On the Robustness of Second-Price Auctions in Prior-Independent Mechanism Design

Jerry Anunrojwong
Santiago Balseiro
Omar Besbes

Columbia Business School

INFORMS Annual Meeting 2022

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
  - posted price
  - second-price auction
  - fixed or random price/reserve
  - first-price auction
  - all-pay auction
  - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Mechanism Design: How To Optimally Sell Things

- Suppose you have an item and n potential buyers but you don't know their willingness-to-pay. What do you do?
- Possible Mechanisms
    - posted price
    - second-price auction
    - fixed or random price/reserve
    - first-price auction
    - all-pay auction
    - many more!
- design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account buyers' incentives

Why Robust Mechanism Design? Wilson Doctrine

- The classical theory assumes the seller **knows the environment perfectly**, and designs the mechanism with that in mind.
  - It assumes a *known common prior*.
  - Often it also assumes Bayes-Nash equilibrium.
- The theory is elegant, but depends too intricately on details:
  - distributional knowledge
  - strategic behavior of bidders
- Two fundamental questions
  - What is an optimal detail-free mechanism?
  - How much can we expect to perform without details relative to the best we could do?

This work is the first to give prescription for an arbitrary number of buyers across a wide spectrum of distribution classes (valuation dependence).

Why Robust Mechanism Design? Wilson Doctrine

- The classical theory assumes the seller **knows the environment perfectly**, and designs the mechanism with that in mind.
  - It assumes a *known common prior*.
  - Often it also assumes Bayes-Nash equilibrium.
- The theory is elegant, but depends too intricately on details:
  - distributional knowledge
  - strategic behavior of bidders
- Two fundamental questions
  - What is an optimal detail-free mechanism?
  - How much can we expect to perform without details relative to the best we could do?

This work is the first to give prescription for an arbitrary number of buyers across a wide spectrum of distribution classes (valuation dependence).

Why Robust Mechanism Design? Wilson Doctrine

- The classical theory assumes the seller **knows the environment perfectly**, and designs the mechanism with that in mind.
  - It assumes a *known common prior*.
  - Often it also assumes Bayes-Nash equilibrium.
- The theory is elegant, but depends too intricately on details:
  - distributional knowledge
  - strategic behavior of bidders
- Two fundamental questions
  - What is an optimal detail-free mechanism?
  - How much can we expect to perform without details relative to the best we could do?

This work is the first to give prescription for an arbitrary number of buyers across a wide spectrum of distribution classes (valuation dependence).

Why Robust Mechanism Design? Wilson Doctrine

- The classical theory assumes the seller **knows the environment perfectly**, and designs the mechanism with that in mind.
  - It assumes a *known common prior*.
  - Often it also assumes Bayes-Nash equilibrium.
- The theory is elegant, but depends too intricately on details:
  - distributional knowledge
  - strategic behavior of bidders
- Two fundamental questions
  - What is an optimal detail-free mechanism?
  - How much can we expect to perform without details relative to the best we could do?

This work is the first to give prescription for an arbitrary number of buyers across a wide spectrum of distribution classes (valuation dependence).

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0, 1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0, 1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
  - IC: each person prefers to report her true value
  - IR: each person prefers to participate rather than the outside option
  - dominant strategy = IC & IR hold for *every valuation*
    (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0, 1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0, 1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
  - IC: each person prefers to report her true value
  - IR: each person prefers to participate rather than the outside option
  - dominant strategy = IC & IR hold for *every valuation*
    (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0, 1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0, 1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
  - IC: each person prefers to report her true value
  - IR: each person prefers to participate rather than the outside option
  - dominant strategy = IC & IR hold for *every valuation*
    (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0, 1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0, 1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
  - IC: each person prefers to report her true value
  - IR: each person prefers to participate rather than the outside option
  - dominant strategy = IC & IR hold for *every valuation*
    (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0, 1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0, 1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
  - IC: each person prefers to report her true value
  - IR: each person prefers to participate rather than the outside option
  - dominant strategy = IC & IR hold for *every valuation*
    (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Mechanism

- Selling one indivisible good to ***n* buyers.**
- Optimize over **direct mechanisms** (x, p).
- each bidder i submits her valuation $v_i \in [0,1]$ truthfully $\Rightarrow \boldsymbol{v} \in [0,1]^n$
- each bidder i is allocated with prob $x_i(\boldsymbol{v})$, pays $p_i(\boldsymbol{v})$
- subject to **dominant strategy** *incentive compatibility* and *individual rationality* constraints
    - IC: each person prefers to report her true value
    - IR: each person prefers to participate rather than the outside option
    - dominant strategy = IC & IR hold for *every valuation*
      (does NOT require: bidders are Bayesian or know about other bidders)

Problem Formulation: Distribution Classes $\mathcal{F}$

- We want the mechanism to perform well against *any* distribution $\boldsymbol{F}$ in a given class $\mathcal{F}$: objective is worst-case over all $\boldsymbol{F} \in \mathcal{F}$.
- Assume we know the **upper bound on the support**, normalized to 1.
- Different distribution classes capture valuation dependency structures: from arbitrary joint distributions to i.i.d.

Problem Formulation: Distribution Classes $\mathcal{F}$

- We want the mechanism to perform well against *any* distribution $\boldsymbol{F}$ in a given class $\mathcal{F}$: objective is worst-case over all $\boldsymbol{F} \in \mathcal{F}$.
- Assume we know the **upper bound on the support**, normalized to 1.
- Different distribution classes capture valuation dependency structures: from arbitrary joint distributions to i.i.d.

Problem Formulation: Distribution Classes $\mathcal{F}$

- We want the mechanism to perform well against *any* distribution $\boldsymbol{F}$ in a given class $\mathcal{F}$: objective is worst-case over all $\boldsymbol{F} \in \mathcal{F}$.
- Assume we know the **upper bound on the support**, normalized to 1.
- Different distribution classes capture valuation dependency structures: from arbitrary joint distributions to i.i.d.

