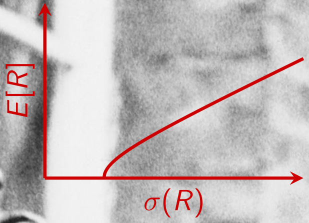

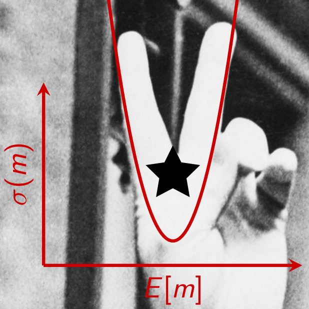

Regularized Conditional Quantile Methods for Portfolio Choice and Stochastic Discount Factor Estimation

Brown Bag Seminar Finance

Johannes Bleher
Joachim Grammig
Johannes Schief

University of Hohenheim
University of Tübingen

Motivation
Pricing kernel restriction

Fundamental asset pricing restriction (excess returns):

$$\mathbb{E}[m_{t+1}R^e_{t+1}] = 0, \qquad \mathbb{E}[m_{t+1}R_{f,t+1}] = 1$$

$R_{f,t+1}$ is the gross risk-free return and the Stochastic Discount Factor m_{t+1} is unobserved.

Why recover m_{t+1}?

- Economic meaning: m_{t+1} prices payoffs across states.
- Risk pricing: summarizes which risks / states receive high state-price weight.
- Hedging: ranks assets by payoff when m_{t+1} is high ("bad times").

Goal: recover a stable, interpretable, state-dependent empirical SDF.

Motivation

Why is recovering m_{t+1} hard in practice?

In finite samples we must solve moment restrictions

$$\frac{1}{T}\sum_{t=1}^{T} m_{t+1} R^{e}_{t+1} \approx 0 \quad \text{with many assets and instruments.}$$

(1) High-dimensional returns

- Many test assets $\Rightarrow$ nearly collinear spans
- Unstable SDFs / prices of risk without shrinkage

(2) State dependence

- Risk pricing varies across macro regimes
- Conditioning raises dimension and can amplify noise

Implication: need a regularized, conditional framework where *m* is still able to price pay-offs.

Motivation
Related literature

Related literature: restrictions, conditioning, and tails

1. **SDF recovery / HJ feasibility:** moment-based SDF proxies and variance bounds; no-arbitrage restrictions. (Hansen and Jagannathan, 1991; Hansen and Jagannathan, 1997)
2. **Conditioning information:** time-varying risk premia via instruments; conditional SDF bounds. (Hansen and Richard, 1987; Lettau and Ludvigson, 2001; Ferson and Siegel, 2003)
3. **Tail-sensitive methods:** direct portfolio policies, quantile/expected-shortfall allocation, and downside-risk pricing. (Brandt, Santa-Clara, and Valkanov, 2009; Koenker, 2005; Rockafellar and Uryasev, 2000; Acerbi and Tasche, 2002; Ang, Chen, and Xing, 2006; Bollerslev, Todorov, and Xu, 2015)

Motivation
Related literature

Related literature: complexity and diagnostics

1. **High-dimensional / ML SDF estimation:** flexible conditioning, shrinkage, and convex-pricing regularization in large cross sections. (Gu, Kelly, and Xiu, 2020; Lettau and Pelger, 2020; Korsaye, Quaini, and Trojani, 2025).
   For a recent overview, see (Sönksen, 2022).
2. **State prices and recovery:** option-implied state-price and recovery approaches. (Ait-Sahalia and Lo, 2000; Ross, 2015; Metaxoglou and Smith, 2017)
3. **Restriction and diagnostic discipline:** simulation-based pricing-kernel validation and theory-vs.-machine-learning risk-premium comparisons. (Grammig and Kuechlin, 2018; Soenksen and Grammig, 2021; Grammig, Hanenberg, Schlag, and Soenksen, 2025)

Motivation

What is new in this work?

This work:

- Estimate a joint conditional quantile surface for portfolio returns under elastic-net regularization.
- Turn the fitted residual ranking into an ES-envelope tail measure and project it onto the pricing set.
- Compare that tail-focused selection to an HJ-type minimum-2^{nd} moment SDF under the same restrictions.
- Study how both kernels vary across macro states.

Key insight: the same fitted conditional return object can be used for

- tail-sensitive portfolio choice,
- ES-envelope tail measure and HJ pricing-kernel recovery,
- and macro-state pricing diagnostics.

