A Distributionally Robust Optimization Framework with Regularization Control

Diego Fonseca[1] – Joint work with Mauricio Junca[2]

[1]Universidad EAFIT, Medellín, Colombia
[2]Universidad de los Andes, Bogotá, Colombia

XVIIth International Conference on Stochastic Programming (ICSP 2025)
Paris, France
August 1, 2025

Outline

Presentation Roadmap

1. **Motivation:** Why do we need a new approach to robustness?
2. **The Proposed Approach:** Adverse Scenario Regularization (ASR-SAA).
3. **Theoretical Foundations:** Equivalence to DRO & Guarantees.
4. **Numerical Study:** Mean-CVaR Portfolio Optimization.
5. **Conclusions.**

Motivation: The Core Problem

Stochastic Programming (SP)

The fundamental problem is to minimize an expected value:

$$\min_{x \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}}[F(x, \xi)]$$

In practice, the true distribution $\mathbb{P}$ is unknown and replaced by an empirical distribution $\hat{\mathbb{P}}_n$ from data $\{\hat{\xi}_1, \ldots, \hat{\xi}_n\}$.

Motivation: The Core Problem

Stochastic Programming (SP)

The fundamental problem is to minimize an expected value:

$$\min_{x \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}}[F(x, \xi)]$$

In practice, the true distribution $\mathbb{P}$ is unknown and replaced by an empirical distribution $\hat{\mathbb{P}}_n$ from data $\{\hat{\xi}_1, \ldots, \hat{\xi}_n\}$.

The Risk-Neutrality Gap

- Minimizing expectation is a **risk-neutral** stance.
- It is insensitive to the dispersion of outcomes.
- An optimal solution in expectation might exhibit high performance variability and poor behavior in **extreme events**.

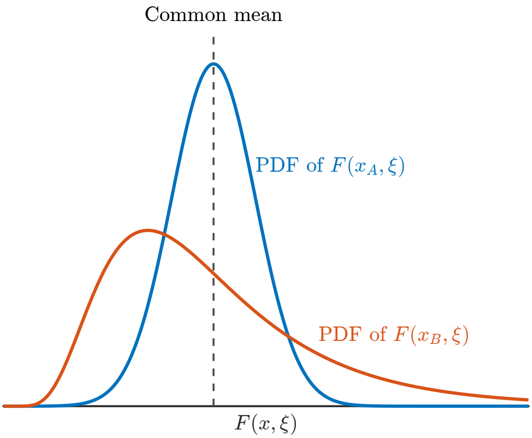

Figure: Same mean, different risk profiles.

Motivation: Two Notions of Robustness

1. Distributional Robustness

- **Goal:** Decisions that perform consistently well *on average* when evaluated out-of-sample.
- **Metric:** Low dispersion of out-of-sample expected performance $\mathbb{E}_{\xi\sim\mathbb{P}}[F(\hat{x},\xi)]$.

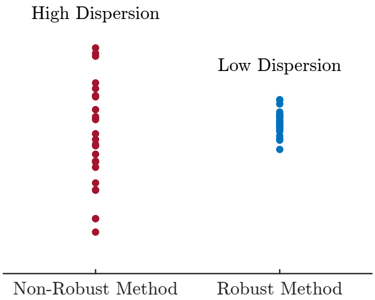

2. Robustness Against Extreme Outcomes

- **Goal:** Mitigate the severity of worst-case realizations of the cost $F(\hat{x},\xi)$.
- **Metric:** Low out-of-sample tail risk, e.g., low $\mathrm{CVaR}_\alpha(F(\hat{x},\xi))$.

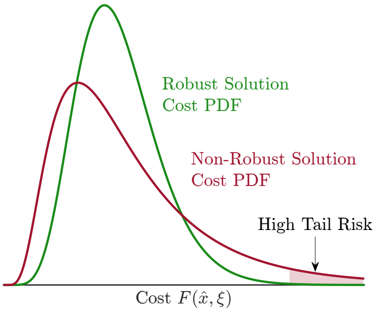

Achieving one type of robustness does not necessarily imply the other.