Problem Formulation: Distribution Classes $\mathcal{F}$

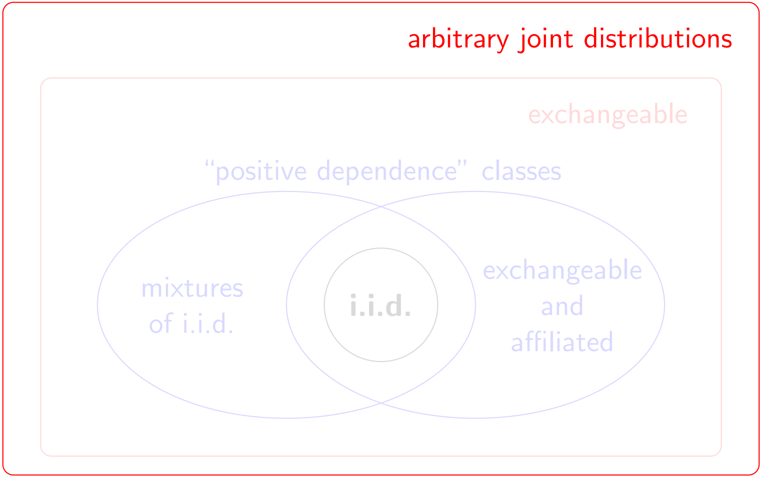

Problem Formulation: Distribution Classes $\mathcal{F}$

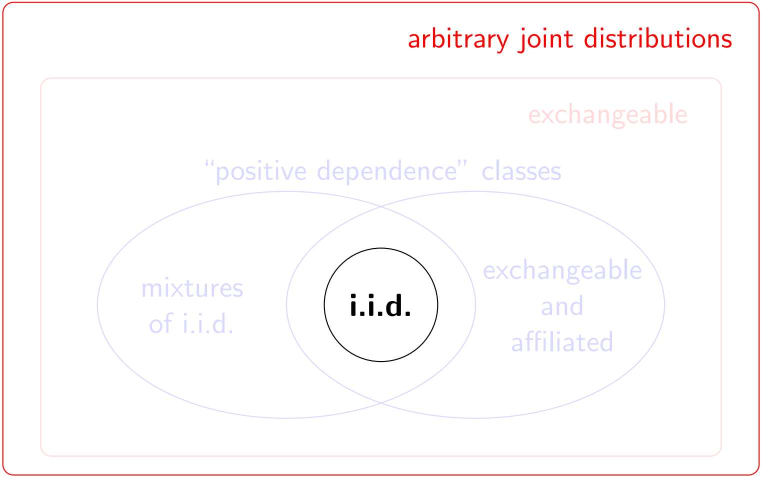

Problem Formulation: Distribution Classes $\mathcal{F}$

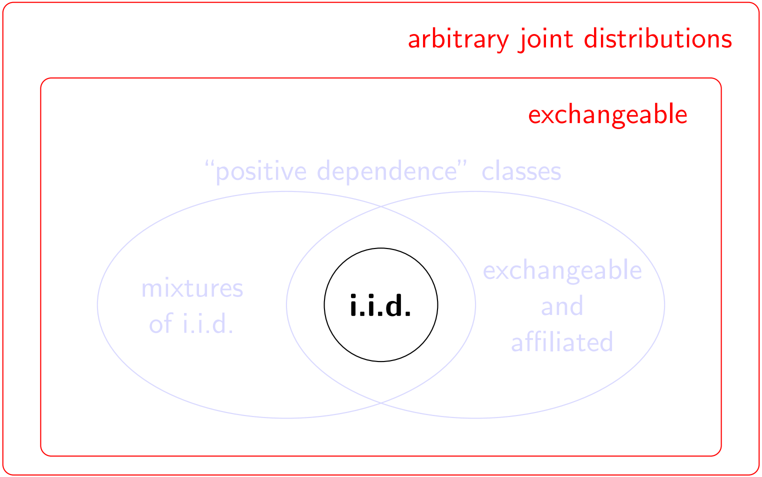

Problem Formulation: Distribution Classes $\mathcal{F}$

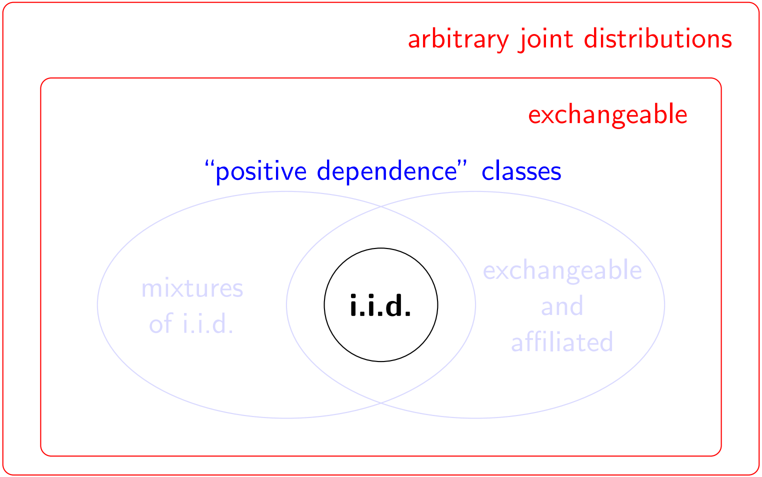

Problem Formulation: Distribution Classes $\mathcal{F}$

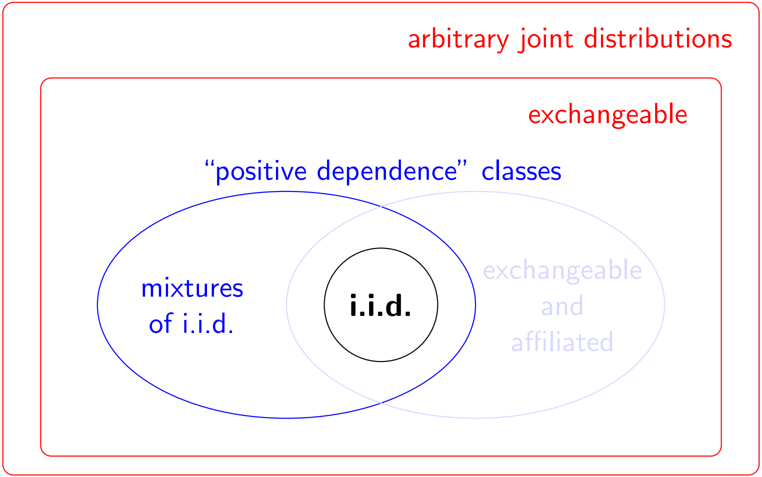

Problem Formulation: Distribution Classes $\mathcal{F}$

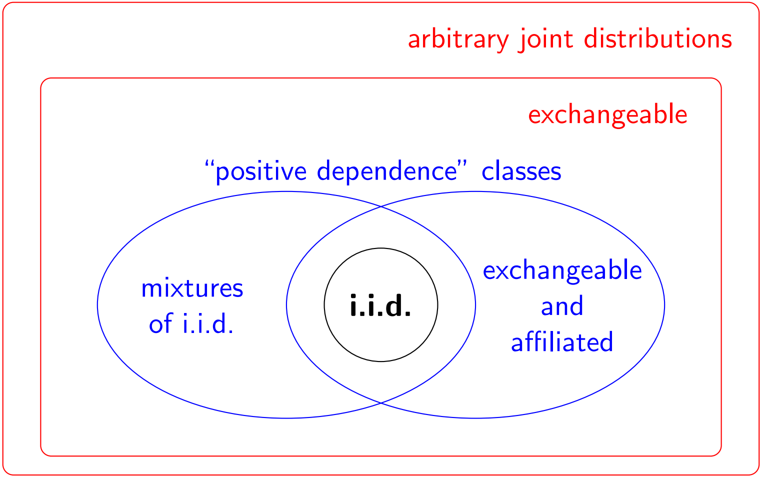

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Distribution Classes $\mathcal{F}$

- **exchangeable** = agents are "symmetric" and can be permuted
  - exchangeable alone allows for arbitrary dependence
- **affiliated** = standard notion for positive dependence in classical (Bayesian) auction theory/mechanism design
  - "Roughly, affiliation means that a high value of one bidder's estimate makes high values of the others' estimates more likely" (Milgrom and Weber, 1982)
  - **We are the first to study affiliation in robust settings**
- **mixtures of i.i.d.** are commonly used in statistics and modeling
  - interpretation: a hidden random type, then i.i.d. conditional on type
  - can also prove: mixtures of i.i.d. implies nonnegative correlation

Problem Formulation: Objective

- The objective is the **regret on revenue**: the difference between the benchmark and the mechanism revenue.

$$\text{Regret} = \text{Benchmark} - \text{Mechanism}$$

- A mechanism m's performance is evaluated by the worst-case regret $\max_{F \in \mathcal{F}} \text{Regret}(m, F)$.
- We focus on the regret because the gap between ideal and actual is an interpretable quantity.
- In contrast, to maximize worst-case revenue, we need additional constraints, e.g. *known mean* o/w worst-case is everyone's value is 0.
- Here, we take the benchmark to be the maximum possible achievable revenue when the valuation is known, i.e. $\max(\boldsymbol{v}) = \max(v_1, \ldots, v_n)$.

Problem Formulation: Objective

- The objective is the **regret on revenue**: the difference between the benchmark and the mechanism revenue.

$$\text{Regret} = \text{Benchmark} - \text{Mechanism}$$

- A mechanism m's performance is evaluated by the worst-case regret $\max_{F \in \mathcal{F}} \text{Regret}(m, F)$.
- We focus on the regret because the gap between ideal and actual is an interpretable quantity.
- In contrast, to maximize worst-case revenue, we need additional constraints, e.g. *known mean* o/w worst-case is everyone's value is 0.
- Here, we take the benchmark to be the maximum possible achievable revenue when the valuation is known, i.e. $\max(\boldsymbol{v}) = \max(v_1, \ldots, v_n)$.

Problem Formulation: Objective

- The objective is the **regret on revenue**: the difference between the benchmark and the mechanism revenue.

$$\text{Regret} = \text{Benchmark} - \text{Mechanism}$$

- A mechanism m's performance is evaluated by the worst-case regret $\max_{F \in \mathcal{F}} \text{Regret}(m, F)$.
- We focus on the regret because the gap between ideal and actual is an interpretable quantity.
- In contrast, to maximize worst-case revenue, we need additional constraints, e.g. *known mean* o/w worst-case is everyone's value is 0.
- Here, we take the benchmark to be the maximum possible achievable revenue when the valuation is known, i.e. $\max(\boldsymbol{v}) = \max(v_1, \ldots, v_n)$.

