A Distribution Optimization Framework for Risk Estimation

Hao Liang, Zhi-Quan Luo

The Chinese University of Hong Kong, Shenzhen

2024 INFORMS Optimization Society Conference
24th, March

Hao Liang, and Zhi-Quan Luo. "**A distribution optimization framework for confidence bounds of risk measures**." *International Conference on Machine Learning*. PMLR, 2023.

Risk-aware Decision-Making

Risk awareness is critical in some applications

Healthcare/Clinical trial

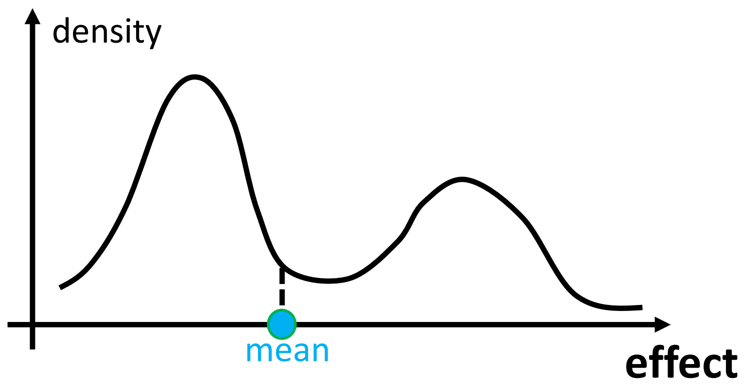

Risk-aware Decision-Making

Risk awareness is critical in some applications

Healthcare/Clinical trial

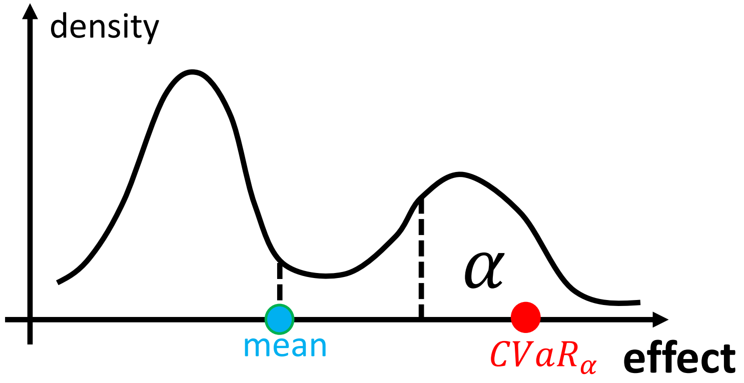

Avoid extreme negative outcomes

Risk-aware Decision-Making

Risk awareness is critical in some applications

Healthcare/Clinical trial

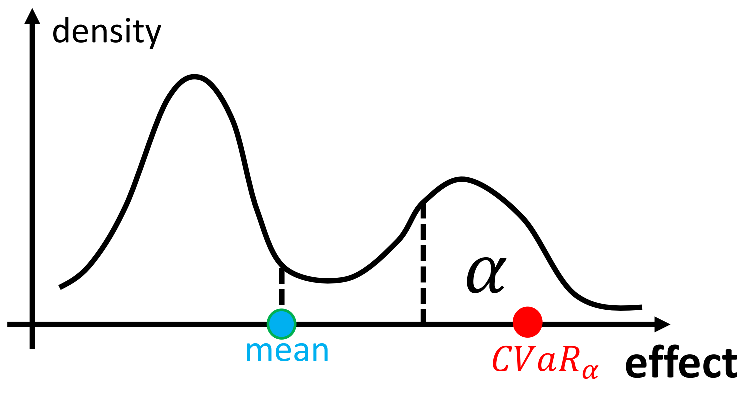

Avoid extreme negative outcomes

Finance

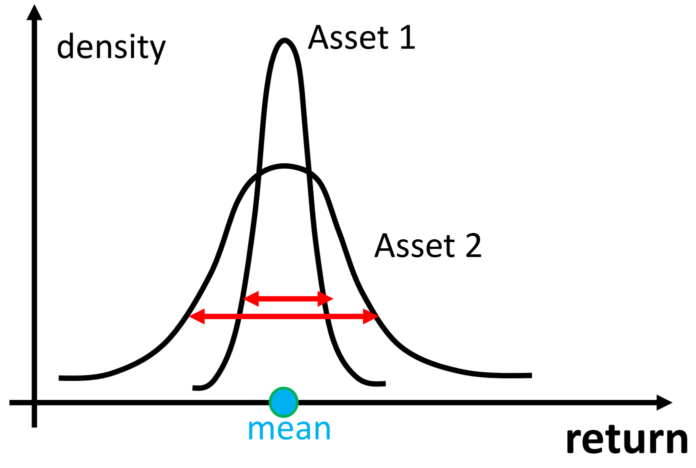

Control violation

Risk-aware Decision-Making

Risk awareness is critical in some applications

Healthcare/Clinical trial

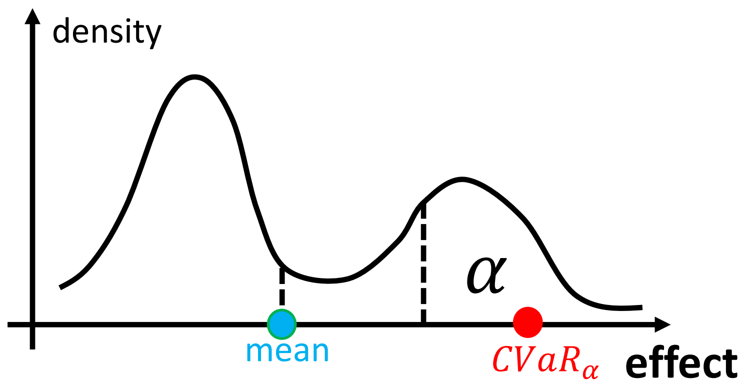

Finance

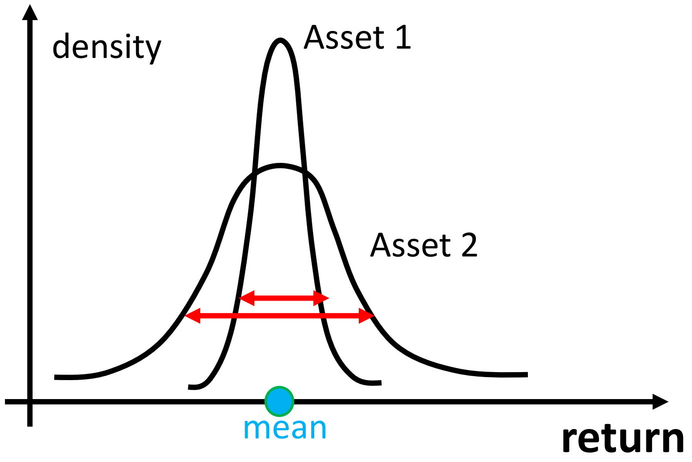

~~Mean~~ Risk measures captures certain distributional characteristics

- Tail mean: extreme negative outcomes
- Higher order moment: violation

Estimation of Risk Measures (RM)

RM reflects the risk preference towards uncertainty

- Entropic risk measure: violation
- Conditional Value at Risk: extreme negative outcomes

Estimation of Risk Measures (RM)

RM reflects the risk preference towards uncertainty

- Entropic risk measure: violation
- Conditional Value at Risk: extreme negative outcomes

Cannot evaluate the risk measure exactly

- Underlying distribution is unknown
- Estimate from finite samples

Estimation of Risk Measures (RM)

RM reflects the risk preference towards uncertainty

- Entropic risk measure: violation
- Conditional Value at Risk: extreme negative outcomes

