Stress Testing Financial Portfolios with Coverage Guarantees

Martin Haugh

Imperial College Business School

Imperial College London

10th Annual Bloomberg-Columbia Machine Learning in Finance Conference

19^{th} **September 2024**

Based on joint work with Zhongze Cai and Xiaocheng Li (Imperial College London)

Factor Models in Finance

- Portfolio value $V_t = v(\mathbf{x}_t)$ a function of risk factors $\mathbf{x}_t$. So time $t+1$ P&L

$$\Delta V_{t+1}(\boldsymbol{\Delta}\mathbf{x}_{t+1}) = v(\mathbf{x}_t + \boldsymbol{\Delta}\mathbf{x}_{t+1}) - v(\mathbf{x}_t) \tag{1}$$

- Typically work with a common factor model, e.g.

$$\boldsymbol{\Delta}\mathbf{x}_{t+1} = \mathbf{B}\,\mathbf{f}_{t+1} + \boldsymbol{\xi}_{t+1}, \qquad t = 0, 1, \ldots \tag{2}$$

  where:

  - $\mathbf{f}_{t+1} \in \mathbb{R}^m$ is the common factor (c.f.) random return vector.
  - $\mathbf{B} = [\mathbf{b}_1 \ldots \mathbf{b}_m] \in \mathbb{R}^{n\times m}$ is the matrix of factor loadings.
  - $\boldsymbol{\xi}_{t+1}$'s $\in \mathbb{R}^n$ are i.i.d. noise vectors.

- Portfolios from many asset classes can be modelled via (1)-(2).

Standard Scenario Analysis (SSA)

- Can define a scenario by jointly stressing any number $k \leq m$ of the c.f. returns.

e.g. Consider a portfolio of futures and options on the S&P 500. A scenario might consist of:

- Shift of -5% to value of S&P 500
- Parallel shift of $+10$ percentage points to implied volatility surface.

Standard Scenario Analysis (SSA)

- Can define a scenario by jointly stressing any number $k \leq m$ of the c.f. returns.

e.g. Consider a portfolio of futures and options on the S&P 500. A scenario might consist of:

  - Shift of -5% to value of S&P 500
  - Parallel shift of $+10$ percentage points to implied volatility surface.

- Risk / portfolio manager often doesn't have an explicit model like (2) at hand.
  - Especially true for portfolios containing derivative securities.

Standard Scenario Analysis (SSA)

- Can define a scenario by jointly stressing any number $k \leq m$ of the c.f. returns.

e.g. Consider a portfolio of futures and options on the S&P 500. A scenario might consist of:

- Shift of -5% to value of S&P 500
- Parallel shift of $+10$ percentage points to implied volatility surface.

- Risk / portfolio manager often doesn't have an explicit model like (2) at hand.
  - Especially true for portfolios containing derivative securities.

- Only a subset of factors - say first $l \leq m$ - are ever considered for stressing.
- In that case SSA works with a "model" of the form

$$\mathbf{\Delta x}_{t+1} = \mathbf{B}_{1:l}\mathbf{f}_{t+1}^{1:l}$$

A Portfolio of Options and Futures on the S&P 500

Underlying	SPX Index

Underlying and Volatility Stress Table

Sum of PnL	Vol Stress								
Underlying Stress	-10	-5	-2	-1	0	1	2	5	10
-20	13,938	11,774	10,631	10,277	9,936	9,608	9,293	8,419	7,183
-10	6,109	4,946	4,436	4,291	4,158	4,035	3,922	3,634	3,296
-5	1,831	1,652	1,637	1,643	1,654	1,670	1,689	1,766	1,946
-2	(314)	89	356	447	539	631	723	1,001	1,461
-1	(920)	(338)	15	132	248	363	478	816	1,361
0	(1,463)	(714)	(280)	(139)	0	137	273	668	1,293
1	(1,939)	(1,035)	(527)	(363)	(203)	(45)	110	559	1,259
2	(2,346)	(1,300)	(723)	(539)	(359)	(182)	(9)	489	1,258
5	(3,125)	(1,744)	(1,003)	(769)	(541)	(318)	(102)	518	1,460
10	(2,921)	(1,297)	(423)	(146)	123	385	641	1,372	2,483
20	2,344	3,559	4,272	4,506	4,738	4,967	5,194	5,860	6,919