Problem Formulation: Objective

- The objective is the **regret on revenue**: the difference between the benchmark and the mechanism revenue.

$$\text{Regret} = \text{Benchmark} - \text{Mechanism}$$

- A mechanism m's performance is evaluated by the worst-case regret $\max_{F \in \mathcal{F}} \text{Regret}(m, F)$.
- We focus on the regret because the gap between ideal and actual is an interpretable quantity.
- In contrast, to maximize worst-case revenue, we need additional constraints, e.g. *known mean* o/w worst-case is everyone's value is 0.
- Here, we take the benchmark to be the maximum possible achievable revenue when the valuation is known, i.e. $\max(\boldsymbol{v}) = \max(v_1, \ldots, v_n)$.

Problem Formulation: Optimization Problem

- Minimax formulation, given a distribution class $\mathcal{F} \subseteq \Delta([0,1]^n)$

$$\min_{\substack{\text{mech } (x,p) \\ \text{IC+IR}}} \max_{\boldsymbol{F} \in \mathcal{F}} \mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\underbrace{\max(\boldsymbol{v})}_{\text{benchmark}} - \underbrace{\sum_{i=1}^{n} p_i(\boldsymbol{v})}_{\text{revenue}} \right]$$

- Mechanism is **prior-independent**
  - the mechanism doesn't need to know $\boldsymbol{F}$ and is independent of $\boldsymbol{F}$
  - performance guarantee over all $\boldsymbol{F} \in \mathcal{F}$
  - "detail-free" (not "fine-tuned") or "robust" to distributional knowledge
- Only consider **dominant strategy** IC & IR mechanisms
  - each buyer need not know other buyers' distributions
  - "robust" to specification of bidder behavior

Problem Formulation: Optimization Problem

- Minimax formulation, given a distribution class $\mathcal{F} \subseteq \Delta([0,1]^n)$

$$\min_{\substack{\text{mech } (x,p) \\ \text{IC+IR}}} \max_{\boldsymbol{F} \in \mathcal{F}} \mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\underbrace{\max(\boldsymbol{v})}_{\text{benchmark}} - \underbrace{\sum_{i=1}^{n} p_i(\boldsymbol{v})}_{\text{revenue}} \right]$$

- Mechanism is **prior-independent**
  - the mechanism doesn't need to know $\boldsymbol{F}$ and is independent of $\boldsymbol{F}$
  - performance guarantee over all $\boldsymbol{F} \in \mathcal{F}$
  - "detail-free" (not "fine-tuned") or "robust" to distributional knowledge
- Only consider **dominant strategy** IC & IR mechanisms
  - each buyer need not know other buyers' distributions
  - "robust" to specification of bidder behavior

Problem Formulation: Optimization Problem

- Minimax formulation, given a distribution class $\mathcal{F} \subseteq \Delta([0,1]^n)$

$$\min_{\substack{\text{mech } (x,p) \\ \text{IC+IR}}} \max_{\boldsymbol{F} \in \mathcal{F}} \mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\underbrace{\max(\boldsymbol{v})}_{\text{benchmark}} - \underbrace{\sum_{i=1}^{n} p_i(\boldsymbol{v})}_{\text{revenue}} \right]$$

- Mechanism is **prior-independent**
  - the mechanism doesn't need to know $\boldsymbol{F}$ and is independent of $\boldsymbol{F}$
  - performance guarantee over all $\boldsymbol{F} \in \mathcal{F}$
  - "detail-free" (not "fine-tuned") or "robust" to distributional knowledge
- Only consider **dominant strategy** IC & IR mechanisms
  - each buyer need not know other buyers' distributions
  - "robust" to specification of bidder behavior

Main Contributions: Preview

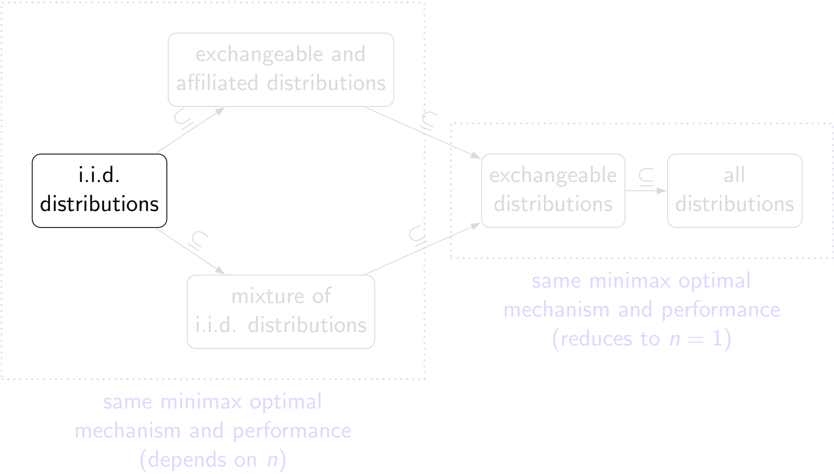

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Main Contributions: Preview

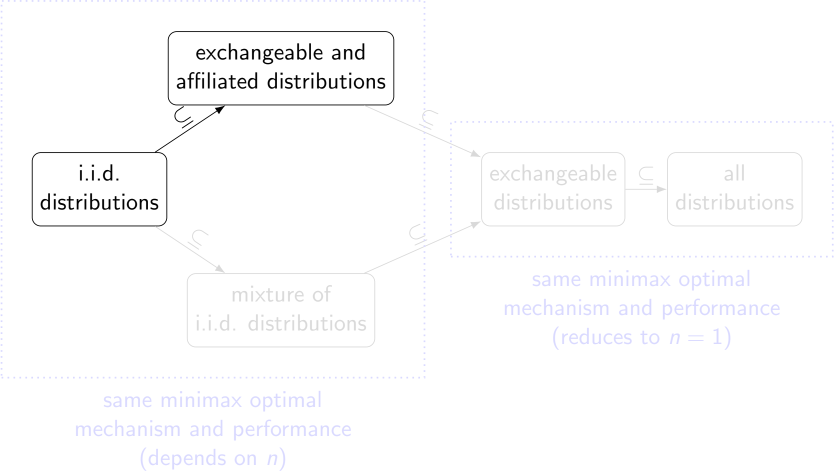

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Main Contributions: Preview

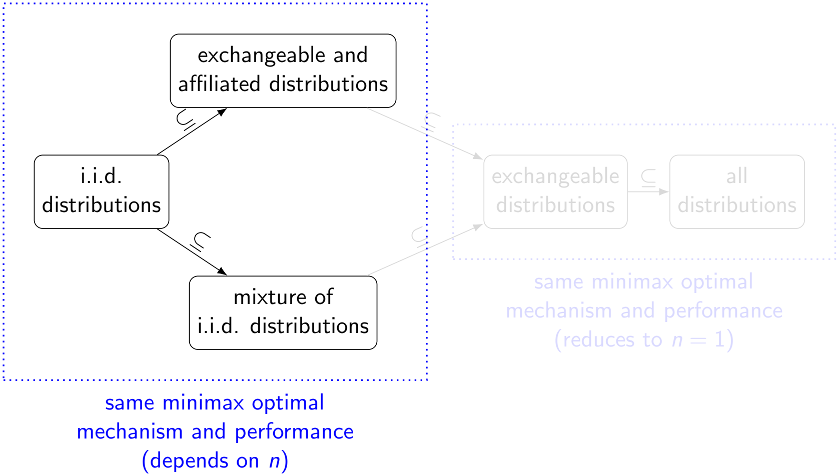

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Main Contributions: Preview

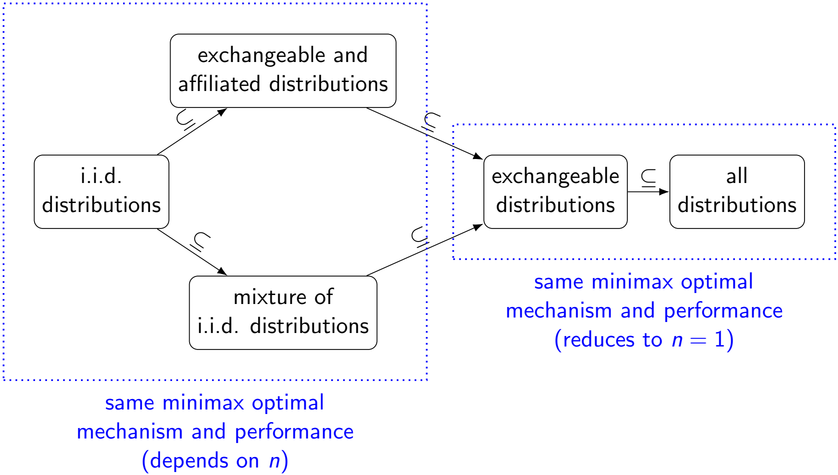

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Main Theorem

Theorem

Under the distribution class of $\{$i.i.d., mixture of i.i.d., exchangeable and affiliated$\}$, the minimax regret admits as an optimal mechanism a second-price auction with random reserve price with cumulative distribution Φ_n^* on $[r_n^*, 1]$ given by

