Robust empirical optimization is almost the same as mean-variance optimization

Jun-ya Gotoh [1] Michael J. Kim [2] Andrew E.B. Lim [3]

[1]Department of Industrial and Systems Engineering, Chuo University, Tokyo, Japan

[2]Department of Mechanical and Industrial Engineering, University of Toronto, Canada

[3]Departments of Decision Sciences and Finance, NUS Business School, Singapore

Workshop on Risk Management Approaches in Engineering Applications, Gainesville, FL, USA

Outline of this talk

1. Basic notions and perspective
    - Empirical Optimization
    - Robust Empirical Optimization (REO)
    - *A REO is almost the same as Mean-Variance Optimization (MVO)*
2. Extension
    - *REOs are generally almost the same as MVO*
3. Selection of Ambiguity Parameter
    - (Standard) Cross-Validation
    - Robust Cross-Validation

Empirical Optimization

We would like to maximize the expected reward

$$\psi(x) := \mathbb{E}_{\mathbb{P}}\left[f(x, Y)\right],$$

over x. We do not know the distribution $\mathbb{P}$ of Y but have n historical data points $(x_1, Y_1), \ldots, (x_n, Y_n)$. We approximate the objective function by the sample mean

$$\hat{\psi}_n(x) := \frac{1}{n}\sum_{i=1}^{n} f(x, Y_i) = \mathbb{E}_{\hat{\mathbb{P}}_n}[f(x, Y)],$$

where $\hat{\mathbb{P}}_n$ is the empirical distribution of Y.

- aka SAA, etc.
- If the distribution $\mathbb{P}$ of Y_i does not depend on the decision x and $Y_1, \ldots, Y_n$ are drawn i.i.d. from $\mathbb{P}$, then $\hat{\psi}_n(x)$ is an unbiased estimate of $\psi(x)$, namely, $\psi(x) = \mathbb{E}_{\mathbb{P}}[\hat{\psi}_n(x)]$.
- many applications in OR/MS, statistics (machine learning), etc.

Two examples from OR/MS

$$\max_x \hat{\psi}_n(x) := \frac{1}{n}\sum_{i=1}^{n} f(x, Y_i) = \mathbb{E}_{\hat{\mathbb{P}}_n}[f(x, Y)]$$

Application 1: Newsvendor Problem

$$f(x, Y) = r \min\{x, Y\} - cx,$$

where
- ▷ $x \in \mathbb{R}$ is the order quantity (decision variable),
- ▷ $Y \in \mathbb{R}$ is the random demand, and
- ▷ r and c are the revenue and cost parameters, $r > c > 0$.

Application 2: Portfolio Selection

$$f(x, R) = -\exp\left(-\gamma R^\top x\right),$$

where
- ▷ $x \in \mathbb{R}^d$ is the portfolio vector (decision variables),
- ▷ $R \in \mathbb{R}^d$ is the vector of random returns, and
- ▷ γ is the risk-aversion parameter.

Robust Empirical Optimization (REO)

$$\max_x \hat{\psi}_n(x) := \frac{1}{n}\sum_{i=1}^{n} f(x, Y_i) = \mathbb{E}_{\hat{\mathbb{P}}_n}[f(x, Y)] \qquad : \text{Empirical optimization}$$

$$\max_x \min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x, Y)\right] + \theta \mathcal{R}(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\}, \qquad : \text{Robust optimization}$$

where $\theta > 0$ is a constant, $\mathbb{Q} \equiv (q_i)$ is a discrete probability distribution with the same support as the empirical distribution $\hat{\mathbb{P}}_n \equiv (\hat{p}_i)$, and

$$\mathcal{R}(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) = \begin{cases} \sum_{i:\, \hat{p}_i > 0} q_i \ln\left(\frac{q_i}{\hat{p}_i}\right), & \text{if } \sum_{i: \hat{p}_i > 0} q_i = 1, q_i \geq 0, \\ +\infty, & \text{otherwise}, \end{cases}$$

is the relative entropy (aka Kullback-Leibler divergence) of $\mathbb{Q}$ relative to $\hat{\mathbb{P}}_n$.