- Table shows an SSA with P&L for simultaneous stresses to S&P 500 and parallel moves in implied volatility surface.
- Other factors could be stressed via steeping / flattening of volatility skew and / or term structure.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

2. SSA implicitly assumes

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t, \mathbf{f}^{\mathbf{s}}_{t+1}) = \mathbf{0}$$

where $\mathbf{f}^{\mathbf{s}}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1}$ denote stressed and unstressed c.f. returns, resp.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

2. SSA implicitly assumes

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1}) = \mathbf{0}$$

   where $\mathbf{f}^{\mathbf{s}}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1}$ denote stressed and unstressed c.f. returns, resp.

3. Could develop a model and set

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} = \mathbb{E}_t[\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid \mathbf{f}^{\mathbf{s}}_{t+1}]. \tag{3}$$

   But (3) ignores uncertainty in $\boldsymbol{\xi}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1})$.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

2. SSA implicitly assumes

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t, \mathbf{f}^{\mathbf{s}}_{t+1}) = \mathbf{0}$$

   where $\mathbf{f}^{\mathbf{s}}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1}$ denote stressed and unstressed c.f. returns, resp.

3. Could develop a model and set

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} = \mathbb{E}_t[\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid \mathbf{f}^{\mathbf{s}}_{t+1}]. \tag{3}$$

   But (3) ignores uncertainty in $\boldsymbol{\xi}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t, \mathbf{f}^{\mathbf{s}}_{t+1})$.

4. SSA generally not back-tested since scenarios (typically) have zero probability.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

2. SSA implicitly assumes

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1}) = \mathbf{0}$$

   where $\mathbf{f}^{\mathbf{s}}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1}$ denote stressed and unstressed c.f. returns, resp.

3. Could develop a model and set

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} = \mathbb{E}_t[\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid \mathbf{f}^{\mathbf{s}}_{t+1}]. \tag{3}$$

   But (3) ignores uncertainty in $\boldsymbol{\xi}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1})$.

4. SSA generally not back-tested since scenarios (typically) have zero probability.

5. SSA not robust to misspecified factors.

Problems with Standard Scenario Analysis

1. SSA only produces a point estimate of scenario loss - typically interpreted as an expectation of portfolio loss in a given scenario.

2. SSA implicitly assumes

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1}) = \mathbf{0}$$

   where $\mathbf{f}^{\mathbf{s}}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1}$ denote stressed and unstressed c.f. returns, resp.

3. Could develop a model and set

$$\mathbf{f}^{\mathbf{s}^c}_{t+1} = \mathbb{E}_t\big[\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid \mathbf{f}^{\mathbf{s}}_{t+1}\big]. \tag{3}$$

   But (3) ignores uncertainty in $\boldsymbol{\xi}_{t+1}$ and $\mathbf{f}^{\mathbf{s}^c}_{t+1} \mid (\mathcal{F}_t,\, \mathbf{f}^{\mathbf{s}}_{t+1})$.

4. SSA generally not back-tested since scenarios (typically) have zero probability.

5. SSA not robust to misspecified factors.

6. SSA not robust to adversarial portfolio selection.

Research Goals / Contributions

- Go beyond SSA by insisting on quantiles and hence prediction intervals (PIs) for scenario losses.
    - PIs are probability forecasts and hence amenable to backtesting via e.g. proper scoring rules.

Research Goals / Contributions

- Go beyond SSA by insisting on quantiles and hence prediction intervals (PIs) for scenario losses.
  - PIs are probability forecasts and hence amenable to backtesting via e.g. proper scoring rules.

- Provide coverage guarantees for the PIs
  - Conditional guarantees via a new non-parametric kernel-based algorithm.
  - Marginal guarantees via conformal prediction methods.

Research Goals / Contributions

- Go beyond SSA by insisting on quantiles and hence prediction intervals (PIs) for scenario losses.
  - PIs are probability forecasts and hence amenable to backtesting via e.g. proper scoring rules.

- Provide coverage guarantees for the PIs
  - Conditional guarantees via a new non-parametric kernel-based algorithm.
  - Marginal guarantees via conformal prediction methods.

- Address issues with SSA.