$$\Phi_n^*(v) = \left(\frac{v}{v - r_n^*}\right)^{n-1} \log\left(\frac{v}{r_n^*}\right) - \sum_{k=1}^{n-1} \frac{1}{k}\left(\frac{v}{v - r_n^*}\right)^{n-1-k},$$

where $r_n^* \in (0, 1/n)$ is the unique solution to

$$(1 - r^*)^{n-1} + \log(r^*) + \sum_{k=1}^{n-1} \frac{(1 - r^*)^k}{k} = 0.$$

Main Theorem

Theorem

Under the distribution class of $\{$i.i.d., mixture of i.i.d., exchangeable and affiliated$\}$, the minimax regret admits as an optimal mechanism a second-price auction with random reserve price with cumulative distribution Φ_n^* on $[r_n^*, 1]$ given by

$$\Phi_n^*(v) = \left(\frac{v}{v - r_n^*}\right)^{n-1} \log\left(\frac{v}{r_n^*}\right) - \sum_{k=1}^{n-1} \frac{1}{k}\left(\frac{v}{v - r_n^*}\right)^{n-1-k},$$

where $r_n^* \in (0, 1/n)$ is the unique solution to

$$(1 - r^*)^{n-1} + \log(r^*) + \sum_{k=1}^{n-1} \frac{(1 - r^*)^k}{k} = 0.$$

Main Theorem

Theorem

Under the distribution class of {i.i.d., mixture of i.i.d., exchangeable and affiliated}, the minimax regret admits as an optimal mechanism a second-price auction with random reserve price with cumulative distribution Φ_n^* on $[r_n^*, 1]$ given by

$$\Phi_n^*(v) = \left(\frac{v}{v - r_n^*}\right)^{n-1} \log\left(\frac{v}{r_n^*}\right) - \sum_{k=1}^{n-1} \frac{1}{k} \left(\frac{v}{v - r_n^*}\right)^{n-1-k},$$

where $r_n^* \in (0, 1/n)$ is the unique solution to

$$(1 - r^*)^{n-1} + \log(r^*) + \sum_{k=1}^{n-1} \frac{(1 - r^*)^k}{k} = 0.$$

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Related Literature (Sample)

- Classical mechanism design, pioneered by Myerson (1981)
- Minimax regret on general distributions with bounded support:
  - Bergemann and Schlag (2008) - 1 buyer, continuous
  - Eren and Maglaras (2010) - 1 buyer, discrete
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Robust benchmark-based mechanism design with i.i.d. values
  - Only 1 or 2 buyers: Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Robust maxmin revenue mechanism design (quite different from us)
  - Bandi and Bertsimas (2014), Kocyigit et al. (2019)
- Different form of robustness: Bergemann and Morris (2005), Che (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)
- No **exact optimality results for general i.i.d. or affiliated buyers**!

Challenges of the Problem

- The space of all mechanisms is large.
- The space of bounded distributions is large.
- **The problem is nonconvex due to class restriction in $\mathcal{F}$ e.g. i.i.d.**

We believe that our methodology is of independent interest.

Our Approach

- We use a **saddle point argument.**
- Let $R(m, \boldsymbol{F}) :=$ expected regret with mechanism m and value dist $\boldsymbol{F}$.
- **Saddle Point Theorem.** If the following saddle inequalities hold then then m^* is an optimal mechanism and F^* a worst-case distribution.

$$\overbrace{R(m^*, \boldsymbol{F}) \leq R(m^*, \boldsymbol{F}^*)}^{\boldsymbol{F}^* \text{ is optimal over all } \boldsymbol{F} \text{ given } m^*} \leq R(m, \boldsymbol{F}^*) \quad \forall m, \boldsymbol{F}$$

$$\underbrace{R(m^*, \boldsymbol{F}^*) \leq R(m, \boldsymbol{F}^*)}_{m^* \text{ is optimal over all } m \text{ given } \boldsymbol{F}^*}$$

- Nash equilibrium (best-response on both sides) of a zero-sum game
  - Seller chooses mechanism m to minimize regret
  - Nature chooses distribution $\boldsymbol{F}$ to maximize regret
- *principled approach* to pin down m^* and F^* (necessary conditions)

Our Approach

- We use a **saddle point argument.**
- Let $R(m, \boldsymbol{F}) :=$ expected regret with mechanism m and value dist $\boldsymbol{F}$.
- **Saddle Point Theorem.** If the following saddle inequalities hold then then m^* is an optimal mechanism and F^* a worst-case distribution.

$$\overbrace{R(m^*, \boldsymbol{F}) \leq R(m^*, \boldsymbol{F}^*)}^{\boldsymbol{F}^* \text{ is optimal over all } \boldsymbol{F} \text{ given } m^*} \leq R(m, \boldsymbol{F}^*) \quad \forall m, \boldsymbol{F}$$

$$\underbrace{R(m^*, \boldsymbol{F}^*) \leq R(m, \boldsymbol{F}^*)}_{m^* \text{ is optimal over all } m \text{ given } \boldsymbol{F}^*}$$

- Nash equilibrium (best-response on both sides) of a zero-sum game
  - Seller chooses mechanism m to minimize regret
  - Nature chooses distribution $\boldsymbol{F}$ to maximize regret
- *principled approach* to pin down m^* and F^* (necessary conditions)

Our Approach

- We use a **saddle point argument.**
- Let $R(m, \boldsymbol{F})$:= expected regret with mechanism m and value dist $\boldsymbol{F}$.
- **Saddle Point Theorem.** If the following saddle inequalities hold then then m^* is an optimal mechanism and F^* a worst-case distribution.

$$\overbrace{R(m^*, \boldsymbol{F}) \leq R(m^*, \boldsymbol{F}^*)}^{\boldsymbol{F}^* \text{ is optimal over all } \boldsymbol{F} \text{ given } m^*} \underbrace{\leq R(m, \boldsymbol{F}^*)}_{m^* \text{ is optimal over all } m \text{ given } \boldsymbol{F}^*} \quad \forall m, \boldsymbol{F}$$

- Nash equilibrium (best-response on both sides) of a zero-sum game
  - Seller chooses mechanism m to minimize regret
  - Nature chooses distribution $\boldsymbol{F}$ to maximize regret
- *principled approach* to pin down m^* and F^* (necessary conditions)

Our Approach

- We use a **saddle point argument.**
- Let $R(m, \boldsymbol{F}) :=$ expected regret with mechanism m and value dist $\boldsymbol{F}$.
- **Saddle Point Theorem.** If the following saddle inequalities hold then then m^* is an optimal mechanism and F^* a worst-case distribution.

$$\overbrace{R(m^*, \boldsymbol{F}) \leq R(m^*, \boldsymbol{F}^*)}^{\boldsymbol{F}^* \text{ is optimal over all } \boldsymbol{F} \text{ given } m^*} \underbrace{\!\!\!\!\! \leq R(m, \boldsymbol{F}^*)}_{m^* \text{ is optimal over all } m \text{ given } \boldsymbol{F}^*} \quad \forall m, \boldsymbol{F}$$

- Nash equilibrium (best-response on both sides) of a zero-sum game
  - Seller chooses mechanism m to minimize regret
  - Nature chooses distribution $\boldsymbol{F}$ to maximize regret
- *principled approach* to pin down m^* and F^* (necessary conditions)

Our Approach

- We use a **saddle point argument.**
- Let $R(m, \boldsymbol{F}) :=$ expected regret with mechanism m and value dist $\boldsymbol{F}$.
- **Saddle Point Theorem.** If the following saddle inequalities hold then then m^* is an optimal mechanism and F^* a worst-case distribution.