- $\mathcal{R}(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n)$ is non-negative and convex in $\mathbb{Q}$, and equal to zero if and only if $\mathbb{Q}$ equals $\hat{\mathbb{P}}_n$.
- θ represents the decision maker's confidence in $\hat{\mathbb{P}}_n$

Literature and our positioning

- Computational tractability [Ben-Tal et al. 13, Bertsimas,Gupta,Kallus 13, Bertsimas,Gupta,Kallus 14, Klaban,Simchi-Levi,Song 13, Wang,Glynn,Ye 13]
- Statistical properties/asymptotics [Ben-Tal et al. 13, Bertsimas,Gupta,Kallus 13, Bertsimas,Gupta,Kallus 14, Wang,Glynn,Ye 13]
- Choice of uncertainty sets [Ben-Tal et al. 13, Bertsimas,Gupta,Kallus 13, Bertsimas,Gupta,Kallus 14, Klaban,Simchi-Levi,Song 13, Wang,Glynn,Ye 13]
- Designing machine learning algorithms robust to data errors [Caramanis,Mannor,Xu 12, El Ghaoui,Lebret 97, Gotoh,Uryasev 13, Xu,Caramanis,Mannor 09, Xu,Caramanis,Mannor 10]

Our research focuses on

- providing perspectives on REO
  - "REO is almost the same as MVO"
- how to determine θ (in a data-driven way)

Fact 1: REO with relative entropy (= a risk min.) = a Mean-Deviation

Proposition (REO objective with relative entropy = a Mean-Deviation objective)

$$g_\theta(x) := \min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x,Y)\right] + \theta \mathcal{R}(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\} = \hat{\psi}_n(x) - \mathcal{D}_\theta\big(f(x,Y) \,|\, \hat{\psi}_n(x)\big),$$

where $\mathcal{D}_\theta\big(f(x,Y) \,|\, \hat{\psi}_n(x)\big) = \theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[\exp\left(-\frac{1}{\theta}\Big(f(x,Y) - \hat{\psi}_n(x)\Big)\right)\right].$

Proof.

$$\begin{aligned} g_\theta(x) &= -\theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[\exp\left(-\tfrac{1}{\theta} f(x,Y)\right)\right] \qquad (\because \text{ well-known duality}) \\ &= \underbrace{\mathbb{E}_{\hat{\mathbb{P}}_n}\Big[f(x,Y)\Big]}_{\hat{\psi}_n(x)} - \theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[\exp\left(-\tfrac{1}{\theta}\Big(f(x,Y) - \mathbb{E}_{\hat{\mathbb{P}}_n}\left[f(x,Y)\right]\Big)\right)\right]. \quad \square \end{aligned}$$

Remark: $\mathcal{D}_\theta$ is a measure of deviation from the mean in the following sense:

1. $\mathcal{D}_\theta(c|\mathbb{E}[c]) = 0$ for any constant $c \in \mathbb{R}$;
2. $\mathcal{D}_\theta(Z|\mathbb{E}[Z]) > 0$ for any (non-constant) random variable Z

Fact 2: Deviation associated with reltive entropy ≈ Variance

$$
\begin{aligned}
&\mathcal{D}_\theta\big(f(x,Y)\,|\,\hat{\psi}_n(x)\big) \\
&= \theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[\exp\left(-\frac{1}{\theta}\Big(f(x,Y)-\hat{\psi}_n(x)\Big)\right)\right] \\
&= \theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[1-\frac{1}{\theta}\Big(f(x,Y)-\hat{\psi}_n(x)\Big)+\frac{1}{2\theta^2}\Big(f(x,Y)-\hat{\psi}_n(x)\Big)^2+\cdots\right] \\
&= \theta \ln\left\{1+\frac{1}{2\theta^2}\mathbb{E}_{\hat{\mathbb{P}}_n}\left[\Big(f(x,Y)-\hat{\psi}_n(x)\Big)^2\right]+\cdots\right\} \\
&= \frac{1}{2\theta}\left\{\frac{1}{n}\sum_{i=1}^{n}\Big(f(x,Y_i)-\hat{\psi}_n(x)\Big)^2\right\}+o(1/\theta) \\
&= \frac{1}{2\theta}\mathbb{V}_{\hat{\mathbb{P}}_n}(f(x,Y))+o(1/\theta)
\end{aligned}
$$

