Carnegie Mellon University

ICML
International Conference
On Machine Learning

Calibrating Decision Robustness via Inverse Conformal Risk Control

Wenbin Zhou[1,2] and Shixiang (Woody) Zhu[1]

[1]*Heinz College of Information Systems and Public Policy*

[2]*Machine Learning Department, School of Computer Science*

Introduction: Robust Optimization

Robust optimization: $\min_{z \in \mathcal{Z}} \max_{y \in U_\lambda} f(z, y)$

Decision variable

Adversarial variable

Robustness parameter

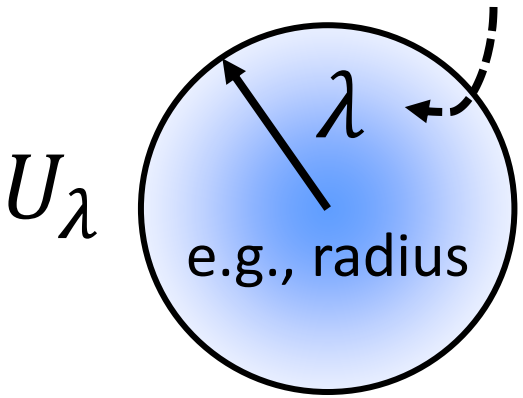

Q: How to select λ? A tradeoff exists:

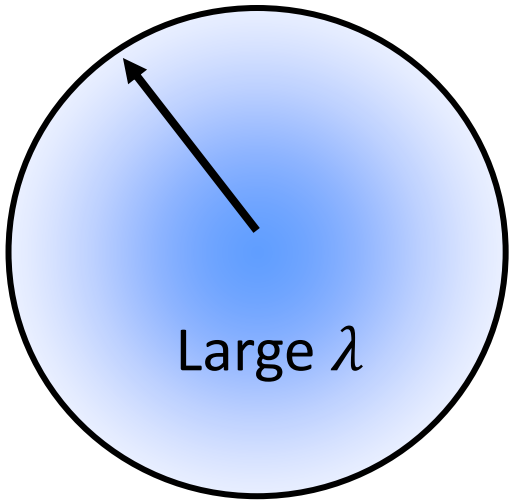

- Covers more plausible adversarial cases (risk ↓)
- May select overly conservative decisions (cost ↑)

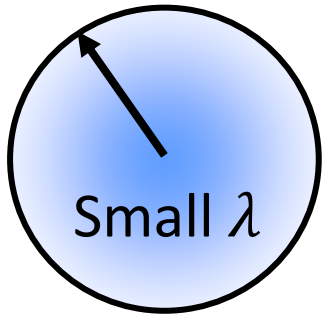

- Covers less plausible adversarial cases (risk ↑)
- Sharper decisions (cost ↓)

Carnegie Mellon University

Introduction: Robust Optimization

How to select λ?