$$\overbrace{R(m^*, \boldsymbol{F}) \leq R(m^*, \boldsymbol{F}^*)}^{\boldsymbol{F}^* \text{ is optimal over all } \boldsymbol{F} \text{ given } m^*} \leq R(m, \boldsymbol{F}^*) \quad \forall m, \boldsymbol{F}$$

$$R(m^*, \boldsymbol{F}) \leq \underbrace{R(m^*, \boldsymbol{F}^*) \leq R(m, \boldsymbol{F}^*)}_{m^* \text{ is optimal over all } m \text{ given } \boldsymbol{F}^*}$$

- Nash equilibrium (best-response on both sides) of a zero-sum game
  - Seller chooses mechanism m to minimize regret
  - Nature chooses distribution $\boldsymbol{F}$ to maximize regret
- *principled approach* to pin down m^* and F^* (necessary conditions)

Pinning down m^* and F^* with necessary conditions

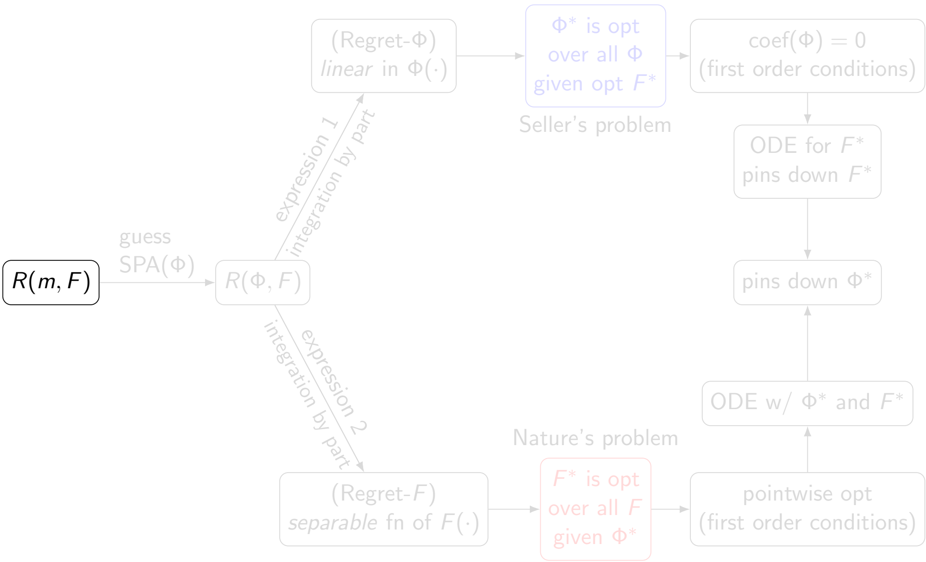

Pinning down m^* and F^* with necessary conditions

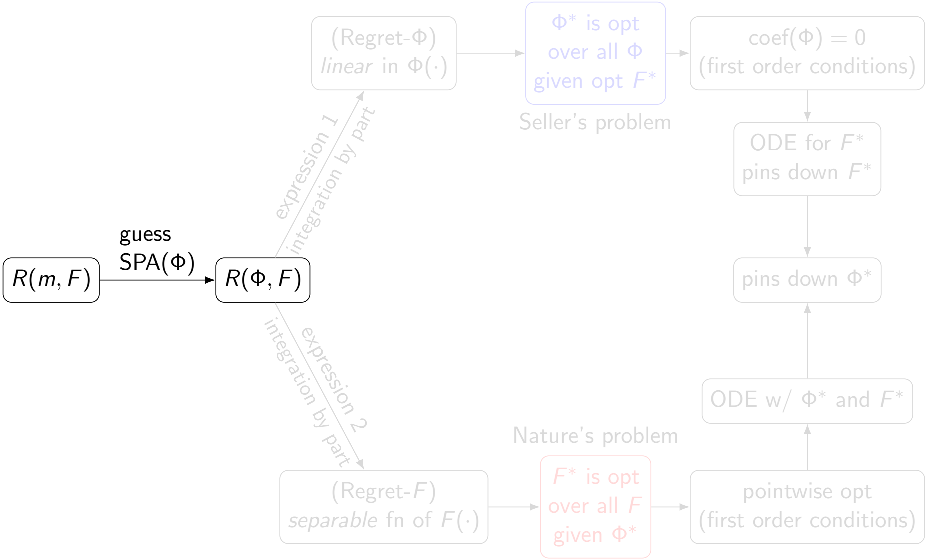

Pinning down m^* and F^* with necessary conditions

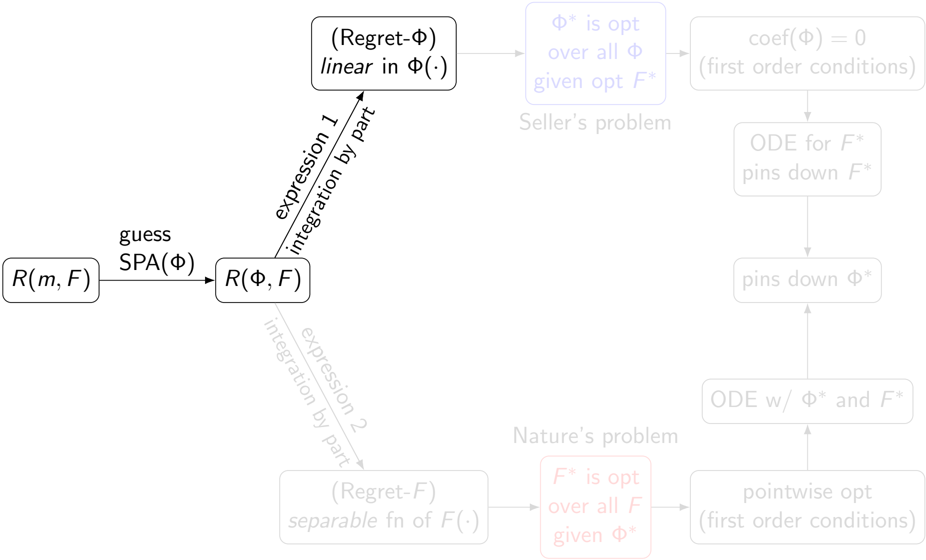

Pinning down m^* and F^* with necessary conditions

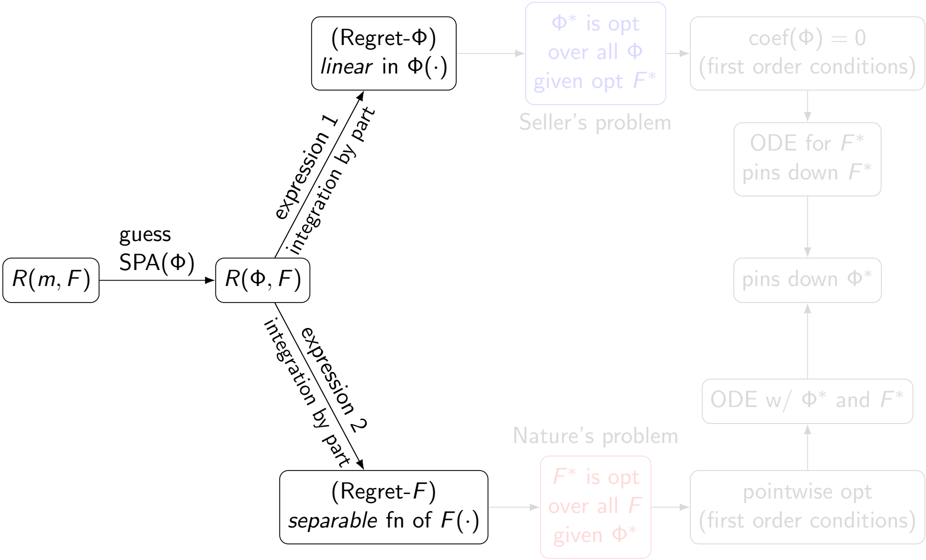

Pinning down m^* and F^* with necessary conditions