REO with relative entropy ≈ Mean-Variance

$$\min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x,Y)\right] + \theta \mathcal{R}(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\} = \mathbb{E}_{\hat{\mathbb{P}}_n}\left[f(x,Y)\right] - \frac{1}{2\theta} \mathbb{V}_{\hat{\mathbb{P}}_n}\left[f(x,Y)\right] + o(1/\theta),$$

Correlation coefficients of out-of-sample rewards by REO and MVO

Newsvendor

θ	$n=10$	$n=50$	$n=100$
5	0.97	0.69	0.66
10	0.92	0.77	0.78
50	0.95	0.96	0.96
100	**0.94**	0.94	0.94
500	0.96	0.92	0.92
1000	0.98	0.96	0.95
10000	0.97	0.99	0.99
∞	1.00	1.00	1.00

Portfolio Selection

θ	$n=50$	$n=100$	$n=150$
0.01	0.692	0.624	0.608
0.05	0.957	**0.959**	0.958
0.10	0.987	0.991	0.993
0.50	0.999	1.000	1.000
1	1.000	1.000	1.000
10	1.000	1.000	1.000
100	1.000	1.000	1.000
∞	1.000	1.000	1.000

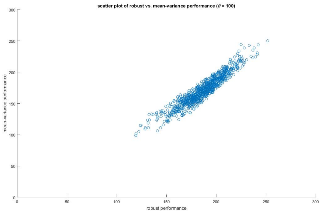

Newsvendor ($\theta = 100$ and $n = 10$)

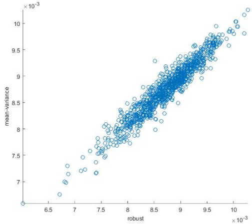

Portfolio Selection ($\theta = 0.05$ and $n = 100$)

Out-of-sample variance reduction (Newsvendor)

$$\mathbb{E}_{\hat{\mathbb{P}}_n}[f(x,Y)] \quad \overset{\text{robustified}}{\rightsquigarrow} \quad \mathbb{E}_{\hat{\mathbb{P}}_n}[f(x,Y)] - \frac{1}{2\theta}\mathbb{V}_{\hat{\mathbb{P}}_n}[f(x,Y)] + o(1/\theta)$$

$\rightsquigarrow$Robustness is achieved by controlling the variability of the reward distribution

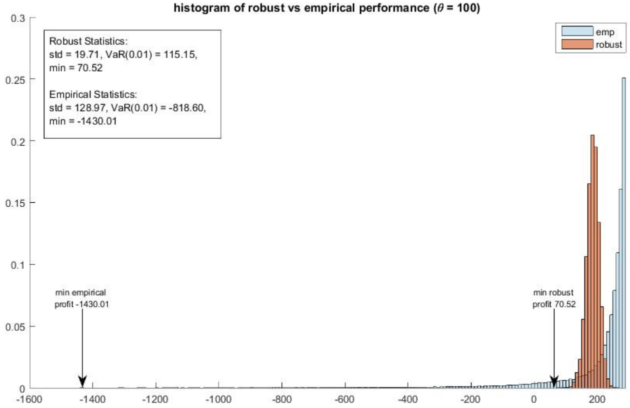

Robust vs. Empirical: out-of-sample profits in the Newsvendor problem with $\theta = 100$