Scenario Analysis with Uncertainty Quantification

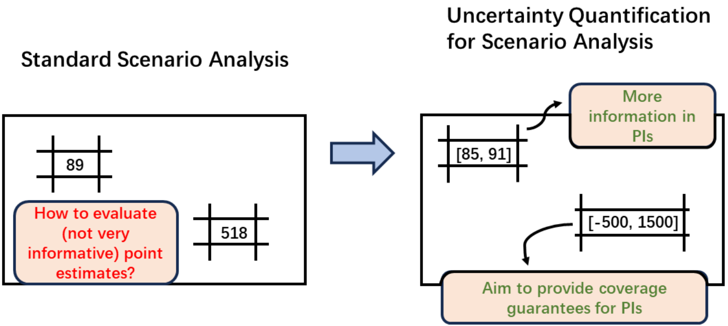

Selected Related Literature

Scenario Analysis

- Breur et al. (2009), Rebonato (2010,2019), Flood and Korenko (2015), Glasserman, Kang, and Kang (2015), H and Ruiz-Lacedelli (2020), Golub, Greenberg, and Ratcliffe (2018), Alfaro and Drehmann (2009), Quagliarello (2009), and Bonti, Kalkbrener, Lotz, and Stahl (2006), Borio,Drehmann, and Tsatsaronis (2012).

Proper Scoring Rules / Elicitability / Monitoring VaR & ES

- Savage (1971), Osband(1985), Lambert et al. (2008), Gneiting(2011), Gneiting and Raftery (2007), Bellini and Bignozzi (2015), Delbaen et al (2016), Fissler and Ziegel (2016), Kou and Peng (2016), Liu and Wang (2021), He, Kou and Peng (2022), Nolde and Ziegel (2017), Hoga and Demetrescu (2023).

Uncertainty Quantification

- Vovk et al. (2005), Shafer and Vovk (2008), Lei et al. (2013), Lei and Wasserman (2014), Lei et al. (2018), Song et al. (2019), Tibshirani et al. (2019), Angelopoulos and Bates (2021), Gibbs and Candès (2021), Bilodeau et al. (2021), Zaffran et al. (2022), Liu et al. (2023), Barber et al. (2023), Gibbs and Candès (2024).

Vast literature on factor models.

The Pinball Loss Function

Definition 1

The pinball loss function $\ell_\beta(\cdot,\cdot)$ *for* $Y, y \in \mathbb{R}$ *and quantile level* $\beta \in [0,1]$ *is*

$$\ell_\beta(Y,y) := \beta \cdot (Y-y)^+ + (1-\beta)\cdot(y-Y)^+.$$

- Pinball loss function is piecewise-linear and convex.
- It is also a proper scoring rule for the β-quantile $Q_\beta(Y)$ of Y so that

$$Q_\beta(Y) = \operatorname*{argmin}_{y\in\mathbb{R}} \mathbb{E}[\ell_\beta(Y,y)].$$

Kernel Scenario Analysis

- Have a history vector $H_T = (\mathbf{x}_T, \mathbf{f}_T)$.
- Feature vector $W_{T+1}(\mathbf{z}) := (H_T, \mathbf{z})$ where $\mathbf{z} := \mathbf{f}^{\mathbf{s}}_{t+1}$ is subset of c.f. returns stressed at time t in given time $t+1$ scenario.
- KSA algorithm takes as input a kernel $\kappa(\cdot,\cdot)$ and expectation predictor φ for scenario loss.

Algorithm Kernel Scenario Analysis (KSA)

Input: Target confidence $1 - \alpha$

for $t = 1, \ldots, T$ **do**
  Set $W_t \leftarrow (H_{t-1}, \mathbf{z}_t)$
  Set $Y_t \leftarrow L_t - \varphi(W_t)$ (loss residual)
end for

Algorithm Kernel Scenario Analysis (KSA)

Input: Target confidence $1-\alpha$

for $t = 1, \dots, T$ **do**

  Set $W_t \leftarrow (H_{t-1}, \mathbf{z}_t)$

  Set $Y_t \leftarrow L_t - \varphi(W_t)$ (loss residual)

end for

For $\beta \in [0,1]$ define empirical pinball loss function

$$\widehat{\ell}_T(y, w; \beta) \leftarrow \frac{\sum_{t=1}^{T} \kappa(W_t, w) \cdot \ell_\beta(Y_t, y)}{\sum_{t=1}^{T} \kappa(W_t, w)}$$