Cannot evaluate the risk measure exactly

- Underlying distribution is unknown
- Estimate from finite samples

Confidence interval (CI) certifies a reliable range in risk-aware context

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
  - CVaR: α fractional tail mean
  $$\boldsymbol{C}_{\alpha}(F) \coloneqq \mathbf{E}[X|X \geq F^{-1}(1-\alpha)]$$
  - ERM:
  $$\boldsymbol{U}_{\boldsymbol{\beta}}(F) \coloneqq \frac{1}{\beta}\log \mathbf{E}_{X\sim F}[\exp(\beta X)]$$
  $$= \mathbf{E}[X] + \frac{\beta}{2}\mathbf{V}[X] + O(|\beta|^2)$$

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
- Given n iid samples $X_1, X_2, \cdots, X_n \sim F, X \in [a, b]$

$$X^n \coloneqq X_1, X_2, \cdots, X_n$$

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
- Given n iid samples $X_1, X_2, \cdots, X_n \sim F, X \in [a, b]$

$$X^n \coloneqq X_1, X_2, \cdots, X_n$$

- Goal: Derive **CI** of $\rho(F)$ given X^n

$$l(\rho) \le \rho(F) \le u(\rho), \text{ w.h.p.}$$

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
- Given n iid samples $X_1, X_2, \cdots, X_n \sim F, X \in [a, b]$

$$X^n := X_1, X_2, \cdots, X_n$$

- Goal: Derive **CI** of $\rho(F)$ given X^n

$$l(\rho) \leq \rho(F) \leq u(\rho), \text{ w.h.p.}$$

Classical concentration bounds on mean

$$|\mu - \hat{\mu}_n| \leq c(\mu) \Longrightarrow$$

$$\hat{\mu}_n - c(\mu) \leq \mu \leq \hat{\mu}_n + c(\mu)$$

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
- Given n iid samples $X_1, X_2, \cdots, X_n \sim F, X \in [a, b]$

$$X^n := X_1, X_2, \cdots, X_n$$

- Goal: Derive **CI** of $\rho(F)$ given X^n

$$l(\rho) \leq \rho(F) \leq u(\rho), \text{ w.h.p.}$$

Classical concentration bounds on mean

$$|\mu - \hat{\mu}_n| \leq c(\mu) \Longrightarrow$$

$$\hat{\mu}_n - c(\mu) \leq \mu \leq \hat{\mu}_n + c(\mu)$$

LCB $l(\mu)$ ↑ **UCB** $u(\mu)$ ↑

Towards Tight Confidence Intervals

Problem setting

- A risk measure ρ assigns risk value $\rho(F)$ to a distribution F
- Given n iid samples $X_1, X_2, \cdots, X_n \sim F, X \in [a, b]$

$$X^n := X_1, X_2, \cdots, X_n$$

- Goal: Derive **CI** of $\rho(F)$ given X^n

$$l(\rho) \leq \rho(F) \leq u(\rho), \text{ w.h.p.}$$

Classical concentration bounds on mean

$$|\mu - \hat{\mu}_n| \leq c(\mu) \Rightarrow$$

$$\hat{\mu}_n - c(\mu) \leq \mu \leq \hat{\mu}_n + c(\mu)$$

LCB $l(\mu)$

UCB $u(\mu)$

→ **Generalization to risk measures?**

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

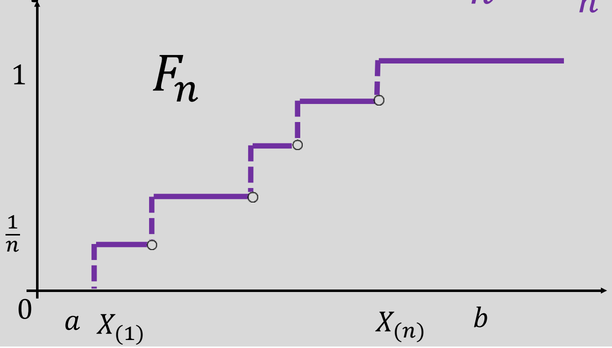

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

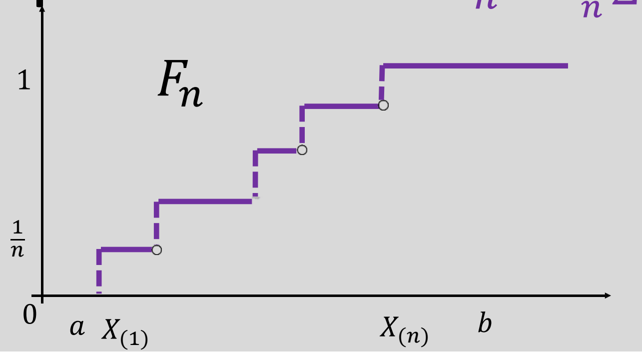

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \leq c_p$$

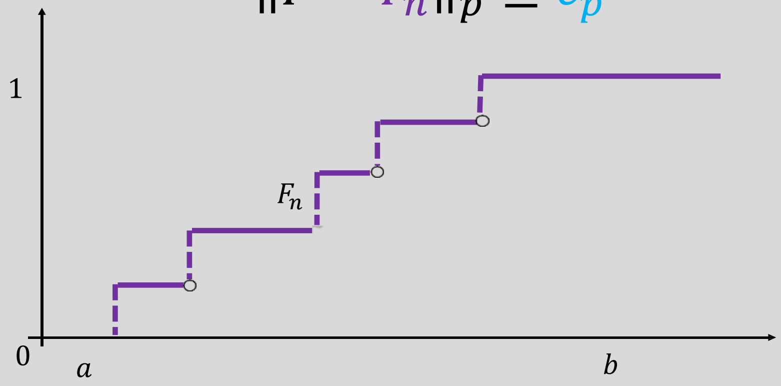

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

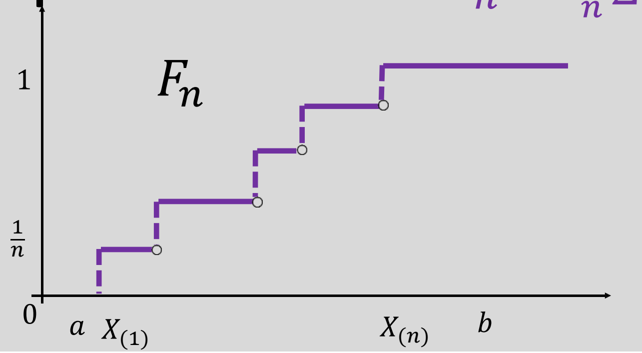

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \leq c_p$$

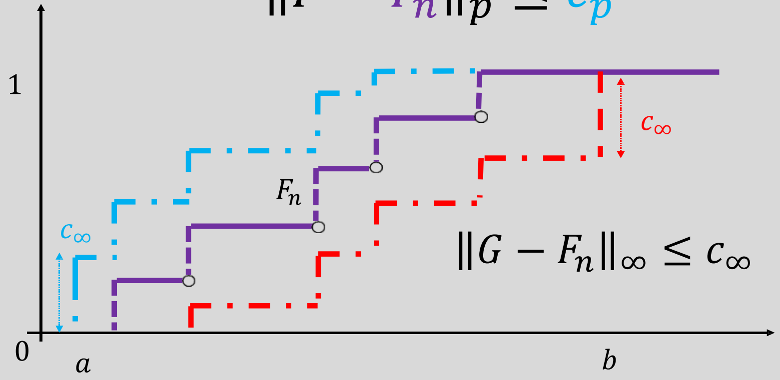

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

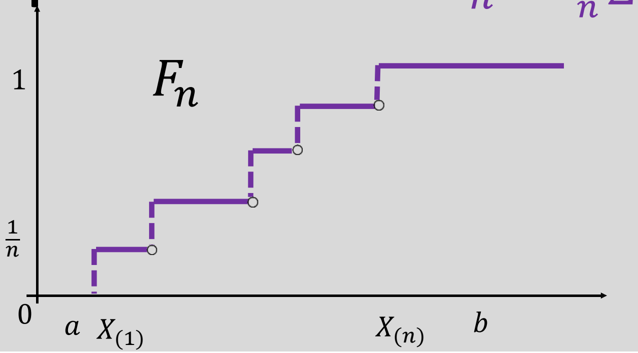

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \leq c_p$$

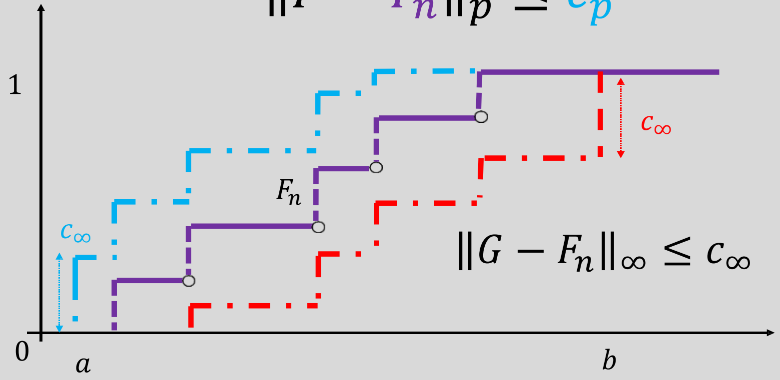

Step 3: Global Lipschitz constant (GLC)