Out-of-sample variance reduction (Portfolio Selection)

$$\mathbb{E}_{\hat{\mathbb{P}}_n}[f(x,Y)] \quad \overset{\text{robustified}}{\rightsquigarrow} \quad \mathbb{E}_{\hat{\mathbb{P}}_n}[f(x,Y)] - \frac{1}{2\theta}\mathbb{V}_{\hat{\mathbb{P}}_n}[f(x,Y)] + o(1/\theta)$$

$\rightsquigarrow$Robustness is achieved by controlling the variability of the reward distribution

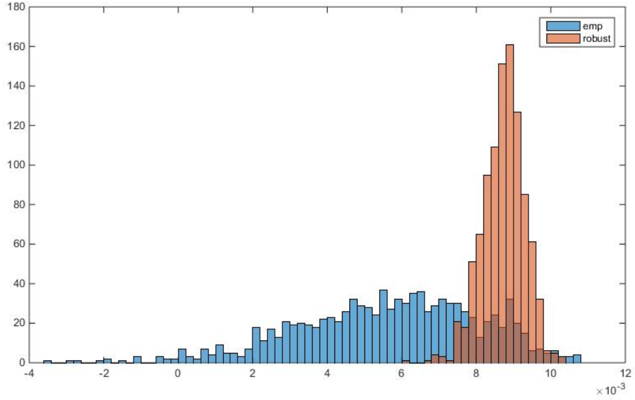

Robust vs. Empirical: out-of-sample certainty equivalent in the Portfolio Selection with $\theta = 0.05$

Extension to REO with Csiszár's ϕ-Divergence

Let ϕ be a closed proper convex function such that $\phi : \mathbb{R} \to \mathbb{R} \cup \{+\infty\}$ and $\phi(z) \geq \phi(1) = 0$ for all z. The ϕ-divergence of $\mathbb{Q}$ relative to $\hat{\mathbb{P}}_n$ is defined by

$$\mathcal{H}_\phi(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) := \begin{cases} \sum_{i:\hat{p}_i>0} \hat{p}_i \phi\left(\frac{q_i}{\hat{p}_i}\right), & \sum_{i:\hat{p}_i>0} q_i = 1, q_i \geq 0, \\ +\infty, & \text{otherwise.} \end{cases}$$

Examples

- relative entropy if $\phi(z) = z \ln z - z + 1$
- χ^2-divergence if $\phi(z) = \frac{1}{z}(z-1)^2$
- modified χ^2-divergence if $\phi(z) = (z-1)^2$
- Hellinger distance if $\phi(z) = (\sqrt{z} - 1)^2$

(See, e.g., [Ben-Tal et al. 13] for other examples.)

The robust objective function associated with the ϕ-divergence penalty:

$$g_{\theta,\phi}(x) := \min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x, Y)\right] + \theta \mathcal{H}_\phi(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\},$$

where constant $\theta > 0$ is the ambiguity parameter.

REO with ϕ-divergence = a Mean-Deviation (= a risk minimization)

Assumption ($\star$)

For every x, the worst case probability measure

$$\mathbb{Q}^* \in \arg\min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x,Y)\right] + \theta \mathcal{H}_\phi(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\}$$

is equivalent to $\hat{\mathbb{P}}_n$; that is, $q_i^* > 0$ if and only if $\hat{p}_i > 0$.

Proposition

If Assumption ($\star$) is satisfied in addition to the assumption of the previous proposition, then the following two objective functions are equal:

1. $g_{\theta,\phi}(x) = \min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}[f(x,Y)] + \theta \mathcal{H}_\phi(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\}$; *(robust objective function)*
2. $\hat{\psi}_n(x) - \mathcal{D}_{\theta,\phi,\hat{\mathbb{P}}_n}(f(x,Y)) \,|\, \hat{\psi}_n(x))$. *(mean-deviation objective function)*

Remark: Maximizations of 1. and 2. are also equal to minimization of a risk measure $\mathcal{R}_{\theta,\phi,\hat{\mathbb{P}}}(x)$.

