The Best of Many Robustness Criteria in Decision Making: Formulation and Application to Robust Pricing

Jerry Anunrojwong*

Columbia Business School

*with Santiago R. Balseiro and Omar Besbes

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria** (i.e. what do we mean when we say robust?) $\Rightarrow$ focus of this work!

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria** (i.e. what do we mean when we say robust?) $\Rightarrow$ focus of this work!

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria** (i.e. what do we mean when we say robust?) $\Rightarrow$ focus of this work!

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria** (i.e. what do we mean when we say robust?) $\Rightarrow$ focus of this work!

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria** (i.e. what do we mean when we say robust?) $\Rightarrow$ focus of this work!

Motivation: Decision Making Under Uncertainty

- If we know everything: stochastic optimization
- More commonly: we know something but not everything
- Incorporate this partial information in decision making while being "robust" to things we don't know ...
- Hence, **robust optimization**!
  - A lot of work on robust optimization under *different environments* (uncertainty descriptions, tractable formulations)
  - Less so on **robust optimality criteria**
    (i.e. what do we mean when we say robust?)
    $\Rightarrow$ focus of this work!

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

What do we mean by "robustly optimal"? 3 Criteria

Definitions often anchored around axioms ...

- Criterion 1: maximin performance (e.g. revenue)
  - Wald (1945), common in the robust OR literature
- Criterion 2: minimax regret
  - Savage (1951), compared to a benchmark, less "conservative"
- Criterion 3: maximin ratio
  - compared to a benchmark, multiplicatively
- In some settings, all criteria are "reasonable" and well-founded.
- In practice, we need one decision. Which one?
  (Do you have to choose?)

Research Questions

- Question 1: How good is a prescription derived from one robustness criterion when evaluated against another robustness criterion?
- Question 2: Does there exist a prescription that performs well under all robustness criteria of interest?

  i.e., can one have **the best of the three focal robustness criteria**?
- Our first step: robust pricing
  (fundamental + well studied **separately** under 3 criteria)

Takeaways: Robust Pricing

Q1: cross-criteria performance not good, Q2: uniformly robust good.

Research Questions

- Question 1: How good is a prescription derived from one robustness criterion when evaluated against another robustness criterion?
- Question 2: Does there exist a prescription that performs well under all robustness criteria of interest?

  i.e., can one have **the best of the three focal robustness criteria**?
- Our first step: robust pricing
  (fundamental + well studied **separately** under 3 criteria)

Takeaways: Robust Pricing

Q1: cross-criteria performance not good, Q2: uniformly robust good.

Research Questions

- Question 1: How good is a prescription derived from one robustness criterion when evaluated against another robustness criterion?
- Question 2: Does there exist a prescription that performs well under all robustness criteria of interest?

  i.e., can one have **the best of the three focal robustness criteria**?
- Our first step: robust pricing
  (fundamental + well studied **separately** under 3 criteria)

Takeaways: Robust Pricing

Q1: cross-criteria performance not good, Q2: uniformly robust good.

Research Questions

- Question 1: How good is a prescription derived from one robustness criterion when evaluated against another robustness criterion?
- Question 2: Does there exist a prescription that performs well under all robustness criteria of interest?

  i.e., can one have **the best of the three focal robustness criteria**?
- Our first step: robust pricing
  (fundamental + well studied **separately** under 3 criteria)

Takeaways: Robust Pricing

Q1: cross-criteria performance not good, Q2: uniformly robust good.

Related Work

- robust pricing
  - Bergemann and Schlag (2008), Eren and Maglaras (2010), Carrasco et al. (2018), Wang et al. (2024), Chen et al. (2022), ...
- robust decision making under uncertainty
  - Wald (1945) for maximin revenue
  - Savage (1951) for minimax regret
  - Borodin and El-Yaniv (2005) for maximin ratio
- robust optimization – multiple objectives
  - Iancu and Trichakis (2014) for Pareto efficiency
  - Armbruster and Delage (2015) for uncertain utilities

Problem Formulation

Robust Pricing

- A seller wants to sell an item to a buyer, valuation $\sim F$.
- If the seller knows F precisely, deterministic posted-price mech is optimal $\mathrm{OPT}(F) = \max_p p\bar{F}(p)$.
- Here, the seller does not know F, only that $F \in \mathscr{F}$
  - support, moments, quantiles

Robust Pricing

- A seller wants to sell an item to a buyer, valuation $\sim F$.
- If the seller knows F precisely, deterministic posted-price mech is optimal $\mathsf{OPT}(F) = \max_p p\bar{F}(p)$.
- Here, the seller does not know F, only that $F \in \mathscr{F}$
  - support, moments, quantiles

Robust Pricing

- A seller wants to sell an item to a buyer, valuation $\sim F$.
- If the seller knows F precisely, deterministic posted-price mech is optimal $\mathsf{OPT}(F) = \max_p p\bar{F}(p)$.
- Here, the seller does not know F, only that $F \in \mathscr{F}$
  - support, moments, quantiles

Robust Criteria

- By Myerson, any incentive-compatible mechanism is a price CDF Φ.
- Given a mechanism Φ and a distribution F, define

$$\text{Revenue}(\Phi, F) := \mathbb{E}_{p\sim\Phi, v\sim F}\left[p\mathbf{1}(v \geq p)\right] = \int\int_{s\leq v} s d\Phi(s) dF(v),$$

$$\text{Regret}(\Phi, F) := \text{OPT}(F) - \text{Revenue}(\Phi, F),$$

$$\text{Ratio}(\Phi, F) := \frac{\text{Revenue}(\Phi, F)}{\text{OPT}(F)}.$$

- We want revenue high, regret low, ratio high. Optimal values are:

$$\theta^*_{\text{Revenue}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Revenue}(\Phi, F),$$

$$\theta^*_{\text{Regret}}(\mathscr{F}) := \min_{\Phi\in\mathscr{M}} \max_{F\in\mathscr{F}} \text{Regret}(\Phi, F),$$

$$\theta^*_{\text{Ratio}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Ratio}(\Phi, F).$$

Robust Criteria

- By Myerson, any incentive-compatible mechanism is a price CDF Φ.
- Given a mechanism Φ and a distribution F, define

$$\text{Revenue}(\Phi, F) := \mathbb{E}_{p\sim\Phi, v\sim F}\left[p\mathbf{1}(v \geq p)\right] = \int \int_{s\leq v} s d\Phi(s) dF(v),$$

$$\text{Regret}(\Phi, F) := \text{OPT}(F) - \text{Revenue}(\Phi, F),$$

$$\text{Ratio}(\Phi, F) := \frac{\text{Revenue}(\Phi, F)}{\text{OPT}(F)}.$$

- We want revenue high, regret low, ratio high. Optimal values are:

$$\theta^*_{\text{Revenue}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Revenue}(\Phi, F),$$

$$\theta^*_{\text{Regret}}(\mathscr{F}) := \min_{\Phi\in\mathscr{M}} \max_{F\in\mathscr{F}} \text{Regret}(\Phi, F),$$

$$\theta^*_{\text{Ratio}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Ratio}(\Phi, F).$$