$$|\rho(F) - \rho(G)| \leq \mathrm{GLC} \cdot \|F - G\|_p, \forall F, G$$

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

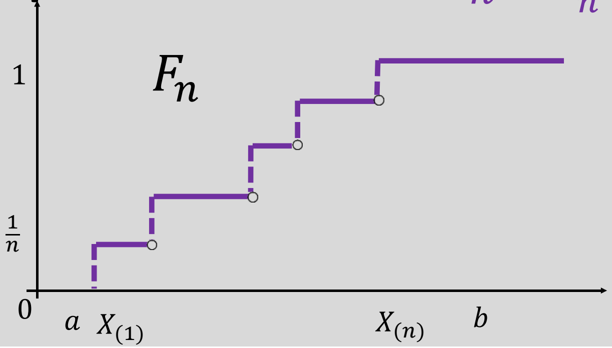

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \leq c_p$$

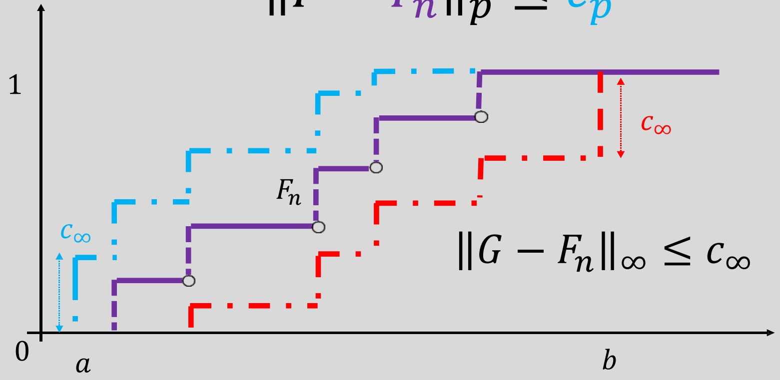

Step 3: Global Lipschitz constant (GLC)

$$|\rho(F) - \rho(G)| \leq \mathrm{GLC} \cdot \|F - G\|_p, \forall F, G$$

Step 4: Global linearization

$$|\rho(F) - \rho(F_n)| \leq \mathrm{GLC} \cdot \|F - F_n\|_p \leq \mathrm{GLC} \cdot c_p$$

$$\rho(F_n) - \mathrm{GLC} \cdot c_p \leq \rho(F) \leq \rho(F_n) + \mathrm{GLC} \cdot c_p$$

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

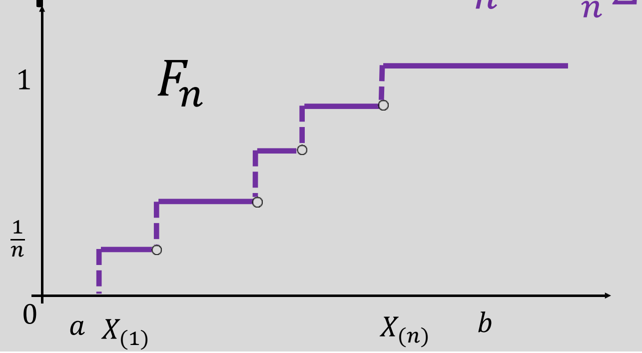

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \leq c_p$$

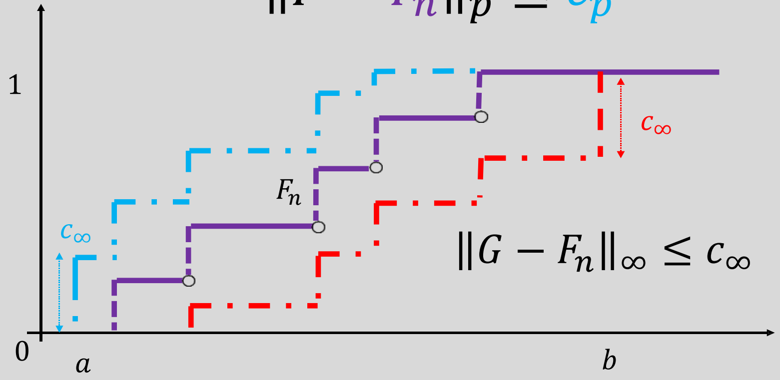

Step 3: Global Lipschitz constant (GLC)

$$|\rho(F) - \rho(G)| \leq \mathrm{GLC} \cdot \|F - G\|_p, \forall F, G$$

Step 4: Global linearization

$$|\rho(F) - \rho(F_n)| \leq \mathrm{GLC} \cdot \|F - F_n\|_p \leq \mathrm{GLC} \cdot c_p$$

$$\rho(F_n) - \mathrm{GLC} \cdot c_p \leq \rho(F) \leq \rho(F_n) + \mathrm{GLC} \cdot c_p$$

LCB $l(\rho)$ **UCB** $u(\rho)$

Existing Method: Global Lipschitz Constant

Step 1: Empirical distribution $F_n := \frac{1}{n}\sum_{i=1} \delta_{X_i}$

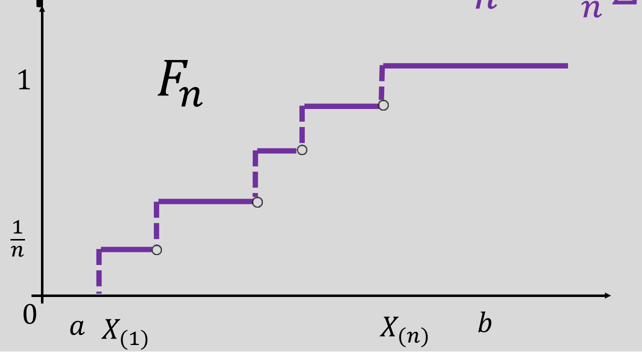

Step 2: Concentration bound on F_n

$$\|F - F_n\|_p \le c_p$$

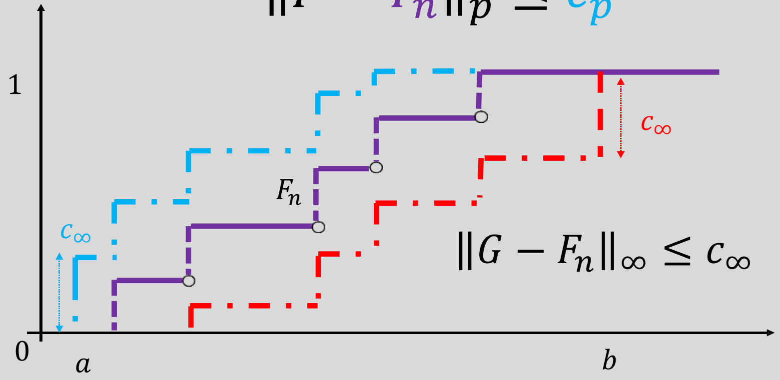

Step 3: Global Lipschitz constant (GLC)

$$|\rho(F) - \rho(G)| \le \mathrm{GLC} \cdot \|F - G\|_p, \forall F, G$$

Step 4: Global linearization

$$|\rho(F) - \rho(F_n)| \le \mathrm{GLC} \cdot \|F - F_n\|_p \le \mathrm{GLC} \cdot c_p$$

$$\rho(F_n) - \mathrm{GLC} \cdot c_p \le \rho(F) \le \rho(F_n) + \mathrm{GLC} \cdot c_p$$