Algorithm Kernel Scenario Analysis (KSA)

Input: Target confidence $1 - \alpha$

for $t = 1, \ldots, T$ **do**
    Set $W_t \leftarrow (H_{t-1}, \mathbf{z}_t)$
    Set $Y_t \leftarrow L_t - \varphi(W_t)$ (loss residual)
end for
For $\beta \in [0,1]$ define empirical pinball loss function

$$\widehat{\ell}_T(y, w; \beta) \leftarrow \frac{\sum_{t=1}^{T} \kappa(W_t, w) \cdot \ell_\beta(Y_t, y)}{\sum_{t=1}^{T} \kappa(W_t, w)}$$

Set $W_{T+1}(\mathbf{z}) \leftarrow (H_T, \mathbf{z})$
For $\beta \in \{\alpha/2, 1 - \alpha/2\}$ compute empirical quantiles

$$\widehat{Q}_T(W_{T+1}(\mathbf{z}); \beta) \leftarrow \operatorname*{argmin}_{y \in \mathbb{R}} \; \widehat{\ell}_T(y, W_{T+1}(\mathbf{z}); \beta)$$

Algorithm Kernel Scenario Analysis (KSA)

Input: Target confidence $1 - \alpha$

for $t = 1, \ldots, T$ **do**

  Set $W_t \leftarrow (H_{t-1}, \mathbf{z}_t)$

  Set $Y_t \leftarrow L_t - \varphi(W_t)$ (loss residual)

end for

For $\beta \in [0, 1]$ define empirical pinball loss function

$$\widehat{\ell}_T(y, w; \beta) \leftarrow \frac{\sum_{t=1}^{T} \kappa(W_t, w) \cdot \ell_\beta(Y_t, y)}{\sum_{t=1}^{T} \kappa(W_t, w)}$$

Set $W_{T+1}(\mathbf{z}) \leftarrow (H_T, \mathbf{z})$

For $\beta \in \{\alpha/2, 1 - \alpha/2\}$ compute empirical quantiles

$$\widehat{Q}_T(W_{T+1}(\mathbf{z}); \beta) \leftarrow \underset{y \in \mathbb{R}}{\operatorname{argmin}} \ \widehat{\ell}_T(y, W_{T+1}(\mathbf{z}); \beta)$$

Return prediction interval (PI)

$$\mathcal{C}_T(\mathbf{z}) \leftarrow \varphi(W_{T+1}(\mathbf{z})) + [\widehat{Q}_T(W_{T+1}(\mathbf{z}); \alpha/2), \quad \widehat{Q}_T(W_{T+1}(\mathbf{z}); 1 - \alpha/2)]$$

Kernel Scenario Analysis

Assumption 1

- $\{W_t\}_{t=0}^{\infty}$ *is a stationary process.*
- *Independence of* $Y_t \mid W_t$*'s.*
- $\{W_t\}_{t=0}^{\infty}$ *ϕ-mixing and time reversed ϕ-mixing.*
- *Kernel function* η *exists so that* $\eta(w,w) = 1$ *and* $0 \le \eta(w,w') \le 1$ *for* $w, w' \in \mathcal{W}$.
- *Also have*

$$D_{\mathrm{TV}}\Big(\pi_{Y|W}(\cdot \mid w)\Big\|\pi_{Y|W}(\cdot \mid w')\Big) \le 1 - \eta(w,w'), \quad \forall w, w' \in \mathcal{W},$$

- *Other (mild) technical conditions*

Theorem 1

Under Assumption 1, for any $\delta > 0$, any scenario $\mathbf{z}$, and for sufficiently large T,

$$\left|\mathbb{P}_{L_{T+1}}\left(L_{T+1} \in \mathcal{C}_T(\mathbf{z}) \mid W_{T+1}(\mathbf{z})\right) - (1-\alpha)\right| \leq \underbrace{2 \cdot \left(1 - \frac{\mathrm{E}_W\left[\kappa(W, W_{T+1}(\mathbf{z})) \cdot \eta(W, W_{T+1}(\mathbf{z}))\right]}{\mathrm{E}_W\left[\kappa(W, W_{T+1}(\mathbf{z}))\right]}\right)}_{\textit{bias}} \underbrace{+ C \cdot \sqrt{\frac{\log T}{T}} + D \cdot \frac{\left(1 + 2\sum_{k=1}^{T} \phi(k)\right)\sqrt{\log \frac{16}{\delta}}}{\sqrt{T}} + \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)}_{\textit{variance}}$$

with probability at least $1 - \delta$.