Decision maker

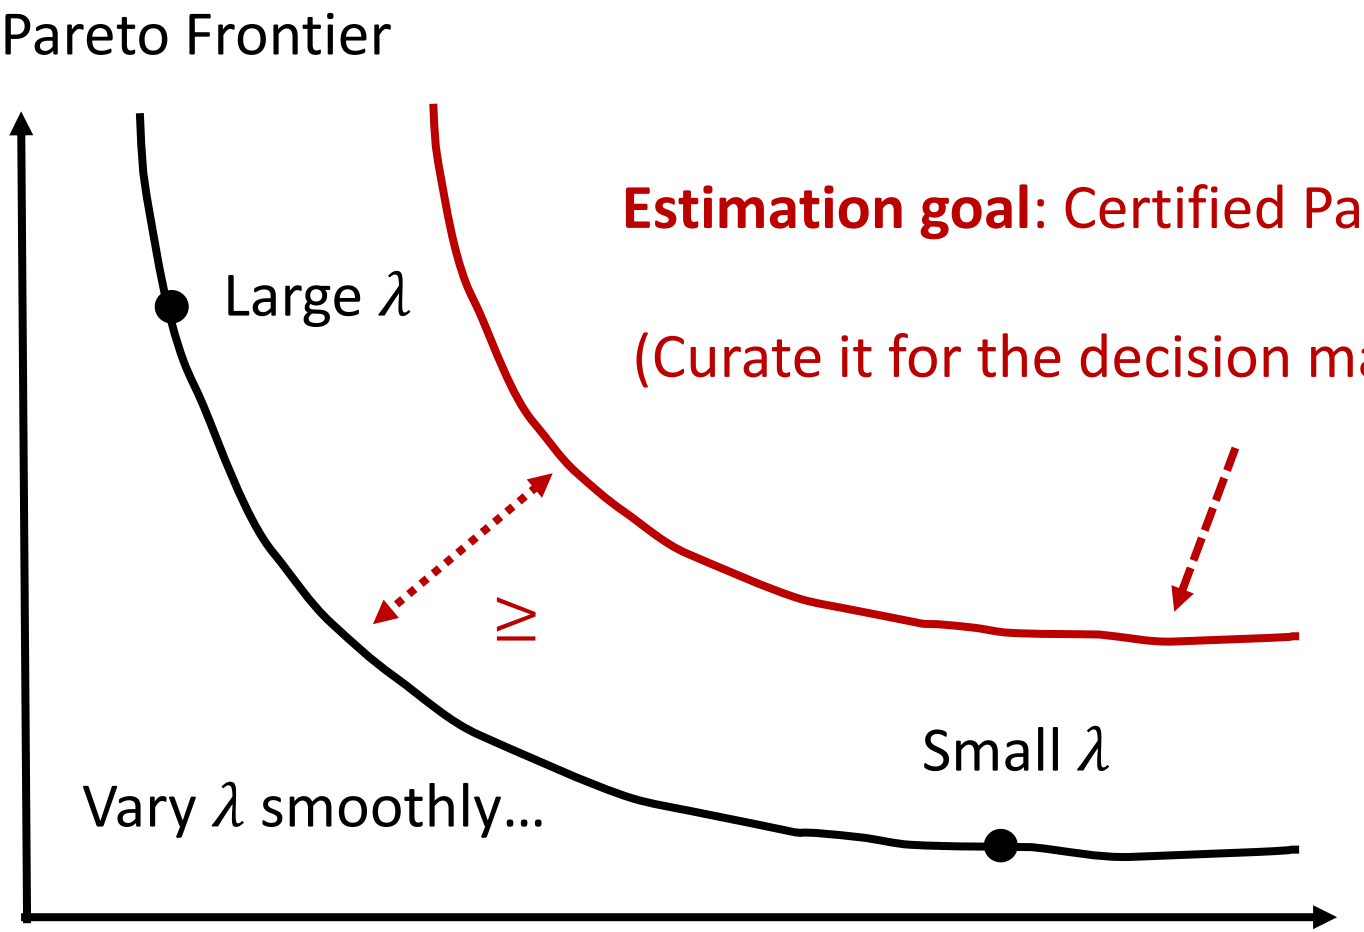

Estimation goal: Certified Pareto Frontier

(Curate it for the decision maker to select λ!)

Problem Setting

Both cost and risk can be formalized by some notion of loss functions:

Cost ← regret = $R_\lambda = f\left(z^{\mathrm{RO}}(\lambda); Y\right) - f(z^\star; Y)$

Risk ← miscoverage = $I_\lambda = \mathbb{1}\{Y \notin U_\lambda\}$

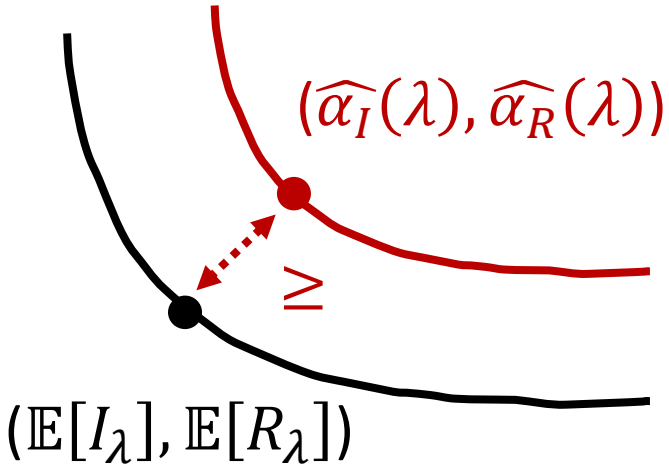

Objective: Certified Pareto Frontier Estimation

For each λ, construct estimators $\widehat{\alpha_I}(\lambda)$ and $\widehat{\alpha_R}(\lambda)$ such that

$$\widehat{\alpha_I}(\lambda) \geq \mathbb{E}[I_\lambda] \quad \text{and} \quad \widehat{\alpha_R}(\lambda) \geq \mathbb{E}[R_\lambda]$$

Carnegie Mellon University

Algorithm: Inverse Conformal Risk Control

Definition: Conformal Risk Control (Angelopoulos et. al., 2022)

Given risk bound α, select loss function parameter λ such that

$$\inf_{\lambda \in \Lambda} \left\{ \lambda : B - \overline{\ell_n}(\lambda) \geq \frac{\lfloor (n+1)(B-\alpha) \rfloor}{n} \right\}$$

Inverted!

Our goal is not to estimate λ

Definition: Inverse Conformal Risk Control (Ours)

Given loss function parameter λ, estimate risk level upper bound:

$$\inf_{\alpha \in [0,1]} \left\{ \alpha : B - \overline{\ell_n}(\lambda) \geq \frac{\lfloor (n+1)(B-\alpha) \rfloor}{n} \right\}$$

Our goal is to curate the Pareto frontier

Theoretical Results

Denote the gap by Δ (either Δ_R or Δ_I):

