Mitigating optimistic bias in entropic risk estimation and optimization with an application to insurance

Utsav Sadana
Université de Montréal

joint work with Angelos Georghiou and Erick Delage

IEMS Seminar, Northwestern University

October 21, 2025

Outline

- What is an entropic risk measure?
- Estimated risk vs true risk
- Optimistic bias in estimating and optimizing entropic risk
  - **Estimation**: Proposed algorithms to mitigate the underestimation of risk
  - **Optimization under distributional ambiguity**: Use the algorithms for calibrating hyperparameters of the ambiguity set for risk minimization
- Case study of a flood insurance pricing problem

What is an Entropic Risk Measure?

Entropic Risk Measure

- Loss is uncertain in many real-world problems
- A risk measure maps the uncertain loss to a real number
- For example, a widely used convex law-invariant risk measure is the entropic risk

$$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) := \begin{cases} \frac{1}{\alpha}\log(\mathbb{E}_{\mathbb{P}}(e^{\alpha\ell(\boldsymbol{\xi})})) & \text{if } \alpha > 0, \\ \mathbb{E}_{\mathbb{P}}[\ell(\boldsymbol{\xi})] & \text{if } \alpha = 0, \end{cases}$$

  $\ell(\boldsymbol{\xi})$ is the loss associated with $\boldsymbol{\xi} \sim \mathbb{P}$ and α is the risk aversion parameter
- Entropic risk is the certainty equivalent of the exponential utility
- For normal loss distribution $\ell(\boldsymbol{\xi}) \sim \mathcal{N}(\mu, \sigma^2)$: $\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) = \mu + \frac{1}{2}\alpha\sigma^2$
- Would you take a fixed loss of μ or a gamble with risk $\mu + \frac{1}{2}\alpha\sigma^2$?

Entropic Risk Measure

- Exponential utility model - Agents' preferences exhibit constant absolute risk aversion (CARA)
- Widespread applications in
  - Risk-sensitive control[1] (entropic risk measure is time-consistent)
  - Portfolio selection[2]
  - Fair and robust decision making[3]
  - Catastrophe insurance pricing[4]
- Diverse communities such as control theory, operations research, economics, and machine learning

---

[1] Howard, R. A., & Matheson, J. E. (1972).Risk-sensitive Markov decision processes. *Management science*, *18*(7), 356–369.

[2] Chen, L., & Sim, M. (2024).Robust CARA optimization [Forthcoming]. *Operations Research*.

[3] Li, T., Beirami, A., Sanjabi, M., & Smith, V. (2023).On Tilted Losses in Machine Learning: Theory and Applications. *Journal of Machine Learning Research*, *24*(142), 1–79.

[4] Bernard, C., Liu, F., & Vanduffel, S. (2020).Optimal insurance in the presence of multiple policyholders. *Journal of Economic Behavior & Organization*, *180*, 638–656.

Literature Review

Three intersecting themes:

- Correcting bias in risk estimators
- Addressing the optimistic bias of SAA (sample average approximation) policy
- Pricing insurance for correlated losses

Correcting Bias in Risk Estimators

- Quantitative risk measurement often relies on precise estimation of risk measures[5]
- Non-parametric bootstrap for correcting bias in Value at Risk (VaR) estimates[6]
- Distributionally robust optimization (DRO) to construct worst-case tail risk bounds[7]
- Extreme value theory (EVT) for unbiased CVaR estimation[8]
- **How to mitigate bias in entropic risk estimation with finite samples?**

---

[5] McNeil, A. J., Frey, R., & Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.

[6] Kim, J. H. T. (2010).Bias correction for estimated distortion risk measure using the bootstrap. *Insurance: Mathematics and Economics*, *47*(2), 198–205.

[7] Lam, H., & Mottet, C. (2017).Tail analysis without parametric models: A worst-case perspective. *Operations Research*, *65*(6), 1696–1711.

[8] Troop, D., Godin, F., & Yu, J. Y. (2021). Bias-corrected peaks-over-threshold estimation of the CVaR. In C. de Campos & M. H. Maathuis (Eds.), *Uncertainty in artificial intelligence* (pp. 1809–1818, Vol. 161). PMLR.

Optimistic Bias in the SAA Policy

- DRO, hold-out, and K-fold cross-validation (CV) address the Optimizer's Curse via hyperparameter tuning[9]
- Estimators of SAA policy performance for Gaussian data in linear optimization problems[10]
- The Optimizer's Information Criterion (OIC) corrects SAA policy bias asymptotically using the loss's influence function[11]
- **Can DRO mitigate the optimistic bias in entropic risk minimization?**

[9] Smith, J. E., & Winkler, R. L. (2006).The Optimizer's Curse: Skepticism and Postdecision Surprise in Decision Analysis. *Management Science*, *52*(3), 311–322.

[10] Gupta, V., Huang, M., & Rusmevichientong, P. (2024).Debiasing in-sample policy performance for small-data, large-scale optimization. *Operations Research*, *72*(2), 848–870.

[11] Iyengar, G., Lam, H., & Wang, T. (2023).Optimizer's information criterion: Dissecting and correcting bias in data-driven optimization. *arXiv preprint arXiv:2306.10081*.

Insurance Pricing

- Seminal paper by Kenneth Arrow in insurance contract design[12]
  - Expected utility-maximizing policyholder will choose full coverage above a deductible
- Several models for the risk-averse insurer and policyholder[13]
- Typical assumption: known loss distribution
- **How to model a risk-averse insurance pricing problem under distributional ambiguity?**

[12] Arrow, K. J. (1963).Uncertainty and the welfare economics of medical care. *American Economic Review*, *53*(5), 941–973.

[13] Bernard, C., Liu, F., & Vanduffel, S. (2020).Optimal insurance in the presence of multiple policyholders. *Journal of Economic Behavior & Organization*, *180*, 638–656.

Optimistic Bias in Entropic Risk Estimation

Estimation of Entropic Risk

- Loss distribution is unknown in real-world applications
- Given: historical data $\mathcal{D} = \{\hat{\boldsymbol{\xi}}_1, \hat{\boldsymbol{\xi}}_2, \cdots, \hat{\boldsymbol{\xi}}_N\}$

$$\hat{\mathbb{P}}_N(\boldsymbol{\xi}) := \frac{1}{N}\sum_{i=1}^{N} \delta_{\hat{\boldsymbol{\xi}}_i}(\boldsymbol{\xi})$$

  where $\delta_{\boldsymbol{\xi}}$ is a Dirac distribution at point $\boldsymbol{\xi}$.