Motivation

Core idea: unified workflow, two kernels

Fit the conditional distribution of portfolio returns via a joint multi-quantile program with elastic-net regularization.

Outputs from the same fitted return object:

- portfolio weights w and conditional tail-risk measures at a state x^*,
- the full conditional quantile surface $q_\tau(x)$,
- two complementary pricing kernels:
  - **HJ-type SDF:** minimum-2^{nd} moment, pricing-feasible kernel,
  - **ES-envelope SDF:** tail-focused kernel projected onto the pricing set.

Takeaway: portfolio choice, tail risk, and state-dependent pricing are linked through one unified estimation workflow.

Roadmap

Method 1: Joint multi-quantile portfolio program

Method 2: Two pricing-feasible SDFs (HJ + projected ES-envelope)

Method 3: Macro-state analysis

Results: Pricing diagnostics & macro patterns

Backup: ES duality and projection argument

Method 1: Joint multi-quantile portfolio program

Data and objects

- Excess returns: $R_t^e \in \mathbb{R}^k$
- State variables: $X_t \in \mathbb{R}^m$
- Portfolio weights: $w \in \mathbb{R}^k$
- Conditional quantile surface:

$$q_\tau(X_t) = X_t' \beta(\tau)$$

Goal: Fit the conditional distribution of $R_t^{e\prime} w$ given X_t.

Method 1: Joint multi-quantile portfolio program

Single-quantile program - Quantile portfolio objective

For a given $\tau \in (0,1)$, define residuals

$$\varepsilon_{\tau,t} = R_t^{e\prime} w - X_t' \beta(\tau).$$

Minimize the empirical check loss:

$$\sum_{t=1}^{T} \left[\tau \, \varepsilon_{\tau,t}^{+} + (1-\tau) \, \varepsilon_{\tau,t}^{-}\right]$$

- Portfolio weights w chosen jointly with quantile fit
- over $\hat{\beta}(\tau)$: Delivers the conditional $\mathrm{VaR}_\tau(X_t)$.
- over w: Minimizes the ES_τ of the portfolio

Method 1: Joint multi-quantile portfolio program

Joint multi-quantile program – Shared portfolio across quantiles

Estimate a grid $\{\tau_j\}_{j=1}^{J}$ jointly:

$$\min_{w,\{\beta(\tau_j)\}} \sum_{j=1}^{J} \sum_{t=1}^{T} \rho_{\tau_j}\left(R_t^{e\prime} w - X_t' \beta(\tau_j)\right)$$

Subject to:

- Budget and box constraints on w
- Target mean constraint
- Optional non-crossing constraints
- Elastic-net penalties:

$$\lambda_w \|w\|_1 + \frac{\eta_w}{2} \|w\|_2^2 \; + \; \sum_{j=1}^{J} \left(\lambda_\beta \|\beta(\tau_j)\|_1 + \frac{\eta_\beta}{2} \|\beta(\tau_j)\|_2^2\right)$$

Idea: estimate the full conditional return distribution of the portfolio while stabilizing both weights and quantile coefficients.

Method 2: Recovery of two pricing-feasible SDFs

Recovery of the HJ SDF – Minimum-2^{nd} moment (HJ-type) SDF

Recover a linear, state-dependent SDF:

$$m_{t+1} = X_t' c - R_t^{e\prime} \gamma$$

Solve:

$$\min_{m,c,\gamma} \sum_{t=1}^{T} m_{t+1}^2$$

Subject to:

$$\frac{1}{T} \sum_{t=1}^{T} m_{t+1} R_{f,t+1} = 1, \qquad \frac{1}{T} \sum_{t=1}^{T} m_{t+1} R_{t+1}^{e} = 0.$$

- Exact pricing by construction
- L^2-optimal on the pricing set
- Optional elastic-net shrinkage on (c, γ)

Method 2: Recovery of two pricing-feasible SDFs

ES-envelope tail measure and projected SDF

Define quantile residuals (at level τ)

$$r_t = R_t^{e\prime} \hat{w} - X_t' \hat{\beta}(\tau).$$

and use the residual ranks from the quantile fit to construct

$$p_{t+1}^{\text{env}} = \bar{p}\, \mathbb{1}\{r_t < a\} + \frac{1 - \bar{p}|B|}{|C|} \mathbb{1}\{r_t = a\}, \qquad m^{\text{raw}} = Tp^{\text{env}},$$

where $\bar{p} = 1/[(1-\tau)T]$, a is the residual cutoff, and B, C collect residuals below and tied at the cutoff.