Theorem: Validity

Under exchangeability, $\mathbb{E}[\Delta] \geq 0$

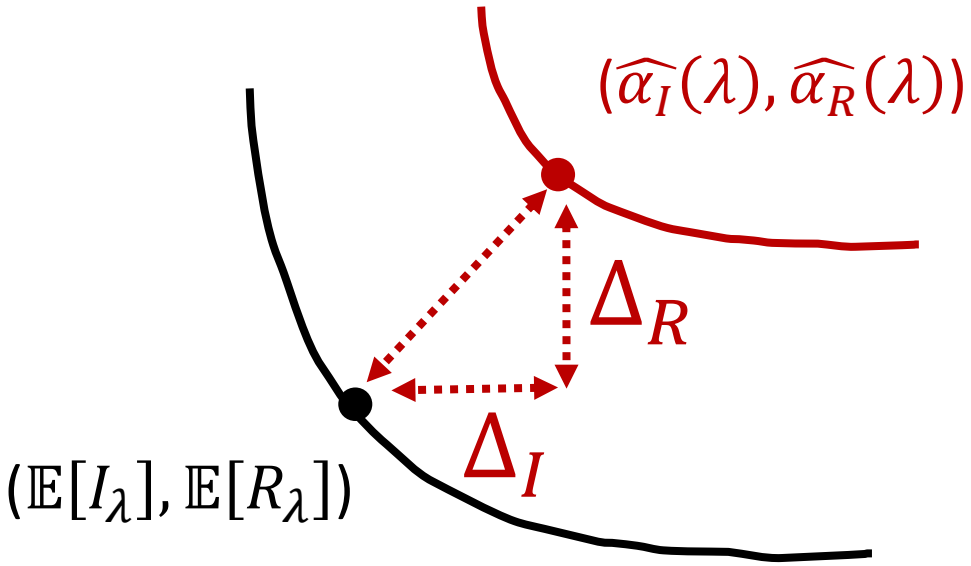

Theorem: Error Bound

Under i.i.d.:

$$|\Delta| \leq \frac{B}{n+1}\left(\sqrt{\frac{n}{2}\log\left(\frac{2}{\delta}\right)} + 2 + \frac{1}{B}\right)$$

Theorem (Informal): Pareto Frontier

Under the majorant consistency assumption, the curve of $(\mathbb{E}[I_\lambda], \mathbb{E}[R_\lambda])$ by varying λ is weakly decreasing.

Algorithm: Recalibration after Post-hoc Selection

During decision-making:

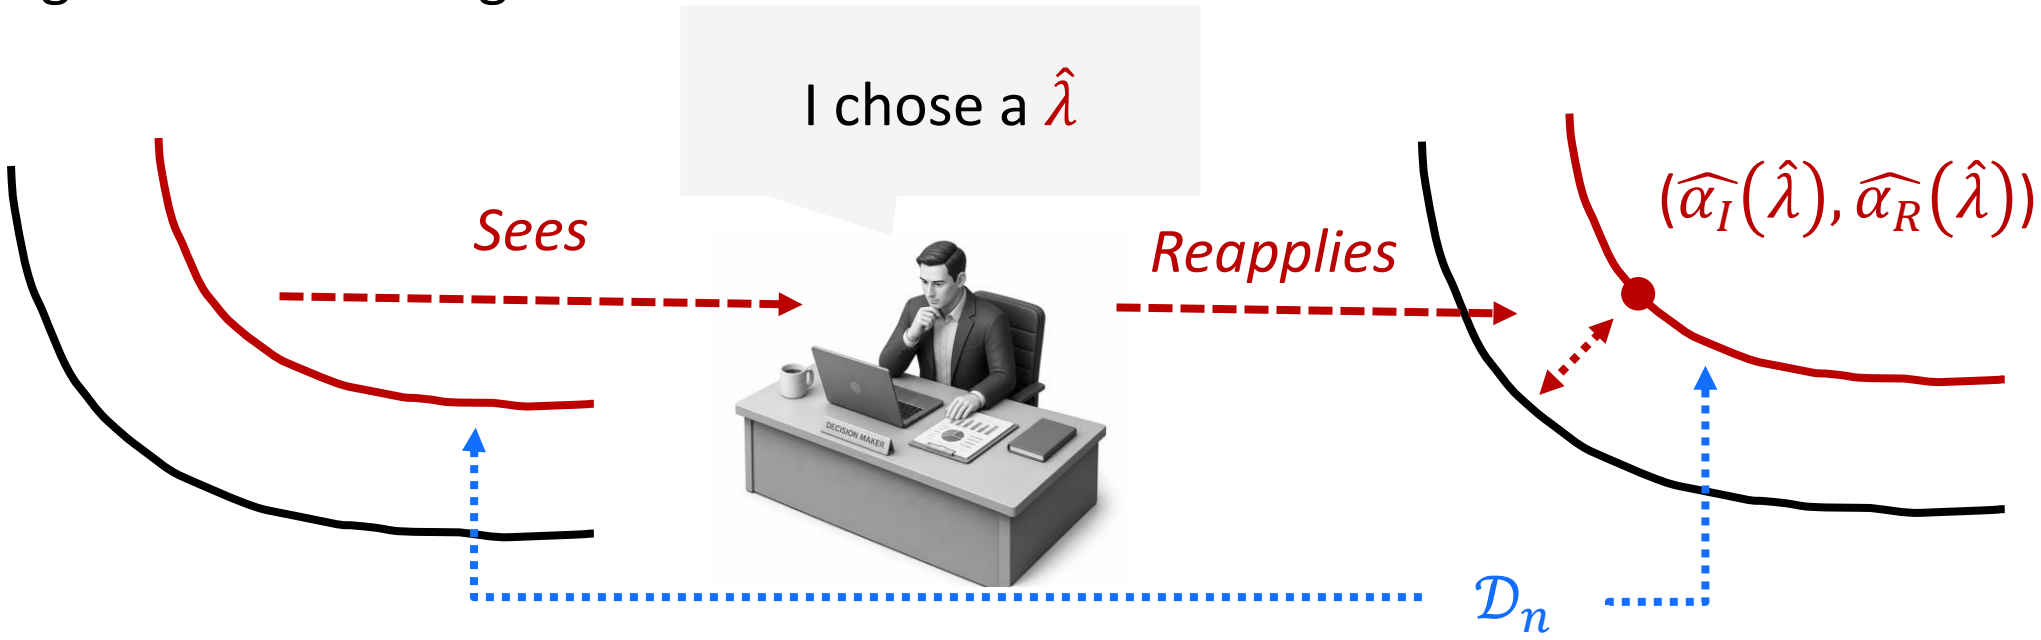

However,

Theorem: Validity

$$\mathbb{E}[\Delta] \geq 0$$

no longer holds due to data sharing.

Carnegie Mellon University

Algorithm: Recalibration after Post-hoc Selection

During decision-making:

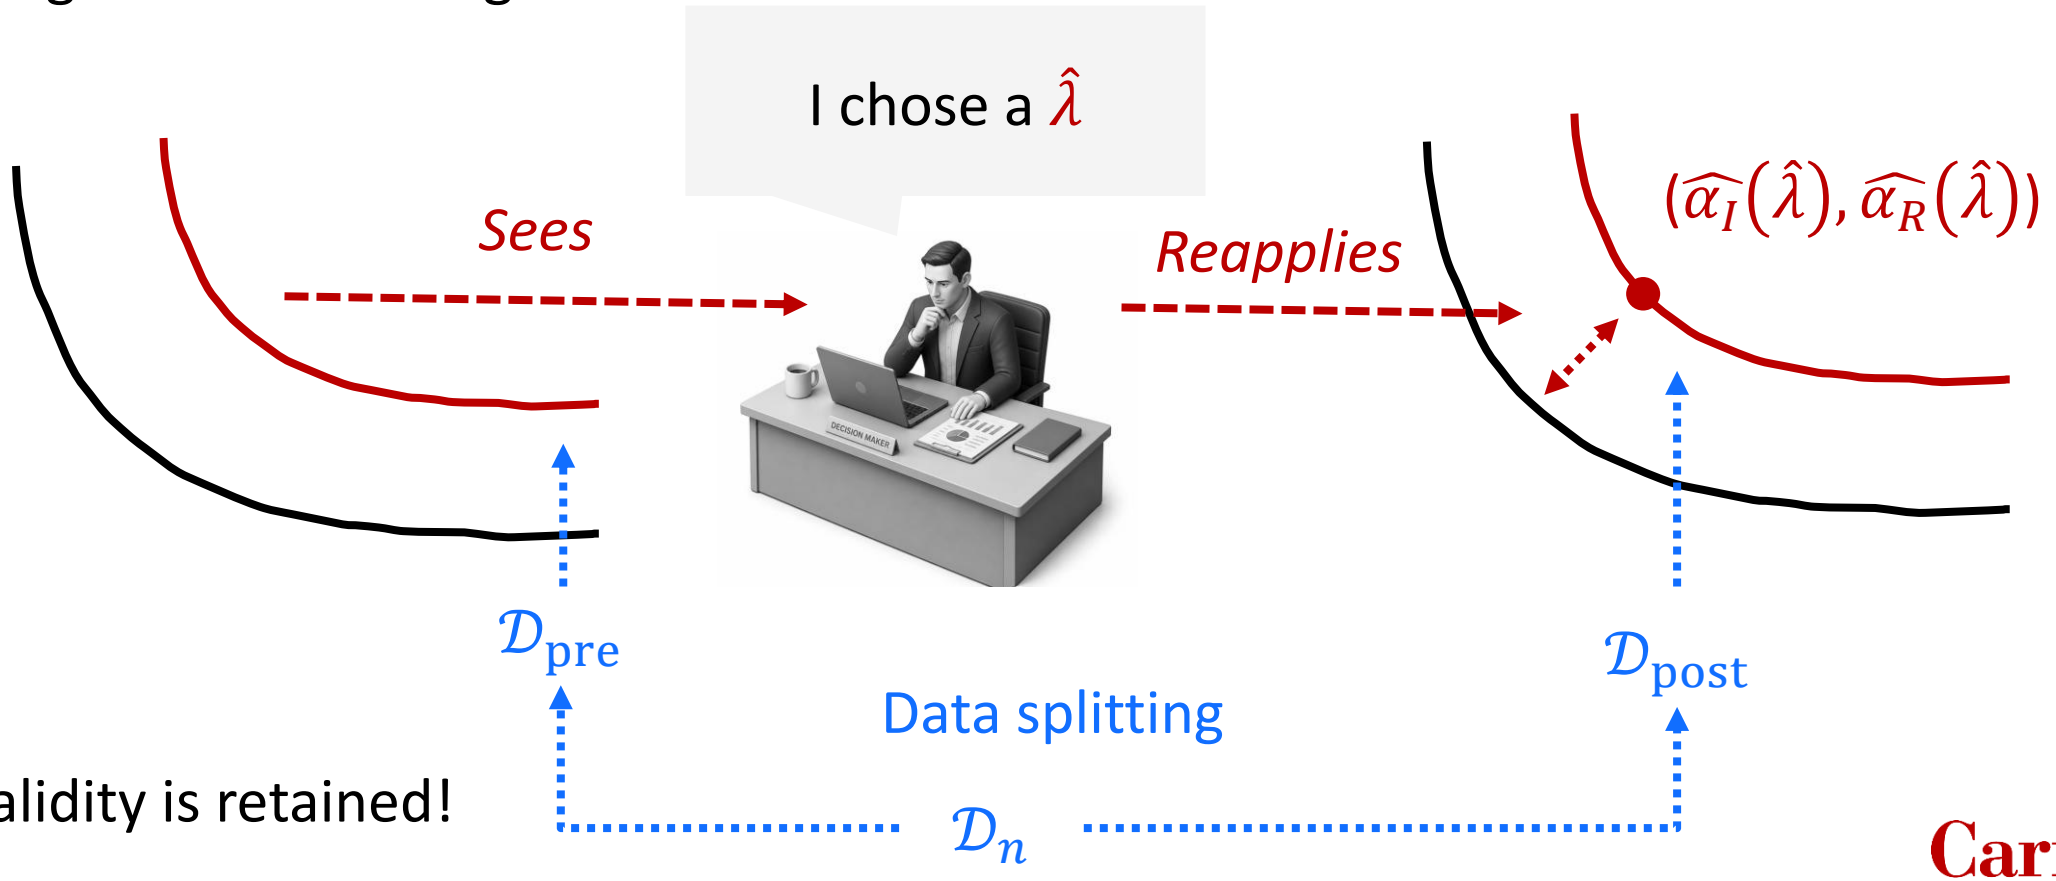

Validity is retained!

Carnegie Mellon University

Experiment: Good Estimation Quality

$\hat{\alpha}(\lambda)$ Evaluation Results (LP)

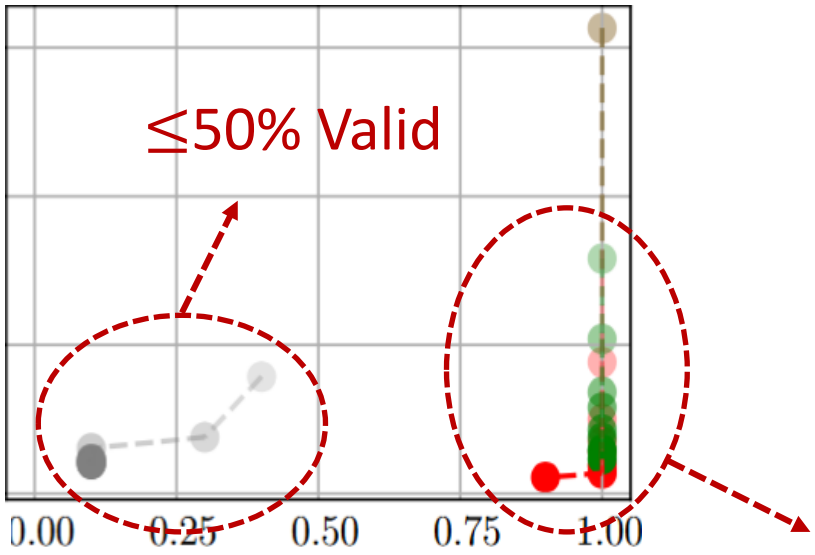

Is $\Delta > 0$?