LCB $l(\rho)$

UCB $u(\rho)$

Limitations

- GLC is not easy to obtain
- Loose due to global linearization of nonlinear risk measure

Optimization Framework: A Distributional Perspective

risk value	→	objective
distribution	→	opt variable
concentration	→	constraint

Optimization Framework: A Distributional Perspective

risk value → objective
distribution → opt variable
concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t.} \ \|G - F_n\|_p \leq c_p$$

Optimization Framework: A Distributional Perspective

risk value → objective

distribution → opt variable

concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t.} \ \|G - F_n\|_p \leq c_p$$

Optimal solutions

- Maximizer $\overline{F_n^p}$
- Minimizer $\underline{F_n^p}$

Optimization Framework: A Distributional Perspective

risk value → objective
distribution → opt variable
concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t.} \ \|G - F_n\|_p \le c_p$$

Optimal solutions

- Maximizer $\overline{F_n^p}$
- Minimizer $\underline{F_n^p}$

Optimal value

- UCB $\rho\left(\overline{F_n^p}\right)$
- LCB $\rho\left(\underline{F_n^p}\right)$

Optimization Framework: A Distributional Perspective

risk value → objective

distribution → opt variable

concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t. } \|G - F_n\|_p \leq c_p$$

Optimal solutions

- Maximizer $\overline{F_n^p}$
- Minimizer $\underline{F_n^p}$

Optimal value

- UCB $\rho\left(\overline{F_n^p}\right)$
- LCB $\rho\left(\underline{F_n^p}\right)$

Tighter CI

$$\rho\left(\overline{F_n^p}\right) \leq \rho(F_n) + \mathrm{LC} \cdot \left\|\overline{F_n^p} - F_n\right\| \leq \rho(F_n) + \mathrm{LC} \cdot c_n^p$$

Optimization Framework: A Distributional Perspective

risk value	→	objective
distribution	→	opt variable
concentration	→	constraint

$$\max \ \rho(G)$$
$$\text{s.t.} \ \|G - F_n\|_p \leq c_p$$

$$\rho\left(\underline{F_n^p}\right) \leq \rho(F) \leq \rho\left(\overline{F_n^p}\right)$$

$$\|G - F_n\|_\infty = \sup_x |G(x) - F_n(x)| \leq c_\infty$$

Optimization Framework: A Distributional Perspective

risk value → objective

distribution → opt variable

concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t. } \|G - F_n\|_p \le c_p$$

$$\rho\left(\underline{F_n^p}\right) \le \rho(F) \le \rho\left(\overline{F_n^p}\right)$$

$$\|G - F_n\|_\infty = \sup_x |G(x) - F_n(x)| \le c_\infty$$

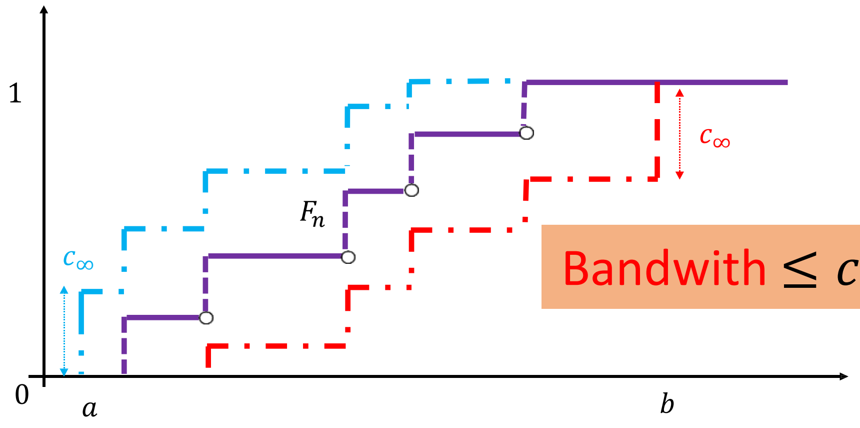

Bandwith $\le c_\infty$

Optimization Framework: A Distributional Perspective

risk value → objective

distribution → opt variable

concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t. } \|G - F_n\|_p \le c_p$$

$$\rho\left(\underline{F_n^p}\right) \le \rho(F) \le \rho\left(\overline{F_n^p}\right)$$

$$\|G - F_n\|_\infty = \sup_x |G(x) - F_n(x)| \le c_\infty$$

$$\|G - F_n\|_1 = \int_a^b |G(x) - F_n(x)| dx \le c_1$$

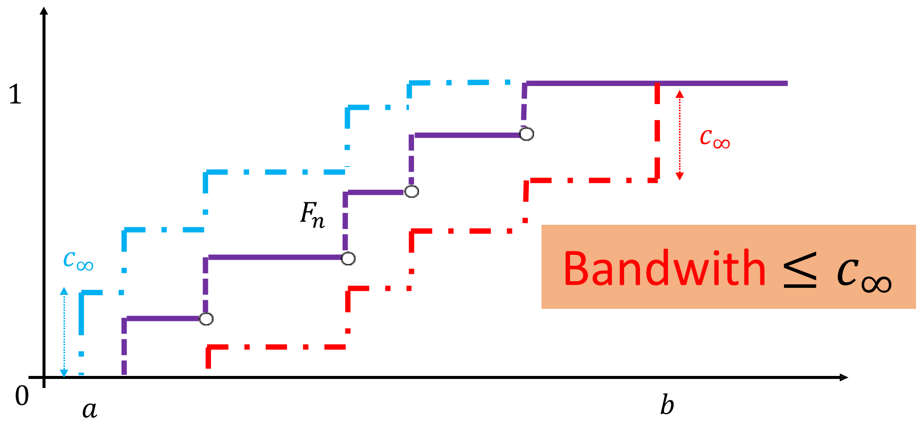

Optimization Framework: A Distributional Perspective

risk value → objective

distribution → opt variable

concentration → constraint

$$\max \ \rho(G)$$
$$\text{s.t. } \|G - F_n\|_p \le c_p$$

$$\rho\left(\underline{F_n^p}\right) \le \rho(F) \le \rho\left(\overline{F_n^p}\right)$$