- ES-envelope probability measure from the Expected Shortfall dual
- Raw mean-one kernel after $m^{\text{raw}} = Tp^{\text{env}}$
- Projected onto the nonnegative pricing-feasible set when exact pricing is required

Idea: Closest pricing-feasible SDF to the pure tail envelope.

Method 2: Recovery of two pricing-feasible SDFs

Economic intuition (tail insurance)

An SDF m_{t+1} prices payoffs across states. Assets that pay when m_{t+1} is high receive high state-price weight and behave like "insurance" assets.

ES-envelope idea:

- Use quantile residuals r_t to identify downside states.
- ES dual weights place mass on the worst $(1-\tau)$ fraction of states, with any cutoff mass split across tied residuals:

$$p_{t+1}^{\text{env}} \propto \mathbb{1}\{r_t < a\} \; + \; \text{boundary mass at } a.$$

Before projection, p^{env} assigns mass only to low-residual states. It therefore provides a transparent tail-state ranking.

Method 2: Recovery of two pricing-feasible SDFs

From tail weights to a pricing kernel

The raw kernel $m^{\text{raw}} = Tp^{\text{env}}$ is **tail-focused** but need not price all test assets exactly:

$$\frac{1}{T}\mathbf{R}^{e\prime} m^{\text{raw}} \neq 0.$$

We therefore compute the **closest nonnegative pricing-feasible** tail kernel when the feasible set is nonempty:

$$m^{\text{proj}} = \arg\min_{m \in \mathcal{A}_+} \|m - m^{\text{raw}}\|_2^2,$$

$$\mathcal{A}_+ = \Bigg\{ m \geq 0 : \frac{1}{T}\sum_{t=1}^{T} m_{t+1} R_{f,t+1} = 1, \qquad \frac{1}{T}\mathbf{R}^{e\prime} m = 0 \Bigg\}.$$

Method 2: Recovery of two pricing-feasible SDFs

From tail weights to a pricing kernel

Subtle but important: tail emphasis and pricing are separate.

- p^{env} selects tail states from residual ranks.
- $m^{\text{raw}} = Tp^{\text{env}}$ is mean-one but does not generally price all assets.
- Projection enforces exact moments when $\mathcal{A}_+$ is feasible.
- The projected kernel inherits its tail emphasis from p^{env}.

Method 2: Recovery of two pricing-feasible SDFs

What is identified?

Objects recovered from the workflow:

- w: tail-sensitive portfolio weights
- $q_\tau(x)$: conditional return quantiles
- m^{HJ}: minimum-2^{nd}-moment pricing-feasible kernel
- m^{ES}: projected tail-envelope kernel

Important comparison: the pricing moments identify an admissible set, not a unique kernel. HJ and ES-envelope are two selections from that set, differing in how they allocate value across states.

Method 3: Macro state analysis

State variables and alignment

Goal: Study how the recovered SDF varies across macro states.

State vector:

$$X_t = (1, Z'_{t-1})'$$

- Z_t = standardized macro variables
- Lagged by one period (no look-ahead)
- Examples:
  - Inflation (YoY)
  - Industrial production growth
  - Unemployment
  - Term spread
  - VIX

Method 3: Macro state analysis

Diagnostics and evaluation

For each recovered SDF m_{t+1}:

$$\overline{mR_f} = \frac{1}{T}\sum_{t=1}^{T} m_{t+1} R_{f,t+1}, \qquad \overline{mR^e} = \frac{1}{T}\mathbf{R}^{e\prime} m$$

- Risk-free pricing error: $\left|\overline{mR_f} - 1\right|$
- Mean absolute risky-asset pricing error:

$$\mathrm{MAE}(m) = \frac{1}{k}\sum_{j=1}^{k}\left|\frac{1}{T}\sum_{t=1}^{T} m_{t+1} R^e_{t+1,j}\right|$$

- Variance proxy $\mathrm{Var}(m)$

Compare:

- HJ kernel (minimum second moment)
- ES-envelope kernel (tail emphasis)

Method 3: Macro state analysis

State dependence of the SDF

Study conditional patterns:

$$E[m_{t+1} \mid X_t^{(j)}]$$

for each macro variable $X_t^{(j)}$.

