Tutorial on Robust Auction Design
Lecture 2

Instructors: Ben Brooks and Songzi Du

Slides @ https://benjaminbrooks.net/
https://econweb.ucsd.edu/~sodu/

EC 2021

Information in auction design

- In lecture 1, we explored various instances of the optimal auction design model
- A core component of each model is a description of the bidders' **information**: What do they know about the value of the good, and what do they know about what others know?
- All of the examples can be viewed as special cases of an **information structure** (also known as a **type space**)

Information in auction design

- In lecture 1, we explored various instances of the optimal auction design model
- A core component of each model is a description of the bidders' **information**: What do they know about the value of the good, and what do they know about what others know?
- All of the examples can be viewed as special cases of an **information structure** (also known as a **type space**)
- More abstractly, fix a set of **payoff-relevant states** Θ
- An **information structure** is a pair $\mathcal{I} = (S, \pi)$, where
  - Each bidder has a set of **signals** S_i (or **types**)
  - There is a function $\pi_i(s_{-i}, \theta | s_i)$ that represents bidder i's **beliefs** about (s_{-i}, θ) (others signals and the state)
- For example, in Vickrey's IPV model, $\Theta = V^n$, $S_i = V$, and $\pi_i(s_{-i}, v | s_i) = \mathbb{I}_{s=v} f_{-i}(s_{-i})$

Information structures and higher-order beliefs

- Harsanyi (1967) famously proposed information structures as an analytically tractable way to represent **higher order beliefs**, i.e.,
  - A first-order belief about the state in $\Delta(\Theta)$
  - A second-order belief about the state and others' first-order beliefs in $\Delta(\Theta \times \Delta(\Theta)^{n-1})$
  - etc.
- Given $\mathcal{I} = (S, \pi)$, there is a natural way to associate each $s_i \in S_i$ with a hierarchy of beliefs
- Mertens and Zamir (1985) later described a sense in which any "reasonable" hierarchy corresponds to a signal in some information structure

Common priors

- A special class of information structures are those in which the beliefs can be derived from a **common prior**
- For example, in the IPV model, we can start with a product distribution $f(v) = \prod_{i=1}^{n} f_i(v_i)$, with marginal distributions f_i
- Private values $\iff$ $S_i = V_i$, and $s_i = v_i$ with probability one
- Each bidder's belief is **derived** from the joint distribution over $S \times V^n$ by Bayesian updating, which due to the independence, gives us $\pi_i(s_{-i}, v|s_i) = \mathbb{I}_{s=v} f_{-i}(s_{-i})$

Common priors

- A special class of information structures are those in which the beliefs can be derived from a **common prior**
- For example, in the IPV model, we can start with a product distribution $f(v) = \prod_{i=1}^{n} f_i(v_i)$, with marginal distributions f_i
- Private values $\iff$ $S_i = V_i$, and $s_i = v_i$ with probability one
- Each bidder's belief is **derived** from the joint distribution over $S \times V^n$ by Bayesian updating, which due to the independence, gives us $\pi_i(s_{-i}, v|s_i) = \mathbb{I}_{s=v} f_{-i}(s_{-i})$
- More broadly, we say that $\pi \in \Delta(S \times \Theta)$ is a **common prior** for $\mathcal{I}$ if $\pi_i(\cdot, \cdot|s_i)$ is obtained by Bayesian updating from π, conditional on s_i
- The **common-prior assumption** (CPA) (i.e., the assumption that there exists a common prior) is somewhat controversial, although it is often made in practice
- Main benefits are (i) tractability and (ii) an integrated view of all agents' (ex ante) welfare (including the mechanism designer)

Common knowledge in mechanism design

- Circling back to the models from Lecture 1, we have made some strong assumptions about what is common knowledge:
  - A common prior from which all agents' beliefs are derived
  - The rules of the game
  - The strategies that are being used

Common knowledge in mechanism design

- Circling back to the models from Lecture 1, we have made some strong assumptions about what is common knowledge:
  - A common prior from which all agents' beliefs are derived
  - The rules of the game
  - The strategies that are being used
- Should we really expect economic agents to agree on all of these things, in a practical setting?
- On top of all of that, we have assumed extremely simple forms for information, e.g., private values, independence, symmetry, regularity
- Should we not expect agents to have information about their own value and also about others' values? This is ruled out in the IPV model: You learn your value exactly, but get no information about others' values (beyond the prior)

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
- In an extreme form, we insist that agents' strategies are optimal, regardless of their higher-order beliefs
- This is referred to as **ex post implementation**