- The empirical entropic risk estimator is given by:

$$\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) := \frac{1}{\alpha}\log\left(\frac{1}{N}\sum_{i=1}^{N} e^{\alpha\ell(\hat{\boldsymbol{\xi}}_i)}\right)$$

- Empirical entropic risk estimator underestimates the true risk (Jensen's inequality)

$$\mathbb{E}\left[\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))\right] < \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$$

Risk Underestimated in Insurance Pricing

- Minimum premium to insure against the loss $\ell(\xi) := \xi$ is given by the entropic risk
- Gamma distributed loss function $\xi \sim \Gamma(10, 0.24)$

$$\pi = \rho_{\mathbb{P}}(\xi)$$
$$= \frac{1}{\alpha} \log\left((1 - 0.24\alpha)^{-10}\right)$$

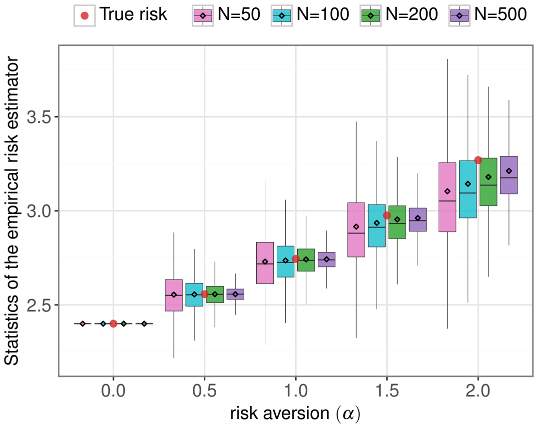

Bias Correction: Mean-Unbiased Estimator

$$\text{Bias:} \quad \delta^{\text{true}} = \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) - \rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))$$

$$\text{Find } \delta(\mathcal{D}_N) \text{ such that } \mathbb{E}[\delta^{\text{true}}] = \mathbb{E}[\delta(\mathcal{D}_N)]$$

$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$

$\mathbb{E}\left[\rho_{\widehat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))\right]$

Bias Correction: Mean-Unbiased Estimator

$$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$$

$$\mathbb{E}\left[\rho_{\widehat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))\right]$$

Non-parametric bootstrap estimator

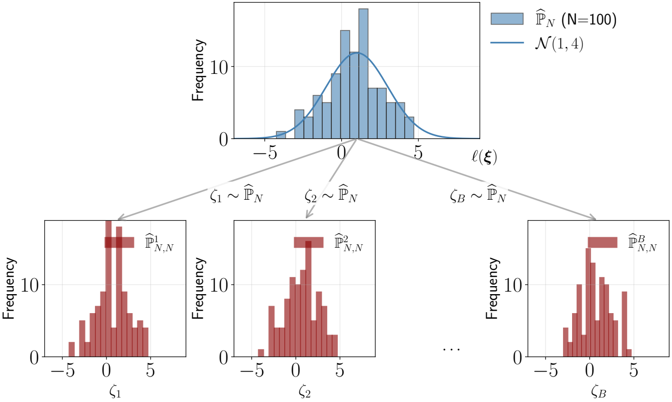

- Calculate bias by resampling from $\hat{\mathbb{P}}_N$ (Efron, 1979)

$$\delta_N(\hat{\mathbb{P}}_N) = \rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) - \mathbb{E}[\rho_{\hat{\mathbb{P}}_{N,N}}(\ell(\boldsymbol{\xi}))]$$

- $\rho_{\mathrm{BS}} = \rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta_N(\hat{\mathbb{P}}_N)$

Bias - True vs Bootstrap

- Suppose $\ell(\xi) \sim \mathcal{N}(1, 4)$ and $\alpha = 2$
- Non-parametric bootstrap underestimates the true correction in finite samples
- Can we further mitigate the underestimation issue?

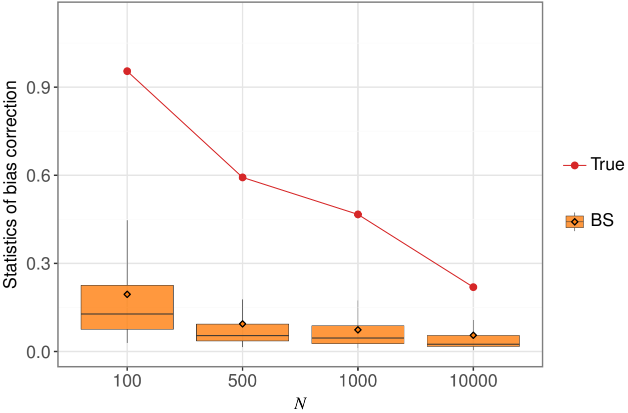

Bias Mitigation - Approximately Median-Unbiased Estimator

$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$

$\mathbb{E}\left[\rho_{\widehat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))\right]$

Proposed Parametric Bootstrap Approach

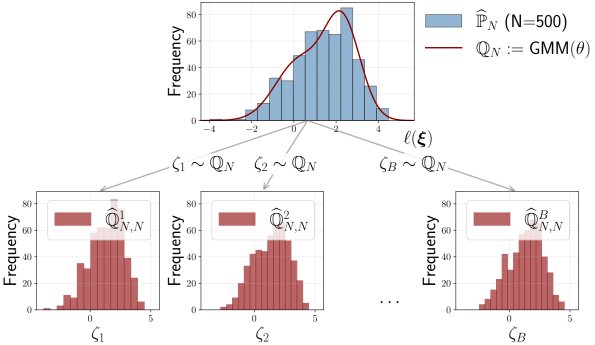

- **Step 1:** Fit $\mathbb{Q}_N$ - Gaussian Mixture Model with parameters $\theta = (\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$
- **Step 2:** Sample $\zeta_i \sim \mathbb{Q}_N$ and compute $\rho_{\hat{\mathbb{Q}}^i_{N,N}}(\zeta)\}$ for each sample
- **Step 3:** Compute $\rho_{\mathbb{Q}_N}(\zeta) = \frac{1}{\alpha} \log \left(\sum_{j=1}^{Y} \pi_j \exp(\alpha \mu_j + \frac{\alpha^2}{2} \sigma_j^2) \right)$
- **Step 4:** Compute $\hat{\delta}_N(\mathbb{Q}_N) = \max(\textbf{median}[\{\rho_{\mathbb{Q}_N}(\zeta) - \rho_{\hat{\mathbb{Q}}^i_{N,N}}(\zeta)\}_{i=1}^{B}], 0)$

Asymptotic consistency

- The bias correction is given by:

$$\hat{\delta}_N(\mathbb{Q}_N) = \max(\textbf{median}[\{\rho_{\mathbb{Q}_N}(\zeta) - \rho_{\hat{\mathbb{Q}}^i_{N,N}}(\zeta)\}_{i=1}^{B}], 0)$$

- The parametric bootstrap estimator is given by:

$$\rho_{\texttt{BS-M}} = \rho_{\hat{\mathbb{P}}_N}(\ell(\xi)) + \hat{\delta}_N(\mathbb{Q}_N)$$

Theorem 1

Under the assumptions

- *the tails of $\ell(\boldsymbol{\xi})$ are exponentially bounded*
- *the tails of $\mathbb{Q}_N$ are exponentially bounded*

the estimator $\rho_{BS\text{-}M}$ is strongly asymptotically consistent.