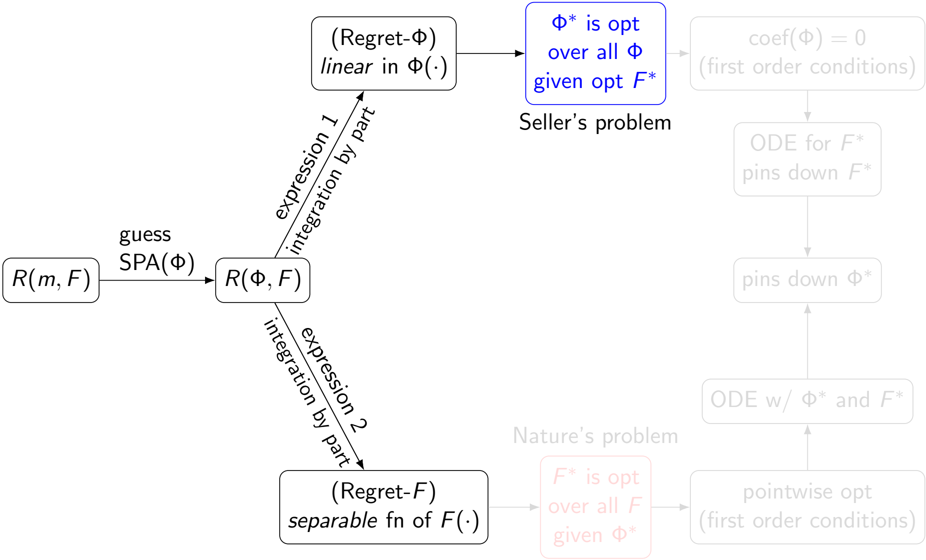

Pinning down m^* and F^* with necessary conditions

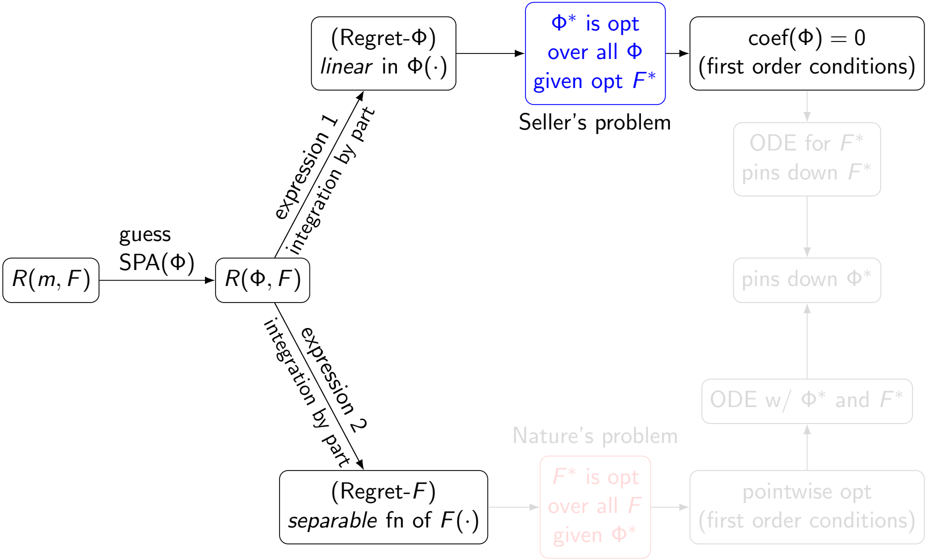

Pinning down m^* and F^* with necessary conditions

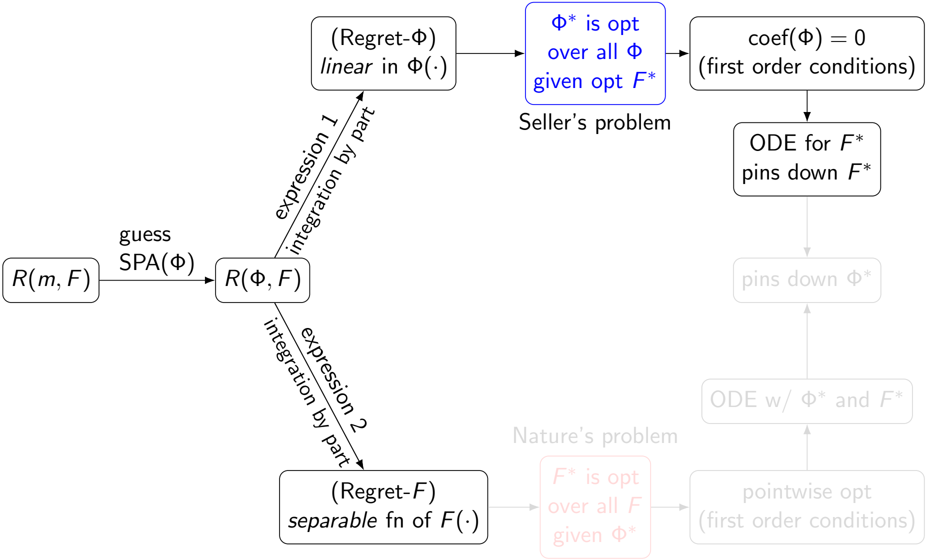

Pinning down m^* and F^* with necessary conditions

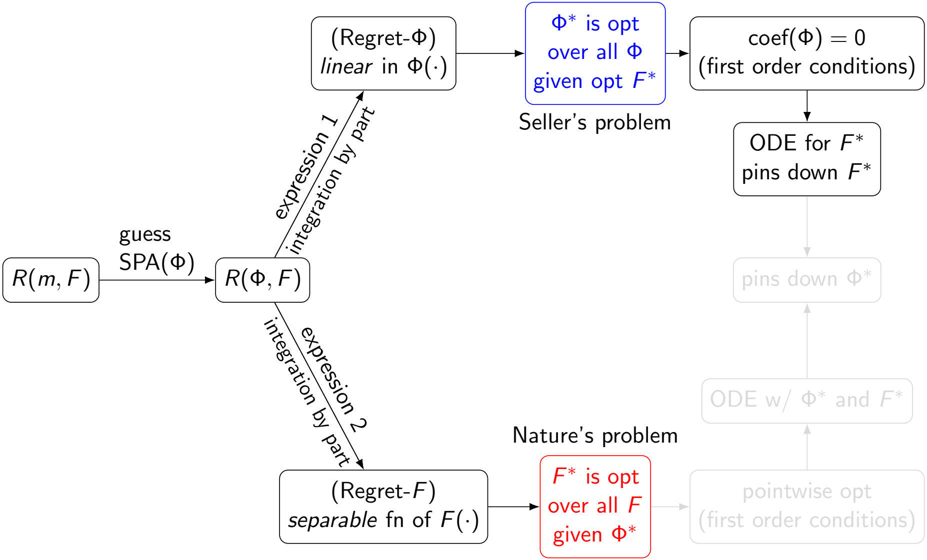

Pinning down m^* and F^* with necessary conditions

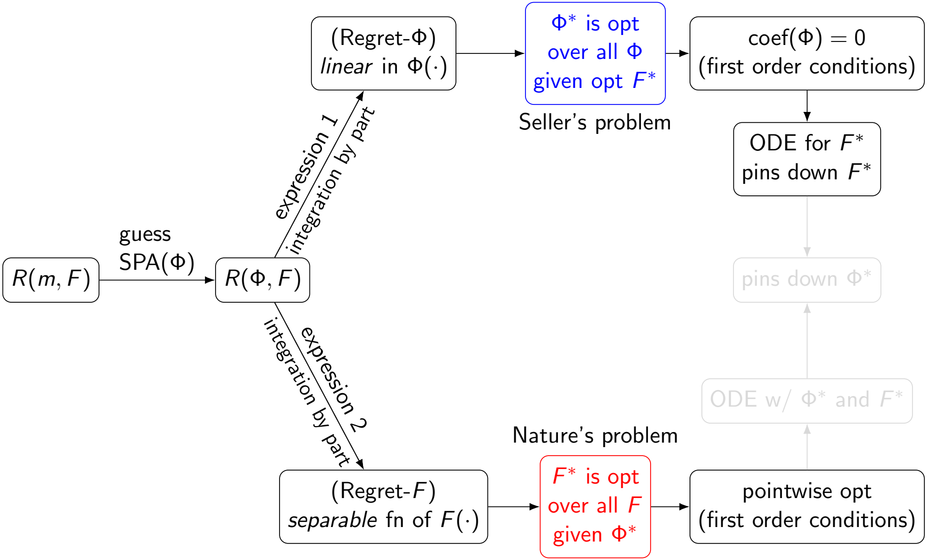

Pinning down m^* and F^* with necessary conditions

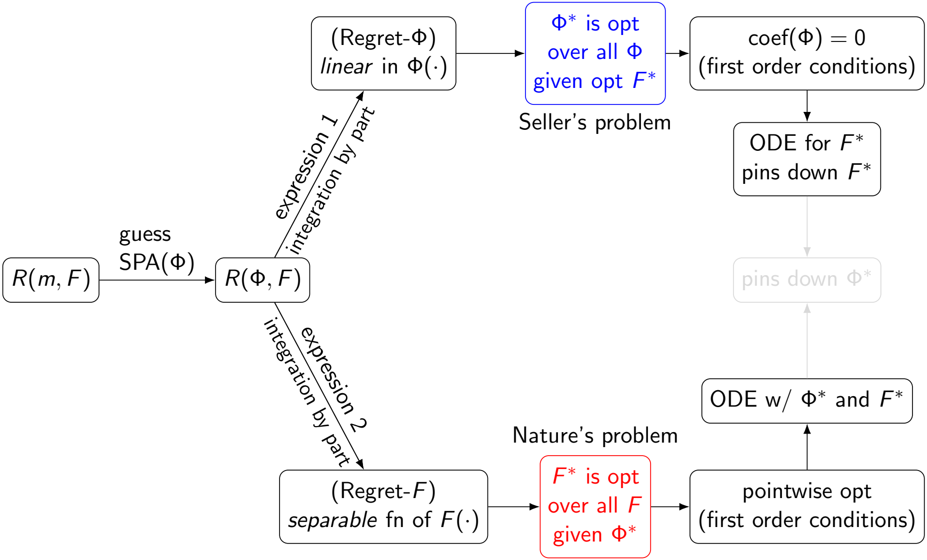

Pinning down m^* and F^* with necessary conditions

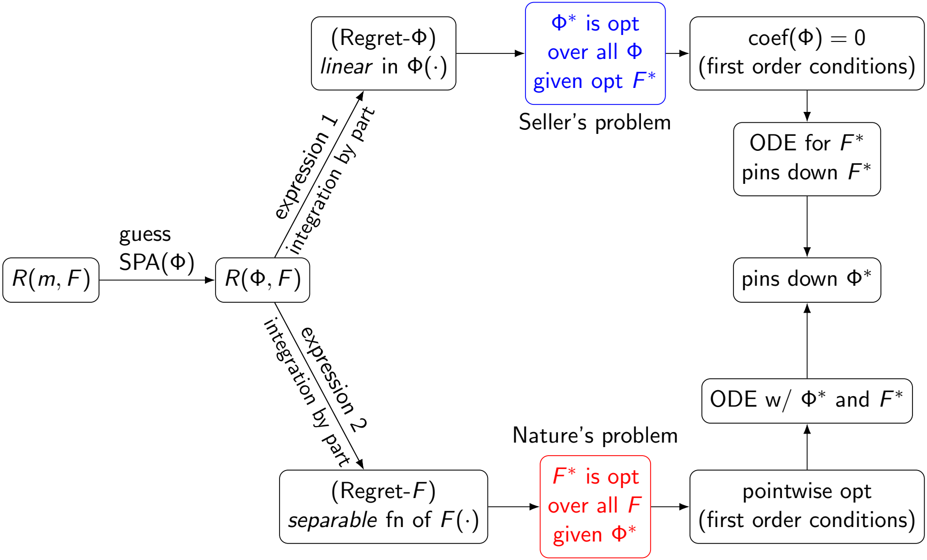

Structure of the Optimal Mechanism

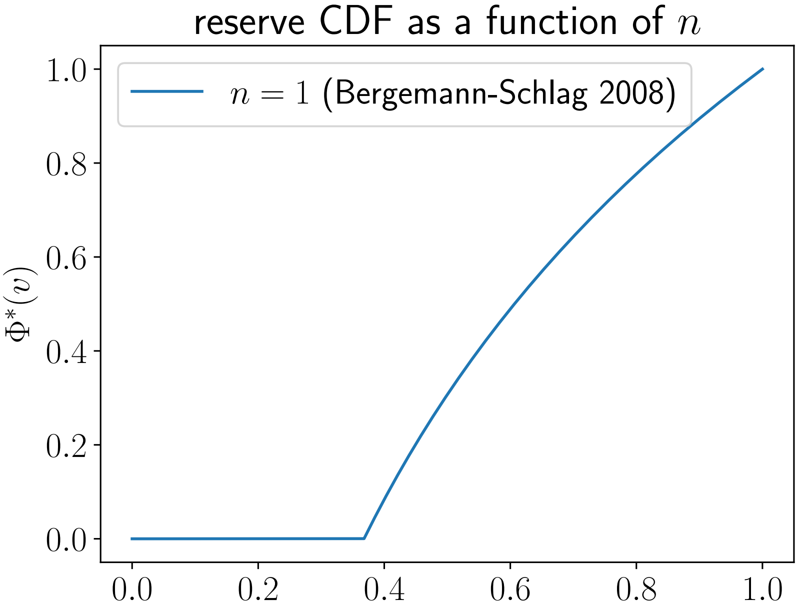

Structure of the Optimal Mechanism

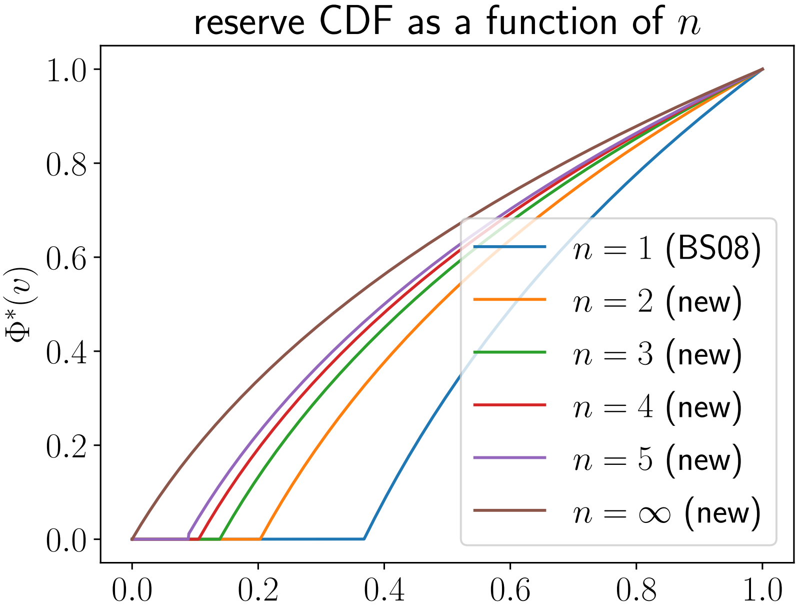

Comparison with Alternative Mechanisms

n	SPA(0)	SPA(r^*)	OPT
1	1.0000	0.5000	0.3679
2	0.5000	0.4444	0.3238
3	0.4444	0.4219	0.3093
4	0.4219	0.4096	0.3021
5	0.4096	0.4019	0.2979
10	0.3874	0.3855	0.2896
25	0.3754	0.3751	0.2847
∞	0.3679	0.3679	0.2815

Table: Worst-case regret for each n.

SPA(0) is the SPA with no reserve.