Two approaches:

- Nonparametric smooths:

$$m_{t+1} \text{ vs } X_t^{(j)}$$

- Tercile splits:

Low / Mid / High states

Question: Does the SDF increase in bad macro states?

Method 3: Macro state analysis

Interpretation: pricing in bad times

Recall:

$$E[m_{t+1} R^e_{t+1}] = 0$$

If m_{t+1} is high in certain states:

- Assets paying in those states are valuable
- They earn lower expected returns
- Those states are "priced"

Comparison:

- HJ kernel $\rightarrow$ smooth mean-variance benchmark
- ES-envelope kernel $\rightarrow$ sharper tail-state sensitivity

Method 3: Macro state analysis

Empirical implementation

Data

- 25 Fama–French size/book-to-market portfolios
- Monthly excess returns
- Lagged macro states:
  - inflation and industrial production growth
  - unemployment, term spread, VIX

Estimator choices

- Baseline grid: $\tau = \{0.50, 0.75, 0.95\}$
- Rolling diagnostics: dense grid $\tau = 0.05, \ldots, 0.95$
- Elastic-net penalties on portfolio weights and quantile coefficients

Outputs: pricing diagnostics, macro-state patterns, and time variation.

Results

Pricing diagnostics

- Both kernels are **pricing-feasible in-sample**:

$$\frac{1}{T}\sum_{t=1}^{T} m_{t+1}R_{f,t+1} = 1, \qquad \frac{1}{T}\sum_{t=1}^{T} m_{t+1}R^{e}_{t+1} = 0.$$

- Key difference: **variance** (HJ is minimum-second-moment; ES-envelope is tail-concentrated).

Metric	HJ	ES-envelope (proj.)
Risk-free pricing error	1.11×10^{-15}	3.33×10^{-16}
Mean absolute pricing error	9.89×10^{-19}	2.28×10^{-18}
ℓ_2 norm of pricing errors	6.02×10^{-18}	1.37×10^{-17}
Variance proxy Var(m)	0.2402	5.0946

$\Rightarrow$ Exact excess-return pricing holds, and the risk-free normalization is satisfied by construction. The ES-envelope kernel has higher variance and places much more weight on tail states.

Rolling holdout: portfolio quantile

Diagnostic	Value
Rolling origins	11
Target hit rate	0.9500
Mean fitted-quantile hit rate	0.8682
Holdout hit-rate range	0.7667–0.9667
Mean holdout lower-tail ES	0.0886

- The fitted 0.95 portfolio-return quantile is not used as a calibrated real-time forecast.
- Its holdout role is to supply the training-window portfolio and residual ordering.
- The SDF diagnostics below test projection and admissibility, not raw quantile forecast accuracy.

Rolling holdout diagnostics: setup and raw failures

- 20-year training window, 5-year holdout window, dense quantile grid.
- Raw ES-envelope does not price risky assets: MAE 1.07×10^{-1}.
- Unconstrained HJ has MAE 1.08×10^{-2} but 50 negative observations.

$\Rightarrow$ Raw tail weights and unconstrained HJ fits are diagnostics, not final admissible kernels.

Rolling holdout diagnostics: projection discipline

- Soft nonnegative projection reduces MAE to 7.80×10^{-5} for HJ and 1.64×10^{-4} for ES.
- Exact nonnegative projection succeeds in 9 of 11 annual origins.
- The two failures share the same empty nonnegative pricing set; separating portfolios are positive in all 60 holdout months.

$\Rightarrow$ The holdout evidence supports the projection discipline, while making infeasible windows visible rather than treating them as numerical failures.