Which model ($\texttt{M}$) should we use to fit the GMM $\mathbb{Q}_N$?

Fitting $\mathbb{Q}_N$

Parametric bootstrap - Maximum Likelihood estimation

- Suppose $\ell(\xi) \sim \mathcal{N}(1, 4)$ and $\alpha = 2$
- BS-MLE - Fit a normal distribution to the data and then bootstrap.
- What if the model is misspecified?

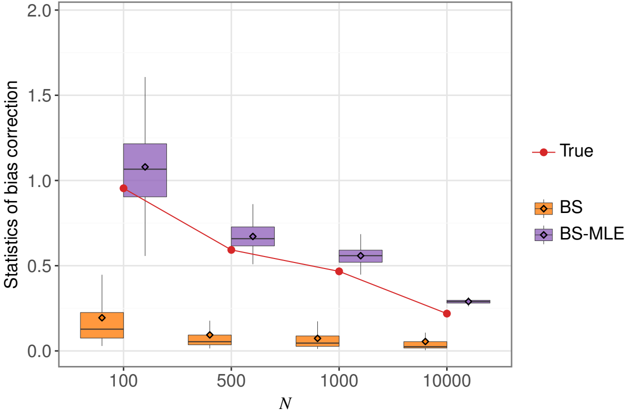

Inaccurate estimation

- Suppose $\xi \sim \text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$, $\boldsymbol{\pi} = [0.7\ 0.3]$, $\boldsymbol{\mu} = [0.5\ 1]$, and $\boldsymbol{\sigma} = [2\ 1]$.
- Slow convergence of the Expectation-Maximization algorithm for overlapping components (Xu & Jordan, 1996)

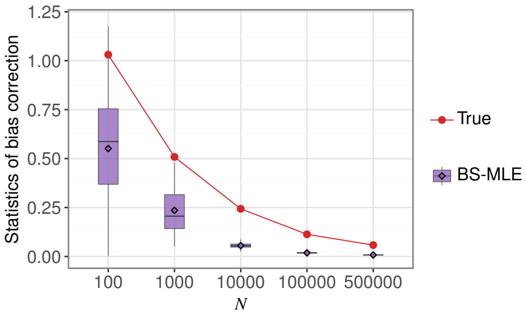

Fitting $\mathbb{Q}_N$ via MLE does not take into account the effect of estimation errors on the bias

Inaccurate estimation

- Suppose $\xi \sim \text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$, $\boldsymbol{\pi} = [0.7\ 0.3]$, $\boldsymbol{\mu} = [0.5\ 1]$, and $\boldsymbol{\sigma} = [2\ 1]$.
- Slow convergence of the Expectation-Maximization algorithm for overlapping components (Xu & Jordan, 1996)

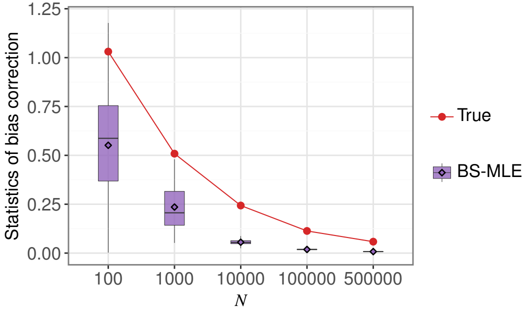

Fitting $\mathbb{Q}_N$ via MLE does not take into account the effect of estimation errors on the bias
Recall, we aim to mimic the true bias $\implies$ Use **"bias-aware" distribution matching**

$\mathbb{Q}_N$ using Entropic Risk Matching

Mimic the bias in the samples

$\mathbb{Q}_N$ using Entropic Risk Matching

Mimic the bias in the samples
Divide samples into $B = \sqrt{N}$ bins

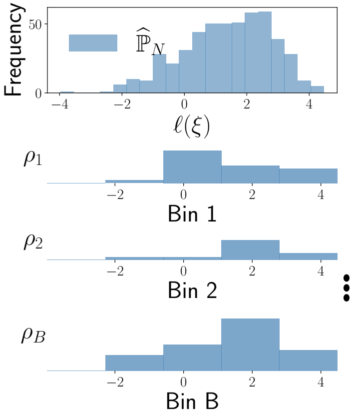

$\mathbb{Q}_N$ using Entropic Risk Matching

Mimic the bias in the samples

Divide samples into $B = \sqrt{N}$ bins

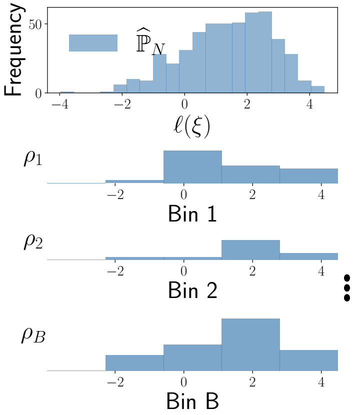

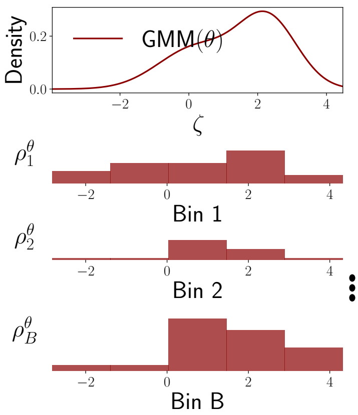

$\mathbb{Q}_N$ using Entropic Risk Matching

Mimic the bias in the samples
Divide samples into $B = \sqrt{N}$ bins

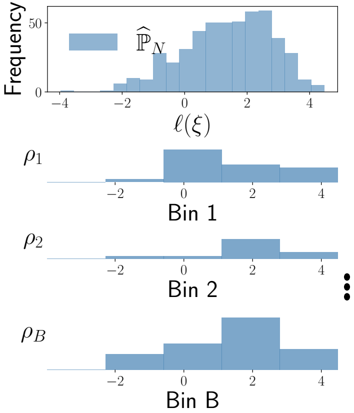

$$\mathcal{R}_{\mathcal{D}_N} = \{\rho_1, \dots, \rho_B\}$$
$$\mathcal{R}_\theta = \{\rho_1^\theta, \dots, \rho_B^\theta\}$$