Robust Criteria

- By Myerson, any incentive-compatible mechanism is a price CDF Φ.
- Given a mechanism Φ and a distribution F, define

$$\text{Revenue}(\Phi, F) := \mathbb{E}_{p\sim\Phi, v\sim F}\left[p\mathbf{1}(v \geq p)\right] = \int \int_{s\leq v} s d\Phi(s) dF(v),$$

$$\text{Regret}(\Phi, F) := \text{OPT}(F) - \text{Revenue}(\Phi, F),$$

$$\text{Ratio}(\Phi, F) := \frac{\text{Revenue}(\Phi, F)}{\text{OPT}(F)}.$$

- We want revenue high, regret low, ratio high. Optimal values are:

$$\theta^*_{\text{Revenue}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Revenue}(\Phi, F),$$

$$\theta^*_{\text{Regret}}(\mathscr{F}) := \min_{\Phi\in\mathscr{M}} \max_{F\in\mathscr{F}} \text{Regret}(\Phi, F),$$

$$\theta^*_{\text{Ratio}}(\mathscr{F}) := \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Ratio}(\Phi, F).$$

Robust Criteria

- By Myerson, any incentive-compatible mechanism is a price CDF Φ.
- Given a mechanism Φ and a distribution F, define

$$\begin{aligned}\text{Revenue}(\Phi, F) &:= \mathbb{E}_{p\sim\Phi, v\sim F}\left[p\mathbf{1}(v \geq p)\right] = \int\int_{s\leq v} s d\Phi(s) dF(v),\\ \text{Regret}(\Phi, F) &:= \text{OPT}(F) - \text{Revenue}(\Phi, F),\\ \text{Ratio}(\Phi, F) &:= \frac{\text{Revenue}(\Phi, F)}{\text{OPT}(F)}.\end{aligned}$$

- We want revenue high, regret low, ratio high. Optimal values are:

$$\begin{aligned}\theta^*_{\text{Revenue}}(\mathscr{F}) &:= \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Revenue}(\Phi, F),\\ \theta^*_{\text{Regret}}(\mathscr{F}) &:= \min_{\Phi\in\mathscr{M}} \max_{F\in\mathscr{F}} \text{Regret}(\Phi, F),\\ \theta^*_{\text{Ratio}}(\mathscr{F}) &:= \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Ratio}(\Phi, F).\end{aligned}$$

Robust Criteria

- By Myerson, any incentive-compatible mechanism is a price CDF Φ.
- Given a mechanism Φ and a distribution F, define

$$\begin{aligned}\text{Revenue}(\Phi, F) &:= \mathbb{E}_{p\sim\Phi, v\sim F}\left[p\mathbf{1}(v \geq p)\right] = \int\int_{s\leq v} s d\Phi(s) dF(v),\\ \text{Regret}(\Phi, F) &:= \text{OPT}(F) - \text{Revenue}(\Phi, F),\\ \text{Ratio}(\Phi, F) &:= \frac{\text{Revenue}(\Phi, F)}{\text{OPT}(F)}.\end{aligned}$$

- We want revenue high, regret low, ratio high. Optimal values are:

$$\begin{aligned}\theta^*_{\text{Revenue}}(\mathscr{F}) &:= \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Revenue}(\Phi, F),\\ \theta^*_{\text{Regret}}(\mathscr{F}) &:= \min_{\Phi\in\mathscr{M}} \max_{F\in\mathscr{F}} \text{Regret}(\Phi, F),\\ \theta^*_{\text{Ratio}}(\mathscr{F}) &:= \max_{\Phi\in\mathscr{M}} \min_{F\in\mathscr{F}} \text{Ratio}(\Phi, F).\end{aligned}$$

Performance Across Robust Criteria

Given that the revenue, regret, ratio all have different scales, we evaluate *relative* performance across criteria.

Fix an uncertainty set $\mathscr{F}$.

The relative performances of a given mechanism Φ are

$$\mathrm{RelPerf}(\Phi, \mathrm{Revenue}, \mathscr{F}) = \frac{\mathrm{WorstRevenue}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Revenue}}},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Regret}, \mathscr{F}) = \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{WorstRegret}(\Phi, \mathscr{F})},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Ratio}, \mathscr{F}) = \frac{\mathrm{WorstRatio}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Ratio}}}.$$

Performance Across Robust Criteria

Given that the revenue, regret, ratio all have different scales, we evaluate *relative* performance across criteria.

Fix an uncertainty set $\mathscr{F}$.

The relative performances of a given mechanism Φ are

$$\mathrm{RelPerf}(\Phi, \mathrm{Revenue}, \mathscr{F}) = \frac{\mathrm{WorstRevenue}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Revenue}}},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Regret}, \mathscr{F}) = \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{WorstRegret}(\Phi, \mathscr{F})},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Ratio}, \mathscr{F}) = \frac{\mathrm{WorstRatio}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Ratio}}}.$$

Performance Across Robust Criteria

Given that the revenue, regret, ratio all have different scales, we evaluate *relative* performance across criteria.

Fix an uncertainty set $\mathscr{F}$.
The relative performances of a given mechanism Φ are

$$\mathrm{RelPerf}(\Phi, \mathrm{Revenue}, \mathscr{F}) = \frac{\mathrm{WorstRevenue}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Revenue}}},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Regret}, \mathscr{F}) = \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{WorstRegret}(\Phi, \mathscr{F})},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Ratio}, \mathscr{F}) = \frac{\mathrm{WorstRatio}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Ratio}}}.$$

Performance Across Robust Criteria

Given that the revenue, regret, ratio all have different scales, we evaluate *relative* performance across criteria.

Fix an uncertainty set $\mathscr{F}$.
The relative performances of a given mechanism Φ are

$$\mathrm{RelPerf}(\Phi, \mathrm{Revenue}, \mathscr{F}) = \frac{\mathrm{WorstRevenue}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Revenue}}},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Regret}, \mathscr{F}) = \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{WorstRegret}(\Phi, \mathscr{F})},$$

$$\mathrm{RelPerf}(\Phi, \mathrm{Ratio}, \mathscr{F}) = \frac{\mathrm{WorstRatio}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Ratio}}}.$$

Performance Across Robust Criteria

Given that the revenue, regret, ratio all have different scales, we evaluate *relative* performance across criteria.