Implementation approach

- An important paper of Bergemann and Morris (2005) formalizes this connection
- Finitely many agents $i = 1, \ldots, N$
- Finite set of **payoff type** profiles: $\Theta = \prod_i \Theta_i$
- Finite set of outcomes Y
- Agents have expected utility preferences over (y, θ) represented by $u_i : Y \times \Theta \to \mathbb{R}$
- The designer's goals are represented by a **social choice correspondence**: $F : \Theta \to 2^Y \setminus \emptyset$
  Interpretation: When the payoff type profile is $\theta \in \Theta$, the mechanism designer wants to implement an outcome in $F(\theta)$

Information structures

- BM specialize to **known-payoff-type** (KPT) information structures $\mathcal{I}$ that can be written in the form
  - Finite sets of types S_i for each i
  - A mapping $\hat{\theta}_i : S_i \to \Theta_i$ for each i
  - A belief function $\hat{\pi}_i : S_i \to \Delta(S_{-i})$
- This corresponds to an assumption on the belief hierarchies, that player i "knows" their payoff type θ_i, and only θ is payoff-relevant

Information structures

- BM specialize to **known-payoff-type** (KPT) information structures $\mathcal{I}$ that can be written in the form
  - Finite sets of types S_i for each i
  - A mapping $\hat{\theta}_i : S_i \to \Theta_i$ for each i
  - A belief function $\hat{\pi}_i : S_i \to \Delta(S_{-i})$
- This corresponds to an assumption on the belief hierarchies, that player i "knows" their payoff type θ_i, and only θ is payoff-relevant
- **Private values**: u_i only depends on θ_i and not θ_{-i}
- **Quasilinear auction model**:
  - $Y = Y_0 \prod_i Y_i$
  - $Y_0 = \Delta(\{0, 1, \dots, N\})$, $Y_i = \mathbb{R}$
  - $u_i(y, \theta) = v(y_0, \theta) - y_i$
    (y_0 is the allocation, y_i is the transfer)

Implementation

- Given $\mathcal{I} = (S, \hat{\theta}, \pi)$, a (direct) mechanism $f : S \to Y$ is **interim incentive compatible (IIC)** if for all i and $s_i, s_i' \in S_i$,

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

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

Implementation

- Given $\mathcal{I} = (S, \hat{\theta}, \pi)$, a (direct) mechanism $f : S \to Y$ is **interim incentive compatible (IIC)** if for all i and $s_i, s_i' \in S_i$,

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

- F is **interim implementable** on $\mathcal{I}$ if there exists an IIC $f : S \to Y$ such that $f(s) \in F(\hat{\theta}(s))$ for all $s \in S$
- $f : \Theta \to Y$ 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 $f : \Theta \to Y$ that is EPIC and $f(\theta) \in F(\theta)$ for all $\theta \in \Theta$

Ex post and dominant strategies

- f is **dominant strategy incentive compatible** if

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

  for all i, θ, θ'_{-i}, and θ'_i
- Clearly implies EPIC, but in general it is weaker
- They coincide in the special case of private values, since then u_i does not depend on θ_{-i}, except through f
- For consistency, I will use the term "ex post" rather than "dominant", even when we specialize to private-value models

Robustness of ex post implementation

Proposition

If F is ex post implementable, then it is implementable on all KPT information structures.

Robustness of ex post implementation

Proposition

If F is ex post implementable, then it is implementable on all KPT information structures.

- Proof: Suppose F is ex post implementable by, say, $f : \Theta \to Y$
- Then F is implementable on $\mathcal{I}$ by the function $f' : S \to Y$ defined by $f'(s) = f(\hat{\theta}(s))$ $\square$

BM's question

- Suppose we want to implement a SCF 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 information structures, even if they are not ex post implementable? Or does the "Wilson doctrine" inevitably lead us to ex post implementation?

Separable environments

- We say that (Θ, Y, u, F) is **separable** if
  - $Y = Y_0 \prod_{i=1}^{N} Y_i$
  - $u_i(y, \theta) = \tilde{u}_i(y_0, y_i, \theta)$
  - There exists $f_0 : \Theta \to Y_0$ and $F_i : \Theta \to 2^{Y_i} \setminus \emptyset$ such that $F(\theta) = f_0(\theta) \prod_i F_i(\theta)$
- Y_0 is the public good component and Y_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, Y, u, F) is **separable** if
  - $Y = Y_0 \prod_{i=1}^{N} Y_i$
  - $u_i(y, \theta) = \tilde{u}_i(y_0, y_i, \theta)$
  - There exists $f_0 : \Theta \to Y_0$ and $F_i : \Theta \to 2^{Y_i} \setminus \emptyset$ such that $F(\theta) = f_0(\theta) \prod_i F_i(\theta)$
- Y_0 is the public good component and Y_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, Y_0 could represent the allocation, Y_i is bidder i's transfer,