(Near) 100% Valid +
Asymptotic consistency

Legend

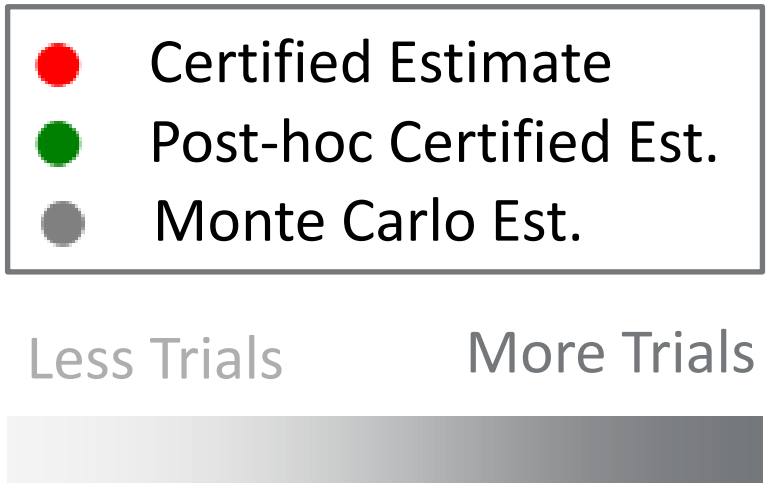

*See the paper for results for more optimization problems

Carnegie Mellon University

Experiment: Informative Constructed Frontier

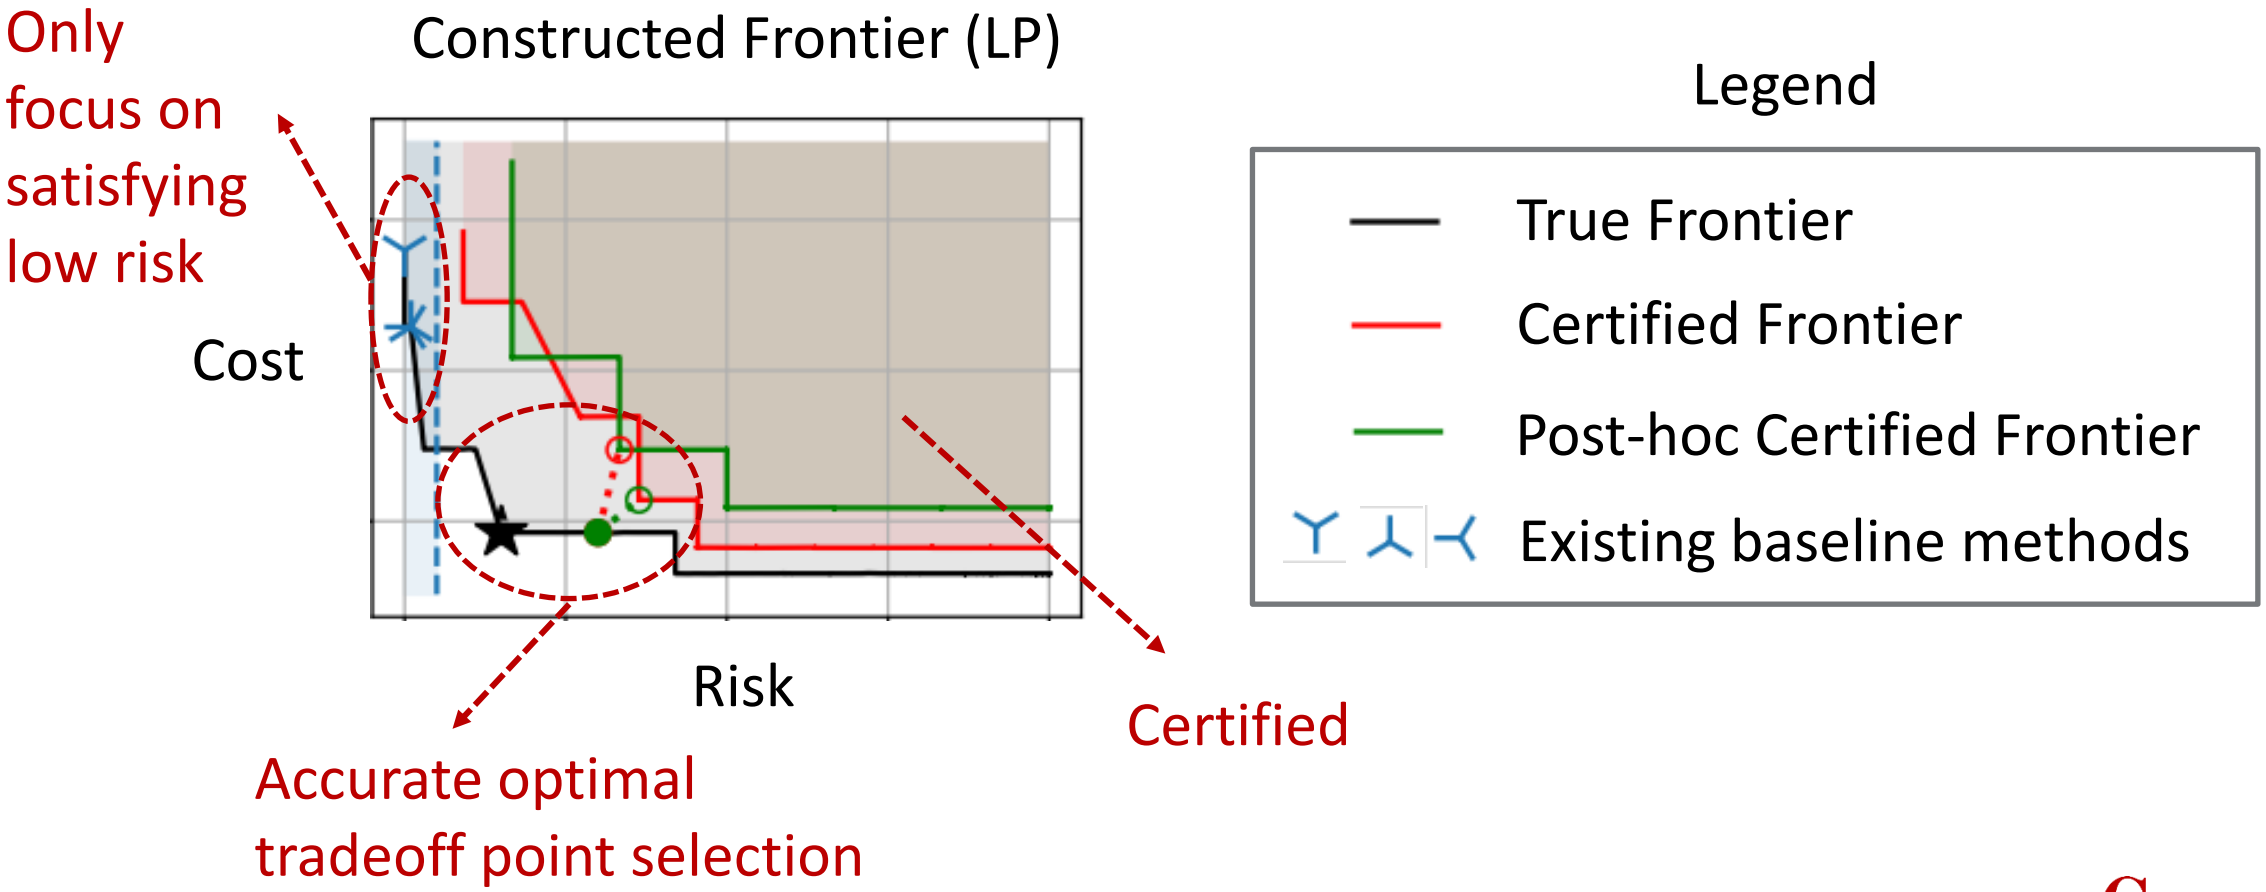

*See the paper for results for more optimization problems

Carnegie Mellon University

Thank you!

arXiv

Contact:

wenbinz2@andrew.cmu.edu

shixiangzhu@cmu.edu

Github codebase

Carnegie Mellon University

Experiment: Good Estimation Quality