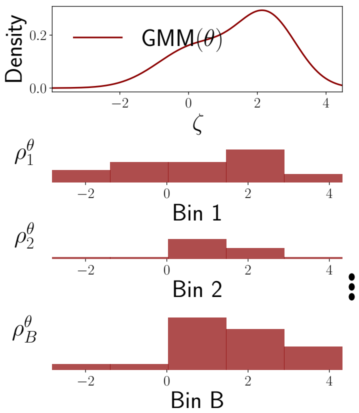

$\mathbb{Q}_N$ using Entropic Risk Matching

Mimic the bias in the samples

Divide samples into $B = \sqrt{N}$ bins

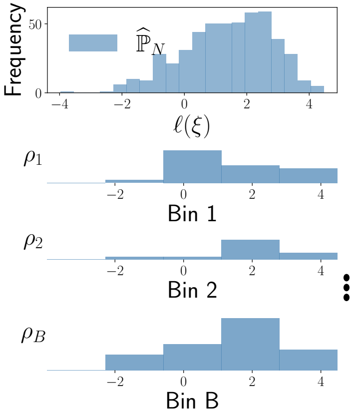

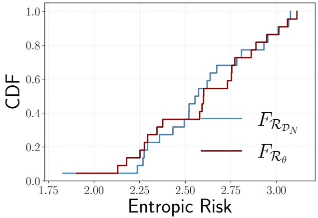

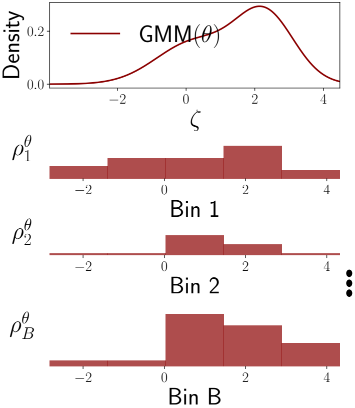

$$\mathcal{R}_{\mathcal{D}_N} = \{\rho_1, \ldots, \rho_B\}$$

$$\mathcal{R}_\theta = \{\rho_1^\theta, \ldots, \rho_B^\theta\}$$

Gradient Descent algorithm

- Find $\boldsymbol{\theta}$

$$\min_{\boldsymbol{\theta}} \mathcal{W}^2\left(\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\boldsymbol{\theta}}}\right) = \left(\int_0^1 |F^{-1}_{\mathcal{R}_{\mathcal{D}_N}}(q) - F^{-1}_{\mathcal{R}_{\boldsymbol{\theta}}}(q)|^2 dq\right)^{1/2}$$

- Optimal $\boldsymbol{\theta}$ is obtained using gradient descent

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \gamma \nabla_{\boldsymbol{\theta}_t} \mathcal{W}^2\left(\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\boldsymbol{\theta}_t}}\right),$$

- Compute gradient by backpropagation
- Making sampling differentiable: For the GMM's discrete component, use Gumbel–max with the Straight-through estimator (STE).
- This is a computationally expensive procedure. Alternative?

$\mathbb{Q}_N$ using Extreme Value Theory

Idea: Entropic risk is sensitive to extreme values. Fit a $\mathbb{Q}_N$ component to model the upper tail of $\mathbb{P}$

Theorem 2 (Fisher–Tippett–Gnedenko Theorem)

Let ζ follow a distribution with cdf $F(\cdot)$. The distribution of $M_n = \max\{\zeta_1, \zeta_2, \cdots, \zeta_n\}$ converges to a non-degenerate distribution G:

$$\lim_{n\to\infty} \mathbb{P}\left(\frac{M_n - b_n}{a_n} \leq x\right) = \lim_{n\to\infty} F(a_n x + b_n)^n \to G(x),$$

where a_n and b_n are normalizing constants

$\implies$ Fit $\Phi^N_{\mu,\sigma}$—distribution of maxima of N i.i.d samples from $\mathcal{N}(\mu,\sigma)$—to the distribution of maxima constructed from data

$\mathbb{Q}_N$ using Extreme Value Theory

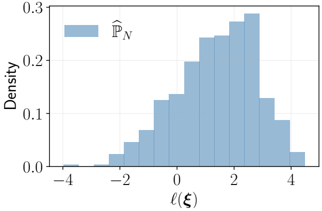

$\mathbb{Q}_N$ using Extreme Value Theory

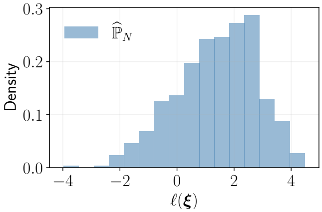

Divide samples into B bins

$\mathbb{Q}_N$ using Extreme Value Theory

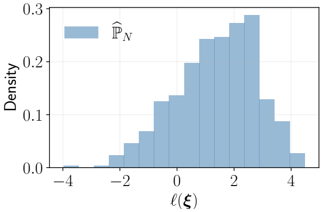

Divide samples into B bins

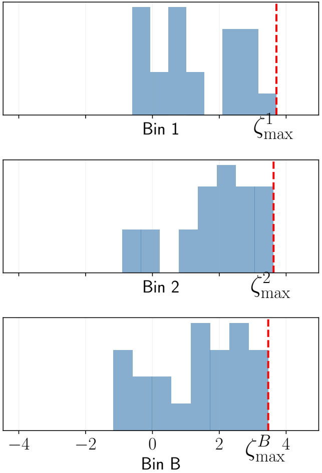

$\mathbb{Q}_N$ using Extreme Value Theory

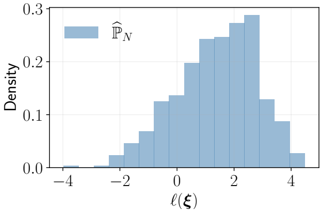

Divide samples into B bins

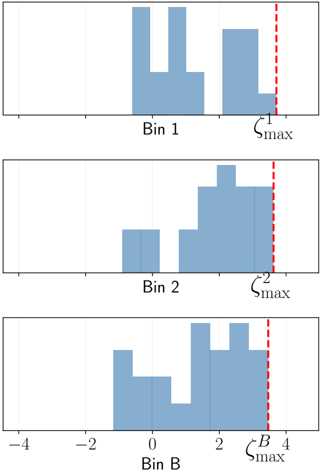

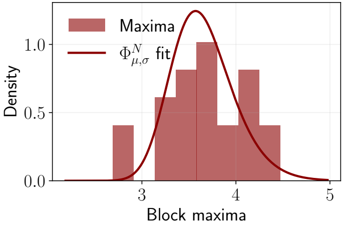

$\mathbb{Q}_N$ using Extreme Value Theory

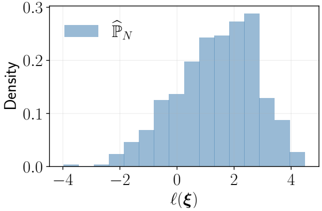

Divide samples into B bins

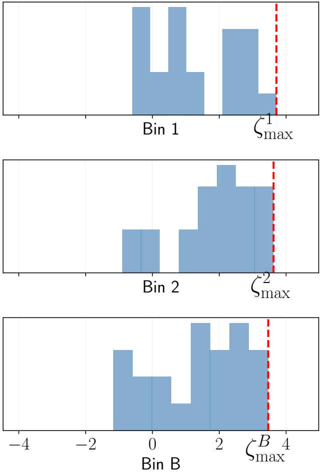

Fit $\Phi^N_{\mu,\sigma}$ to the maxima (analytic).