Results

Pricing diagnostics

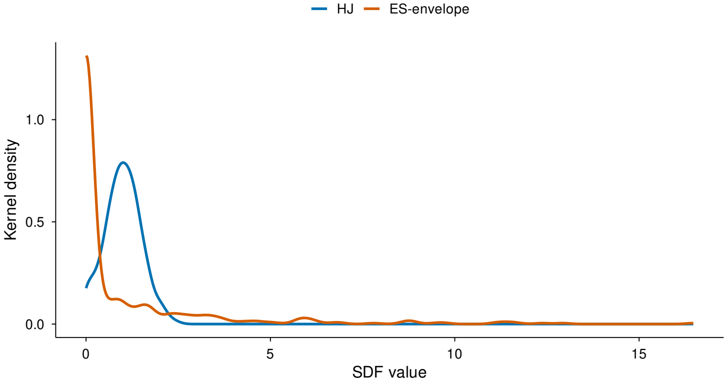

Results

Hansen–Jagannathan bound

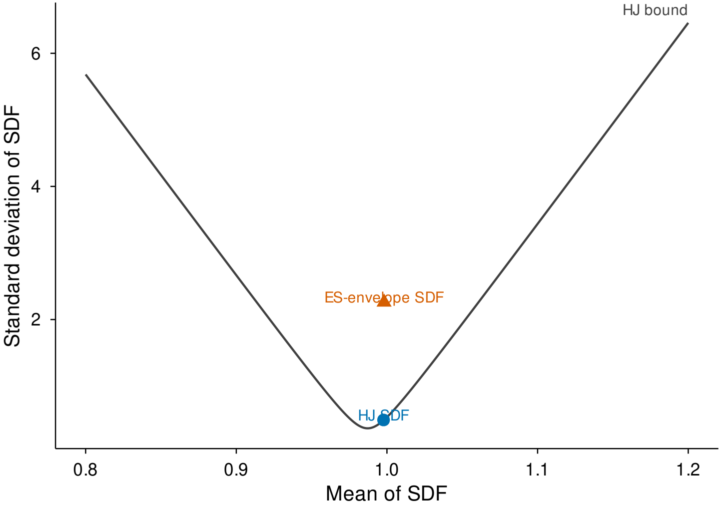

Results

Macro-state dependence

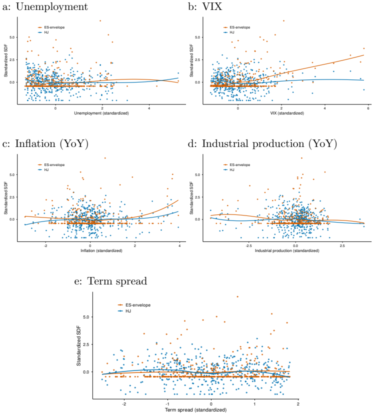

Macro-state takeaways

- ES-envelope–VIX is the clearest positive state pattern.
- HJ variation is smaller; its more reliable spread is low term spread.
- Labor-market evidence is mixed; it is not a stable positive tail-kernel relation.
- Inflation is non-monotone in the current application.

Results

Implied macro state prices

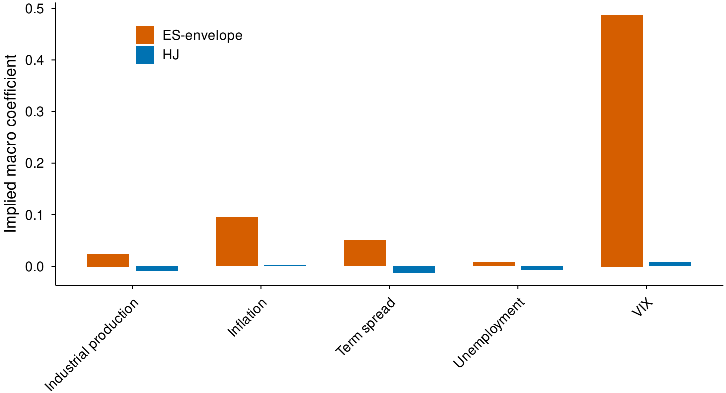

- Once payoff-space restrictions are imposed, the HJ kernel produces moderate, interpretable macro loadings.
- The ES-envelope SDF tilts more toward volatility and drawdown-related states under the same restrictions.

Results
Time variation

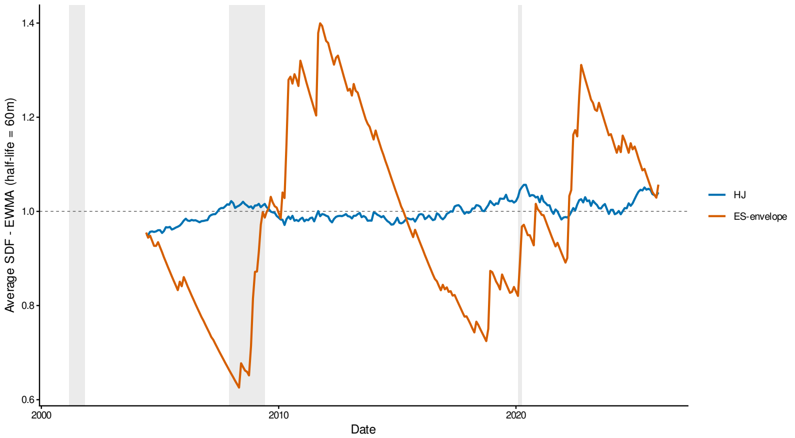

- EWMA half-life: 60 months. Recent states receive more weight.
- Use this view for timing: ES-envelope spikes more sharply around stress episodes.