Fix an uncertainty set $\mathscr{F}$.
The relative performances of a given mechanism Φ are

$$\begin{aligned}
\mathrm{RelPerf}(\Phi, \mathrm{Revenue}, \mathscr{F}) &= \frac{\mathrm{WorstRevenue}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Revenue}}}, \\
\mathrm{RelPerf}(\Phi, \mathrm{Regret}, \mathscr{F}) &= \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{WorstRegret}(\Phi, \mathscr{F})}, \\
\mathrm{RelPerf}(\Phi, \mathrm{Ratio}, \mathscr{F}) &= \frac{\mathrm{WorstRatio}(\Phi, \mathscr{F})}{\theta^*_{\mathrm{Ratio}}}.
\end{aligned}$$

Best of Many Robustness Criteria

Rather than fixing the mechanisms to be one of the focal mechanisms, we directly optimize over all mechanisms.

The relative performance of Φ over all criteria is

$$\begin{aligned}
&\mathsf{RelPerf}(\Phi, \mathsf{All}, \mathscr{F}) \\
&= \min_{F \in \mathscr{F}} \min \left\{ \frac{\mathsf{Revenue}(\Phi, F)}{\theta^*_{\mathsf{Revenue}}}, \frac{\theta^*_{\mathsf{Regret}}}{\mathsf{Regret}(\Phi, F)}, \frac{\mathsf{Ratio}(\Phi, F)}{\theta^*_{\mathsf{Ratio}}} \right\},
\end{aligned}$$

We solve

$$c^*(\mathscr{F}) = \max_{\Phi \in \mathscr{M}} \mathsf{RelPerf}(\Phi, \mathsf{All}, \mathscr{F}).$$

Best of Many Robustness Criteria

Rather than fixing the mechanisms to be one of the focal mechanisms, we directly optimize over all mechanisms.

The relative performance of Φ over all criteria is

$$\begin{aligned}
&\mathsf{RelPerf}(\Phi, \mathsf{All}, \mathscr{F}) \\
&= \min_{F \in \mathscr{F}} \min \left\{ \frac{\mathsf{Revenue}(\Phi, F)}{\theta^*_{\mathsf{Revenue}}}, \frac{\theta^*_{\mathsf{Regret}}}{\mathsf{Regret}(\Phi, F)}, \frac{\mathsf{Ratio}(\Phi, F)}{\theta^*_{\mathsf{Ratio}}} \right\},
\end{aligned}$$

We solve

$$c^*(\mathscr{F}) = \max_{\Phi \in \mathscr{M}} \mathsf{RelPerf}(\Phi, \mathsf{All}, \mathscr{F}).$$

Best of Many Robustness Criteria

Rather than fixing the mechanisms to be one of the focal mechanisms, we directly optimize over all mechanisms.

The relative performance of Φ over all criteria is

$$\begin{aligned}&\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathscr{F})\\&= \min_{F \in \mathscr{F}} \min \left\{ \frac{\mathrm{Revenue}(\Phi, F)}{\theta^*_{\mathrm{Revenue}}}, \frac{\theta^*_{\mathrm{Regret}}}{\mathrm{Regret}(\Phi, F)}, \frac{\mathrm{Ratio}(\Phi, F)}{\theta^*_{\mathrm{Ratio}}} \right\},\end{aligned}$$

We solve

$$c^*(\mathscr{F}) = \max_{\Phi \in \mathscr{M}} \mathrm{RelPerf}(\Phi, \mathrm{All}, \mathscr{F}).$$

Reformulation via Linear Programs

Unifying revenue, regret, and ratio with λ-regret

$$\text{worst case } \lambda\text{-regret of } \Phi \text{ is } = \max_{F \in \mathscr{F}} \left[\lambda \mathrm{OPT}(F) - \mathrm{Revenue}(\Phi, F)\right]$$

$$\text{minimax } \lambda\text{-regret is } = \min_{\Phi \in \mathscr{M}} \max_{F \in \mathscr{F}} \left[\lambda \mathrm{OPT}(F) - \mathrm{Revenue}(\Phi, F)\right]$$

Prop. Revenue is $\lambda = 0$. Regret is $\lambda = 1$. Ratio is λ such that minimax λ-regret is zero.

Unifying revenue, regret, and ratio with λ-regret

$$\text{worst case } \lambda\text{-regret of } \Phi \text{ is } = \max_{F \in \mathscr{F}} \left[\lambda \mathrm{OPT}(F) - \mathrm{Revenue}(\Phi, F)\right]$$

$$\text{minimax } \lambda\text{-regret is } = \min_{\Phi \in \mathscr{M}} \max_{F \in \mathscr{F}} \left[\lambda \mathrm{OPT}(F) - \mathrm{Revenue}(\Phi, F)\right]$$

Prop. Revenue is $\lambda = 0$. Regret is $\lambda = 1$. Ratio is λ such that minimax λ-regret is zero.

λ-regret Linear Program

Fix a mechanism Φ and the uncertainty set to be some known moments m_i and quantiles (q_j, r_j) on a known grid $\mathscr{G} \subseteq [0,1]$:

$$\mathscr{F} = \left\{ F \in \Delta(\mathscr{G}) : \int v^i dF(v) = m_i \ \forall i,\ \bar{F}(r_j) = q_j \ \forall j \right\}$$

Then the worst case λ-regret of Φ is from this LP:

$$\min_{\theta, \alpha(\cdot), \beta(\cdot)} \quad \theta$$

$$\text{s.t. } \theta \geq \sum_{i \in \mathscr{I}} \alpha_i(p) m_i + \sum_{j \in \mathscr{J}} \beta_j(p) q_j \quad \forall p \in \mathscr{G}$$

$$\lambda p \mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_i(p) v^i - \sum_{j \in \mathscr{J}} \beta_j(p) \mathbf{1}(v \geq r_j) \leq 0$$

λ-regret Linear Program

Fix a mechanism Φ and the uncertainty set to be some known moments m_i and quantiles (q_j, r_j) on a known grid $\mathscr{G} \subseteq [0,1]$:

$$\mathscr{F} = \left\{ F \in \Delta(\mathscr{G}) : \int v^i dF(v) = m_i \ \forall i,\ \bar{F}(r_j) = q_j \ \forall j \right\}$$

Then the worst case λ-regret of Φ is from this LP:

$$\begin{aligned}
\min_{\theta,\alpha(\cdot),\beta(\cdot)} \quad & \theta \\
\text{s.t.} \quad & \theta \geq \sum_{i \in \mathscr{I}} \alpha_i(p) m_i + \sum_{j \in \mathscr{J}} \beta_j(p) q_j \quad \forall p \in \mathscr{G} \\
& \lambda p \mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_i(p) v^i - \sum_{j \in \mathscr{J}} \beta_j(p) \mathbf{1}(v \geq r_j) \leq 0
\end{aligned}$$

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

$$\min_{F \in \mathscr{F}} \quad \lambda \mathsf{OPT}(F) - \mathsf{Revenue}(\Phi, F)$$

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

$$\min_{F \in \mathscr{F}} \quad \lambda \quad \underbrace{\mathsf{OPT}(F)}_{\max_p p \int \mathbf{1}(v \geq p) dF(v)} \quad - \quad \underbrace{\mathsf{Revenue}(\Phi, F)}_{\int \int_{s \leq v} s d\Phi(s) dF(v)}$$