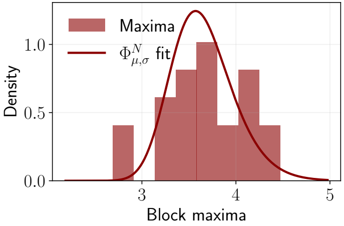

$\mathbb{Q}_N$ using Extreme Value Theory

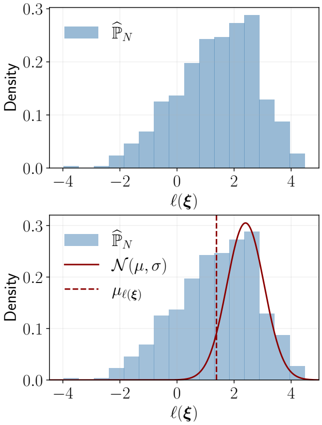

Divide samples into B bins

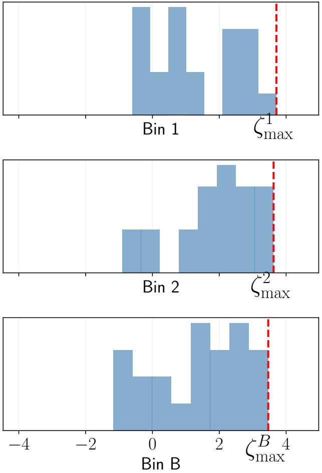

Fit $\Phi^N_{\mu,\sigma}$ to the maxima (analytic).

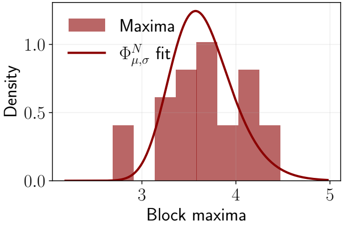

$\mathbb{Q}_N$ using Extreme Value Theory

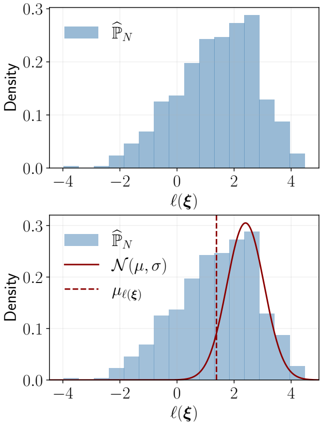

GMM $=$
$50\% \mu_{\ell(\xi)} + 50\% \mathcal{N}(\mu, \sigma)$

Divide samples into B bins

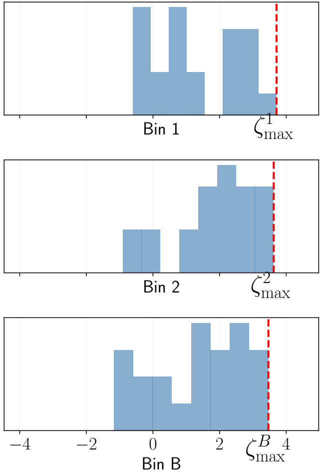

Fit $\Phi^N_{\mu,\sigma}$ to the maxima (analytic).

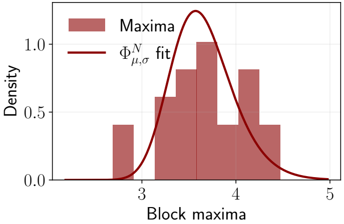

Effect of Distribution Fitting

Recall $\xi \sim \text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$, $\boldsymbol{\pi} = [0.7\ 0.3]$, $\boldsymbol{\mu} = [0.5\ 1]$, and $\boldsymbol{\sigma} = [2\ 1]$.

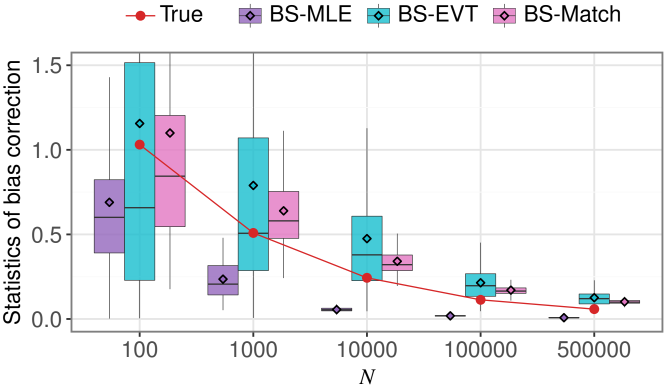

Proposed approaches typically **overestimate bias, but stay close to true bias**

Comparing With Other Methods

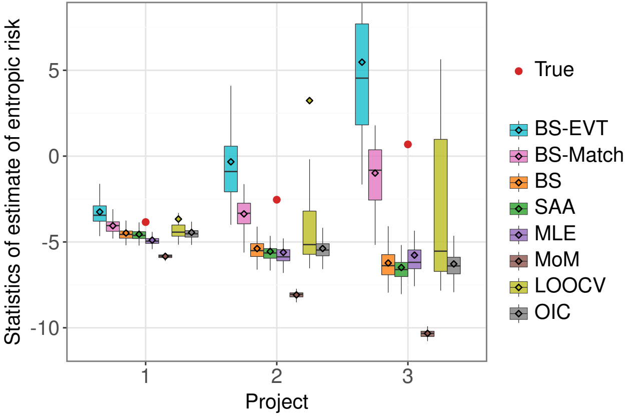

- Project 1, 2, 3 have respective losses 0.4ξ, 0.6ξ and 0.8ξ with $\xi \sim$ GMM: $\mu_\xi = -18.6$ and $\sigma_\xi = 2.9$
- Most methods underestimate risk while **proposed methods are close to true risk**

Optimization

Distributionally robust optimization (DRO)

- Entropic risk minimization with distribution $\mathbb{P}$

$$\rho^* = \min_{\boldsymbol{z} \in \mathcal{Z}} \rho_{\mathbb{P}}(\ell(\boldsymbol{z}, \boldsymbol{\xi})) := \frac{1}{\alpha} \log \left(\mathbb{E}_{\mathbb{P}}[e^{\alpha \ell(\boldsymbol{z}, \boldsymbol{\xi})}] \right)$$