Regular Measure of Deviation

$$\mathcal{D}_{\theta,\phi,\mathbb{P}}(Z \mid \mathbb{E}_{\mathbb{P}}[Z]) := \inf_{c} \left\{ c + \theta \mathbb{E}_{\mathbb{P}} \left[\phi^* \left(\frac{\mathbb{E}_{\mathbb{P}}[Z] - Z - c}{\theta} \right) \right] \right\}$$

where ϕ^* denotes the conjugate of ϕ, i.e., $\phi^*(\zeta) = \sup_z \{z\zeta - \phi(z)\}$.

Definition (Regular Measure of Deviation [Rockafellar, Uryasev 13])

Given any probability space $(\Omega, \mathcal{F}, \mathbb{P})$, let $\mathcal{L}^2(\Omega)$ denote the space of square-integrable random variables, i.e., $\mathbb{E}\left[X^2\right] < \infty$. A functional $\mathcal{D} : \mathcal{L}^2(\Omega) \to [0, \infty]$ is said to be *a regular measure of deviation* if it is closed convex and satisfies

1. $\mathcal{D}(c) = 0$ for any constant $c \in \mathbb{R}$.
2. $\mathcal{D}(Z) > 0$ for any (non-constant) random variable $Z \in \mathcal{L}^2(\Omega)$.

Proposition

Let ϕ be a closed proper convex function such that $\phi(1) = 0$ and $\phi'(1) = 0$. Then, for any random variable $Z \in \mathcal{L}^2(\Omega)$, $\mathcal{D}(Z) = \mathcal{D}_{\theta,\phi,\mathbb{P}}(Z \mid \mathbb{E}_{\mathbb{P}}[Z])$ is a regular measure of deviation.

REO with ϕ-divergence $\approx$ MVO

Proposition

Suppose that ϕ is convex, twice continuously differentiable, and that $\phi(1) = 0$, $\phi'(1) = 0$ and $\phi''(1) > 0$. Then the deviation measure $\mathcal{D}_{\theta,\phi,\hat{\mathbb{P}}_n}(f(x,Y) \,|\, \hat{\psi}_n(x))$ satisfies

$$\mathcal{D}_{\theta,\phi,\hat{\mathbb{P}}_n}(f(x,Y) \,|\, \hat{\psi}_n(x)) \quad = \quad \frac{1}{2\theta\phi''(1)}\hat{\sigma}_n^2(x) + o(1/\theta),$$

where $\hat{\sigma}_n^2(x) := \mathbb{V}_{\hat{\mathbb{P}}_n}(f(x,Y))$.

Consequently,

$$\min_{\mathbb{Q}} \left\{ \mathbb{E}_{\mathbb{Q}}\left[f(x,Y)\right] + \theta\mathcal{H}_\phi(\mathbb{Q} \,|\, \hat{\mathbb{P}}_n) \right\} = \hat{\psi}_n(x) - \frac{1}{2\theta\phi''(1)}\hat{\sigma}_n^2(x) + o(1/\theta),$$

$\rightsquigarrow$Robustness is achieved by controlling the variability of the reward distribution

Remark: Mean-Variance Optimization may not be a convex optimization even when the reward $f(x,Y)$ is concave in x, while Robust Optimization is convex whenever $f(x,Y)$ is concave.

1. Basic notions and perspective
   - Empirical Optimization
   - Robust Empirical Optimization (REO)
   - *A REO is almost the same as Mean-Variance Optimization (MVO)*

2. Extension
   - *REOs are generally almost the same as MVO*

3. Selection of Ambiguity Parameter
   - (Standard) Cross-Validation
   - Robust Cross-Validation

Selection of Ambiguity Parameter θ

- With relative entropy, REO results in a convex optimization problem:

$$\max_x -\theta \ln \mathbb{E}_{\hat{\mathbb{P}}_n}\left[\exp\left(-\frac{1}{\theta} f(x,Y)\right)\right]$$

- How to select θ? $\rightsquigarrow$ Cross-Validation

1. Let $\Theta := \{\theta_1, ..., \theta_H\}$ be a set of candidates of θ.
2. Divide the whole data samples available into $k = 10$ subsets, I,II,..., X.
3. For each $\theta \in \Theta$, do the procedure on the right.
4. Pick up the best θ so that the mean of the CV distribution is maximized.