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

$$\min_{F \in \mathscr{F}} \max_{p} \quad \lambda p \mathbf{1}(v \geq p) dF(v) - \int \int_{s \leq v} s d\Phi(s) dF(v)$$

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

$$\min_{\theta} \quad \theta \text{ s.t.}$$

$$\theta \geq \max_{F \in \mathscr{F}} \lambda p \mathbf{1}(v \geq p) dF(v) - \int \int_{s \leq v} s d\Phi(s) dF(v) \quad \forall p$$

Proof of λ-regret Linear Program

LP duality for each price p of OPT, then epigraph.

$$\min_{F \in \mathscr{F}, \theta} \quad \theta \text{ s.t.}$$

$$\theta \geq \max_{F \in \mathscr{F}} \lambda p \mathbf{1}(v \geq p) dF(v) - \int \int_{s \leq v} s d\Phi(s) dF(v) \quad \forall p$$

Because $\mathscr{F}$ is linear, we dualize the LP of $\max_{F \in \mathscr{F}}$ for each p with dual variables $\alpha_i(p)$ for moments, $\beta_j(p)$ for quantiles, giving our results.

Cross-Criterion Performance

How does a mechanism Φ that is robustly optimized to one criterion (e.g. regret) performs under another (e.g. ratio)?

We can solve an LP to get a specific Φ for the first criterion, and solve another LP to evaluate it under the second. But there might be many possible optimal solutions ...

We take the *worst case approach*. Because all 3 criteria can be written as λ-regret, we solve the *cross regret* problem:

$$R^*_{\lambda_{\mathsf{new}}}(\mathscr{F}, (\lambda_{\mathsf{old}}, r_{\mathsf{old}})) = \min_{\Phi \in \mathscr{M}_{\mathsf{old}}} R^{\Phi}_{\lambda_{\mathsf{new}}}(\mathscr{F})$$

$$\text{where } \mathscr{M}_{\mathsf{old}} = \left\{ \Phi : R^{\Phi}_{\lambda_{\mathsf{old}}}(\mathscr{F}) \le r_{\mathsf{old}} \right\}.$$

Using the LP formulation for λ-regret we have before, this is also an LP.

Cross-Criterion Performance

How does a mechanism Φ that is robustly optimized to one criterion (e.g. regret) performs under another (e.g. ratio)?

We can solve an LP to get a specific Φ for the first criterion, and solve another LP to evaluate it under the second. But there might be many possible optimal solutions ...

We take the *worst case approach*. Because all 3 criteria can be written as λ-regret, we solve the *cross regret* problem:

$$R^*_{\lambda_{\text{new}}}(\mathscr{F}, (\lambda_{\text{old}}, r_{\text{old}})) = \min_{\Phi \in \mathscr{M}_{\text{old}}} R^{\Phi}_{\lambda_{\text{new}}}(\mathscr{F})$$
$$\text{where } \mathscr{M}_{\text{old}} = \left\{ \Phi : R^{\Phi}_{\lambda_{\text{old}}}(\mathscr{F}) \le r_{\text{old}} \right\}.$$

Using the LP formulation for λ-regret we have before, this is also an LP.

Cross-Criterion Performance

How does a mechanism Φ that is robustly optimized to one criterion (e.g. regret) performs under another (e.g. ratio)?

We can solve an LP to get a specific Φ for the first criterion, and solve another LP to evaluate it under the second. But there might be many possible optimal solutions ...

We take the *worst case approach*. Because all 3 criteria can be written as λ-regret, we solve the *cross regret* problem:

$$R^*_{\lambda_{\text{new}}}(\mathscr{F}, (\lambda_{\text{old}}, r_{\text{old}})) = \min_{\Phi \in \mathscr{M}_{\text{old}}} R^{\Phi}_{\lambda_{\text{new}}}(\mathscr{F})$$

$$\text{where } \mathscr{M}_{\text{old}} = \left\{ \Phi : R^{\Phi}_{\lambda_{\text{old}}}(\mathscr{F}) \leq r_{\text{old}} \right\}.$$

Using the LP formulation for λ-regret we have before, this is also an LP.

Cross-Criterion Performance

How does a mechanism Φ that is robustly optimized to one criterion (e.g. regret) performs under another (e.g. ratio)?

We can solve an LP to get a specific Φ for the first criterion, and solve another LP to evaluate it under the second. But there might be many possible optimal solutions ...

We take the *worst case approach*. Because all 3 criteria can be written as λ-regret, we solve the *cross regret* problem:

$$R^*_{\lambda_{\text{new}}}(\mathscr{F}, (\lambda_{\text{old}}, r_{\text{old}})) = \min_{\Phi \in \mathscr{M}_{\text{old}}} R^{\Phi}_{\lambda_{\text{new}}}(\mathscr{F})$$
$$\text{where } \mathscr{M}_{\text{old}} = \left\{\Phi : R^{\Phi}_{\lambda_{\text{old}}}(\mathscr{F}) \le r_{\text{old}}\right\}.$$

Using the LP formulation for λ-regret we have before, this is also an LP.

Cross-Criterion Performance

How does a mechanism Φ that is robustly optimized to one criterion (e.g. regret) performs under another (e.g. ratio)?

We can solve an LP to get a specific Φ for the first criterion, and solve another LP to evaluate it under the second. But there might be many possible optimal solutions ...

We take the *worst case approach*. Because all 3 criteria can be written as λ-regret, we solve the *cross regret* problem:

$$R^*_{\lambda_{\mathsf{new}}}(\mathscr{F}, (\lambda_{\mathsf{old}}, r_{\mathsf{old}})) = \min_{\Phi \in \mathscr{M}_{\mathsf{old}}} R^{\Phi}_{\lambda_{\mathsf{new}}}(\mathscr{F})$$

$$\text{where } \mathscr{M}_{\mathsf{old}} = \left\{ \Phi : R^{\Phi}_{\lambda_{\mathsf{old}}}(\mathscr{F}) \leq r_{\mathsf{old}} \right\}.$$

Using the LP formulation for λ-regret we have before, this is also an LP.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0, 1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$

$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$

$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0,1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$

$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$

$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0,1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$

$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$

$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0,1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$

$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$

$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0,1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$

$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$

$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0, 1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$
$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$
$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. ($\theta^*_{\text{Regret}} =$ optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

LP for the Uniformly Robust Mechanism

We want the highest $c \in [0,1]$ s.t. there exists a mechanism Φ that is a factor of c from optimal θ^* for all 3 criteria:

$$\min_{F \in \mathscr{F}} \text{Revenue}(\Phi, F) \geq c \cdot \theta^*_{\text{Revenue}}$$
$$\max_{F \in \mathscr{F}} \text{Regret}(\Phi, F) \leq \theta^*_{\text{Regret}}/c$$
$$\min_{F \in \mathscr{F}} \text{Ratio}(\Phi, F) \geq c \cdot \theta^*_{\text{Ratio}}.$$

Fix c. Compute the θ^*'s first. (θ^*_{Regret} = optimal minimax regret, etc.)