Motivation: The WDRO Approach

WDRO Formulation

WDRO hedges against misspecification by optimizing for the worst case within a Wasserstein ball $\mathcal{D} = \mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)$ of radius ε:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)} \mathbb{E}_{\xi \sim \mathbb{Q}}[F(x, \xi)]$$

Motivation: The WDRO Approach

WDRO Formulation

WDRO hedges against misspecification by optimizing for the worst case within a Wasserstein ball $\mathcal{D} = \mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)$ of radius ε:

$$\min_{x\in\mathcal{X}} \sup_{\mathbb{Q}\in\mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)} \mathbb{E}_{\xi\sim\mathbb{Q}}[F(x,\xi)]$$

The Bridge to Regularization: An Asymptotic View (Kuhn et al. 2024; Gao et al. 2022)

Under technical conditions (mainly **smoothness** of $F(x,\cdot)$ and some **integrability** assumptions), for a small radius ε, the WDRO objective can be approximated by:

$$\sup_{\mathbb{Q}\in\mathcal{B}_\varepsilon(\hat{\mathbb{P}}_n)} \mathbb{E}_{\mathbb{Q}}[F(x,\xi)] = \mathbb{E}_{\hat{\mathbb{P}}_n}[F(x,\xi)] + \varepsilon\cdot\left(\mathbb{E}_{\hat{\mathbb{P}}_n}[\|\nabla_\xi F(x,\xi)\|_*^q]\right)^{1/q} + o(\varepsilon)$$

The exponent q in the gradient norm is the **Hölder conjugate** of the Wasserstein order p.

Implication: The choice of p dictates the penalty.

- **1-WDRO** ($p=1 \implies q=\infty$): Penalizes the *maximum* gradient norm over the sample (a robust, worst-case penalty).

Motivation: The WDRO Approach

WDRO Formulation

WDRO hedges against misspecification by optimizing for the worst case within a Wasserstein ball $\mathcal{D} = \mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)$ of radius ε:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{B}_\varepsilon(\widehat{\mathbb{P}}_n)} \mathbb{E}_{\xi \sim \mathbb{Q}}[F(x, \xi)]$$

The Bridge to Regularization: An Asymptotic View (Kuhn et al. 2024; Gao et al. 2022)

Under technical conditions (mainly **smoothness** of $F(x, \cdot)$ and some **integrability** assumptions), for a small radius ε, the WDRO objective can be approximated by:

$$\sup_{\mathbb{Q} \in \mathcal{B}_\varepsilon(\hat{\mathbb{P}}_n)} \mathbb{E}_{\mathbb{Q}}[F(x, \xi)] = \mathbb{E}_{\hat{\mathbb{P}}_n}[F(x, \xi)] + \varepsilon \cdot \left(\mathbb{E}_{\hat{\mathbb{P}}_n}[\|\nabla_\xi F(x, \xi)\|_*^q]\right)^{1/q} + o(\varepsilon)$$

The exponent q in the gradient norm is the **Hölder conjugate** of the Wasserstein order p.

Implication: The choice of p dictates the penalty.

- **1-WDRO** ($p = 1 \implies q = \infty$): Penalizes the *maximum* gradient norm over the sample (a robust, worst-case penalty).
- **2-WDRO** ($p = 2 \implies q = 2$): Penalizes the *root-mean-square* of gradient norms (an average-type penalty).

A Motivating Example

We construct a problem to specifically compare 1-WDRO and 2-WDRO.

- **Question:** Does the choice of the Wasserstein order p affect robustness against extreme outcomes?
- **Finding:** Yes. The choice of p can lead to markedly different tail risk behaviors, even when the average performance is similar.

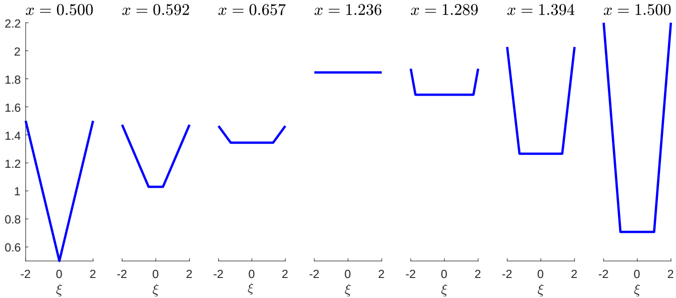

Figure: First, we design a cost function $F(x, \xi)$ whose shape transitions as the decision x varies, allowing different methods to find solutions with distinct risk profiles.

