optimal multipliers for (P)
- Recall that under the regularity hypothesis, only local downward IC and IR for the lowest type are binding
- The optimal multiplier for (i, v) is $\alpha_i(v) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$
- Basic fact about linear programs: The value remains the same if a subset of the binding constraints are replaced by a weighted sum of those constraints, with weights that are proportional to the optimal multipliers
- As a result, the value of (P) remains the same if we replace

$$v_i q_i(v_i, v_{-i}) - t_i(v_i, v_{-i}) \geq v_i q_i(v_i - \Delta, v_{-i}) - t_i(v_i - \Delta, v_{-i}) \; \forall v_{-i}$$

with the weighted sum, for any constant $C_i(v_i)$:

$$\sum_{v_{-i}} C_i(v_i)\alpha_i(v_i, v_{-i}) \left[v_i q_i(v_i, v_{-i}) - t_i(v_i, v_{-i})\right]$$
$$\geq \sum_{v_{-i}} C_i(v_i)\alpha_i(v_i, v_{-i}) \left[v_i q_i(v_i - \Delta, v_{-i}) - t_i(v_i - \Delta, v_{-i})\right]$$

Reinterpretation as beliefs

- If we take $C_i(v_i) = 1/\sum_{v_{-i}} \alpha(v_i, v_{-i})$, then $\pi_i^*(v_{-i}|v_i) := C_i(v_i)\alpha(v_i, v_{-i})$ is a belief!
- Hence, the aggregated constraint is just a Bayesian local-downward IC constraint for the beliefs π^*
- We can do the same thing with the EPIR constraints for the lowest type, and aggregate them into a Bayesian IR constraint with the beliefs π^*, all without changing the value of (P)
- Finally, since the other ex post constraints are slack at the optimal solution to (P), we can aggregate them however we want without changing the value
- Thus, (P′) with the beliefs π^* has the same value as (P) $\square$

Where does this depend on regularity?

- With regularity, for each (v_i, v_{-i}), there is only one binding constraint
  (local down IC or IR at the bottom)
- We can take the multiplier on this single binding constraint to be proportional to $\pi_i(v_{-i}|v_i)$
- But in the irregular case, there will be non-local downward constraints that bind, in particular local-downward and local upward

A minor paradox

- Chung and Ely show that these minmax beliefs violate the common prior assumption

A minor paradox

- Chung and Ely show that these minmax beliefs violate the common prior assumption
- At the same time, Morris's theorem, which we discussed in lecture 2, shows that beliefs satisfy the CPA if and only if there is no trade
- So, there should be a feasible and admissible trade for this type space, i.e., transfers such that every agent has an interim expected non-negative net payoff, and some agent has a strictly positive net payoff

A minor paradox

- Chung and Ely show that these minmax beliefs violate the common prior assumption
- At the same time, Morris's theorem, which we discussed in lecture 2, shows that beliefs satisfy the CPA if and only if there is no trade
- So, there should be a feasible and admissible trade for this type space, i.e., transfers such that every agent has an interim expected non-negative net payoff, and some agent has a strictly positive net payoff
- But then, why can't the seller charge a small fee to the agent with strictly positive payoff, and scale up the transfers arbitrarily to use them as a money pump? How could optimal revenue be finite???

A minor paradox

- Chung and Ely show that these minmax beliefs violate the common prior assumption
- At the same time, Morris's theorem, which we discussed in lecture 2, shows that beliefs satisfy the CPA if and only if there is no trade
- So, there should be a feasible and admissible trade for this type space, i.e., transfers such that every agent has an interim expected non-negative net payoff, and some agent has a strictly positive net payoff
- But then, why can't the seller charge a small fee to the agent with strictly positive payoff, and scale up the transfers arbitrarily to use them as a money pump? How could optimal revenue be finite???
- The answer is IC: In our previous no trade result, we didn't allow agents to misreport
- With such truthtelling constraints, there may still be no trade

No trade without CPA but with IC

- To see how this can happen, consider the following example: There are two states $\Theta = \{0, 1\}$, equally likely, player 1 has one type, and $T_2 = \{0, 1\}$
- Player 1 thinks their signal is correct with probability $p > 1/2$, and player 2 thinks it's correct with probability $q \neq p$, $q > 1/2$
- Without IC, there would be a feasible and acceptable trade, where player 2 pays player 1
  - when $t_2 = \theta$ if $p > q$
  - when $t_2 \neq \theta$ if $p < q$
- But if we add IC, then when $q < p$, player 2 makes money when they are wrong, so they would be better off misreporting and maximize the probability of being incorrect!