For each constraint, rewrite it as λ-regret and use LP duality as before $\rightarrow$ LP feasibility problem.

Find largest feasible c by line search.

Results

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ

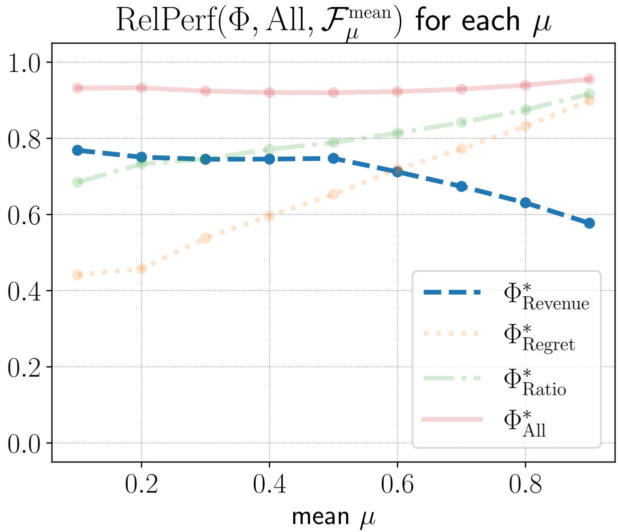

$\mathscr{F}$ = KNOWN SUPPORT $[0, 1]$, MEAN μ

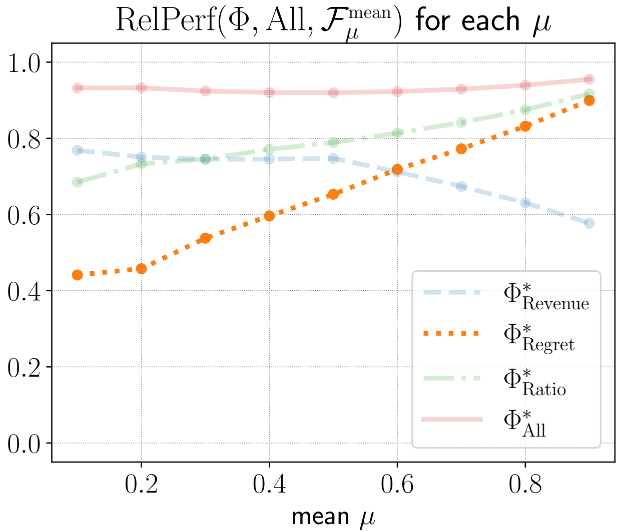

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ

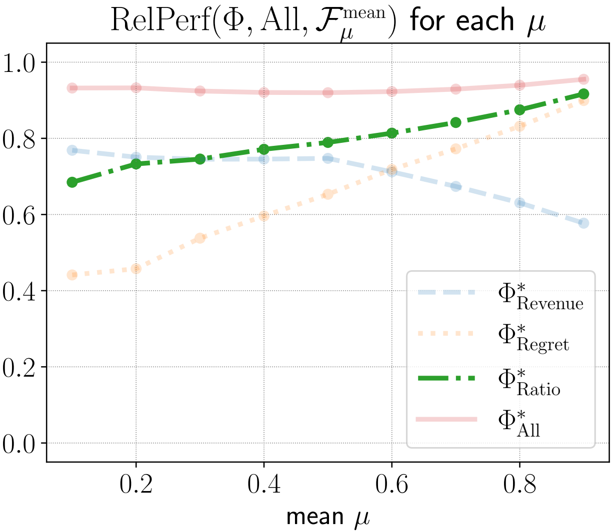

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ

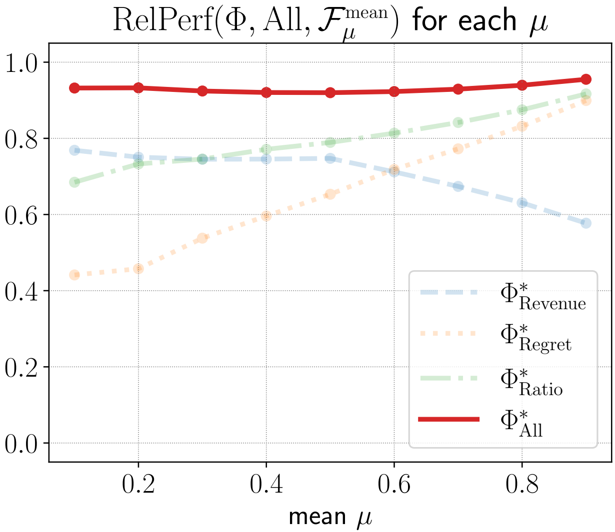

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ

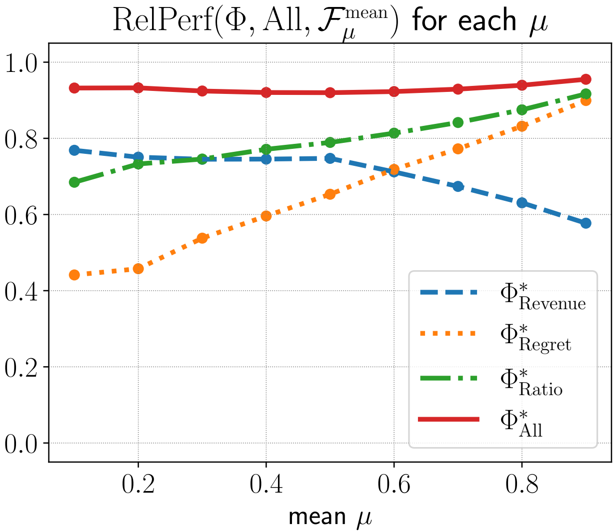

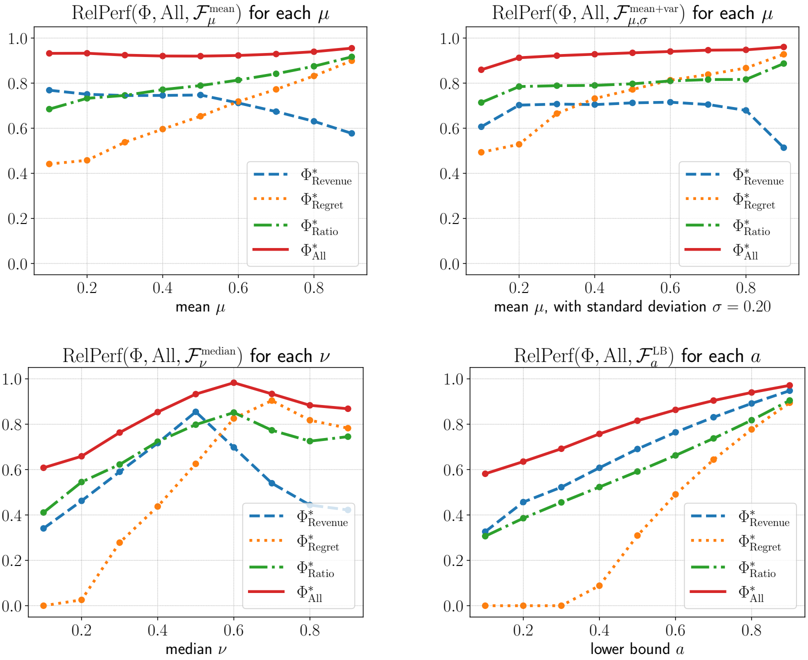

$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\mu}^{\mathrm{mean}})$ for each μ
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
mean μ
$\Phi^*_{\mathrm{Revenue}}$
Φ^*_{Regret}
Φ^*_{Ratio}
Φ^*_{All}
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\mu,\sigma}^{\mathrm{mean+var}})$ for each μ
mean μ, with standard deviation $\sigma = 0.20$
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\nu}^{\mathrm{median}})$ for each ν
median ν
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{a}^{\mathrm{LB}})$ for each a
lower bound a

Price CDFs (Mechanisms) for mean $\mu = 0.5$

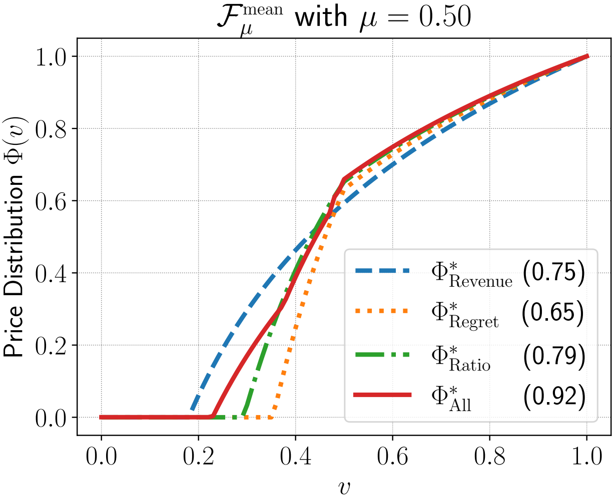

Worst case relative performances across instances

Additional Information	Uniformly Robust Mechanism	Focal Mechanisms		
		revenue	regret	ratio
mean	**92%**	58%	44%	**68%**
mean and variance	**86%**	51%	49%	**71%**
median	**61%**	34%	0%	**41%**
lower bound	**58%**	**33%**	0%	31%

Worst case relative performances across instances

Additional Information	Uniformly Robust Mechanism	Focal Mechanisms		
		revenue	regret	ratio
mean	**92%**	58%	44%	**68%**
mean and variance	**86%**	51%	49%	**71%**
median	**61%**	34%	0%	**41%**
lower bound	**58%**	**33%**	0%	31%

Worst case relative performances across instances

Additional Information	Uniformly Robust Mechanism	Focal Mechanisms		
		revenue	regret	ratio
mean	**92%**	58%	44%	**68%**
mean and variance	**86%**	51%	49%	**71%**
median	**61%**	34%	0%	**41%**
lower bound	**58%**	**33%**	0%	31%

Worst case relative performances across instances

Additional Information	Uniformly Robust Mechanism	Focal Mechanisms		
		revenue	regret	ratio
mean	**92%**	58%	44%	**68%**
mean and variance	**86%**	51%	49%	**71%**
median	**61%**	34%	0%	**41%**
lower bound	**58%**	**33%**	0%	31%

Conclusion

- We propose general formulations to investigate **criteria-overfitting** in robust decision making.
- We analyze the case of **robust pricing** – nontrivial LPs.
- Mechanisms robust for one criterion can perform badly under another criterion.
- There exist "uniformly robust" mechanisms that perform almost as good as the best across criteria.
- Results depend on the structure of the problem
  – investigating these systematically is future work!

Conclusion

- We propose general formulations to investigate **criteria-overfitting** in robust decision making.