Adaptive Conformal Scenario Analysis

- Conformal prediction algorithms developed to provide prediction intervals (regression) or prediction sets (classification) with valid coverage guarantees and minimal assumptions.
- Often applied to both probabilistic methods, e.g. logistic regression, and non-probabilistic methods, e.g. deep learning, random forests.
- Main focus on IID data but algorithms recently developed for time-series data.
- We adopt algorithm of Gibbs and Candès (2021) to scenario analysis setting
  - Model agnostic and takes as input a quantile predictor g_t.

Algorithm Adaptive Conformal Scenario Analysis (ACSA)

Input: Target confidence $1 - \alpha$, step size γ, quantile predictor g_1

Set $\beta_1 = \alpha$

for $t = 1, \ldots, T$ **do**

For any stress scenario $\mathbf{z}$, construct prediction interval

$$\mathcal{C}_t(\mathbf{z}) \leftarrow [g_t(W_{t+1}(\mathbf{z}); \beta_t/2), \quad g_t(W_{t+1}(\mathbf{z}); 1 - \beta_t/2)]$$

Observe L_{t+1} and $\mathbf{z}_{t+1}$, and compute

$$\mathrm{err}_{t+1} \leftarrow \mathbb{1}\{L_{t+1} \notin \mathcal{C}_t(\mathbf{z}_{t+1})\}$$

Update

$$\beta_{t+1} \leftarrow \beta_t + \gamma \cdot (\alpha - \mathrm{err}_{t+1})$$

end for

Assumption 2

Quantile predictor $g_t(W_{t+1}(\mathbf{z}); \beta)$ *satisfies*

$$\mathbb{P}(L_{t+1} \in [g_t(W_{t+1}(\mathbf{z}); 0), g_t(W_{t+1}(\mathbf{z}); 1)] \mid W_{t+1}(\mathbf{z})) = 1.$$

Assumption 2

Quantile predictor $g_t(W_{t+1}(\mathbf{z});\beta)$ *satisfies*

$$\mathbb{P}(L_{t+1} \in [g_t(W_{t+1}(\mathbf{z});0), g_t(W_{t+1}(\mathbf{z});1)] \mid W_{t+1}(\mathbf{z})) = 1.$$

- Extend g_t for $\beta \notin [0,1]$ via

$$g_t(W_{t+1}(\mathbf{z});\beta) = \begin{cases} g_t(W_{t+1}(\mathbf{z});0), & \text{for } \beta < 0 \\ g_t(W_{t+1}(\mathbf{z});1), & \text{for } \beta > 1. \end{cases}$$

Assumption 2

Quantile predictor $g_t(W_{t+1}(\mathbf{z});\beta)$ *satisfies*

$$\mathbb{P}(L_{t+1} \in [g_t(W_{t+1}(\mathbf{z});0), g_t(W_{t+1}(\mathbf{z});1)] \mid W_{t+1}(\mathbf{z})) = 1.$$

- Extend g_t for $\beta \notin [0,1]$ via

$$g_t(W_{t+1}(\mathbf{z});\beta) = \begin{cases} g_t(W_{t+1}(\mathbf{z});0), & \text{for } \beta < 0 \\ g_t(W_{t+1}(\mathbf{z});1), & \text{for } \beta > 1. \end{cases}$$

Theorem 2 (Gibbs and Candès 2021)

Under Assumption 2, the prediction interval $\mathcal{C}_t(\cdot)$ *generated by Algorithm 5 satisfies*

$$\left| \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{L_{t+1} \in \mathcal{C}_t(\mathbf{z}_{t+1})\} - (1-\alpha) \right| \leq \frac{\max\{\alpha, 1-\alpha\} + \gamma}{T\gamma} \quad \textit{a.s.}$$

In particular, as $T \to \infty$, *the empirical coverage rate converges to* $1-\alpha$ *a.s.*

Adaptive Conformal Scenario Analysis

- Theorem provides a calibration / coverage guarantee.
- It says nothing about sharpness / precision, i.e. width of PIs, which depends on quality of g_t.