- Sample average approximation:

$$\rho_{\mathtt{SAA}} = \min_{\boldsymbol{z} \in \mathcal{Z}} \rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{z}, \boldsymbol{\xi})) := \frac{1}{\alpha} \log \left(\mathbb{E}_{\hat{\mathbb{P}}_N}[e^{\alpha \ell(\boldsymbol{z}, \boldsymbol{\xi})}] \right).$$

  Often results in the Optimizer's Curse or overfitting

- Distributionally robust optimization - protection against distributional ambiguity

$$\rho_{\mathtt{DRO}} = \min_{\boldsymbol{z} \in \mathcal{Z}} \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon(\hat{\mathbb{P}}_N)} \frac{1}{\alpha} \log \left(\mathbb{E}_{\mathbb{Q}}[e^{\alpha \ell(\boldsymbol{z}, \boldsymbol{\xi})}] \right).$$

  where $\mathcal{B}_\epsilon(\hat{\mathbb{P}}_N)$ is the ambiguity set of all distributions at a "distance" ϵ from $\hat{\mathbb{P}}_N$

Ambiguity Set

- **KL-divergence** ambiguity set is ill-suited
  - $\mathrm{KL}(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \epsilon$ - cannot reposition scenarios
  - $\mathrm{KL}(\hat{\mathbb{P}}_N, \mathbb{Q}) \leq \epsilon$ - worst case loss is ∞
- **p-Wasserstein** ambiguity set - unbounded worst-case loss if $p < \infty$

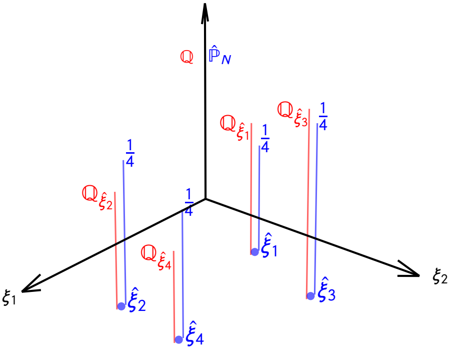

Type-∞ Wasserstein Ambiguity Set

$$\mathcal{B}_\epsilon(\hat{\mathbb{P}}_N) := \left\{ \mathbb{Q} \in \mathcal{M}(\Xi) \mid \exists\, \boldsymbol{\xi}_i \in \Xi \text{ s.t. } \|\boldsymbol{\xi}_i - \hat{\boldsymbol{\xi}}_i\|_\infty \leq \epsilon,\ \forall i \in [N], \mathbb{Q} = \frac{1}{N}\sum_{i=1}^{N} \delta_{\boldsymbol{\xi}_i} \right\}.$$

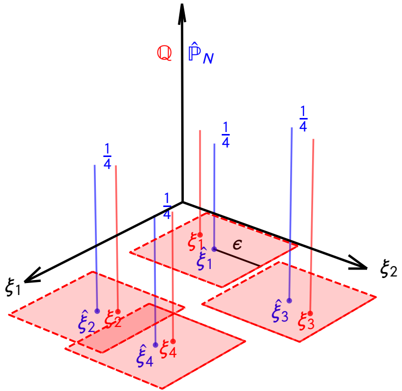

Distributionally robust optimization

Theorem 3

For a linear loss function, DRO problem becomes a regularized exponential cone program:

$$\min_{\boldsymbol{z} \in \mathcal{Z}} \frac{1}{\alpha} \log \left(\frac{1}{N} \sum_{i=1}^{N} e^{\alpha \boldsymbol{z}^\top \hat{\boldsymbol{\xi}}_i} \right) + \epsilon \|\boldsymbol{z}\|_*.$$

- For piecewise concave loss function in $\boldsymbol{\xi} \in \Xi$, DRO problem can be reformulated as a convex optimization problem using Fenchel duality
- We solve the exponential cone program using the MOSEK solver
- ϵ is typically chosen by K-fold cross validation

Theorem 4

$\rho_{DRO} \to \rho^*$ *and* $\rho_{SAA} \to \rho^*$ *in probability at the rate* $\mathcal{O}(1/\sqrt{N})$

Underestimation of the Risk of In-Sample Decisions

Proposition 5

The entropic risk estimator based on the K-fold CV underestimates the entropic risk of the policy constructed using $N(1 - \frac{1}{K})$ data points.

To calibrate ϵ, we will use the proposed bias-mitigation approaches.

Case Study

Insurance Pricing

Insurance contract for household h: coverage (z_h) and premium (π_h)

Insurer

risk aversion α_0

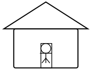

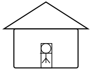

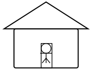

household 1	household h	household M
loss ξ_1	loss ξ_h	loss ξ_M
risk aversion α_1	risk aversion α_h	risk aversion α_M

Insurance Pricing

Insurance contract for household h: coverage (z_h) and premium (π_h)

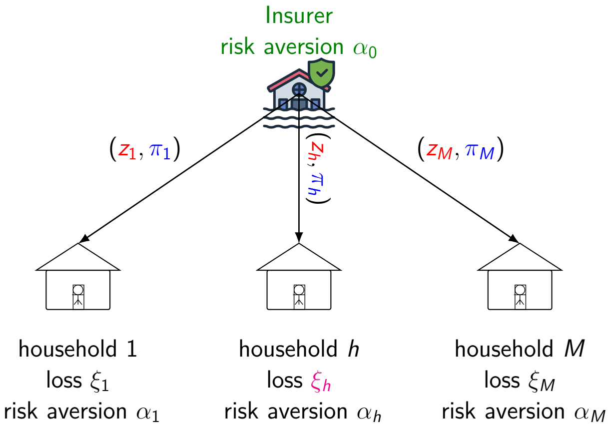

Insurance Pricing

Insurance contract for household h: coverage (z_h) and premium (π_h)

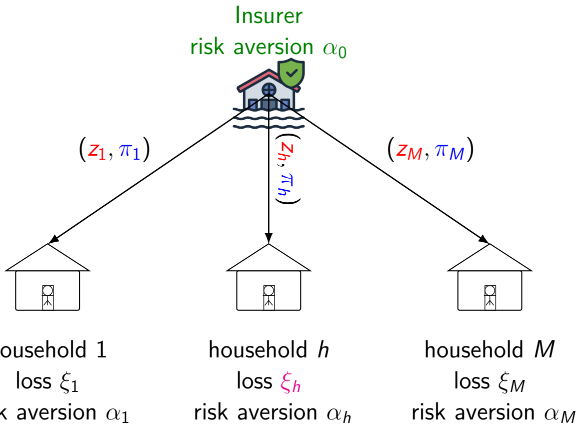

$$\underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\pi_h + (1 - z_h)\xi_h)}_{\text{risk when insured}} \leq \underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\xi_h)}_{\text{risk when uninsured}} \quad \forall\, h \in [M]$$

