Robust Mechanism Design with Support Information

Jerry Anunrojwong
Santiago R. Balseiro
Omar Besbes

Columbia Business School

May 15, 2023

Problem Motivation

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?
- Many possible mechanisms: posted price, second-price auction, etc.
- Mechanism design: design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account 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?
- Many possible mechanisms: posted price, second-price auction, etc.
- Mechanism design: design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account 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?
- Many possible mechanisms: posted price, second-price auction, etc.
- Mechanism design: design the *rules of the game* (mechanism) to optimize an objective (e.g. maximize revenue) while taking into account incentives

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?
What features of the robust mechanism are important?
What is the value of support ("scale") information?

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?
What features of the robust mechanism are important?
What is the value of support ("scale") information?

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?
What features of the robust mechanism are important?
What is the value of support ("scale") information?

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?

What features of the robust mechanism are important?

What is the value of support ("scale") information?

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?
What features of the robust mechanism are important?
What is the value of support ("scale") information?

Robust Mechanism Design with *Partial Information*

- Why **robust** mechanism design?
  - Classical theory assumes the seller *knows the environment perfectly*.
  - Elegant theory, but **strong assumptions** on knowledge
- Why **partial information**?
  - Robust mechanism design protects against the **worst-case**.
  - But the *really* worst case is often too extreme.
  - The resulting mechanism is *detail-free* but sometimes *too conservative*.
- We may know **something**!
  - Here, we know the **scale**: lower & upper bounds $[a, b]$
  - e.g. cost-per-click in search ads $\sim \$2 - 4$.

How does the mechanism change depending on information?
What features of the robust mechanism are important?
What is the value of support ("scale") information?

Problem Formulation

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms. (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

Problem Formulation: DSIC Mechanism Class

- Selling one indivisible good to ***n* buyers.**
- Optimize over mechanisms $(x, p) \in \mathcal{M}$.
- each bidder i submits her bid $v_i \in [a, b] \Rightarrow \boldsymbol{v} \in [a, b]^n$
- 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)
- We will mostly focus on $\mathcal{M}_{\text{all}}$, all DSIC mechanisms.
  (Later will look at other mechanism classes too.)

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 $\boldsymbol{F} \in \mathcal{F}$.
- We know the **lower bound** a and **upper bound** b of the support.
- We consider different distribution classes but "positively dependent" distribution classes turn out to be equally powerful as n **i.i.d. bidders**; will focus on $\mathcal{F}_{\text{iid}}$.

a/b = scale information

a/b low = less information
a/b high = more information
$a = 0$: minimal information

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 $\boldsymbol{F} \in \mathcal{F}$.
- We know the **lower bound** a and **upper bound** b of the support.
- We consider different distribution classes but "positively dependent" distribution classes turn out to be equally powerful as n **i.i.d. bidders**; will focus on $\mathcal{F}_{\text{iid}}$.

a/b = scale information

a/b low = less information
a/b high = more information
$a = 0$: minimal information

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 $\boldsymbol{F} \in \mathcal{F}$.
- We know the **lower bound** a and **upper bound** b of the support.
- We consider different distribution classes but "positively dependent" distribution classes turn out to be equally powerful as n **i.i.d. bidders**; will focus on $\mathcal{F}_{\text{iid}}$.

a/b = scale information

a/b low = less information
a/b high = more information
$a = 0$: minimal information

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 $\boldsymbol{F} \in \mathcal{F}$.
- We know the **lower bound** a and **upper bound** b of the support.
- We consider different distribution classes but "positively dependent" distribution classes turn out to be equally powerful as n **i.i.d. bidders**; will focus on $\mathcal{F}_{\text{iid}}$.

a/b = scale information

$$a/b \text{ low} = \text{less information}$$
$$a/b \text{ high} = \text{more information}$$
$$a = 0\text{: minimal information}$$

Problem Formulation: Objective

Compare revenue $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^n p_i(v)\right]$ to benchmark $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}[\max(\boldsymbol{v})]$.