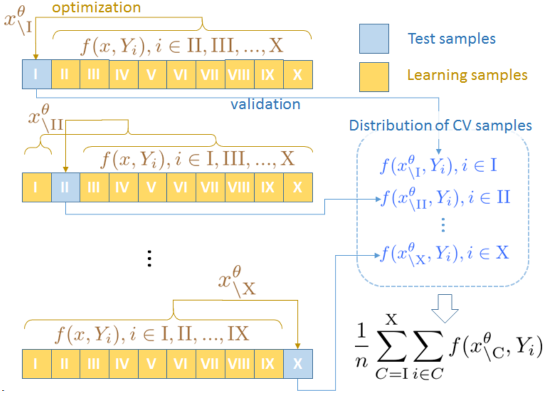

Cross-Validation for Newsvendor problem

Table: 10-Fold Cross-Validation Results for robust optimization and empirical optimization for the newsvendor problem

		Mean		Standard Deviation		1% VaR	
n	Choice of θ	robust	emp	robust	emp	robust	emp
10	2460	271.73	215.77	24.64	128.40	191.58	-818.60
50	17180	273.73	272.25	16.87	24.42	223.40	54.49
100	61990	278.84	277.68	12.36	14.94	232.96	149.62

- CV is conducted for $\theta \in \{10, 50, 100, 500, 1000, 5000, 10000\}$.
- Stats are computed over 100,000 data sets drawn from a mixture of two exponential distributions

Cross-Validation for Portfolio Selection problem

Table: 10-Fold Cross-Validation Results for Robust Optimization and Empirical Optimization for Portfolio Selection Problem (in units of return $[\times 10^{-2}\%]$)

n	Choice of θ	Mean		Standard Deviation		1% VaR	
		robust	emp	robust	emp	robust	emp
50	0.300	65.72	46.55	27.52	26.80	-18.93	-19.37
100	0.300	72.03	55.90	23.51	24.48	-4.53	-9.46
150	0.255	76.91	62.36	19.59	22.45	5.30	-0.78

The second column reports the geometric mean of the chosen θ's over the 1000 data sets.

- CV is conducted for $\theta \in \{0.01, 0.05, 0.1, 0.5, 1, 5, 10, 50, 100\}$.
- stats are computed over 1000 data sets drawn from a normal distribution

In-sample Distribution vs. Out-of-sample Distribution for Portfolio Selection

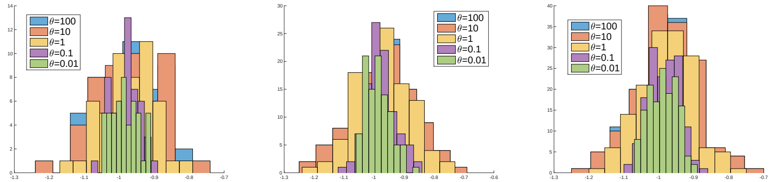

(i) $n = 50$ (ii) $n = 100$ (iii) $n = 150$

In-sample utility for the first data set (out of 1000 data sets)