Insurance Pricing

Insurance contract for household h: coverage (z_h) and premium (π_h)

Insurer
risk aversion α_0

$$\underbrace{\rho^{\alpha_0}_{\hat{\mathbb{P}}_N}\left(\sum_{h=1}^{M} z_h \xi_h - \sum_{h=1}^{M} \pi_h\right)}_{\text{insurer's risk}}$$

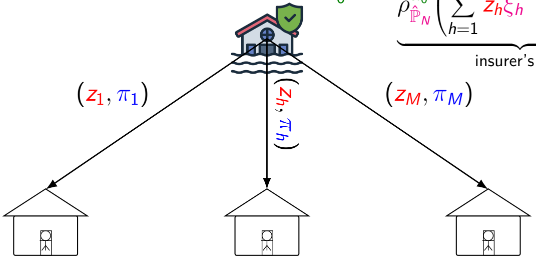

household 1	household h	household M
loss ξ_1	loss ξ_h	loss ξ_M
risk aversion α_1	risk aversion α_h	risk aversion α_M

$$\underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\pi_h + (1 - z_h)\xi_h)}_{\text{risk when insured}} \leq \underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\xi_h)}_{\text{risk when uninsured}} \quad \forall h \in [M]$$

Insurance Pricing

Insurance contract for household h: coverage (z_h) and premium (π_h)

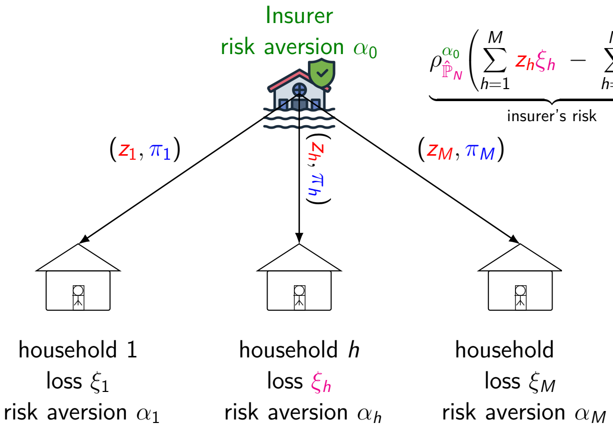

$$\underbrace{\rho^{\alpha_0}_{\hat{\mathbb{P}}_N}\left(\sum_{h=1}^{M} z_h \xi_h - \sum_{h=1}^{M} \pi_h\right)}_{\text{insurer's risk}}$$

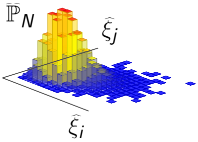

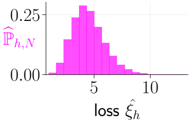

$$\underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\pi_h + (1 - z_h)\xi_h)}_{\text{risk when insured}} \leq \underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}(\xi_h)}_{\text{risk when uninsured}} \quad \forall\, h \in [M]$$

Insurance Pricing

Sample Average Approximation

$$\rho_{\mathtt{SAA}} = \min \underbrace{\rho^{\alpha_0}_{\hat{\mathbb{P}}_N}\left(\sum_{h=1}^{M} z_h \xi_h - \sum_{h=1}^{M} \pi_h\right)}_{\text{risk of insurer}}$$

$$\text{s.t.} \boldsymbol{\pi} \in \mathbb{R}^M_+,\ \boldsymbol{z} \in [0,1]^M$$

$$\underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\pi_h + (1 - z_h)\xi_h\right)}_{\text{risk when insured}} \leq \underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\xi_h\right)}_{\text{risk when uninsured}} \qquad \forall h \in [M]$$

Insurance Pricing

Sample Average Approximation

$$\rho_{\mathtt{SAA}} = \min \underbrace{\rho^{\alpha_0}_{\hat{\mathbb{P}}_N}\left(\sum_{h=1}^{M} z_h \xi_h - \sum_{h=1}^{M} \pi_h\right)}_{\text{risk of insurer}}$$

$$\text{s.t.}\boldsymbol{\pi} \in \mathbb{R}^M_+,\ \boldsymbol{z} \in [0,1]^M$$

$$\underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\pi_h + (1 - z_h)\xi_h\right)}_{\text{risk when insured}} \leq \underbrace{\rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\xi_h\right)}_{\text{risk when uninsured}} \qquad \forall h \in [M]$$

Distributionally Robust Optimization

$$\begin{aligned} \rho_{\mathtt{DRO}} = \min \quad & \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon(\hat{\mathbb{P}}_N)} \rho^{\alpha_0}_{\mathbb{Q}}\left(\sum_{h=1}^{M} z_h \xi_h - \sum_{h=1}^{M} \pi_h\right) \\ \text{s.t.} \quad & \boldsymbol{\pi} \in \mathbb{R}^M_+,\ \boldsymbol{z} \in [0,1]^M \\ & \rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\pi_h + (1 - z_h)\xi_h\right) \leq \rho^{\alpha_h}_{\hat{\mathbb{P}}_{h,N}}\left(\xi_h\right) \quad \forall h \in [M] \end{aligned}$$

Calibrating ϵ

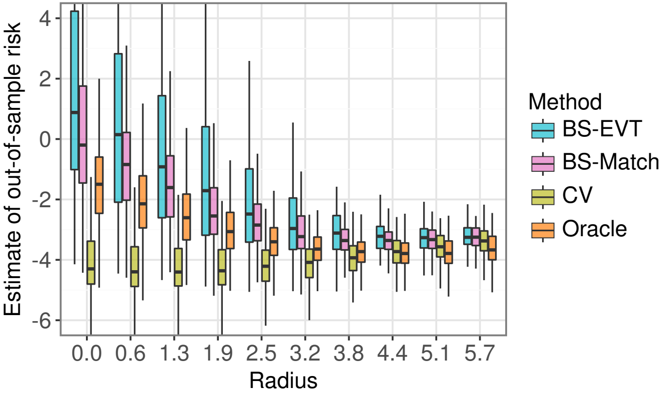

Bootstrapping applied on Cross Validation data

- Vanilla CV systematically chooses smaller ϵ
- Proposed methods choose similar ϵ with oracle