$$\text{absolute gap} = \text{Regret}(m, \boldsymbol{F}) = \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}[\max(\boldsymbol{v})] - \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^n p_i(v)\right]$$

$$\text{relative gap} = \text{Ratio}(m, \boldsymbol{F}) = \frac{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^n p_i(v)\right]}{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}[\max(\boldsymbol{v})]}$$

$$\text{MinimaxRegret}(\mathcal{M}, \mathcal{F}) := \inf_{m\in\mathcal{M}} \sup_{\boldsymbol{F}\in\mathcal{F}} \text{Regret}(m, \boldsymbol{F})$$

$$\text{MaximinRatio}(\mathcal{M}, \mathcal{F}) := \sup_{m\in\mathcal{M}} \inf_{\boldsymbol{F}\in\mathcal{F}} \text{Ratio}(m, \boldsymbol{F})$$

Problem Formulation: Objective

Compare revenue $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]$ to benchmark $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right]$.

$$\text{absolute gap} = \text{Regret}(m, \boldsymbol{F}) = \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right] - \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]$$

$$\text{relative gap} = \text{Ratio}(m, \boldsymbol{F}) = \frac{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]}{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right]}$$

$$\text{MinimaxRegret}(\mathcal{M}, \mathcal{F}) := \inf_{m\in\mathcal{M}} \sup_{\boldsymbol{F}\in\mathcal{F}} \text{Regret}(m, \boldsymbol{F})$$

$$\text{MaximinRatio}(\mathcal{M}, \mathcal{F}) := \sup_{m\in\mathcal{M}} \inf_{\boldsymbol{F}\in\mathcal{F}} \text{Ratio}(m, \boldsymbol{F})$$

Problem Formulation: Objective

Compare revenue $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]$ to benchmark $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right]$.

$$\text{absolute gap} = \text{Regret}(m, \boldsymbol{F}) = \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right] - \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]$$

$$\text{relative gap} = \text{Ratio}(m, \boldsymbol{F}) = \frac{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(v)\right]}{\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\max(\boldsymbol{v})\right]}$$

$$\text{MinimaxRegret}(\mathcal{M}, \mathcal{F}) := \inf_{m\in\mathcal{M}} \sup_{\boldsymbol{F}\in\mathcal{F}} \text{Regret}(m, \boldsymbol{F})$$

$$\text{MaximinRatio}(\mathcal{M}, \mathcal{F}) := \sup_{m\in\mathcal{M}} \inf_{\boldsymbol{F}\in\mathcal{F}} \text{Ratio}(m, \boldsymbol{F})$$

Related Literature

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Related Literature

Problem	Scale Info	Regret	Ratio
pricing ($n = 1$)	any level	Bergemann-Schlag'08	Eren-Maglaras'10
auctions (any n)	no info	**Our previous work**	0
auctions (any n)	**any level**	**This work**	**This work**

- Minimax regret on general distributions with bounded support:
  - Caldentey et al. (2017) - 1 buyer (pricing), multiple time periods
  - Kocyigit et al. (2021) - n correlated buyers (reduce to 1 buyer)
- Classical mechanism design, pioneered by Myerson (1981)
- One or two i.i.d. bidders with benchmark:
  - Dhangwatnotai et al. (2015), Fu et al. (2015), Allouah and Besbes (2020), Hartline and Johnsen (2021)
- Maximizing worst-case revenue:
  - Bandi and Bertsimas (2014), Bachrach and Talgam-Cohen (2022)
- Approximation ratio results for many buyers:
  - Hartline and Roughgarden (2009), Talgam-Cohen (2021)

Main Results

Unifying Regret and Ratio with λ-Regret

For $\lambda \in (0, 1]$ constant:

$$R_\lambda(\mathcal{M}, \mathcal{F}) := \inf_{m \in \mathcal{M}} \sup_{\boldsymbol{F} \in \mathcal{F}} R_\lambda(m, \boldsymbol{F}) := \mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\lambda \max(\boldsymbol{v}) - \sum_{i=1}^{n} p_i(\boldsymbol{v}) \right]$$

Regret is $\lambda = 1$. Ratio is the largest λ such that $R_\lambda(\mathcal{M}, \mathcal{F}) \leq 0$.

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

The need to go beyond SPA.

- Typically robust DSIC mechanism design gives SPA.
- **But SPA with random reserve is not optimal when a/b is large!**
- Imagine a/b is very close to 1. So we know the scale very well.
- Among SPAs, setting any reserve is *risky*: guaranteed payoff too high.
- So among SPAs, no reserve is optimal. Anything better?
- What if ... rather than not allocating below the threshold, we **sometimes allocate to the non-highest bidder**?