A Motivating Example

We construct a problem to specifically compare 1-WDRO and 2-WDRO.

- **Question:** Does the choice of the Wasserstein order p affect robustness against extreme outcomes?
- **Finding:** Yes. The choice of p can lead to markedly different tail risk behaviors, even when the average performance is similar.

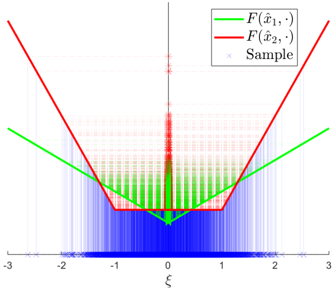

Figure: (a) Out-of-sample scatter plot

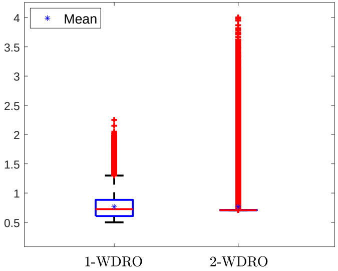

Figure: (b) Out-of-sample boxplot

Figure: For a **single sample**, both methods yield similar average costs. However, the boxplot reveals that the 2-WDRO solution allows for significantly larger extreme outcomes.

Motivation: The Choice of p Matters

A Motivating Example

We construct a problem to specifically compare 1-WDRO and 2-WDRO.

- **Question:** Does the choice of the Wasserstein order p affect robustness against extreme outcomes?
- **Finding:** Yes. The choice of p can lead to markedly different tail risk behaviors, even when the average performance is similar.

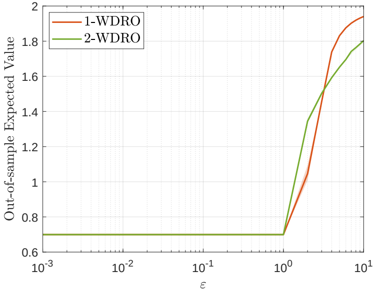

Figure: (a) Out-of-sample Expected Value

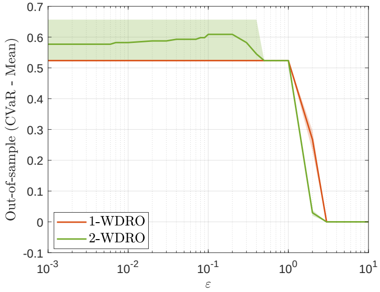

Figure: (b) Out-of-sample (CVaR - Mean)

Figure: This is not an isolated finding. **Simulating over 150 samples confirms** that, for this problem, 2-WDRO consistently leads to solutions with a higher tail risk premium.

Our Approach: From Sample Points to Adverse Scenarios

Where WDRO's Regularization Focuses

- As we saw, WDRO implicitly regularizes by penalizing the cost function's sensitivity (the gradient norm $\|\nabla_\xi F\|$) at the **sample points** $\hat{\xi}_i$.
- **A critical observation:** The most impactful, high-cost events are often triggered by **adverse scenarios** which may be underrepresented or absent in a finite sample.

Our Approach: From Sample Points to Adverse Scenarios

Where WDRO's Regularization Focuses

- As we saw, WDRO implicitly regularizes by penalizing the cost function's sensitivity (the gradient norm $\|\nabla_\xi F\|$) at the **sample points** $\hat{\xi}_i$.
- **A critical observation:** The most impactful, high-cost events are often triggered by **adverse scenarios** which may be underrepresented or absent in a finite sample.

Our Core Idea

Let's design a regularizer that explicitly targets robustness where it matters most: at a pre-defined set of known or anticipated **adverse scenarios**.

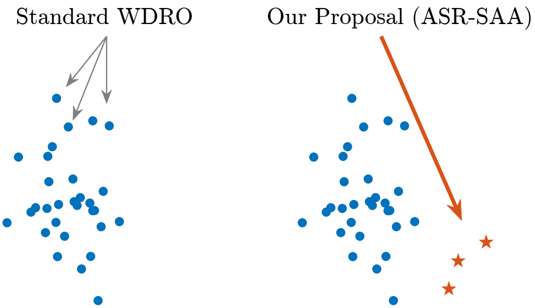

Figure: Shifting the regularization's focus.