Aside on no-trade with IC

- Indeed, we can enrich our analysis of no-trade and the common prior to include IC
- Given a Harsanyi type space (T, π), recall that a trade $\gamma : T \to \mathbb{R}^n$ is **feasible** if $\sum_i \gamma_i(t) \leq 0$
- It is **acceptable** if

$$\sum_{t_{-i},\theta} \pi_i(t_{-i}, \theta | t_i)\gamma_i(t_i, t_{-i}, \theta) \geq 0,$$

  for all i, t_i, and with some strict inequality
- Now we introduce: A trade is **incentive compatible** if

$$\sum_{t_{-i},\theta} \pi_i(t_{-i}, \theta | t_i) \left[\gamma_i(t_i, t_{-i}, \theta) - \gamma_i(t_i', t_{-i}, \theta)\right] \geq 0,$$

  for all i, t_i, t_i'

Farkas Variant II

Lemma (Morris' theorem)

Given $A \in \mathbb{R}^{m \times n}$ and a set S of columns, exactly one of the following is true:

(i) *There exists an $x \in \mathbb{R}^n$ such that $x \geq 0$ and $Ax = 0$ and $x_j > 0$ for all $j \in S$*

(ii) *There exists $j \in S$ and $y \in \mathbb{R}^m$ such that $yA \geq 0$ for all j, with strict inequality for column j*

Farkas dual of trade with IC

Theorem (Morris, 1994)

Exactly one of the following is true:

- *There exists a feasible, acceptable, and IC trade*
- *Beliefs are* noisy concordant*: There exist multipliers* $\pi(\theta, t)$*,* $\lambda_i(t_i)$*, and* $\mu_i(t_i, t_i')$ *such that*

$$\pi(\theta, t) = \lambda_i(t_i)\pi_i(t_{-i}, \theta | t_i) + \sum_{t_i'} \left[\mu_i(t_i, t_i')\pi_i(t_{-i}, \theta | t_i) - \mu_i(t_i', t_i)\pi_i(t_{-i}, \theta | t_i')\right]$$

with $\lambda_i(t_i)$ *not all zero*

Farkas dual of trade with IC

Theorem (Morris, 1994)

Exactly one of the following is true:

- *There exists a feasible, acceptable, and IC trade*
- *Beliefs are* noisy concordant*: There exist multipliers* $\pi(\theta, t)$, $\lambda_i(t_i)$, *and* $\mu_i(t_i, t_i')$ *such that*

$$\pi(\theta, t) = \lambda_i(t_i)\pi_i(t_{-i}, \theta|t_i) + \sum_{t_i'} \left[\mu_i(t_i, t_i')\pi_i(t_{-i}, \theta|t_i) - \mu_i(t_i', t_i)\pi_i(t_{-i}, \theta|t_i')\right]$$

  with $\lambda_i(t_i)$ *not all zero*

- $\pi \geq 0$ is the multiplier on feasibility, $\lambda_i \geq 0$ is on acceptability, and $\mu_i \geq 0$ is on IC
- NB We could rescale the multipliers so that π integrates to 1

Noisy concordant beliefs

- Interpretation: π is a "baseline" common prior $\pi(\theta, t)$
- Moreover, summing this equation across (t_{-i}, θ), we would get

$$\sum_{\theta, t_{-i}} \pi(\theta, t_i, t_{-i}) = \lambda_i(t_i) + \sum_{t_i'} \left[\mu_i(t_i, t_i') - \mu_i(t_i', t_i)\right]$$

- So, it is as if the distribution of t_i is jumbled: Some types t_i' have the belief of t_i and some types t_i have the belief of t_i'
- We can view this as a big Markov matrix, describing the transition from some types to others, and μ_i describes the net flow in and out of t_i, and λ_i describes the net number of types staying the same
- Indeed, we can rearrange the system to

$$\pi_i(t_{-i}, \theta | t_i) = \frac{\pi(\theta, t) + \mu_i(t_i', t_i)\pi_i(t_{-i}, \theta | t_i')}{\lambda_i(t_i) + \sum_{t_i'} \mu_i(t_i, t_i')}$$

so player i's beliefs are a weighted average of their beliefs obtained via Bayes rule under the prior and others' beliefs

Noisy concordance in the maxmin beliefs

- Clearly, it must be that the beliefs derived by Chung and Ely are noisy concordant
- How can we see this? Recall that the belief is

$$\pi^*(v_{-i}|v_i) = \frac{\sum_{v_i' \geq v_i} f(v_i', v_{-i})}{\sum_{v_i' \geq v_i, v_{-i}'} f(v_i', v_{-i})}$$

- The obvious solution is all multipliers zero except:

$$\lambda_i(0) = \sum_{v_i' \geq 0, v_{-i}'} f(v_i', v_{-i}), \quad \mu_i(v_i, v_i - \Delta) = \sum_{v_i' \geq v_i, v_{-i}'} f(v_i', v_{-i}),$$

i.e., local downward IC and IR at the bottom!