Separable environments

- We say that (Θ, Y, u, F) is **separable** if
  - $Y = Y_0 \prod_{i=1}^{N} Y_i$
  - $u_i(y, \theta) = \tilde{u}_i(y_0, y_i, \theta)$
  - There exists $f_0 : \Theta \to Y_0$ and $F_i : \Theta \to 2^{Y_i} \setminus \emptyset$ such that $F(\theta) = f_0(\theta) \prod_i F_i(\theta)$
- Y_0 is the public good component and Y_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, Y_0 could represent the allocation, Y_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, Y, u, F) *is separable and* F *is implementable on all KPT information structures, then* F *is ex post implementable.*

BM's main result

Proposition

If (Θ, Y, u, F) *is separable and* F *is implementable on all KPT information structures, then* F *is ex post implementable.*

- For some i and θ_{-i}, look at the information structure where $S_i = \Theta_i$ and $S_j = \{\theta_j\}$

BM's main result

Proposition

If (Θ, Y, u, F) is separable and F is implementable on all KPT information structures, then F is ex post implementable.

- For some i and θ_{-i}, look at the information structure where $S_i = \Theta_i$ and $S_j = \{\theta_j\}$
- Separability and the fact that F is interim implementable $\Longrightarrow$ there exist $g_i^{i,\theta_{-i}} : \Theta \to Y_i$ such that

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

BM's main result

Proposition

If (Θ, Y, u, F) is separable and F is implementable on all KPT information structures, then F is ex post implementable.

- For some i and θ_{-i}, look at the information structure where $S_i = \Theta_i$ and $S_j = \{\theta_j\}$
- Separability and the fact that F is interim implementable $\implies$ there exist $g_i^{i,\theta_{-i}} : \Theta \to Y_i$ such that

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

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

Necessity of separability

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

Necessity of separability

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

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

Auctions with ex post implementation

- Let's return to the auction model with n bidders, values $v_i \in V = \{v^1, \ldots, v^M\}$, where $v^k - v^{k-1} = \gamma$
- 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 information structures 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 information structures, 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:

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

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

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

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

- NB: Very different if we use interim participation!

An upper bound on revenue

- As in the independent case, we can derive an upper bound on optimal revenue in terms of **virtual values**
- Consider the following weighted sum of the EPIC/EPIR constraints:

$$\sum_{m=1}^{M} \alpha_i(v^m, v_{-i})[v^m(q_i(v^m, v_{-i}) - q_i(v^{m-1}, v_{-i})) - (t_i(v^m, v_{-i}) - t_i(v^{m-1}, v_{-i}))],$$

  where we interpret $q_i(v^0, v_{-i}) = t_i(v^0, v_{-i}) = 0$
- We can then rearrange to

$$\sum_{m=1}^{M} q_i(v^m, v_{-i})(v^m \alpha_i(v^m, v_{-i}) - v^{m+1}\alpha_i(v^{m+1}, v_{-i})) - \sum_{m=1}^{M} t_i(v^m, v_{-i})(\alpha_i(v^m, v_{-i}) - \alpha_i(v^{m+1}, v_{-i}))$$

Virtual values and regularity

- If we set $\alpha_i(v) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$, then this sum reduces to

$$\sum_{m=1}^{M} f(v^m, v_{-i}) \left[q_i(v^m, v_{-i}) \underbrace{\left(v^m + \gamma \frac{\sum_{l=m+1}^{M} f(v^l, v_{-i})}{f(v^m, v_{-i})} \right)}_{\equiv \phi_i(v^m, v_{-i})} - t_i(v^m, v_{-i}) \right]$$

- $\phi_i(v)$ represents a generalized virtual value
- Since $\alpha_i \geq 0$, the sum is non-negative for any EPIC/IR mechanism

Virtual values and regularity

- If we set $\alpha_i(v) = \sum_{v_i' \geq v_i} f(v_i', v_{-i})$, then this sum reduces to

$$\sum_{m=1}^{M} f(v^m, v_{-i}) \left[q_i(v^m, v_{-i}) \underbrace{\left(v^m + \gamma \frac{\sum_{l=m+1}^{M} f(v^l, v_{-i})}{f(v^m, v_{-i})} \right)}_{\equiv \phi_i(v^m, v_{-i})} - t_i(v^m, v_{-i}) \right]$$

- $\phi_i(v)$ represents a generalized virtual value
- Since $\alpha_i \geq 0$, the sum is non-negative for any EPIC/IR mechanism
- If we add the sums to expected revenue, we get

$$\sum_v \sum_i f(v) t_i(v) \leq \sum_v \sum_i f(v) \phi_i(v) q_i(v)$$

Generalized regularity

- f is **regular** if ϕ_i is non-decreasing in v_i for all i, v_{-i}
- In this case, the upper bound is maximized by a mechanism that allocates the good to a bidder with the highest $\phi_i(v)$
- This allocation can be implemented by an auction in which the high-virtual value bidder has to pay $\min\{v_i | \phi_i(v_i, v_{-i}) \geq 0\}$, so the upper bound is attained
- Moreover, $\alpha_i(v)$ must be optimal Lagrange multiplier on local downward EPIC and EPIR for the lowest value
- (The other multipliers are all zero)

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 information structures
- 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 information structure, 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 information structure?