$$\|G - F_n\|_\infty = \sup_x |G(x) - F_n(x)| \le c_\infty$$

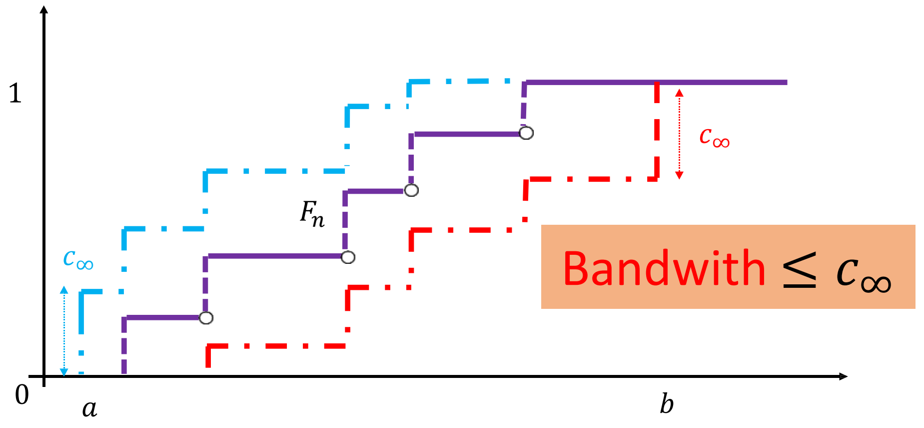

$$\|G - F_n\|_1 = \int_a^b |G(x) - F_n(x)|dx \le c_1$$

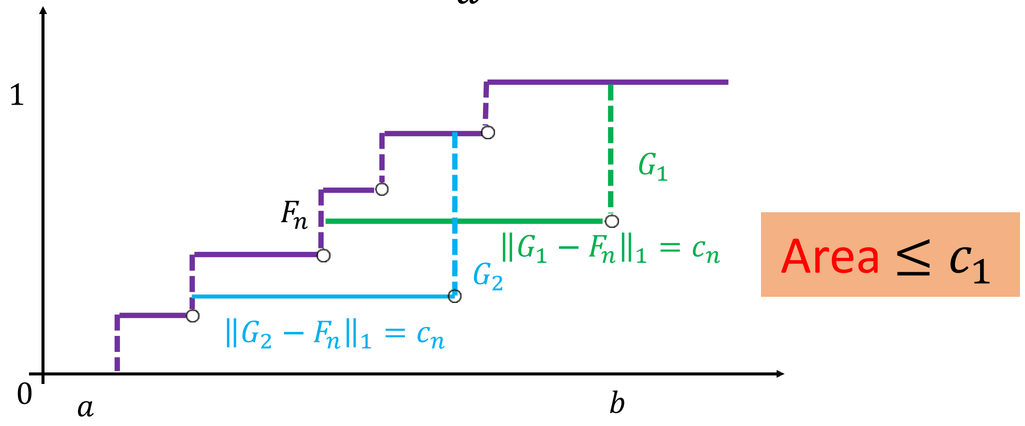

Closed-form Solution

Challenges

- Infinite-dimensional CDF
- Norm ball constraint
- Diverse and nonlinear risk measures

Closed-form Solution

Challenges

- Infinite-dimensional CDF
- Norm ball constraint
- Diverse and nonlinear risk measures

→ **How to obtain optimal solution?**

Closed-form Solution

Challenges

- Infinite-dimensional CDF
- Norm ball constraint
- Diverse and nonlinear risk measures

→ **How to obtain optimal solution?**

Closed form as **transformation** of F_n, computationally efficient

Closed-form Solution

Challenges
- Infinite-dimensional CDF
- Norm ball constraint
- Diverse and nonlinear risk measures

→ **How to obtain optimal solution?**

Closed form as **transformation** of F_n, computationally efficient

$$\|G - F_n\|_\infty = \sup_x |G(x) - F_n(x)| \leq c_\infty$$

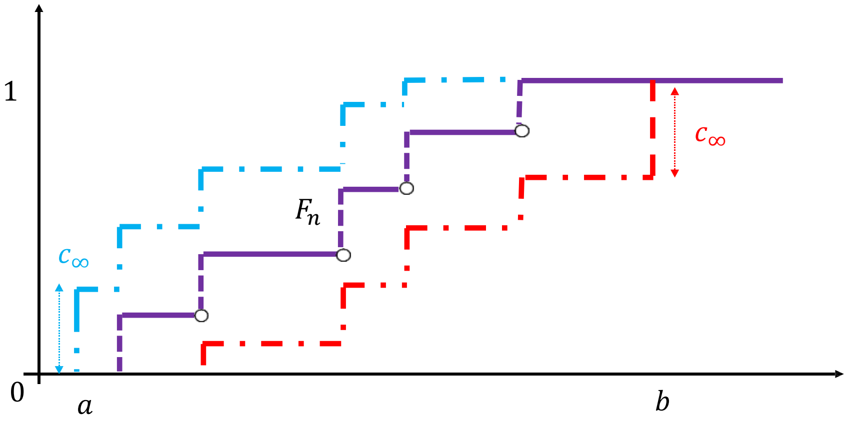

Envelope

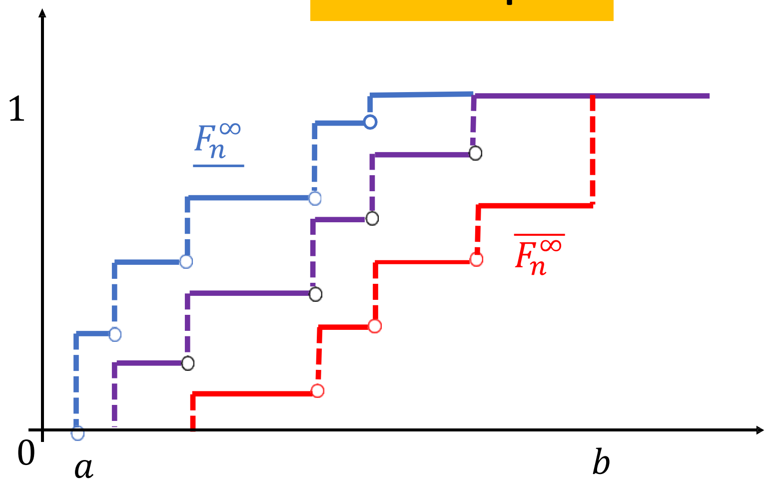

Closed-form Solution

Challenges

- Infinite-dimensional CDF
- Norm ball constraint
- Diverse and nonlinear risk measures

→ **How to obtain optimal solution?**

Closed form as **transformation** of F_n, computationally efficient

$$\|G - F_n\|_1 = \int_a^b |G(x) - F_n(x)|dx \leq c_1$$

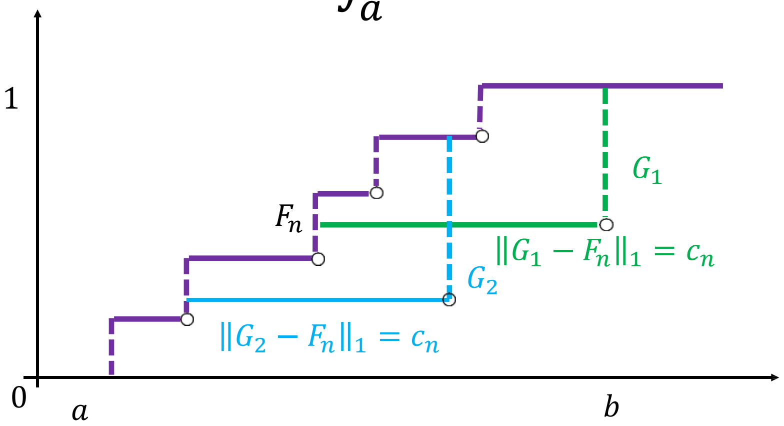

Water-filling

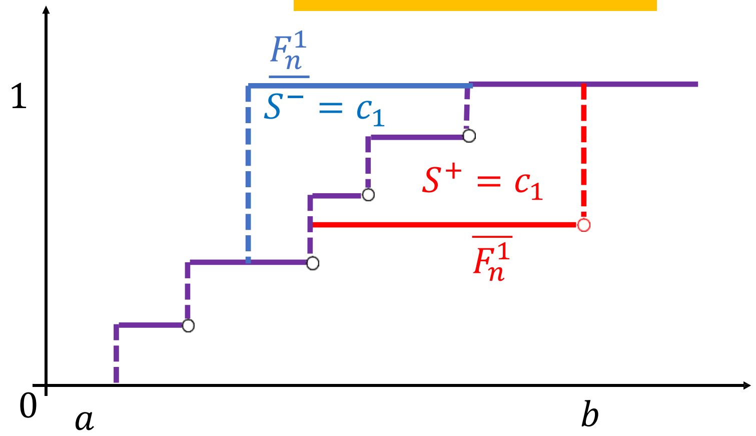

Intrinsic Optimality

Ours vs. GLC

$$\rho\left(\overline{F_n^p}\right) < \rho(F_n) + \text{GLC} \cdot c_n^p$$

$$\text{GLC} = \sup_{G,G' \in D(a,b)} \frac{\rho(G) - \rho(G')}{\|G - G'\|}$$

→

$$\text{LLC} = \sup_{G,G' \in B(F,c)} \frac{\rho(G) - \rho(G')}{\|G - G'\|}$$

Intrinsic Optimality

Ours vs. LLC vs. GLC

$$\rho\left(\overline{F_n^p}\right) < \rho(F_n) + \mathrm{LLC} \cdot c_n^p < \rho(F_n) + \mathrm{GLC} \cdot c_n^p$$

$$\mathrm{GLC} = \sup_{G,G' \in D(a,b)} \frac{\rho(G) - \rho(G')}{\|G - G'\|} \longrightarrow \mathrm{LLC} = \sup_{G,G' \in B(F,c)} \frac{\rho(G) - \rho(G')}{\|G - G'\|}$$