New Mechanism Class

SPA is not optimal.

Contribution: a new building block for mechanisms.

Define the *threshold mechanism* with threshold r and default player i:

- If the highest is above r, allocate to the highest
- Otherwise, allocate to the default player i (can be $\emptyset, 1, 2, \ldots, n$)

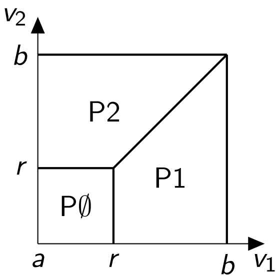

Figure: default player ∅

Figure: default player 1

Figure: default player 2

New Mechanism Class

SPA is not optimal.

Contribution: a new building block for mechanisms.

Define the *threshold mechanism* with threshold r and default player i:

- If the highest is above r, allocate to the highest
- Otherwise, allocate to the default player i (can be $\emptyset, 1, 2, \ldots, n$)

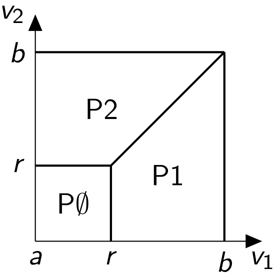

Figure: default player $\emptyset$

Figure: default player 1

Figure: default player 2

New Mechanism Class

SPA is not optimal.
Contribution: a new building block for mechanisms.

Define the *threshold mechanism* with threshold r and default player i:

- If the highest is above r, allocate to the highest
- Otherwise, allocate to the default player i (can be $\emptyset, 1, 2, \ldots, n$)

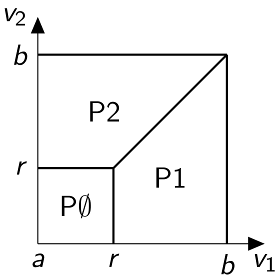

Figure: default player ∅

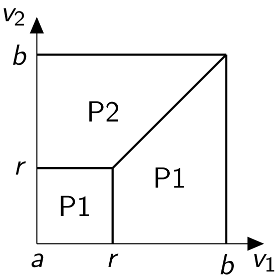

Figure: default player 1

Figure: default player 2

New Mechanism Class

SPA is not optimal.

Contribution: a new building block for mechanisms.

Define the *threshold mechanism* with threshold r and default player i:

- If the highest is above r, allocate to the highest
- Otherwise, allocate to the default player i (can be $\emptyset, 1, 2, \ldots, n$)

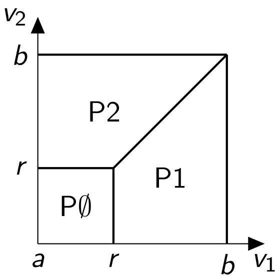

Figure: default player $\emptyset$

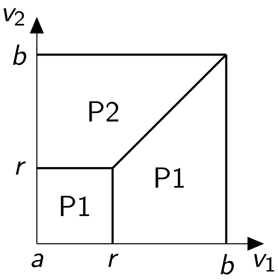

Figure: default player 1

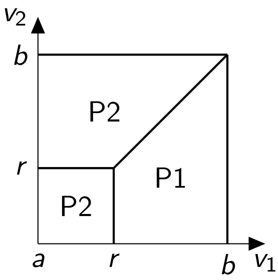

Figure: default player 2

SPA(r) and POOL(τ) mechanism families

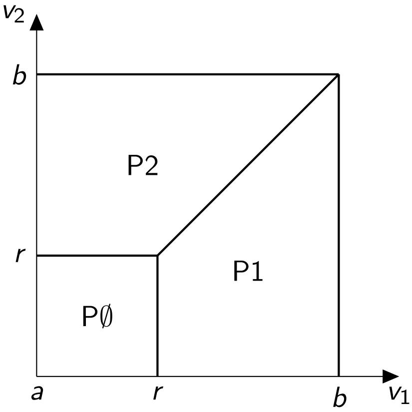

Figure: SPA(r)