- In fact, concordance is exactly the same as transfer neutrality:

$$f(v) = \underbrace{\lambda_i(v_i)}_{=\beta_i(v_i)} \pi_i(v_{-i}|v_i) + \sum_{v_i'} \left[\underbrace{\mu_i(v_i, v_i')}_{=\alpha_i(v_i, v_i')} \pi_i(v_{-i}|v_i) - \mu_i(v_i', v_i)\pi_i(v_{-i}|v_i') \right],$$

except that we have solved for beliefs as a function of multipliers, rather than multipliers as a function of beliefs!

Retrospective on maxmin ex post mechanisms

- Chung and Ely have an incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the rather strong regularity assumption, as well as private values

Retrospective on maxmin ex post mechanisms

- Chung and Ely have an incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the rather strong regularity assumption, as well as private values
- Recently, some papers have been extending the theory beyond Chung and Ely (Yamashita and Zhu 2020, Chen and Li 2018)
- Loosely speaking, the result goes through when only local downward IC binds at the optimum
- All of these results rely on non-common prior beliefs on the part of the bidders

Retrospective on maxmin ex post mechanisms

- Chung and Ely have an incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the rather strong regularity assumption, as well as private values
- Recently, some papers have been extending the theory beyond Chung and Ely (Yamashita and Zhu 2020, Chen and Li 2018)
- Loosely speaking, the result goes through when only local downward IC binds at the optimum
- All of these results rely on non-common prior beliefs on the part of the bidders
- Chung and Ely show by examples that relaxing regularity and imposing the CPA both break the result
- It remains an open question what are maxmin mechanisms without regularity and/or with the CPA

Retrospective on maxmin ex post mechanisms

- Chung and Ely have an incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the rather strong regularity assumption, as well as private values
- Recently, some papers have been extending the theory beyond Chung and Ely (Yamashita and Zhu 2020, Chen and Li 2018)
- Loosely speaking, the result goes through when only local downward IC binds at the optimum
- All of these results rely on non-common prior beliefs on the part of the bidders
- Chung and Ely show by examples that relaxing regularity and imposing the CPA both break the result
- It remains an open question what are maxmin mechanisms without regularity and/or with the CPA
- Moreover, even if the optimal ex post mechanism solves the maxmin problem, there are other mechanisms that do just as well on the worst case, and improve elsewhere (Borgers, 2013)

Problems with ex post implementation

- Bergemann and Morris (2005) show that ex post is equivalent to robustness to beliefs only in the special case of "separable" environments, which is quite strong
- Similarly, Chung and Ely (2007) give a maxmin foundation for ex post mechanisms in the specific context of auction design only when the distribution is regular

Problems with ex post implementation

- Bergemann and Morris (2005) show that ex post is equivalent to robustness to beliefs only in the special case of "separable" environments, which is quite strong
- Similarly, Chung and Ely (2007) give a maxmin foundation for ex post mechanisms in the specific context of auction design only when the distribution is regular
- Another problem with ex post implementation:
  In general, not many SCFs are ex post implementable
- Jehiel, Meyer-ter-Vehn, Moldovanu, and Zame (2006) show in an interdependent value and quasilinear environment with multidimensional signals, that generically (on preferences), the only ex post implementable SCFs are constant, i.e., they choose the same alternative independent of the type profile

The setting of JMMZ

- Suppose there are just two agents and two alternatives
- Each agent has a two-dimensional signal $s_i \in [0,1]^L$
- Preferences $v_{i,k}(s_1, s_2)$, assumed to be continuous
- The mechanism specifies an outcome $q(s)$ and transfers $t_i(s)$
- If (q, t) is ex post IC, then the transfer cannot depend on s_i, except through the outcome that gets implemented (otherwise i would misreport whichever signal minimizes
- So really, we can write $t_{i,k}(s_{-i})$ (allowed arbitrary dependence on s_{-i}