Adaptive Conformal Scenario Analysis

- Theorem provides a calibration / coverage guarantee.
- It says nothing about sharpness / precision, i.e. width of PIs, which depends on quality of g_t.

e.g. Suppose g_t returns a constant C for any $\beta \in (0,1)$ and L_t a continuous r.var.

- Then β_t's deterministic and

$$\mathcal{C}_t(\mathbf{z}_{t+1}) = \begin{cases} [C, C], & \text{if } \beta_t \in (0,1) \\ \text{Supp}(L_{t+1}), & \text{otherwise.} \end{cases}$$

Adaptive Conformal Scenario Analysis

- Theorem provides a calibration / coverage guarantee.
- It says nothing about sharpness / precision, i.e. width of PIs, which depends on quality of g_t.

e.g. Suppose g_t returns a constant C for any $\beta \in (0,1)$ and L_t a continuous r.var.

- Then β_t's deterministic and

$$\mathcal{C}_t(\mathbf{z}_{t+1}) = \begin{cases} [C, C], & \text{if } \beta_t \in (0,1) \\ \text{Supp}(L_{t+1}), & \text{otherwise.} \end{cases}$$

- $\mathcal{C}_t(\mathbf{z}_{t+1}) = \text{Supp}(L_{t+1})$ at least $100 \times (1-\alpha)\%$ of the time.

Adaptive Conformal Scenario Analysis

- Theorem provides a calibration / coverage guarantee.
- It says nothing about sharpness / precision, i.e. width of PIs, which depends on quality of g_t.

e.g. Suppose g_t returns a constant C for any $\beta \in (0,1)$ and L_t a continuous r.var.

- Then β_t's deterministic and

$$\mathcal{C}_t(\mathbf{z}_{t+1}) = \begin{cases} [C, C], & \text{if } \beta_t \in (0,1) \\ \text{Supp}(L_{t+1}), & \text{otherwise.} \end{cases}$$

- $\mathcal{C}_t(\mathbf{z}_{t+1}) = \text{Supp}(L_{t+1})$ at least $100 \times (1-\alpha)\%$ of the time.

- Better g_t's generally producer sharper PIs.

But ACSA Only Provides a Marginal Guarantee

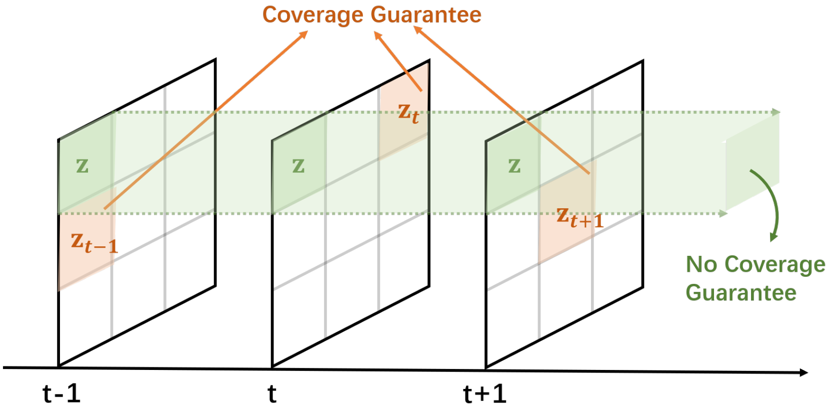

Numerical Experiments

- Portfolio consists of equal positions on approx 90 European options on S&P 500 and S&P 500 itself.
- Ground truth is

$$\mathbf{\Delta x}_{t+1} = \mathbf{B}\mathbf{f}_{t+1} + \boldsymbol{\xi}_{t+1}.$$

- Common factor return $\mathbf{f}_{t+1} \in \mathbb{R}^4$
  - First component $\mathbf{f}_{t+1}^{(1)}$ is log return of S&P 500 with daily vol ~ Garch$(1,1)$
  - Other components are implied volatility factors and satisfy

$$\mathbf{f}_{t+1}^{(2:4)} = \mathbf{G}\mathbf{f}_t^{(2:4)} + \boldsymbol{\epsilon}_{t+1}.$$