Chung and Ely (2007)

- We endow the seller with a prior f over bidders' values
- Chung and Ely (2007): If f is regular, then the answer is no
- In particular, for any mechanism, there is a KPT information structure such that revenue is no greater than Π^*
- In fact, there is a "worst case" information structure $\mathcal{I}^*$ such that maximum revenue across all Bayesian mechanisms is Π^*
- Thus, an optimal ex post mechanism $\mathcal{M}^*$ and $\mathcal{I}^*$ are a "saddle point", in the sense that $\mathcal{M}^*$ maximizes revenue on $\mathcal{I}^*$, and $\mathcal{I}^*$ minimizes revenue on $\mathcal{M}^*$
- We will subsequently return to this notion of a saddle point

A worst-case belief structure

- In $\mathcal{I}^*$, 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:

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

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

$$v^m q_i(v^m, v_{-i}) - t_i(v^m, v_{-i}) \geq v_i q_i(v^{m-1}, v_{-i}) - t_i(v^{m-1}, v_{-i}) \ \forall v_{-i}$$

with the weighted sum, for any constant $C_i(v^m)$:

$$\sum_{v_{-i}} C_i(v^m)\alpha_i(v^m, v_{-i}) \left[v^m q_i(v^m, v_{-i}) - t_i(v^m, v_{-i})\right]$$
$$\geq \sum_{v_{-i}} C_i(v^m)\alpha_i(v^m, v_{-i}) \left[v^m q_i(v^{m-1}, v_{-i}) - t_i(v^{m-1}, v_{-i})\right]$$

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)

Retrospective on maxmin ex post mechanisms

- An incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the somewhat strong regularity assumption, as well as private values

Retrospective on maxmin ex post mechanisms

- An incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the somewhat 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)
- All of these results rely on non-common prior beliefs on the part of the bidders

Retrospective on maxmin ex post mechanisms

- An incredibly beautiful result
- Further justifies ex post implementation in auctions
- But, still relies on the somewhat 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)
- 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
- 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)

Ex post implementation versus maxmin

- Bergemann and Morris (2005) give us a nice characterization of when ex post implementation captures the desire for robustness embodied in the Wilson critique
- Similarly, Chung and Ely (2007) give a maxmin foundation for ex post mechanisms in the specific context of auction design

Ex post implementation versus maxmin

- Bergemann and Morris (2005) give us a nice characterization of when ex post implementation captures the desire for robustness embodied in the Wilson critique
- Similarly, Chung and Ely (2007) give a maxmin foundation for ex post mechanisms in the specific context of auction design
- 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 multiple signals, that generically, the only ex post implementable SCFs are constant, i.e., they choose the same alternative independent of the type profile

Ex post implementation versus maxmin

- Bergemann and Morris (2005) give us a nice characterization of when ex post implementation captures the desire for robustness embodied in the Wilson critique
- Similarly, Chung and Ely (2007) give a maxmin foundation for ex post mechanisms in the specific context of auction design
- 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 multiple signals, that generically, the only ex post implementable SCFs are constant, i.e., they choose the same alternative independent of the type profile
- 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), maxmin may be more fruitful

Informational vs. other kinds of robustness

- It's important to distinguish the Wilson critique with other robustness considerations
- In AGT, we often assume a simple form for information (e.g., IPV), and look for "robustness" to fundamentals, like the prior distribution of values
- The evaluation criterion generally takes the form of constant factor approximation of some interesting benchmark

Informational vs. other kinds of robustness

- It's important to distinguish the Wilson critique with other robustness considerations
- In AGT, we often assume a simple form for information (e.g., IPV), and look for "robustness" to fundamentals, like the prior distribution of values
- The evaluation criterion generally takes the form of constant factor approximation of some interesting benchmark
- In robust auction design, we may be willing to fix a prior on fundamentals (at least for the seller), while we look for robustness with respect to information/beliefs (although some other notions of robustness are sometimes studied)
- We also typically use implementation theoretic/optimization criteria
- There are obvious areas of overlap that remain largely unexplored, namely, constant factor approximations with rich information