Results

Time variation

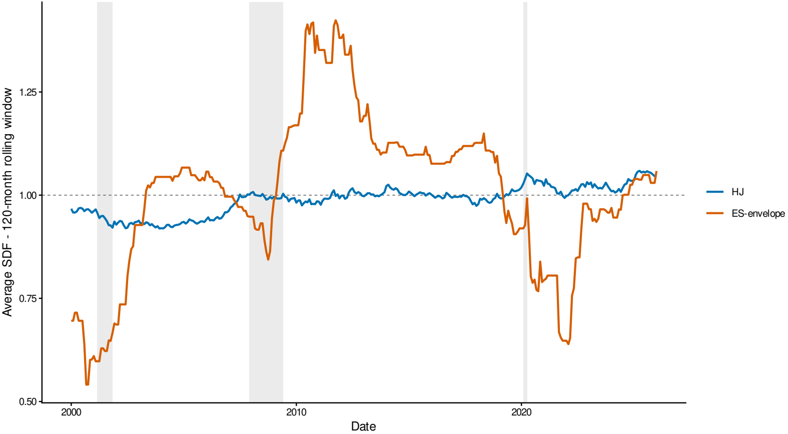

Takeaway: the 120-month rolling mean is a slower level diagnostic. It should confirm broad tail amplification, not match the EWMA turning points month by month.

Results
Conclusion

HJ kernel (minimum 2^{nd}-moment)

- Pricing-feasible by construction
- Smooth weights across states
- Captures “average” market pricing

ES-envelope kernel (tail emphasis)

- Tail-ranking from quantile residuals
- Projection enforces exact pricing when feasible
- Concentrates on downside residual states

Bottom line: reported projected kernels price the cross-section when feasible, but they encode *different* state-pricing (m smooth vs m tail-concentrated).

Implications

1. **Pricing selections vary with macro states**
   - ES-envelope–VIX is the stable positive pattern.
   - HJ evidence is cleaner for low term-spread states.
   - Downside residual states receive larger kernel weights in the tail-focused selection.
2. **Two economically distinct pricing selections**
   - HJ: smooth, mean-variance benchmark.
   - ES-envelope: tail-focused selection from residual ranks.
3. **Applications**
   - Tail-hedge portfolio construction.
   - Macro stress indicators.
   - Connecting tail kernels to preference models.
   - Cross-asset tail pricing.

Contribution

Contribution 1: one optimization, three objects

- Portfolio weights w
- Full conditional quantile surface $q_\tau(x)$
- Tail measure plus projected stochastic discount factor m_{t+1}

Tail risk and asset pricing are linked by residual ranks

$$\text{Quantile fit} \quad \Longrightarrow \quad \text{Residual ranking} \quad \Longrightarrow \quad \text{ES-envelope SDF}$$

Contribution

Contribution 2: two selections from one pricing set

- HJ SDF = minimum-2^{nd} moment projection.
- ES-envelope SDF = tail-focused selection after projection.
- Exact restrictions hold when the nonnegative pricing set is feasible.
- Otherwise, soft projection is a stress-test approximation, not an exact SDF.

Core insight:
One joint quantile program identifies tail states; projection determines when the resulting object is an admissible pricing kernel.

Backup

Boundary states flag stress, not clean extractable payoffs

What zero state prices mean

- Projection hits the boundary of the admissible pricing set.
- The fitted kernel assigns no marginal value to observed stress states.
- This is not model-free evidence of arbitrage.

What the monthly refit found

Signal	Net	DD
HJ	1.0%	−4.2%
ES	1.1%	−4.3%
Mean LP	1.6%	−4.1%

181 out-of-sample months; unit gross exposure; 10bp turnover cost.

Useful as risk timing, not as a clean-payoff machine.

Interpretation. A tradable claim needs recurring states, positive lower-tail payoffs, ETF implementability, and robustness to costs, lags, and short-sale limits.