Calibrating ϵ

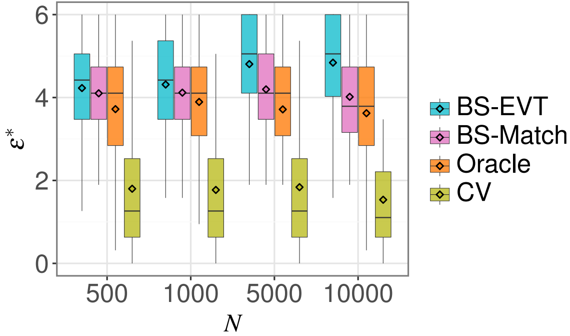

Bootstrapping applied on Cross Validation data

- Vanilla CV systematically chooses smaller ϵ
- Proposed method chooses similar ϵ with oracle

Effect of Sample Size N

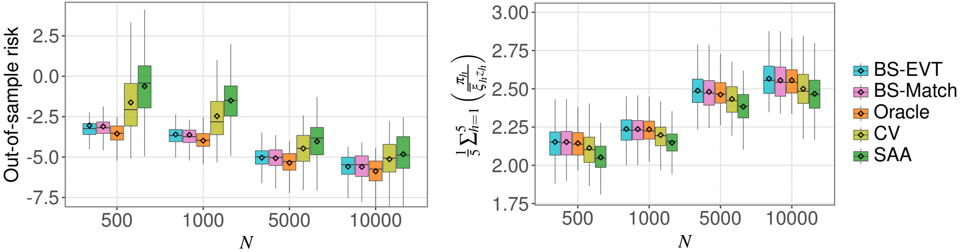

As N increases:

- both insurer and households are more capable of estimating their risk
- insurer extracts higher premiums for the same coverage level

Effect of Correlation

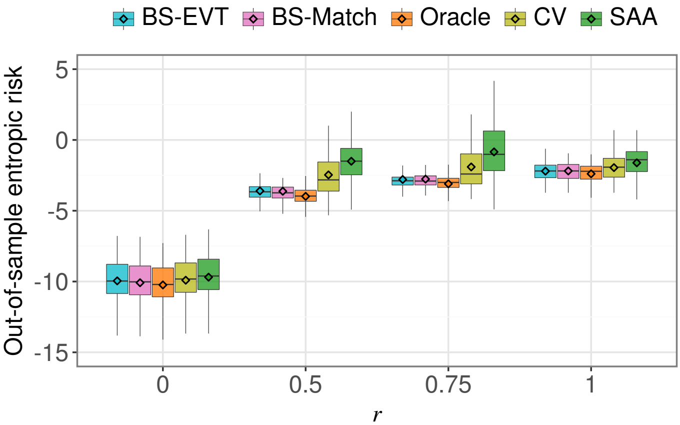

As r increases

- extreme loss events are more likely to occur simultaneously
- insurer's risk exposure increases

Effect of Correlation

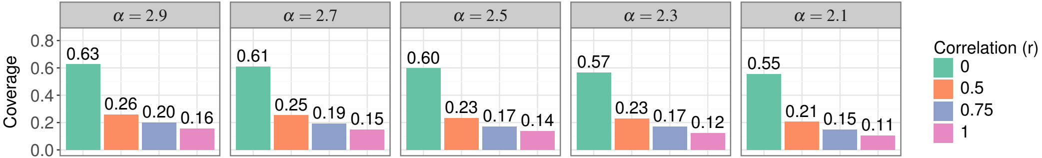

(a) `BS-Match`.

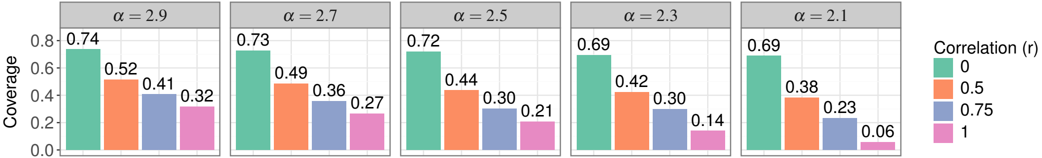

(b) `SAA`.

- r increases $\Longrightarrow$ risk pooling diminishes $\Longrightarrow$ insurer reduces coverage
- **`SAA` offers higher coverage due to underestimation of risk**

Key Takeaways

Estimating risk using sample averages can make us highly optimistic!

- Bias mitigation can partially address this issue
- Future research could aim to reduce the variance of the SAA estimator

DRO techniques can alleviate optimistic bias in decision-making

- still requires debiasing techniques for calibration
- effective calibration techniques need to be cheap
  - number of folds in K-Fold CV
  - computational complexity of bias correction

utsav.sadana@umontreal.ca
https://utsav19.github.io/

References

Arrow, K. J. (1963).Uncertainty and the welfare economics of medical care. *American Economic Review*, *53*(5), 941–973.

Bernard, C., Liu, F., & Vanduffel, S. (2020).Optimal insurance in the presence of multiple policyholders. *Journal of Economic Behavior & Organization*, *180*, 638–656.

Chen, L., & Sim, M. (2024).Robust CARA optimization [Forthcoming]. *Operations Research*.

Efron, B. (1979).Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, *7*(1), 1–26.

Gupta, V., Huang, M., & Rusmevichientong, P. (2024).Debiasing in-sample policy performance for small-data, large-scale optimization. *Operations Research*, *72*(2), 848–870.

Howard, R. A., & Matheson, J. E. (1972).Risk-sensitive Markov decision processes. *Management science*, *18*(7), 356–369.

Iyengar, G., Lam, H., & Wang, T. (2023).Optimizer's information criterion: Dissecting and correcting bias in data-driven optimization. *arXiv preprint arXiv:2306.10081*.

Kim, J. H. T. (2010).Bias correction for estimated distortion risk measure using the bootstrap. *Insurance: Mathematics and Economics*, *47*(2), 198–205.

Lam, H., & Mottet, C. (2017).Tail analysis without parametric models: A worst-case perspective. *Operations Research*, *65*(6), 1696–1711.

Li, T., Beirami, A., Sanjabi, M., & Smith, V. (2023).On Tilted Losses in Machine Learning: Theory and Applications. *Journal of Machine Learning Research*, *24*(142), 1–79.

References (cont.)

McNeil, A. J., Frey, R., & Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.

Smith, J. E., & Winkler, R. L. (2006).The Optimizer's Curse: Skepticism and Postdecision Surprise in Decision Analysis. *Management Science*, *52*(3), 311–322.

Troop, D., Godin, F., & Yu, J. Y. (2021). Bias-corrected peaks-over-threshold estimation of the CVaR. In C. de Campos & M. H. Maathuis (Eds.), *Uncertainty in artificial intelligence* (pp. 1809–1818, Vol. 161). PMLR.

Xu, L., & Jordan, M. I. (1996).On Convergence Properties of the EM Algorithm for Gaussian Mixtures. *Neural Computation*, *8*(1), 129–151.