Using indifference

- If q is non-constant, then there is some set $X \subset [0,1]^{2L}$ where $q = 1$ and otherwise $q = 2$
- Suppose the boundary is nice and transfers are differentiable
- At the boundary, both agents must be indifferent between the two alternatives, so we get

$$\underbrace{v_{i,1}(s) - v_{i,2}(s)}_{\equiv \mu_1(s)} - \underbrace{t_{i,1}(s_{-i}) - t_{i,2}(s_{-i})}_{\equiv \tau_i(s_{-i})} = 0$$

$$\iff \nu_i(s) - \tau_i(s_{-i}) = 0$$

- But since this equation holds everywhere on the boundary, it must be that the normal vector to the boundary is the gradient

$$(\nabla_{s_i}\mu_i(s), \nabla_{s_{-i}}\mu_i(s) - \tau_i(s_{-i}))$$

The differential equation

- But we could repeat the same analysis for agent $-i$, and find that the gradient is also

$$(\nabla_{s_i}\mu_{-i}(s) - \tau_{-i}(s_i), \nabla_{s_{-i}}\mu_{-i}(s))$$

- These two gradients must be proportional: for some $\alpha > 0$

$$\begin{pmatrix} \nabla_{s_i}\mu_i(s) \\ \nabla_{s_{-i}}\mu_i(s) - \tau_i(s_{-i}) \end{pmatrix} = \alpha \begin{pmatrix} \nabla_{s_i}\mu_{-i}(s) - \tau_{-i}(s_i) \\ \nabla_{s_{-i}}\mu_{-i}(s) \end{pmatrix}$$

- In particular, we have that given s_i, we must have for $\alpha(s) > 0$,

$$\nabla_{s_i}\mu_i(s_i, s_{-i}) = \alpha(s)(\nabla_{s_i}(\mu_{-i}(s_i, s_{-i}) - \nabla_{s_i}\tau_{-i}(s_i))$$

- This is a PDE in s_{-i} that has to be satisfied by preferences, and is quite restrictive...

A parametric example

- Suppose $L = 2$ and

$$v_{i,k}(s) = s_{i,k}(a_{i,k} + b_{i,k}s_{-i,k})$$

- Then

$$\mu_i(s) = a_{i,k}s_{i,k} - a_{i,l}s_{i,l} + b_{i,k}s_{i,k}s_{-i,k} - b_{i,l}s_{i,l}s_{-i,l}$$

  and hence

$$\nabla_{s_i}\mu_i(s) = \alpha(s)(\nabla_{s_i}(\mu_{-i}(s_i, s_{-i}) - \nabla_{s_i}\tau_{-i}(s_i))$$
$$\iff \begin{pmatrix} a_{i,k} + b_{i,k}s_{-i,k} \\ -a_{i,l} - b_{i,l}s_{-i,l} \end{pmatrix} = \alpha(s)\begin{pmatrix} b_{-i,k}s_{-i,k} - \tau_{-i,k}(s_i) \\ b_{-i,l}s_{-i,l} - \tau_{-i,l}(s_i) \end{pmatrix}$$

- But this can only hold if $b_{i,k} = \alpha(s)b_{-i,k}$ and $b_{i,l} = \alpha(s)b_{-i,l}$ for $\alpha(s) > 0$, i.e.,

$$b_{i,k}b_{-i,l} = b_{i,l}b_{-i,k}$$

  and this condition is “non-generic” on $\mathbb{R}^4$

Final thoughts on ex post implementation

- JMMZ generalize this idea beyond the parametric example, and relaxing the assumptions of differentiability/smoothness of the boundary
- They show that the PDE must be satisfied more generally, and give give genericity conditions on v so that it cannot be satisfied
- Thus, for most valuations, only constant SCFs can be implemented

Final thoughts on ex post implementation

- JMMZ generalize this idea beyond the parametric example, and relaxing the assumptions of differentiability/smoothness of the boundary
- They show that the PDE must be satisfied more generally, and give give genericity conditions on v so that it cannot be satisfied
- Thus, for most valuations, only constant SCFs can be implemented
- Together with the strong sufficient conditions of BM and CE for ex post to be without loss, this further dampens our hope of building a general theory based on ex post implementation

Final thoughts on ex post implementation

- JMMZ generalize this idea beyond the parametric example, and relaxing the assumptions of differentiability/smoothness of the boundary
- They show that the PDE must be satisfied more generally, and give give genericity conditions on v so that it cannot be satisfied
- Thus, for most valuations, only constant SCFs can be implemented
- Together with the strong sufficient conditions of BM and CE for ex post to be without loss, this further dampens our hope of building a general theory based on ex post implementation
- Of course, the maxmin criterion can still be applied, even when not many SCFs are ex post implementable
- This suggests that in more general environments (e.g., those without KPT), we may want to push beyond ex post implementation, and consider more general mechanisms