Our Approach: The ASR-SAA Framework

Adverse Scenario Regularized SAA (ASR-SAA)

We propose to solve the following regularized optimization problem:

$$\min_{x \in \mathcal{X}} \underbrace{\frac{1}{n} \sum_{i=1}^{n} F(x, \widehat{\xi}_i)}_{\text{Empirical Performance}} + \underbrace{\varepsilon R(x)}_{\text{Adverse Regularization}}$$

The parameter $\varepsilon \geq 0$ controls the trade-off between average performance and robustness.

Our Approach: The ASR-SAA Framework

Adverse Scenario Regularized SAA (ASR-SAA)

We propose to solve the following regularized optimization problem:

$$\min_{x\in\mathcal{X}} \underbrace{\frac{1}{n}\sum_{i=1}^{n} F(x,\widehat{\xi}_i)}_{\text{Empirical Performance}} + \underbrace{\varepsilon\, R(x)}_{\text{Adverse Regularization}}$$

The parameter $\varepsilon \geq 0$ controls the trade-off between average performance and robustness.

The Adverse Regularizer $R(x)$

The regularizer directly penalizes the cost function's sensitivity at adverse points:

$$R(x) = \sum_{j=1}^{m} r_j \left\| \nabla_\xi F(x,\zeta_j) \right\|$$

- $\zeta_j \in \Xi_{\mathrm{adv}}$: An **adverse scenario** (e.g., from historical crises, expert knowledge).
- r_j: A weight representing the **relative importance** of scenario ζ_j.
- $\|\nabla_\xi F(x,\zeta_j)\|$: Measures how **sensitive** the cost is to small perturbations around that adverse scenario. A smaller norm implies a "flatter", more robust response.

Theory: A Bridge to Decision-Dependent DRO

A General Theoretical Connection

A key insight of our work is that this approach is not an ad-hoc heuristic. The theoretical results apply to a **broad class of regularized SAA problems**, with our ASR-SAA being a practical instance.

Theory: A Bridge to Decision-Dependent DRO

A General Theoretical Connection

A key insight of our work is that this approach is not an ad-hoc heuristic. The theoretical results apply to a **broad class of regularized SAA problems**, with our ASR-SAA being a practical instance.

Theorem (Equivalence to Decision-Dependent WDRO)

Under mild technical conditions, any regularized problem of the form

$$\min_{x \in \mathcal{X}} \left(\mathbb{E}_{\hat{\mathbb{P}}_n}[F(x, \xi)] + \varepsilon R(x) \right)$$

is equivalent to a novel **decision-dependent** *WDRO problem:*

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{B}_{\varepsilon R(x)}\left(F(x, \cdot)_{\#}\widehat{\mathbb{P}}_n\right)} \mathbb{E}_{\varsigma \sim \mathbb{Q}}[\varsigma]$$

Here, the ambiguity set is defined on the space of outcomes, and its center and radius both change with the decision x.

Theory: Reliability Guarantees

Main Assumptions (Informal)

Our theoretical guarantees hold under standard assumptions in the field:

- **Lipschitz Continuity:** The cost function $F(x, \cdot)$ is Lipschitz continuous.
- **Regularizer Control:** There exists $\alpha \geq 0$ such that $R(x) + \alpha$ upper bounds the Lipschitz constant of $F(x, \cdot)$ for all $x \in \mathcal{X}$.
- **Light Tails:** The true distribution $\mathbb{P}$ has sufficiently light tails (e.g., finite exponential moments).

Main Assumptions (Informal)

Our theoretical guarantees hold under standard assumptions in the field:

- **Lipschitz Continuity:** The cost function $F(x, \cdot)$ is Lipschitz continuous.
- **Regularizer Control:** There exists $\alpha \geq 0$ such that $R(x) + \alpha$ upper bounds the Lipschitz constant of $F(x, \cdot)$ for all $x \in \mathcal{X}$.
- **Light Tails:** The true distribution $\mathbb{P}$ has sufficiently light tails (e.g., finite exponential moments).