Dual Representation and the SDF

Expected Shortfall from quantile residuals

Define quantile residuals as

$$r_t = R_t^{e\prime}\hat{w} - X_t'\hat{\beta}(\tau).$$

Lower-tail Expected Shortfall at level τ as

$$\mathrm{ES}_\tau = \min_{\alpha,z}\ \alpha + \frac{1}{(1-\tau)T}\sum_{t=1}^{T} z_t$$

subject to

$$z_t \geq -r_t - \alpha, \qquad z_t \geq 0.$$

- This is again a convex linear program (cp. Rockafellar and Uryasev, 2000; Rockafellar and Uryasev, 2002).
- $\alpha =$ empirical VaR_τ.
- z_t capture tail exceedances.

Dual Representation and the SDF

Dual representation with residuals

For losses $\ell_t = -r_t$, the empirical ES dual is

$$\max_p \sum_{t=1}^{T} p_{t+1}\ell_t \quad \text{s.t.} \quad 0 \leq p_{t+1} \leq \frac{1}{(1-\tau)T}, \qquad \sum_{t=1}^{T} p_{t+1} = 1.$$

Optimal mass is placed on the largest losses, equivalently on the lowest residuals. With $h = (1-\tau)T$, $\bar{p} = 1/h$, $a = r_{(\lceil h \rceil)}$, $B = \{s : r_s < a\}$, and $C = \{s : r_s = a\}$, we use

$$p_{t+1}^{\text{env}} = \bar{p}\,\mathbb{1}\{r_t < a\} + \frac{1-\bar{p}|B|}{|C|}\,\mathbb{1}\{r_t = a\}.$$

(cf. Rockafellar and Uryasev, 2000; Rockafellar and Uryasev, 2002; Acerbi and Tasche, 2002)

Dual Representation and the SDF

Tail-envelope probabilities

Using residual ranks, the ES dual gives a tail probability vector p^{env}. The mean-one raw kernel is

$$m_{t+1}^{\text{raw}} = T p_{t+1}^{\text{env}}.$$

- p^{env} is nonnegative, sums to one, and respects the ES dual cap.
- m^{raw} is mean-one and nonnegative.
- m^{raw} is an SDF only after pricing projection, unless it already satisfies the pricing moments.

Dual Representation and the SDF

Projection onto pricing set

Nonnegative pricing set

$$\mathcal{A}_+ = \left\{ m \geq 0 : \frac{1}{T}\mathbf{e}'m = 1, \quad \frac{1}{T}\sum_{t=1}^{T} m_{t+1} R_t^e = 0 \right\}.$$

projected envelope

$$m^{\text{proj}} = \arg\min_{m \in \mathcal{A}_+} \|m - m^{\text{raw}}\|_2^2.$$

Smallest L^2 adjustment ensuring exact pricing and nonnegative state prices within the feasible set.

Dual Representation and the SDF

When is the ES-envelope an SDF?

Let $r_t = R_t'\hat{w} - X_t'\hat{\beta}(\tau)$, $h = (1-\tau)T$, $\bar{p} = 1/h$, $a = r_{(\lceil h \rceil)}$, $B = \{s : r_s < a\}$, and $C = \{s : r_s = a\}$. From the ES dual, we use

$$p_{t+1}^{\text{env}} = \bar{p}\,\mathbb{1}\{r_t < a\} + \frac{1-\bar{p}|B|}{|C|}\,\mathbb{1}\{r_t = a\}.$$

- p_{t+1}^{env} are dual weights on tail losses.
- The dual cap is assigned below the cutoff; remaining mass is split across cutoff ties.
- They define a probability measure concentrating on residual tail states.
- The raw kernel $m^{\text{raw}} = Tp^{\text{env}}$ is nonnegative but need not price risky assets.
- After projection onto

$$\frac{1}{T}\mathbf{e}'m = 1, \quad \frac{1}{T}\sum m_{t+1}R_t^e = 0,$$

  the measure satisfies the asset-pricing moment conditions.

Hence: the projected m_{t+1}^{proj} is a pricing kernel when the feasible set is nonempty.

Dual Representation and the SDF

Connection to the quantile program

Key observation:
The SDF is constructed from the same residual ordering that solves

$$\min_{w,\beta(\tau)} \sum_{t=1}^{T} \rho_\tau(\varepsilon_{\tau,t})$$

with $\rho_\tau(u) = u\,(\tau - \mathbb{1}\{u < 0\})$.

- The quantile fit determines the state ranking.
- The ES dual assigns mass to the worst-ranked states.
- The pricing kernel is obtained after adding the no-arbitrage projection.