Numerical Experiments

- Portfolio consists of equal positions on approx 90 European options on S&P 500 and S&P 500 itself.
- Ground truth is

$$\mathbf{\Delta x}_{t+1} = \mathbf{B}\mathbf{f}_{t+1} + \boldsymbol{\xi}_{t+1}.$$

- Common factor return $\mathbf{f}_{t+1} \in \mathbb{R}^4$
  - First component $\mathbf{f}_{t+1}^{(1)}$ is log return of S&P 500 with daily vol $\sim$ Garch$(1, 1)$
  - Other components are implied volatility factors and satisfy

$$\mathbf{f}_{t+1}^{(2:4)} = \mathbf{G}\mathbf{f}_{t}^{(2:4)} + \boldsymbol{\epsilon}_{t+1}.$$

- Scenarios are joint stresses to $\mathbf{f}_{t+1}^{(1:2)}$, i.e S&P 500 return and first implied volatility factor.

Numerical Experiments

Algorithm	Variants	Coverage	PI Length ↓	CRPS ↓
Oracle		90.0%	2.56	0.452

Numerical Experiments

Algorithm	Variants	Coverage	PI Length ↓	CRPS ↓
Oracle		90.0%	2.56	0.452
Empirical quantile	Vanilla	72.4%	15.32	0.653
	SSA	77.1%	12.38	0.631

Numerical Experiments

Algorithm	Variants	Coverage	PI Length ↓	CRPS ↓
Oracle		90.0%	2.56	0.452
Empirical quantile	Vanilla	72.4%	15.32	0.653
	SSA	77.1%	12.38	0.631
ACSA	Linear-g_t	91.3%	8.63	0.572
	NN-g_t	90.9%	4.14	0.519
	Oracle-g_t	90.4%	2.73	0.464

Numerical Experiments

Algorithm	Variants	Coverage	PI Length ↓	CRPS ↓
Oracle		90.0%	2.56	0.452
Empirical quantile	Vanilla	72.4%	15.32	0.653
	SSA	77.1%	12.38	0.631
ACSA	Linear-g_t	91.3%	8.63	0.572
	NN-g_t	90.9%	4.14	0.519
	Oracle-g_t	90.4%	2.73	0.464
KSA $H_t = \mathbf{f}_t$	$\varphi_t = 0$	84.5%	4.17	0.482
	SSA-φ_t	87.2%	3.31	0.474
	NN-φ_t	89.6%	3.14	0.469
	Oracle-φ_t	89.7%	3.09	0.461

Numerical Experiments

Algorithm	Variants	Coverage	PI Length ↓	CRPS ↓
Oracle		90.0%	2.56	0.452
Empirical quantile	Vanilla	72.4%	15.32	0.653
	SSA	77.1%	12.38	0.631
ACSA	Linear-g_t	91.3%	8.63	0.572
	NN-g_t	90.9%	4.14	0.519
	Oracle-g_t	90.4%	2.73	0.464
KSA $H_t = \mathbf{f}_t$	$\varphi_t = 0$	84.5%	4.17	0.482
	SSA-φ_t	87.2%	3.31	0.474
	NN-φ_t	89.6%	3.14	0.469
	Oracle-φ_t	89.7%	3.09	0.461
KSA NN-φ_t misspecified	$H_t = \mathbf{f}_t^{(1:3)}$	88.7%	3.43	0.471
	$H_t = \mathbf{f}_t^{(1:2)}$	84.2%	3.17	0.498

Consistency of ACSA and KSA

	Parallel Shift						
Mkt	-1.5	-1.0	-0.5	0	0.5	1.0	1.5
-3	-5.1	-4.3	-2.3	-0.7	1.1	2.6	5.4
-2	-4.6	-3.7	-1.9	-0.5	0.5	1.7	4.1
-1	-2.5	-1.7	-0.4	0.4	0.7	1.5	3.6
0	-1.3	-0.5	0.3	0.9	1.3	1.7	2.9
1	0.7	1.1	1.1	1.4	2.3	2.4	2.7
2	2.3	2.4	2.4	2.6	3.1	4.4	5.3
3	4.2	4.6	4.9	5.4	6.2	6.8	7.5