Finite-Sample Guarantee

For any confidence level $1 - \beta$, we define a radius[a] $\varepsilon_n(\beta)$ that vanishes as n increases.
With probability at least $1 - \beta$, the true out-of-sample cost is bounded:

$$\mathbb{E}_{\mathbb{P}}[F(\hat{x}_n, \xi)] \leq \widehat{J}_n + \varepsilon_n(\beta)\alpha$$

[a] The rate of convergence for $\varepsilon_n(\beta)$ is derived from the celebrated bounds on the empirical measure's Wasserstein distance by Fournier & Guillin (2015).

Asymptotic Consistency

As the sample size $n \to \infty$ (and thus $\varepsilon_n \to 0$):

- **Value Convergence:** The optimal value of our problem converges to the true optimal value.
- **Solution Convergence:** The solutions $\{\hat{x}_n\}$ converge to true optimal solutions.

Implication: Our method is statistically sound and reliable.

Application: Mean-CVaR Portfolio Optimization

Problem: The Mean-CVaR Portfolio

A cornerstone of modern risk management. Find a portfolio weight vector w to solve:

$$\min_{w \in \mathcal{W}} \underbrace{\mathbb{E}_{\xi}[-w^\top \xi]}_{\text{Minimize Expected Loss}} + \rho \cdot \underbrace{\mathrm{CVaR}_{\alpha}(-w^\top \xi)}_{\text{Control Tail Risk}}$$

This can be cast into our general framework $\min \mathbb{E}[F(x, \xi)]$ where $x = (w, \tau)$ and ξ are the asset returns.

Application: Mean-CVaR Portfolio Optimization

Problem: The Mean-CVaR Portfolio

A cornerstone of modern risk management. Find a portfolio weight vector w to solve:

$$\min_{w \in \mathcal{W}} \underbrace{\mathbb{E}_\xi[-w^\top \xi]}_{\text{Minimize Expected Loss}} + \rho \cdot \underbrace{\mathrm{CVaR}_\alpha(-w^\top \xi)}_{\text{Control Tail Risk}}$$

This can be cast into our general framework $\min \mathbb{E}[F(x, \xi)]$ where $x = (w, \tau)$ and ξ are the asset returns.

Applying the ASR-SAA Framework

- The cost function $F(w, \tau, \xi)$ becomes:

$$F(w, \tau, \xi) = -w^\top \xi + \rho \left(\tau + \frac{1}{\alpha} \Big(-w^\top \xi - \tau \Big)_+ \right)$$

- We can compute its subgradient $\nabla_\xi F(w, \tau, \xi)$ to build our regularizer $R(w, \tau)$.
- **Adverse scenarios** ζ_j are naturally defined, e.g., vectors of asset returns during historical market crashes.

Application: Mean-CVaR Portfolio Optimization

Problem: The Mean-CVaR Portfolio

A cornerstone of modern risk management. Find a portfolio weight vector w to solve:

$$\min_{w \in \mathcal{W}} \underbrace{\mathbb{E}_{\xi}[-w^\top \xi]}_{\text{Minimize Expected Loss}} + \rho \cdot \underbrace{\mathrm{CVaR}_{\alpha}(-w^\top \xi)}_{\text{Control Tail Risk}}$$

This can be cast into our general framework $\min \mathbb{E}[F(x, \xi)]$ where $x = (w, \tau)$ and ξ are the asset returns.

Applying the ASR-SAA Framework

- The cost function $F(w, \tau, \xi)$ becomes:

$$F(w, \tau, \xi) = -w^\top \xi + \rho \left(\tau + \frac{1}{\alpha} \Big(-w^\top \xi - \tau \Big)_+ \right)$$

- We can compute its subgradient $\nabla_{\xi} F(w, \tau, \xi)$ to build our regularizer $R(w, \tau)$.
- **Adverse scenarios** ζ_j are naturally defined, e.g., vectors of asset returns during historical market crashes.

Computational Tractability

Our ASR-SAA problem for this application can be reformulated and solved efficiently as a **Mixed-Integer Second-Order Cone Program (MISOCP)**.

Numerical Study: Case Study 1 - Setup

Objective

To test the performance of ASR-SAA when **relevant** historical crisis data is available to inform the choice of adverse scenarios.

Experimental Design

- **Asset Universe:** A portfolio of **23 top S&P 500 companies** by market capitalization as of 2022.