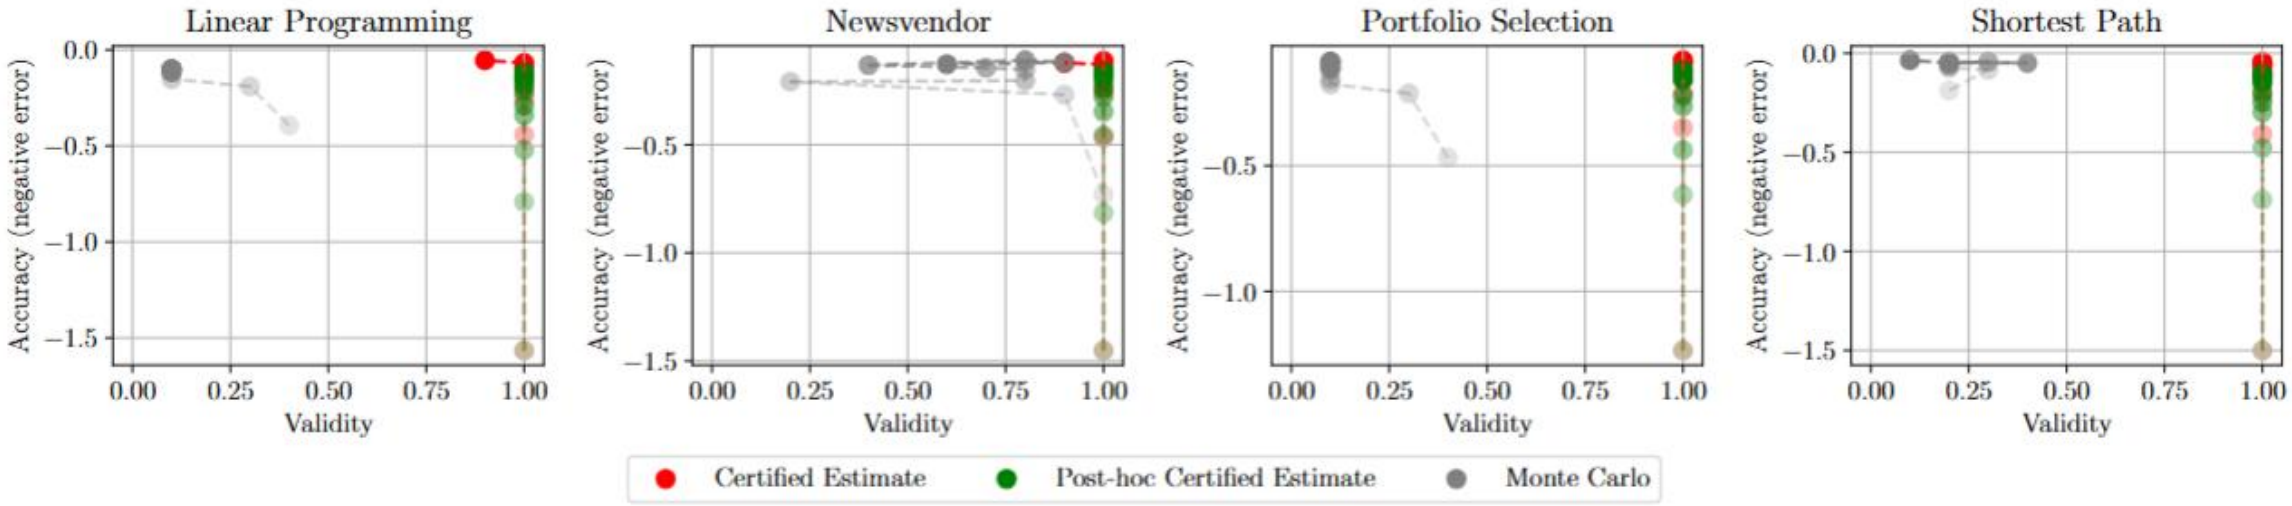

Figure 3. Validity-accuracy tradeoff curves under four optimization settings. Connected dots trace estimator performance as the number of calibration samples n increases from 1 to 50, with greater opacity indicating smaller n. The proposed method is shown in red, and the baseline Monte Carlo estimator in gray. Both axes represent metrics where higher values indicate better performance (↑), so methods appearing closer to the upper-right corner are more desirable.