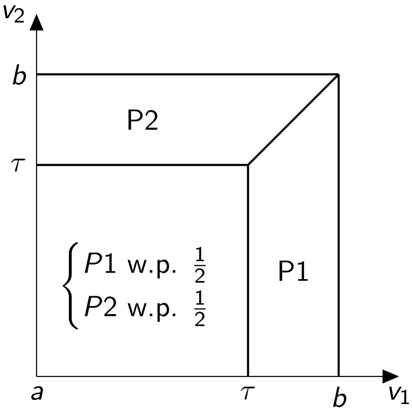

Figure: IRON(τ)

Main Theorem

Theorem (Minimax λ-regret mechanism m^* over $\mathcal{M}_{\text{all}}$ against $\mathcal{F}_{\text{iid}}$)

There are thresholds a_l, a_h such that

- *For $a/b \leq a_l$ (low information), $m^* = \mathrm{SPA}(\Phi)$ with Φ on $[a, b]$.*
- *For $a/b \geq a_h$ (high information), $m^* = \mathrm{POOL}(\Psi)$ with Ψ on $[a, b]$*
- *For $a_l \leq a/b \leq a_h$ (moderate information),*
  m^* is a convex combination of $\mathrm{SPA}(\Phi)$ for Φ in $[a, v^*]$
  and $\mathrm{POOL}(\Psi)$ for Ψ in $[v^*, b]$ for some v^*.

We can characterize Φ and Ψ in closed form.

Our Approach

- We use a **saddle point argument.**
- **Saddle Point Theorem.** If the following saddle inequalities hold then m^* is an optimal mechanism and F^* a worst-case distribution.

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

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

- 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
- Pin down m^* and F^* with (necessary) first-order conditions
- Prove that $(m^*, \boldsymbol{F}^*)$ are actually saddle points beyond FOCs.

Our Approach

- We use a **saddle point argument.**
- **Saddle Point Theorem.** If the following saddle inequalities hold then m^* is an optimal mechanism and F^* a worst-case distribution.

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

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

- 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
- Pin down m^* and F^* with (necessary) first-order conditions
- Prove that $(m^*, \boldsymbol{F}^*)$ are actually saddle points beyond FOCs.

Our Approach

- We use a **saddle point argument.**
- **Saddle Point Theorem.** If the following saddle inequalities hold then m^* is an optimal mechanism and F^* a worst-case distribution.

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

- 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
- Pin down m^* and F^* with (necessary) first-order conditions
- Prove that $(m^*, \boldsymbol{F}^*)$ are actually saddle points beyond FOCs.

Our Approach

- We use a **saddle point argument.**
- **Saddle Point Theorem.** If the following saddle inequalities hold then m^* is an optimal mechanism and F^* a worst-case distribution.

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

- 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
- Pin down m^* and F^* with (necessary) first-order conditions
- Prove that $(m^*, \boldsymbol{F}^*)$ are actually saddle points beyond FOCs.

General Insights

Quantifying the value of scale information and competition

more scale information

more competition

a/b	0.01	0.10	0.25	0.50	0.75	0.99
$n = 1$	0.1784	0.3028	0.4191	0.5906	0.7766	0.9900
$n = 2$	0.2158	0.4038	0.5660	0.7463	0.8841	0.9957
$n = 3$	0.2406	0.4529	0.6110	0.7779	0.9001	0.9963
$n = 5$	0.2686	0.4864	0.6408	0.7981	0.9102	0.9967
$n = 8$	0.2836	0.5035	0.6556	0.8080	0.9150	0.9969

Table: Maximin ratio as a function of a/b for each number of buyers n.

Takeaway: *just a little* scale info gives very good guarantees!
With 2 agents, $a/b = 0.1 \Rightarrow\sim 40\%$, $a/b = 0.5 \Rightarrow\sim 75\%$.

Quantifying importance of mechanism features

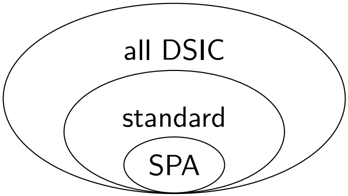

We show **strict separation** between mechanism classes.
Important to **allocate to the non-highest**.
Gaps capture the **cost of simplicity**.