Numerical Study: Case Study 1 - Setup

Objective

To test the performance of ASR-SAA when **relevant** historical crisis data is available to inform the choice of adverse scenarios.

Experimental Design

- **Asset Universe:** A portfolio of **23 top S&P 500 companies** by market capitalization as of 2022.
- **Adverse Scenarios:** Daily returns from **2020** (COVID-19 crisis).
- **Training Data:** Daily returns from **2021** (market recovery).
- **Out-of-Sample Test:** Daily returns from **2022** (geopolitical/economic stress).

Numerical Study: Case Study 1 - Setup

Objective

To test the performance of ASR-SAA when **relevant** historical crisis data is available to inform the choice of adverse scenarios.

Experimental Design

- **Asset Universe:** A portfolio of **23 top S&P 500 companies** by market capitalization as of 2022.
- **Adverse Scenarios:** Daily returns from **2020** (COVID-19 crisis).
- **Training Data:** Daily returns from **2021** (market recovery).
- **Out-of-Sample Test:** Daily returns from **2022** (geopolitical/economic stress).

Methodology

- We define adverse scenarios as the asset returns on days where the S&P 500 index dropped by more than a certain threshold (e.g., -2%, -3.5%, -5%).
- We compare our ASR-SAA variants against SAA, 1-WDRO, and 2-WDRO.

Numerical Study: Case Study 1 - Setup

Objective

To test the performance of ASR-SAA when **relevant** historical crisis data is available to inform the choice of adverse scenarios.

Experimental Design

- **Asset Universe:** A portfolio of **23 top S&P 500 companies** by market capitalization as of 2022.
- **Adverse Scenarios:** Daily returns from **2020** (COVID-19 crisis).
- **Training Data:** Daily returns from **2021** (market recovery).
- **Out-of-Sample Test:** Daily returns from **2022** (geopolitical/economic stress).

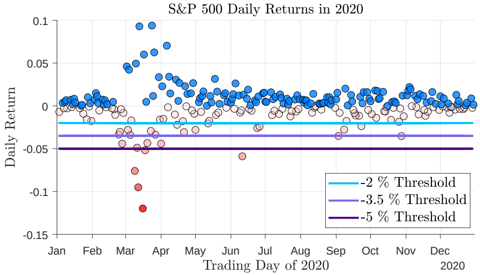

Figure: S&P 500 daily returns in 2020, showing thresholds for adverse scenario selection.

Methodology

- We define adverse scenarios as the asset returns on days where the S&P 500 index dropped by more than a certain threshold (e.g., -2%, -3.5%, -5%).
- We compare our ASR-SAA variants against SAA, 1-WDRO, and 2-WDRO.

Numerical Study: Case Study 1 - Results

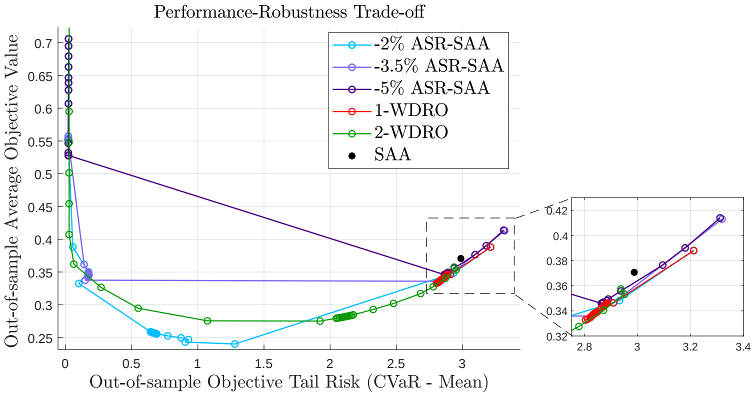

Figure: Out-of-sample performance-risk trade-off. Lower on the y-axis (Expected Cost) is better. Left on the x-axis (Cost Risk Premium) is better.