Carnegie Mellon University

Experiment: Informative Constructed Frontier

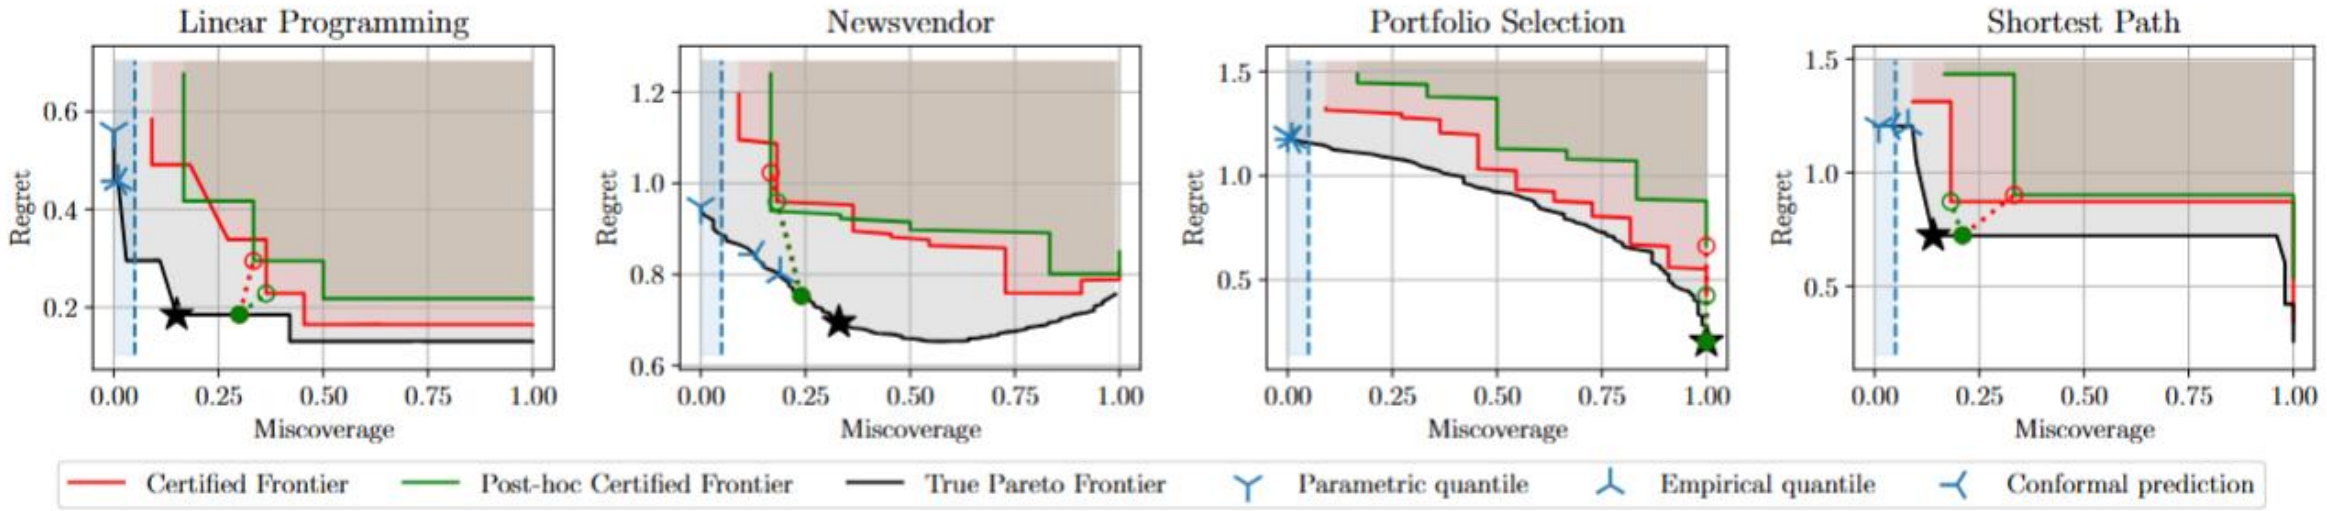

Figure 4. Miscoverage-regret tradeoff Pareto frontiers. Given some prespecified preference function, there is an optimal tradeoff point (black star) that can be computed from the true Pareto frontier (black line). We can obtain a certified tradeoff point (red dot) and a post-hoc tradeoff point (green dot) by using CREME. Blue markers represent tradeoff points derived from the three baseline methods that identify λ values ensuring the miscoverage rate remains below 5% (vertical dashed blue line).

Carnegie Mellon University

Remark: Inverse Conformal Risk Control

The proposed estimator $\widehat{\alpha_\ell}(\lambda)$ can be lower bounded by the following proxy estimator:

$$\widetilde{\alpha_\ell}(\lambda) = \frac{n}{n+1}\overline{\ell_n}(\lambda) + \frac{B}{n+1}$$

Which is efficient to compute and easy to implement in practice.