- We analyze the case of **robust pricing** – nontrivial LPs.
- Mechanisms robust for one criterion can perform badly under another criterion.
- There exist "uniformly robust" mechanisms that perform almost as good as the best across criteria.
- Results depend on the structure of the problem – investigating these systematically is future work!

Conclusion

- We propose general formulations to investigate **criteria-overfitting** in robust decision making.
- We analyze the case of **robust pricing** – nontrivial LPs.
- Mechanisms robust for one criterion can perform badly under another criterion.
- There exist "uniformly robust" mechanisms that perform almost as good as the best across criteria.
- Results depend on the structure of the problem – investigating these systematically is future work!

Conclusion

- We propose general formulations to investigate **criteria-overfitting** in robust decision making.
- We analyze the case of **robust pricing** – nontrivial LPs.
- Mechanisms robust for one criterion can perform badly under another criterion.
- There exist "uniformly robust" mechanisms that perform almost as good as the best across criteria.
- Results depend on the structure of the problem
  – investigating these systematically is future work!

Conclusion

- We propose general formulations to investigate **criteria-overfitting** in robust decision making.
- We analyze the case of **robust pricing** – nontrivial LPs.
- Mechanisms robust for one criterion can perform badly under another criterion.
- There exist "uniformly robust" mechanisms that perform almost as good as the best across criteria.
- Results depend on the structure of the problem – investigating these systematically is future work!

Appendix

Worst Case Cross-Criterion Performances

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ

Mean		Evaluated Criterion: revenue	regret	ratio
Mechanism Criterion	revenue	1.00	0.58	0.75
	regret	0.45	1.00	0.44
	ratio	0.68	0.93	1.00

Table: The uncertainty set is known mean information $\mathscr{F}_{\mu}^{\text{mean}}$. Each cell is $\min_{\mu \in \mathbb{G}} \text{RelPerf}(\Phi^*_{\text{MechanismCriterion}}(\mathscr{F}_{\mu}^{\text{mean}}), \text{EvaluatedMetric}, \mathscr{F}_{\mu}^{\text{mean}})$, the performance of a mechanism optimized for one criterion (row) when evaluated under another criterion (column).

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEAN μ, STDEV $\sigma = 0.20$

Mean & Variance		Evaluated Criterion		
		revenue	regret	ratio
Mechanism Criterion	revenue	1.00	0.51	0.67
	regret	0.49	1.00	0.58
	ratio	0.78	0.71	1.00

Table: The uncertainty set is known first two moments (mean and variance) information $\mathscr{F}_{\mu,\sigma}^{\text{mean+var}}$. Each cell is $\min_{\mu,\sigma\in\mathbb{G}} \text{RelPerf}(\Phi^*_{\text{MechanismCriterion}}(\mathscr{F}_{\mu,\sigma}^{\text{mean+var}}), \text{EvaluatedMetric}, \mathscr{F}_{\mu,\sigma}^{\text{mean+var}})$, the performance of a mechanism optimized for one criterion (row) when evaluated under another criterion (column).