Maximin ratio over different $\mathcal{M}$ for $n = 4$

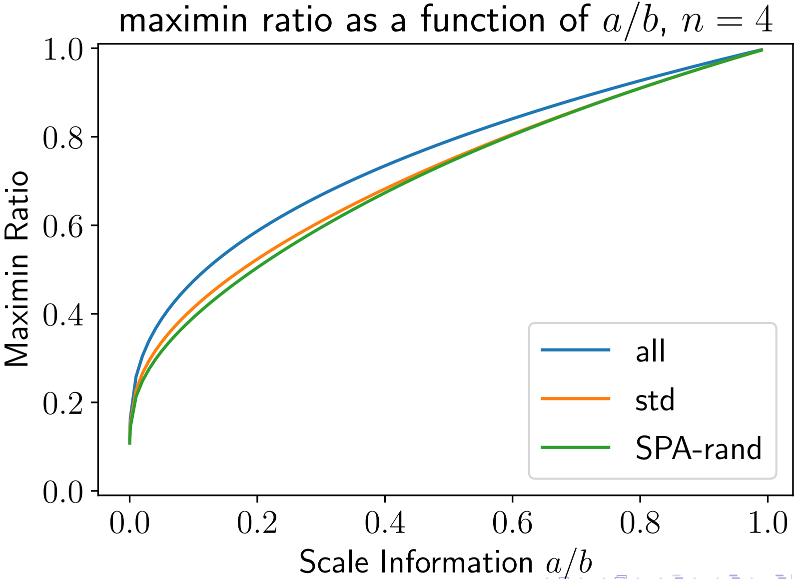

Same results hold for other $\mathcal{F}$ with **positive dependence**!

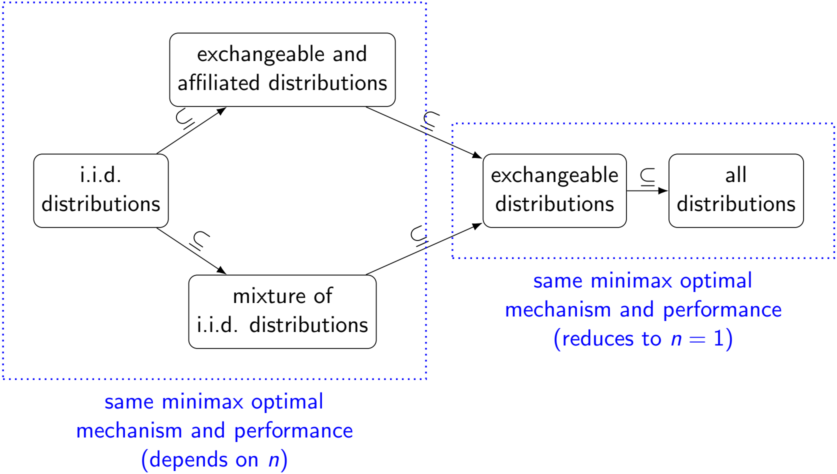

The same mechanism (SPA **and** POOL**)** is minimax optimal across many distribution classes.

Conclusion

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Conclusion

- Closed-form characterization of a minimax optimal mechanism, knowing only the **lower** & **upper** bounds on the support.
- General framework: n agents, several distribution classes $\mathcal{F}$, several mechanism classes $\mathcal{M}$, both regret and ratio objectives.
- Propose new mechanism classes with bases SPA(r) and POOL(τ)
- Quantify value of scale information and competition.
- Distribution classes don't matter but mechanism classes do matter!
- Biggest gap is between all versus standard mechanisms: **in robust settings, should sometimes allocate to the non-highest!**
- Broader agenda: **robust mechanism design with partial information**