Intrinsic Optimality

Our bounds improve the **tightest** Local Lipschitz Constant (LLC)!

$$\rho\left(\overline{F_n^p}\right) < \rho(F_n) + \mathrm{LLC} \cdot c_n^p < \rho(F_n) + \mathrm{GLC} \cdot c_n^p$$

LLC vs. GLC

RM	LLC ($p = \infty$)	GLC ($p = \infty$)	Improvement
CVaR	$\frac{b - F_n^{-1}((1-\alpha-c)^+)}{\alpha}$	$\frac{b-a}{\alpha}$	✓
SRM	$\left\lVert \phi(\underline{F_n^\infty}) \right\rVert_1$	$(b-a)\phi(1)$	✓
DRM	$\left\lVert g'(1 - \underline{F_n^\infty}) \right\rVert_1$	$(b-a)\,\lVert g' \rVert_\infty$	✓
ERM	$\frac{\exp(\beta b) - \exp(\beta a)}{\beta \int_a^b \exp(\beta x) d\overline{\underline{F_n^\infty}}(x))}$	$\frac{\exp(\beta(b-a)) - 1}{\beta}$	✓
RDEU	$\left\lVert w'(\underline{F_n^\infty}) v' \right\rVert_1$	$\lVert w' \rVert_\infty \lVert v' \rVert_1$	✓

LLC vs. Ours

RM	CVaR	SRM	DRM	ERM	RDEU
LLC	$\frac{b - F^{-1}(1-\alpha-c)}{\alpha}$	$\lVert \phi(\underline{F^\infty}) \rVert_1$	$\lVert g'(1 - \underline{F^\infty}) \rVert_1$	$\frac{\exp(\beta b) - \exp(\beta a)}{\int_a^b \exp(\beta x) d\underline{F^\infty}(x)}$	$\lVert w'(\underline{F^\infty}) v' \rVert_1$
$\frac{\mathbf{T}(\overline{F^\infty}) - \mathbf{T}(F)}{c}$	$\frac{b - F^{-1}(1-\alpha)}{\alpha}$	$\lVert \phi(F) \rVert_1$	$\lVert g'(1 - F) \rVert_1$	$\frac{\exp(\beta b) - \exp(\beta a)}{\beta \int_a^b \exp(\beta x) dF(x)}$	$\lVert w'(F) v' \rVert_1$
Improvement	✓	✓	✓	✓	✓

Numerical Experiments

Comparisons of CIs for **CVaR** and **ERM** with varying sample sizes

CVaR UCB w/ $\|\cdot\|_\infty$

ERM UCB w/ $\|\cdot\|_1$

CVaR LCB w/ $\|\cdot\|_\infty$

ERM LCB w/ $\|\cdot\|_1$

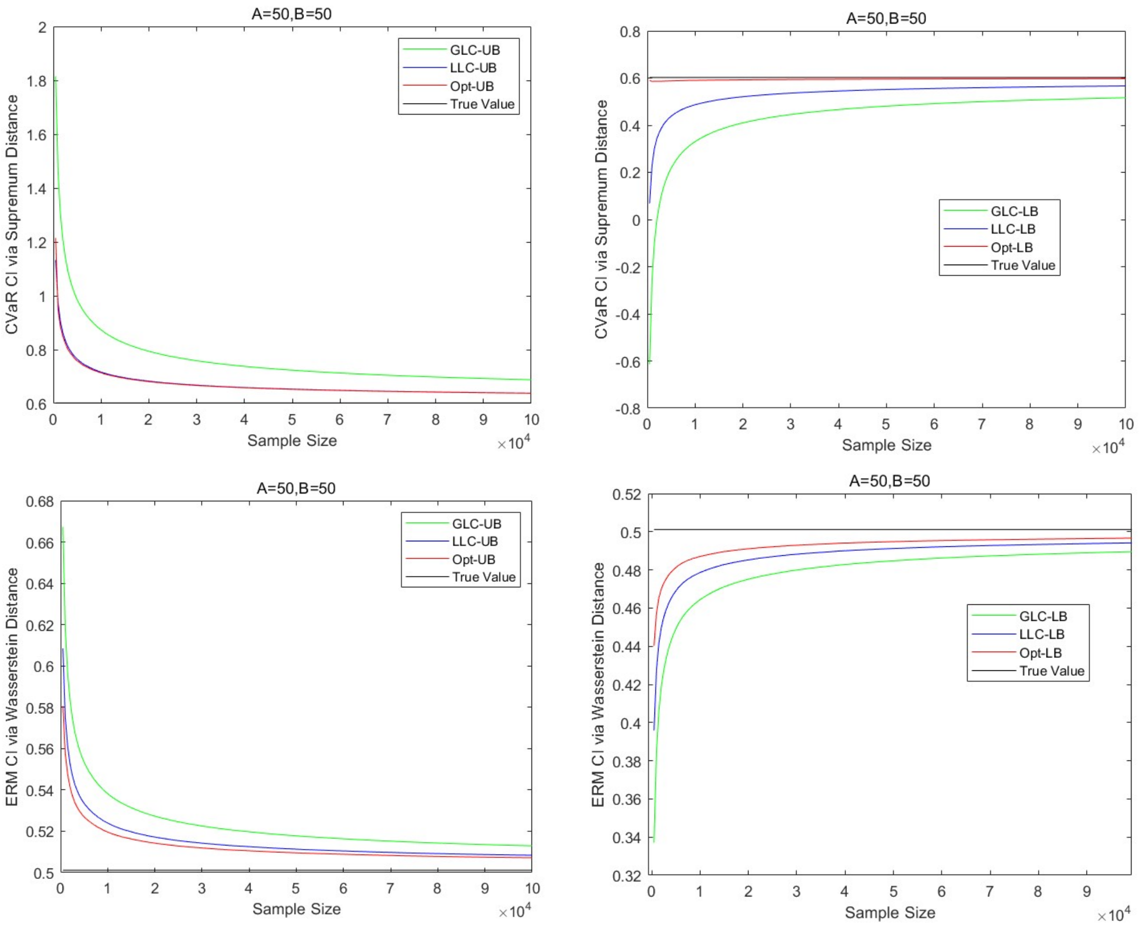

Summary

- **Tight** CI
- Apply to **broad** classes of RMs
  - **Spectral risk measure**, including **Conditional value at risk**
  - **Distortion risk measure**
  - **Certainty equivalent**, including **Entropic risk measure**
  - **Rank-dependent expected utility**
- **LC-free** for specific RM

  Common optimal solutions for different RMs
- Computationally **efficient**
- Motivate improved risk-aware bandit algorithms