Numerical Study: Case Study 1 - Results

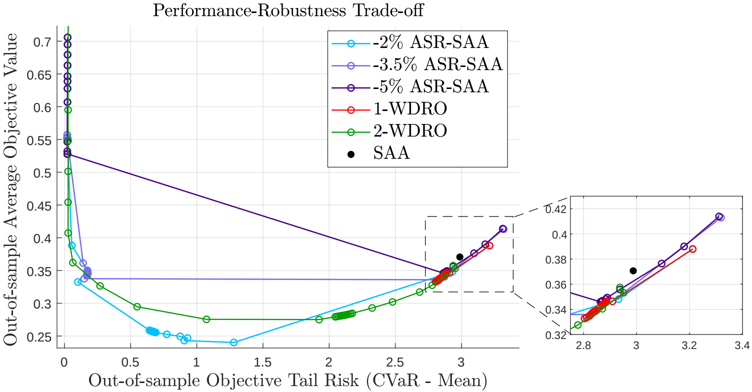

Figure: Out-of-sample performance-risk trade-off. Lower on the y-axis (Expected Cost) is better. Left on the x-axis (Cost Risk Premium) is better.

Key Finding

When informed by relevant crisis data, the **-2% ASR-SAA** variant achieves a superior trade-off frontier, outperforming all benchmarks, including 2-WDRO.

Numerical Study: Case Study 2 - A Critical Test

Objective

To investigate the method's sensitivity to the **quality** of adverse scenarios. What happens if the scenarios are not representative of the future crisis?

Experimental Design

- **Asset Universe:** A portfolio of **23 major companies** reflecting the market composition at the start of 2008.

Numerical Study: Case Study 2 - A Critical Test

Objective

To investigate the method's sensitivity to the **quality** of adverse scenarios. What happens if the scenarios are not representative of the future crisis?

Experimental Design

- **Asset Universe:** A portfolio of **23 major companies** reflecting the market composition at the start of 2008.
- **Adverse Scenarios:** Daily returns from **2006** (a stable, pre-crisis period).
- **Training Data:** Daily returns from **2007** (pre-crisis).
- **Out-of-Sample Test:** Daily returns from **2008** (Global Financial Crisis).

Numerical Study: Case Study 2 - A Critical Test

Objective

To investigate the method's sensitivity to the **quality** of adverse scenarios. What happens if the scenarios are not representative of the future crisis?

Experimental Design

- **Asset Universe:** A portfolio of **23 major companies** reflecting the market composition at the start of 2008.
- **Adverse Scenarios:** Daily returns from **2006** (a stable, pre-crisis period).
- **Training Data:** Daily returns from **2007** (pre-crisis).
- **Out-of-Sample Test:** Daily returns from **2008** (Global Financial Crisis).

The Challenge

The "adverse" days from 2006 are mild and not reflective of the extreme events of 2008. This is a deliberate stress test of our framework's main assumption.

Objective

To investigate the method's sensitivity to the **quality** of adverse scenarios. What happens if the scenarios are not representative of the future crisis?

Experimental Design

- **Asset Universe:** A portfolio of **23 major companies** reflecting the market composition at the start of 2008.
- **Adverse Scenarios:** Daily returns from **2006** (a stable, pre-crisis period).
- **Training Data:** Daily returns from **2007** (pre-crisis).
- **Out-of-Sample Test:** Daily returns from **2008** (Global Financial Crisis).

The Challenge

The "adverse" days from 2006 are mild and not reflective of the extreme events of 2008. This is a deliberate stress test of our framework's main assumption.

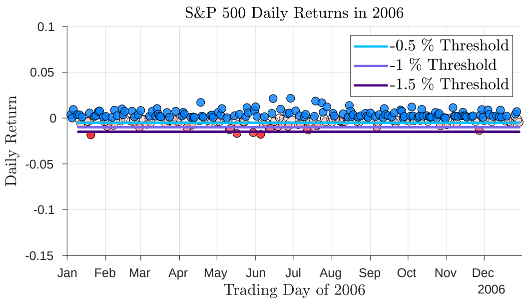

Figure: S&P 500 daily returns in 2006. Note the much lower volatility compared to 2020.