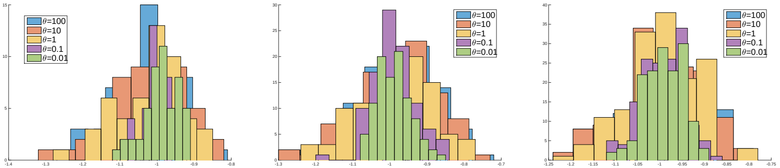

(i) $n = 50$ (ii) $n = 100$ (iii) $n = 150$

Utility with validation samples of 10-fold CV for the first data set

In-sample vs. Out-of-sample Mean/StDev. for Portfolio Selection

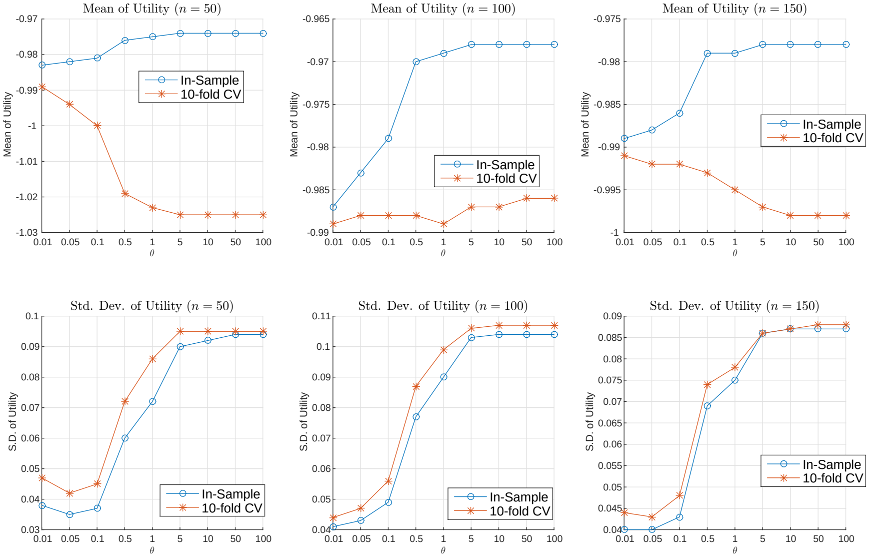

Robust Cross-Validation

Table: Improved 10-Fold Cross-Validation Results (in units of return [$\times 10^{-2}$%])

(1) Minimum-variance cross-validation

n	Choice of θ	Mean		Standard Deviation		1% VaR	
		robust	emp	robust	emp	robust	emp
50	0.013	83.59	46.55	8.02	26.80	60.83	-19.37
100	0.019	86.26	55.90	5.29	24.48	72.11	-9.46
150	0.025	87.21	62.36	4.29	22.45	76.24	-0.78

CV is conducted for $\theta \in \{0.01, 0.05, 0.1, 0.5, 1, 5, 10, 50, 100\}$.

(2) Mean-variance cross-validation: Optimizing mean $-c\times$ std.devi. ($c = 3.09/\sqrt{n-1}$)

n	Choice of θ	Mean		Standard Deviation		1% VaR	
		robust	emp	robust	emp	robust	emp
50	0.026	80.54	46.55	14.21	26.80	20.79	-19.37
100	0.027	84.68	55.90	11.09	24.48	29.85	-9.46
150	0.029	86.79	62.36	6.84	22.45	61.96	-0.78

CV is conducted for $\theta \in \{0.01, 0.05, 0.1, 0.5, 1, 5, 10, 50, 100\}$.

(3) Constrained cross-validation: Optimizing θ over $[0.01, 0.1]$

n	Choice of θ	Mean		Standard Deviation		1% VaR	
		robust	emp	robust	emp	robust	emp
50	0.030	81.36	46.55	11.50	26.80	44.93	-19.37
100	0.032	86.21	55.90	7.44	24.48	63.99	-9.46
150	0.032	87.65	62.36	6.03	22.45	71.03	-0.78

CV is conducted for $\theta \in \{0.01, 0.02, ..., 0.09, 0.1\}$.

Concluding remarks

1. REO $\approx$ MVO

$\rightsquigarrow$ Robustness is achieved by controlling the variability of the reward distribution

2. Parameter selection via CV should also be robust,
   i.e., we should take into account the variability in the CV distribution

   We have demonstrated three robust CVs:

   2.1 *Minimum-variance CV*: Choose θ that minimizes the variance of utility in cross-validation;
   2.2 *Mean-variance CV*: Choose θ that optimizes a mean-variance objective in cross-validation;
   2.3 *Constrained CV*: Optimize θ over a range where the sample variance of utility, and the difference between in-sample and out-of-sample expectd utilities, is small.

   One can imagine more sophisticated versions where the range is chosen adaptively as a function of the data set.