$\mathscr{F}$ = KNOWN SUPPORT $[0,1]$, MEDIAN ν

Median		Evaluated Criterion: revenue	regret	ratio
Mechanism Criterion	revenue	1.00	0.41	0.34
	regret	0.00	1.00	0.00
	ratio	0.41	0.52	1.00

Table: The uncertainty set is known median information $\mathscr{F}_\nu^{\text{median}}$. Each cell is $\min_{\nu \in \mathbb{G}} \text{RelPerf}(\Phi^*_{\text{MechanismCriterion}}(\mathscr{F}_\nu^{\text{median}}), \text{EvaluatedMetric}, \mathscr{F}_\nu^{\text{median}})$, the performance of a mechanism optimized for one criterion (row) when evaluated under another criterion (column).

$\mathscr{F}$ = KNOWN SUPPORT $[a, 1]$

Lower Bound		Evaluated Criterion: revenue	regret	ratio
Mechanism Criterion	revenue	1.00	0.41	0.33
	regret	0.00	1.00	0.00
	ratio	0.31	0.53	1.00

Table: The uncertainty set is known lower bound information $\mathscr{F}_a^{\mathsf{LB}}$. Each cell is $\min_{a \in \mathbb{G}} \mathsf{RelPerf}(\Phi^*_{\mathsf{MechanismCriterion}}(\mathscr{F}_a^{\mathsf{LB}}), \mathsf{EvaluatedMetric}, \mathscr{F}_a^{\mathsf{LB}})$, the performance of a mechanism optimized for one criterion (row) when evaluated under another criterion (column).

Randomized Mechanisms

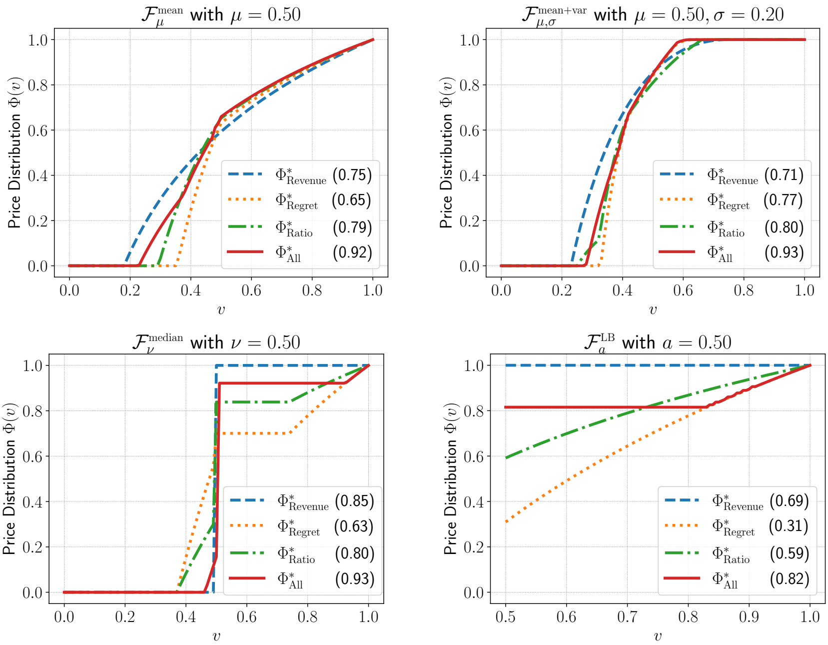
$\mathcal{F}_{\mu}^{\text{mean}}$ with $\mu = 0.50$
$\mathcal{F}_{\mu,\sigma}^{\text{mean+var}}$ with $\mu = 0.50, \sigma = 0.20$
$\mathcal{F}_{\nu}^{\text{median}}$ with $\nu = 0.50$
$\mathcal{F}_{a}^{\text{LB}}$ with $a = 0.50$
Price Distribution $\Phi(v)$
v
Φ^*_{Revenue} (0.75)
Φ^*_{Regret} (0.65)
Φ^*_{Ratio} (0.79)
Φ^*_{All} (0.92)
Φ^*_{Revenue} (0.71)
Φ^*_{Regret} (0.77)
Φ^*_{Ratio} (0.80)
Φ^*_{All} (0.93)
Φ^*_{Revenue} (0.85)
Φ^*_{Regret} (0.63)
Φ^*_{Ratio} (0.80)
Φ^*_{All} (0.93)
Φ^*_{Revenue} (0.69)
Φ^*_{Regret} (0.31)
Φ^*_{Ratio} (0.59)
Φ^*_{All} (0.82)

Deterministic Pricing Results
(Relative Performances Across Parameter Values)

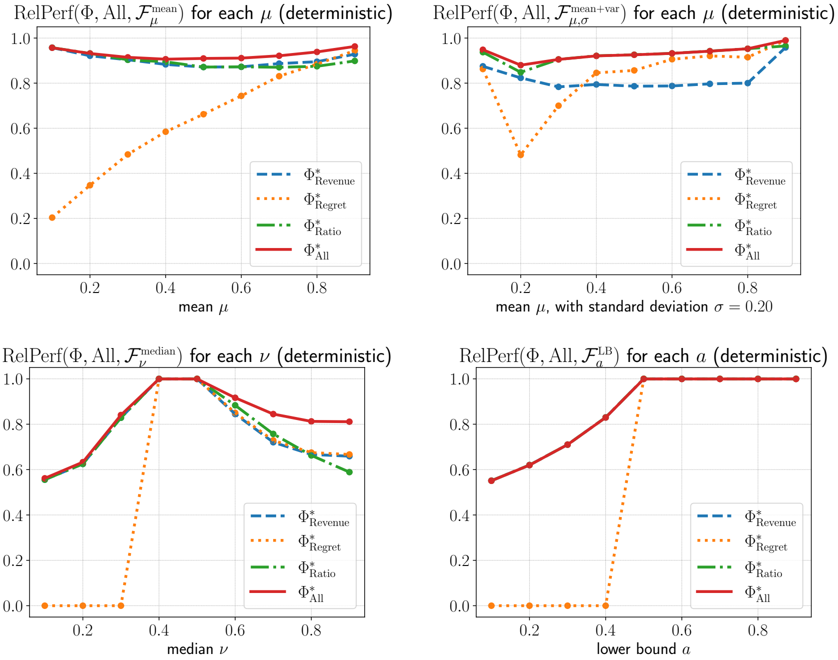

$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\mu}^{\mathrm{mean}})$ for each μ (deterministic)
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
mean μ
$\Phi^*_{\mathrm{Revenue}}$
Φ^*_{Regret}
Φ^*_{Ratio}
Φ^*_{All}
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\mu,\sigma}^{\mathrm{mean+var}})$ for each μ (deterministic)
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
mean μ, with standard deviation $\sigma = 0.20$
$\Phi^*_{\mathrm{Revenue}}$
Φ^*_{Regret}
Φ^*_{Ratio}
Φ^*_{All}
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{\nu}^{\mathrm{median}})$ for each ν (deterministic)
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
median ν
$\Phi^*_{\mathrm{Revenue}}$
Φ^*_{Regret}
Φ^*_{Ratio}
Φ^*_{All}
$\mathrm{RelPerf}(\Phi, \mathrm{All}, \mathcal{F}_{a}^{\mathrm{LB}})$ for each a (deterministic)
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
lower bound a
$\Phi^*_{\mathrm{Revenue}}$
Φ^*_{Regret}
Φ^*_{Ratio}
Φ^*_{All}