ACSA with Oracle g_t

	Parallel Shift						
Mkt	-1.5	-1.0	-0.5	0	0.5	1.0	1.5
-3	-5.0	-4.4	-2.2	-0.9	0.8	2.5	5.1
-2	-4.5	-3.5	-2.2	-0.8	0.2	1.5	4.2
-1	-2.7	-1.6	-0.8	0.5	0.9	1.5	3.9
0	-1.1	-0.5	-0.1	0.8	1.1	1.4	2.8
1	0.8	1.3	1.4	1.5	2.0	2.2	2.7
2	2.0	2.3	2.4	2.8	2.9	3.9	4.8
3	4.4	4.5	4.8	5.2	6.1	6.9	7.2

KSA with Oracle φ_t

ACSA with Oracle vs SSA

	Parallel Shift						
Mkt	-1.5	-1.0	-0.5	0	0.5	1.0	1.5
-3	-5.1	-4.3	-2.3	-0.7	1.1	2.6	5.4
-2	-4.6	-3.7	-1.9	-0.5	0.5	1.7	4.1
-1	-2.5	-1.7	-0.4	0.4	0.7	1.5	3.6
0	-1.3	-0.5	0.3	0.9	1.3	1.7	2.9
1	0.7	1.1	1.1	1.4	2.3	2.4	2.7
2	2.3	2.4	2.4	2.6	3.1	4.4	5.3
3	4.2	4.6	4.9	5.4	6.2	6.8	7.5

ACSA with Oracle g_t

	Parallel Shift						
Mkt	-1.5	-1.0	-0.5	0	0.5	1.0	1.5
-3	-2.8	-2.3	-1.6	-0.8	0.3	1.5	3.0
-2	-2.2	-1.9	-1.4	-0.8	0.0	1.0	2.1
-1	-1.5	-1.3	-1.0	-0.5	0.0	0.8	1.7
0	-0.6	-0.5	-0.3	0.0	0.4	0.9	1.6
1	0.4	0.5	0.6	0.7	1.0	1.3	1.8
2	1.5	1.5	1.5	1.6	1.8	2.0	2.3
3	2.5	2.6	2.6	2.6	2.7	2.8	3.0

Standard Scenario Analysis

Conclusions and Further Research

Conclusions

Two algorithms for scenario analysis prediction intervals:

1. ACAS algorithm - simple application of Gibbs and Candès (2021) from conformal prediction literature
   - Model agnostic. Marginal coverage guarantees.
2. KSA algorithm using ideas from non-parametric statistics and uncertainty calibration literatures
   - Conditional coverage guarantees. Not good with extreme scenarios.

Conclusions and Further Research

Conclusions

Two algorithms for scenario analysis prediction intervals:

1. ACAS algorithm - simple application of Gibbs and Candès (2021) from conformal prediction literature
   - Model agnostic. Marginal coverage guarantees.
2. KSA algorithm using ideas from non-parametric statistics and uncertainty calibration literatures
   - Conditional coverage guarantees. Not good with extreme scenarios.
3. Adaptive KSA algorithm. Convergence rate of $O\left(t^{-\frac{1}{d_W+2}}\right)$ for coverage guarantee when:
   - Stationary distribution of $\{W_t\}_{t=0}^{\infty}$ is Gaussian.
   - κ_{h_t} and η RBF kernels with h_t decreasing in t.

Conclusions and Further Research

Conclusions

Two algorithms for scenario analysis prediction intervals:

1. ACAS algorithm - simple application of Gibbs and Candès (2021) from conformal prediction literature
   - Model agnostic. Marginal coverage guarantees.
2. KSA algorithm using ideas from non-parametric statistics and uncertainty calibration literatures
   - Conditional coverage guarantees. Not good with extreme scenarios.
3. Adaptive KSA algorithm. Convergence rate of $O\left(t^{-\frac{1}{d_W+2}}\right)$ for coverage guarantee when:
   - Stationary distribution of $\{W_t\}_{t=0}^{\infty}$ is Gaussian.
   - κ_{h_t} and η RBF kernels with h_t decreasing in t.

Ongoing Research

- More extensive numerical experiments.
- Combining ACSA with KSA.

Thank you!