Numerical Study: Case Study 2 - Results & Lesson Learned

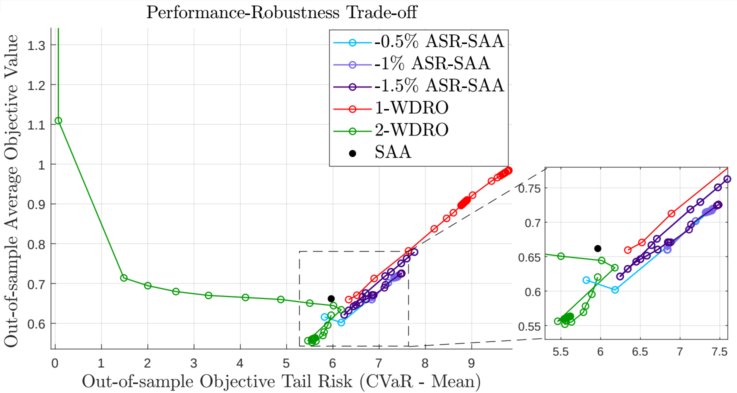

Figure: Out-of-sample performance during the 2008 crisis.

Numerical Study: Case Study 2 - Results & Lesson Learned

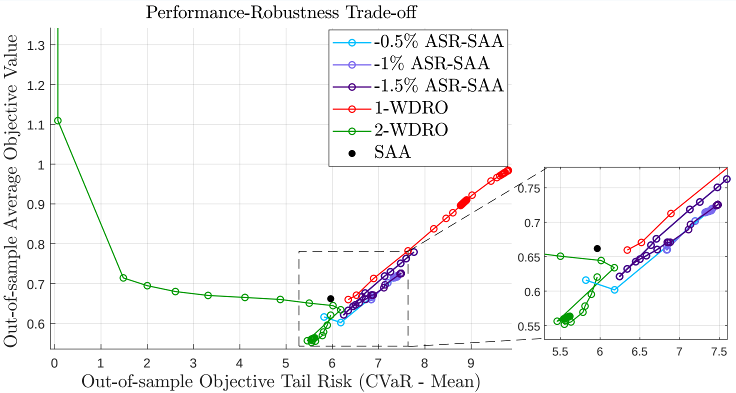

Figure: Out-of-sample performance during the 2008 crisis.

Key Finding & Lesson Learned

In this case, **2-WDRO** provides the best performance.

This highlights a crucial feature of our ASR-SAA framework: its effectiveness is contingent on the availability of **relevant adverse scenarios**.

Case 1: Relevant Scenarios (2020-2022)

ASR-SAA $>$ 2-WDRO

When provided with relevant crisis data, ASR-SAA can achieve a better risk-return profile than benchmark WDRO.

Numerical Study: Synthesis of Findings

Case 1: Relevant Scenarios (2020-2022)

ASR-SAA $>$ 2-WDRO

When provided with relevant crisis data, ASR-SAA can achieve a better risk-return profile than benchmark WDRO.

Case 2: Unrepresentative Scenarios (2006-2008)

2-WDRO $>$ ASR-SAA

If scenarios are not "adverse enough," data-driven methods like 2-WDRO can be more robust.

Case 1: Relevant Scenarios (2020-2022)

$$\text{ASR-SAA} > \text{2-WDRO}$$

When provided with relevant crisis data, ASR-SAA can achieve a better risk-return profile than benchmark WDRO.

Case 2: Unrepresentative Scenarios (2006-2008)

$$\text{2-WDRO} > \text{ASR-SAA}$$

If scenarios are not "adverse enough," data-driven methods like 2-WDRO can be more robust.

Overall Implication

ASR-SAA is not a universal replacement for WDRO, but a tractable tool for decision-makers to **directly incorporate** expert knowledge or historical stress-event data into the optimization process. Even with weak scenarios, it still improved upon SAA and 1-WDRO.

Conclusions

Summary of Contributions

1. We proposed **ASR-SAA**, a regularized framework that directly targets robustness against pre-defined adverse scenarios, bridging a gap in standard WDRO.
2. We established a novel theoretical connection, showing that a general class of regularized SAA methods are equivalent to a **decision-dependent WDRO** problem.
3. We provided rigorous **finite-sample guarantees** and proofs of **asymptotic consistency**, ensuring the statistical reliability of the framework.
4. Through numerical experiments, we demonstrated that ASR-SAA can offer a **practical advantage** over state-of-the-art methods when relevant adverse information is available, providing a valuable tool for risk-aware decision-making.

Thank You

Questions?

Diego Fonseca
diegofonseca@eafit.edu.co