Applications to Risk-aware Multi-armed Bandits

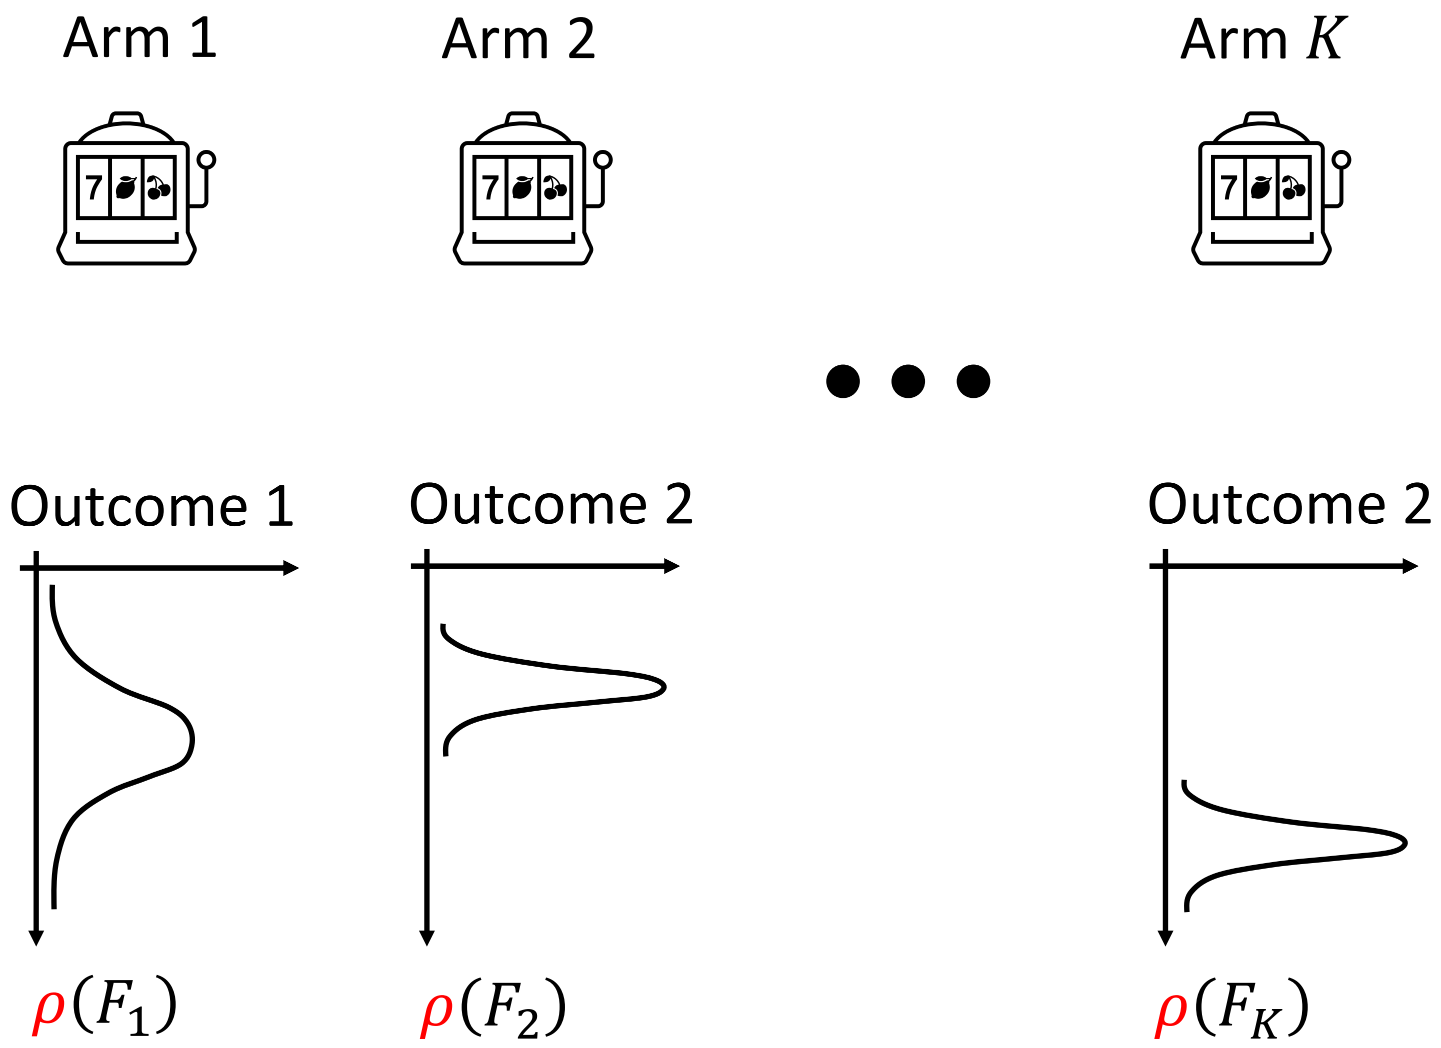

Applications to Risk-aware Multi-armed Bandits

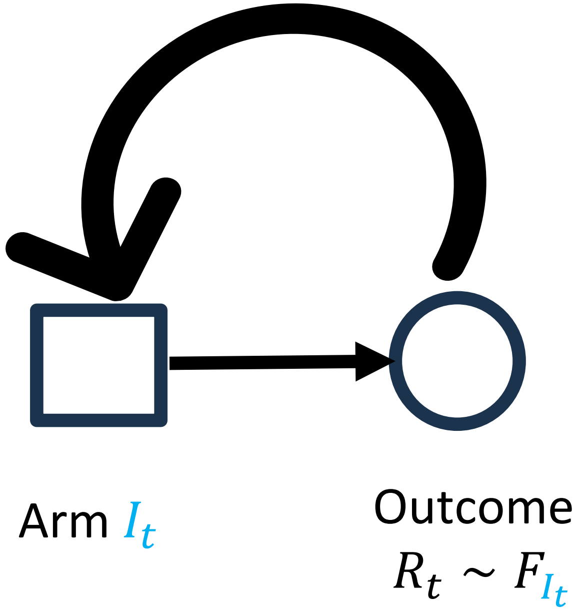

Applications to Risk-sensitive Multi-armed Bandits

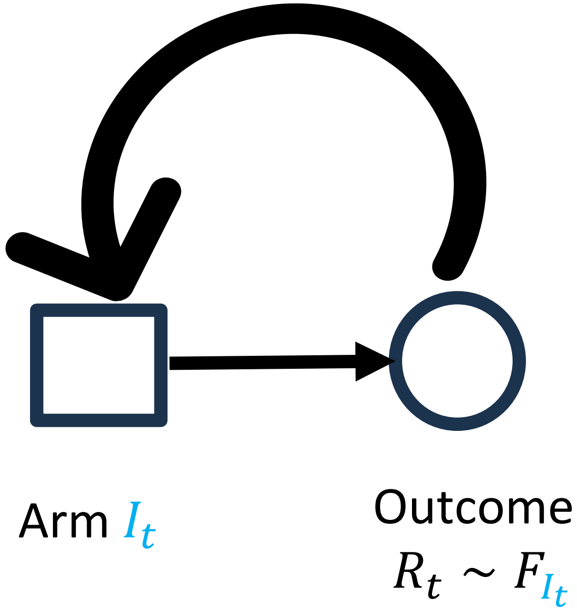

Maximize cumulative value

$$\sum_{t=1}^{N} \rho(F_{I_t})$$

Applications to Risk-sensitive Multi-armed Bandits

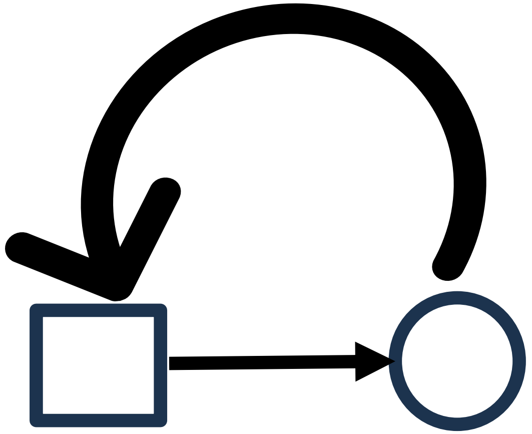

Arm I_t

Outcome

$R_t \sim F_{I_t}$

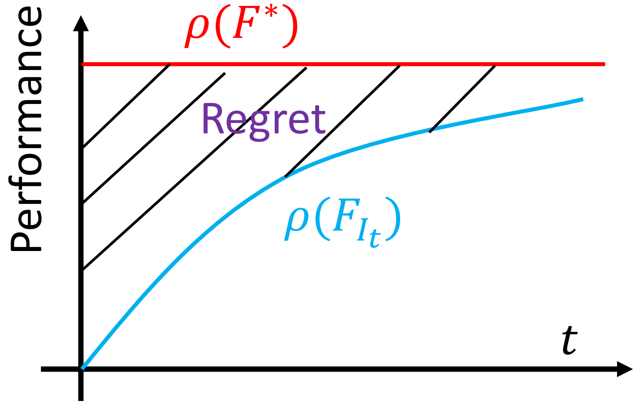

Maximize cumulative value

$$\sum_{t=1}^{N} \rho(F_{I_t})$$

Applications to **Risk-sensitive** Multi-armed Bandits

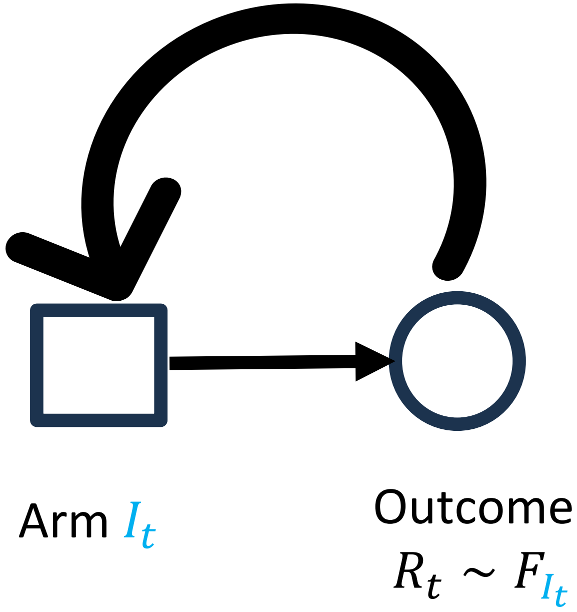

Maximize cumulative value

$$\sum_{t=1}^{N} \rho(F_{I_t})$$

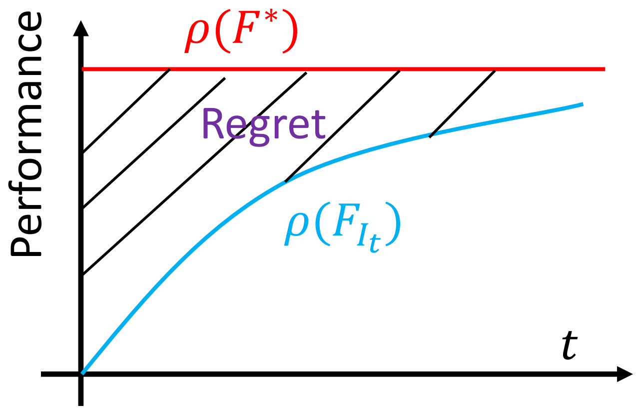

Low regret

↓

High sample efficiency

Optimism in Face of Uncertainty

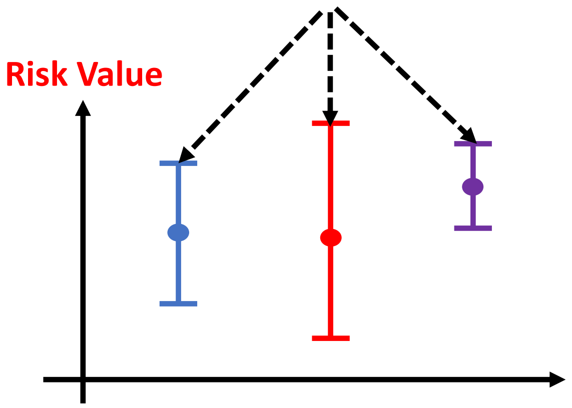

- Act greedily w.r.t. **Upper Confidence Bound** of **Risk Value**

- Encourages exploring actions with the best possible outcomes

Tighter UCB ⟶ Higher efficiency

A Meta Bandit Algorithm for Generic Risk Measures

Lower Confidence Band

For $t = 1:N$

- Maintain EDF for each arm $\widehat{F_{i,t}}$
- Compute $\overline{F_{i,t}}$ for each $\widehat{F_{i,t}}$
- Choose action

$$I_t = \mathrm{argmax}_{i \in [K]}\, \rho\left(\overline{F_{i,t}}\right)$$

- Observe $R_t \sim F_{I_t}$

- Applicable to generic risk measures
- Reduce to CVaR-UCB in [1]