Appendix

Maximin ratio over different $\mathcal{M}$ for $n = 4$

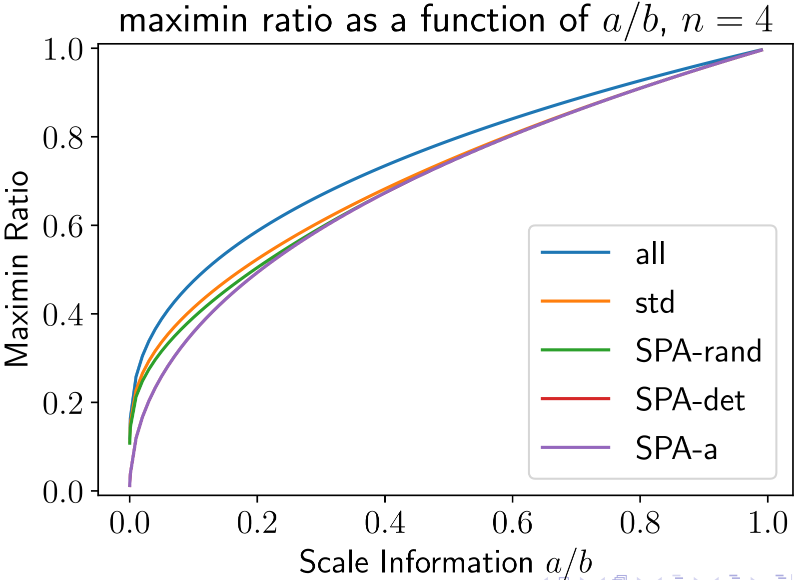

Maximin ratio over different $\mathcal{M}$ for $n = 2$

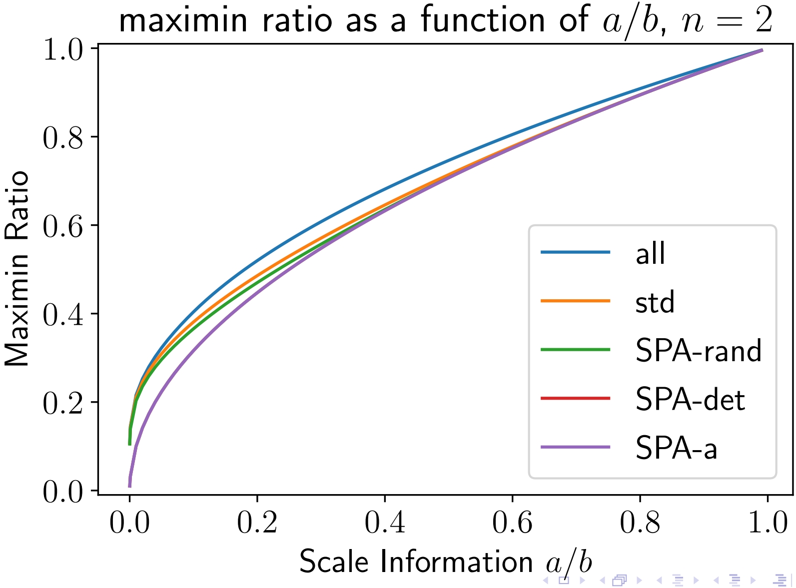

Maximin ratio over different $\mathcal{M}$ for $n = 3$

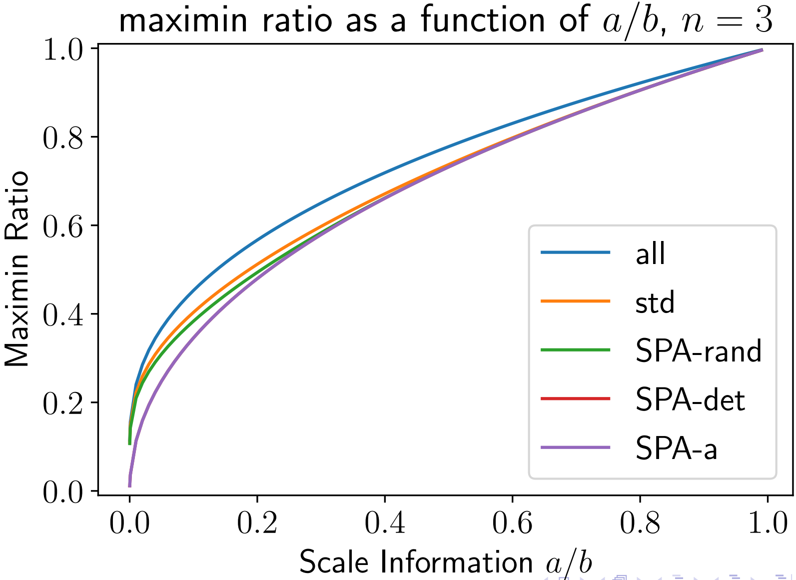

Maximin ratio over different $\mathcal{M}$ for $n = 1$