SPA(r^*) is the SPA with optimal deterministic reserve.

OPT is the optimal mechanism.

OPT is a significant improvement compared to SPA(0) and SPA(r^*).

Comparison with Alternative Mechanisms

n	SPA(0)	SPA(r^*)	OPT
1	1.0000	0.5000	0.3679
2	0.5000	0.4444	0.3238
3	0.4444	0.4219	0.3093
4	0.4219	0.4096	0.3021
5	0.4096	0.4019	0.2979
10	0.3874	0.3855	0.2896
25	0.3754	0.3751	0.2847
∞	0.3679	0.3679	0.2815

Table: Worst-case regret for each n.

SPA(0) is the SPA with no reserve.

SPA(r^*) is the SPA with optimal deterministic reserve.

OPT is the optimal mechanism.

OPT is a significant improvement compared to SPA(0) and SPA(r^*).

Comparison with Alternative Mechanisms

n	SPA(0)	SPA(r^*)	OPT
1	1.0000	0.5000	0.3679
2	0.5000	0.4444	0.3238
3	0.4444	0.4219	0.3093
4	0.4219	0.4096	0.3021
5	0.4096	0.4019	0.2979
10	0.3874	0.3855	0.2896
25	0.3754	0.3751	0.2847
∞	0.3679	0.3679	0.2815

Table: Worst-case regret for each n.
SPA(0) is the SPA with no reserve.
SPA(r^*) is the SPA with optimal deterministic reserve.
OPT is the optimal mechanism.
OPT is a significant improvement compared to SPA(0) and SPA(r^*).