- Provable gain over GLC-UCB in [2]

[1] Tamkin, A., Keramati, R., Dann, C., and Brunskill, E. Distributionally-aware exploration for cvar bandits.
In NeurIPS 2019 Workshop on Safety and Robustness on Decision Making, 2019.
[2] Cassel, A., Mannor, S., and Zeevi, A. A general approach to multi-armed bandits under risk criteria.
In Conference On Learning Theory, pp. 1295–1306. PMLR, 2018.

A Meta Bandit Algorithm for Generic Risk Measures

Upper Confidence Band

For $t = 1 : N$

- Maintain EDF for each arm $\widehat{F_{i,t}}$
- Compute $\overline{F_{i,t}}$ for each $\widehat{F_{i,t}}$
- Choose action

  $$I_t = \mathrm{argmax}_{i\in[K]}\, \rho\left(\overline{F_{i,t}}\right)$$
- Observe $R_t \sim F_{I_t}$

Regret gain

$$\frac{\sum_{i>1}\left(b - F_i^{-1}(1-\alpha-2c_i)\right)^2/\Delta_i^2}{\sum_{i>1}(b-a)^2/\Delta_i^2}$$

Numerical experiments

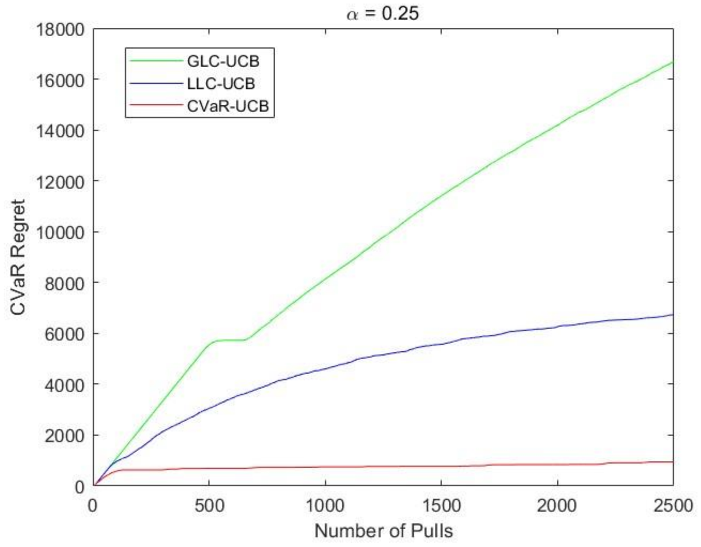

- Applicable to generic risk measures
- Reduce to CVaR-UCB in [1]
- Provable gain over GLC-UCB in [2]

[1] Tamkin, A., Keramati, R., Dann, C., and Brunskill, E. Distributionally-aware exploration for cvar bandits. *In NeurIPS 2019 Workshop on Safety and Robustness on Decision Making, 2019.*

[2] Cassel, A., Mannor, S., and Zeevi, A. A general approach to multi-armed bandits under risk criteria. *In Conference On Learning Theory, pp. 1295–1306. PMLR, 2018.*

Thank You

Now seeking postdoc and research positions!
Feel free to contact via Email
haoliang1@link.cuhk.edu.cn