Deterministic Pricing Results
(Optimal Prices Across Parameter Values)

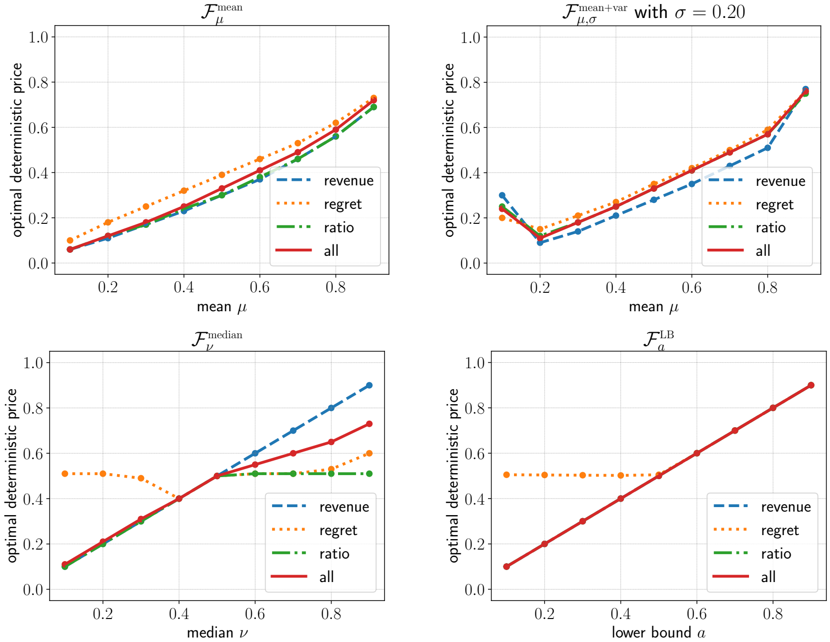

$\mathcal{F}_{\mu}^{\text{mean}}$
$\mathcal{F}_{\mu,\sigma}^{\text{mean+var}}$ with $\sigma = 0.20$
$\mathcal{F}_{\nu}^{\text{median}}$
$\mathcal{F}_{a}^{\text{LB}}$
optimal deterministic price
mean μ
median ν
lower bound a
revenue
regret
ratio
all

Omitted LP Formulations

Cross-Regret LP

$$
\begin{aligned}
\min_{\substack{\Phi,\theta \\ \alpha_{\text{new}},\beta_{\text{new}} \\ \alpha_{\text{old}},\beta_{\text{old}}}} \quad & \theta \\
\text{s.t.} \quad & \theta \geq \sum_{i \in \mathscr{I}} \alpha_{\text{new},i}(p) m_i + \sum_{j \in \mathscr{J}} \beta_{\text{new},j}(p) q_j \quad \forall p \in \mathscr{G} \\
& \lambda_{\text{new}} p \mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_{\text{new},i}(p) v^i - \sum_{j \in \mathscr{J}} \beta_{\text{new},j}(p) \mathbf{1}(v \geq r_j) \leq 0 \quad \forall v, p \in \mathscr{G} \\
& r_{\text{old}} \geq \sum_{i \in \mathscr{I}} \alpha_{\text{old},i}(p) m_i + \sum_{j \in \mathscr{J}} \beta_{\text{old},j}(p) q_j \quad \forall p \in \mathscr{G} \\
& \lambda_{\text{old}} p \mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_{\text{old},i}(p) v^i - \sum_{j \in \mathscr{J}} \beta_{\text{old},j}(p) \mathbf{1}(v \geq r_j) \leq 0 \quad \forall v, p \in \mathscr{G} \\
& \Phi \text{ is a CDF.}
\end{aligned}
$$

Uniformly Robust LP Feasibility

There exists a mechanism Φ that achieves $(\theta_{\text{Revenue}}, \theta_{\text{Regret}}, \theta_{\text{Ratio}})$ if and only if the following linear problem in variables $\alpha_{\text{Revenue}}, \alpha_{\text{Regret}}, \alpha_{\text{Ratio}}, \beta_{\text{Revenue}}, \beta_{\text{Regret}}, \beta_{\text{Ratio}}, \Phi$ is feasible:

$$\sum_{i \in \mathscr{I}} \alpha_{\text{Revenue},i}(p) m_i + \sum_{j \in \mathscr{J}} \beta_{\text{Revenue},j}(p) q_j \leq -\theta_{\text{Revenue}} \quad \forall p \in \mathscr{G}$$

$$\sum_{i \in \mathscr{I}} \alpha_{\text{Regret},i}(p) m_i + \sum_{j \in \mathscr{J}} \beta_{\text{Regret},j}(p) q_j \leq \theta_{\text{Regret}} \quad \forall p \in \mathscr{G}$$

$$\sum_{i \in \mathscr{I}} \alpha_{\text{Ratio},i}(p) m_i + \sum_{j \in \mathscr{J}} \beta_{\text{Ratio},j}(p) q_j \leq 0 \quad \forall p \in \mathscr{G}$$

$$-\int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_{\text{Revenue},i}(p) v^i - \sum_{j \in \mathscr{J}} \beta_{\text{Revenue},j}(p) \mathbf{1}(v \leq r_j) \leq 0 \quad \forall v, p \in \mathscr{G}$$

$$p\mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_{\text{Regret},i}(p) v^i - \sum_{j \in \mathscr{J}} \beta_{\text{Regret},j}(p) \mathbf{1}(v \leq r_j) \leq 0 \quad \forall v, p \in \mathscr{G}$$

$$\theta_{\text{Ratio}} p\mathbf{1}(v \geq p) - \int_{s \leq v} s d\Phi(s) - \sum_{i \in \mathscr{I}} \alpha_{\text{Ratio},i}(p) v^i - \sum_{j \in \mathscr{J}} \beta_{\text{Ratio},j}(p) \mathbf{1}(v \leq r_j) \leq 0 \quad \forall v, p \in \mathscr{G}$$

Φ is a CDF,