Results apply to many distribution classes

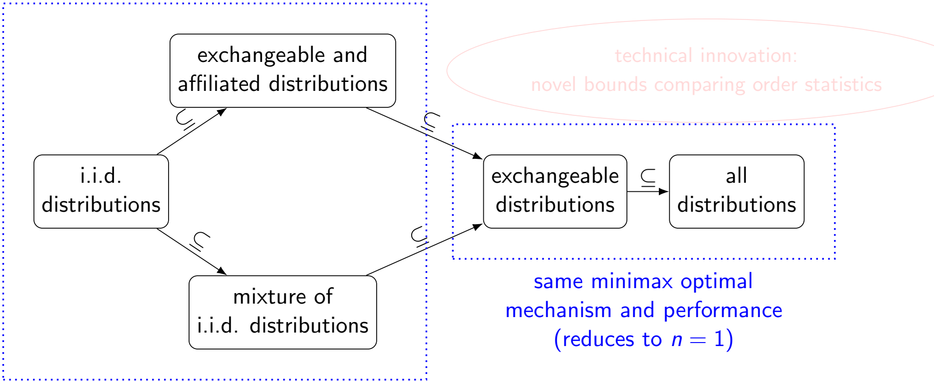

same minimax optimal
mechanism and performance
(reduces to $n = 1$)

same minimax optimal
mechanism and performance
(depends on n)

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Results apply to many distribution classes

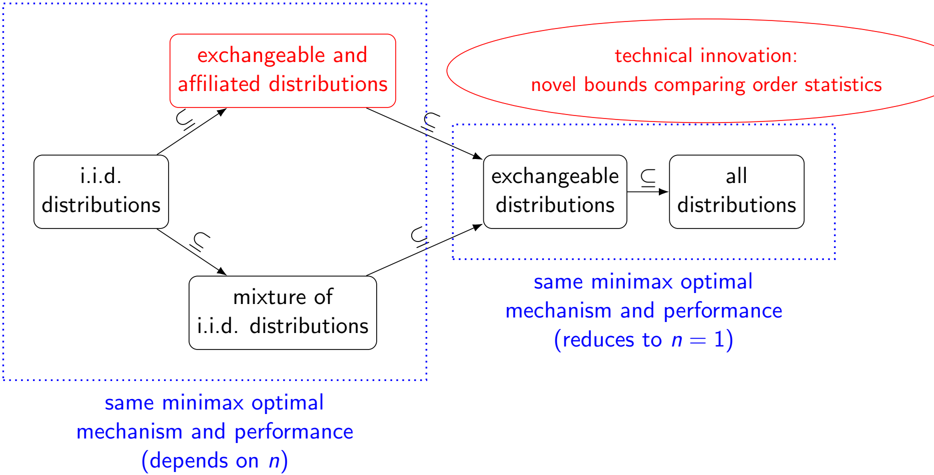

Second Price Auction with Random Reserve is minimax optimal across many distribution classes!

Value of Competition: Positive vs General Dependence

- $\mathcal{F}_{\text{iid}}, \mathcal{F}_{\text{aff}}, \mathcal{F}_{\text{mix}}$ (positive dependence) – regret $\downarrow 0.2815$ as $n \to \infty$
- $\mathcal{F}_{\text{all}}, \mathcal{F}_{\text{exc}}$ (general dependence) – regret is always $1/e \approx 0.3679$

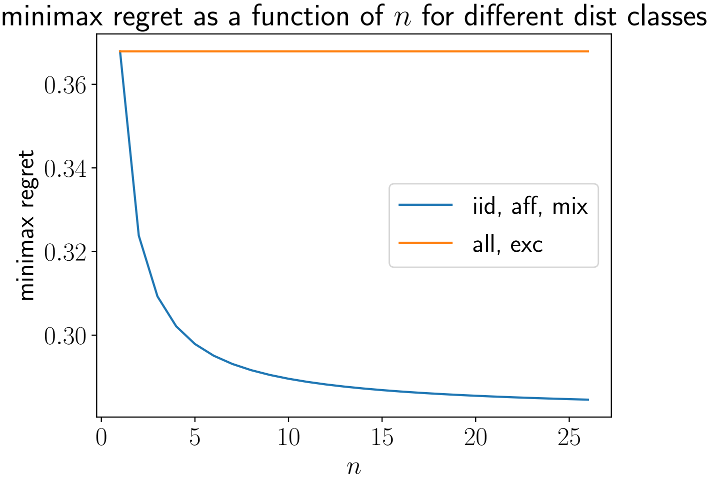

Value of Competition: Positive vs General Dependence

- $\mathcal{F}_{\text{iid}}, \mathcal{F}_{\text{aff}}, \mathcal{F}_{\text{mix}}$ (positive dependence) – regret $\downarrow 0.2815$ as $n \to \infty$
- $\mathcal{F}_{\text{all}}, \mathcal{F}_{\text{exc}}$ (general dependence) – regret is always $1/e \approx 0.3679$

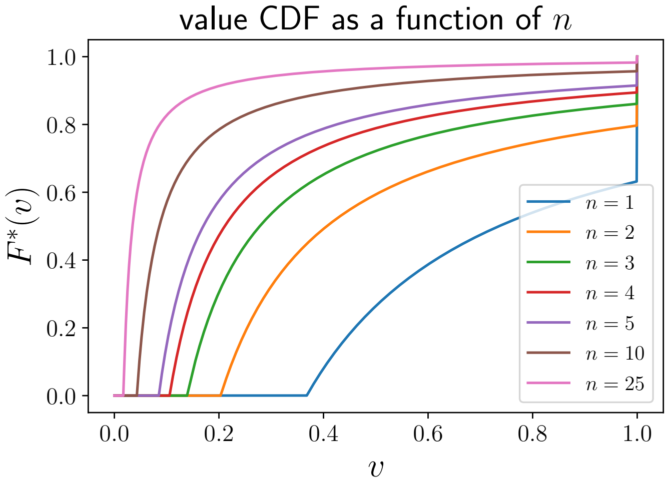

Conclusion

- Systematic study of prior-independent mechanism design
- Closed-form characterization of a minimax optimal mechanism, knowing only the upper bounds on the support
- General framework: n agents, several distribution classes (i.i.d., mixtures of i.i.d., exchangeable and affiliated, exchangeable, all).
- Our results show the strength (or lack thereof) of different distributional class assumptions and quantify the value of competition.
- Related questions:
  - alternative support assumptions
  - alternative benchmarks
  - different notions of robustness
  - value of additional distributional information

Conclusion

- Systematic study of prior-independent mechanism design
- Closed-form characterization of a minimax optimal mechanism, knowing only the upper bounds on the support
- General framework: n agents, several distribution classes (i.i.d., mixtures of i.i.d., exchangeable and affiliated, exchangeable, all).
- Our results show the strength (or lack thereof) of different distributional class assumptions and quantify the value of competition.
- Related questions:
  - alternative support assumptions
  - alternative benchmarks
  - different notions of robustness
  - value of additional distributional information

Conclusion

- Systematic study of prior-independent mechanism design
- Closed-form characterization of a minimax optimal mechanism, knowing only the upper bounds on the support
- General framework: n agents, several distribution classes (i.i.d., mixtures of i.i.d., exchangeable and affiliated, exchangeable, all).
- Our results show the strength (or lack thereof) of different distributional class assumptions and quantify the value of competition.
- Related questions:
  - alternative support assumptions
  - alternative benchmarks
  - different notions of robustness
  - value of additional distributional information

Conclusion

- Systematic study of prior-independent mechanism design
- Closed-form characterization of a minimax optimal mechanism, knowing only the upper bounds on the support
- General framework: n agents, several distribution classes (i.i.d., mixtures of i.i.d., exchangeable and affiliated, exchangeable, all).
- Our results show the strength (or lack thereof) of different distributional class assumptions and quantify the value of competition.
- Related questions:
  - alternative support assumptions
  - alternative benchmarks
  - different notions of robustness
  - value of additional distributional information

Conclusion

- Systematic study of prior-independent mechanism design
- Closed-form characterization of a minimax optimal mechanism, knowing only the upper bounds on the support
- General framework: n agents, several distribution classes (i.i.d., mixtures of i.i.d., exchangeable and affiliated, exchangeable, all).
- Our results show the strength (or lack thereof) of different distributional class assumptions and quantify the value of competition.
- Related questions:
  - alternative support assumptions
  - alternative benchmarks
  - different notions of robustness
  - value of additional distributional information