Implication: the projection turns the tail object into an admissible pricing kernel.

Dual Representation and the SDF

HJ SDF as minimum-2^{nd} moment projection

Solve:

$$\min_{m} \frac{1}{T} \sum m_{t+1}^2$$

subject to

$$\frac{1}{T} \sum m_{t+1} = 1, \qquad \frac{1}{T} \sum m_{t+1} R_t^e = 0.$$

Geometric interpretation:

- Affine pricing hyperplane in $\mathbb{R}^T$
- HJ SDF $= L^2$-projection of origin onto that hyperplane
- Minimum possible 2^{nd} moment among pricing-feasible kernels

Dual Representation and the SDF

Geometry: HJ vs ES

HJ SDF:

- Quadratic program
- L^2-projection
- Smooth weights across states

ES-envelope SDF:

- Linear program (dual)
- Extreme-point solution
- Concentrated mass in tail states

Different risk emphasis, both are projected to pricing set.

References

Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking and Finance*, 26(7):1487–1503, 2002. doi: 10.1016/S0378-4266(02)00283-2.

Yacine Ait-Sahalia and Andrew W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1–2):9–51, 2000. doi: 10.1016/S0304-4076(99)00016-0.

Andrew Ang, Joseph Chen, and Yuhang Xing. Downside risk. *The Review of Financial Studies*, 19(4):1191–1239, 2006. doi: 10.1093/rfs/hhj035.

Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1): 113–134, 2015. doi: 10.1016/j.jfineco.2015.02.010.

References

Michael W. Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *The Review of Financial Studies*, 22(9):3411–3447, 2009. doi: 10.1093/rfs/hhp003.

Wayne E. Ferson and Andrew F. Siegel. Stochastic discount factor bounds with conditioning information. *The Review of Financial Studies*, 16(2):567–595, 2003. doi: 10.1093/rfs/hhg004.

Joachim Grammig and Eva-Maria Kuechlin. A two-step indirect inference approach to estimate the long-run risk asset pricing model. *Journal of Econometrics*, 205(1):6–33, 2018. doi: 10.1016/j.jeconom.2018.03.003.

References

Joachim Grammig, Constantin Hanenberg, Christian Schlag, and Jantje Soenksen. Diverging roads: Theory-based vs. machine-learning-implied stock risk premia. *Journal of Financial Econometrics*, 23(2):1–55, 2025. doi: 10.1093/jjfinec/nbaf005.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5): 2223–2273, 2020. doi: 10.1093/rfs/hhaa009.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991. doi: 10.1086/261750.

References

Lars Peter Hansen and Ravi Jagannathan. Assessing specification errors in stochastic discount factor models. *The Journal of Finance*, 52(2):557–590, 1997. doi: 10.1111/j.1540-6261.1997.tb04813.x.

Lars Peter Hansen and Scott F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, 1987. doi: 10.2307/1913601.

Roger Koenker. *Quantile Regression*. Cambridge University Press, 2005.

Sofonias Alemu Korsaye, Alberto Quaini, and Fabio Trojani. Smart stochastic discount factors. *Management Science*, 2025. doi: 10.1287/mnsc.2024.05750. Articles in Advance.

References

Martin Lettau and Sydney Ludvigson. Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, 56 (3):815–849, 2001. doi: 10.1111/0022-1082.00347.

Martin Lettau and Markus Pelger. Factors that fit the time series and cross-section of stock returns. *The Review of Financial Studies*, 33(5):2274–2325, 2020. doi: 10.1093/rfs/hhaa020.

Konstantinos Metaxoglou and Aaron Smith. State prices of conditional quantiles: New evidence on time variation in the pricing kernel. *Journal of Applied Econometrics*, 32(1):192–217, 2017. doi: 10.1002/jae.2515.

R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–42, 2000.

References

R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471, 2002. doi: 10.1016/S0378-4266(02)00271-6.

Stephen A. Ross. The recovery theorem. *The Journal of Finance*, 70(2):615–648, 2015. doi: 10.1111/jofi.12092.

Jantje Soenksen and Joachim Grammig. Empirical asset pricing with multi-period disaster risk: A simulation-based approach. *Journal of Econometrics*, 222(1):805–832, 2021. doi: 10.1016/j.jeconom.2020.08.001.

References

Jantje Sönksen. Machine learning for asset pricing. In Felix Chan and László Mátyás, editors, *Econometrics with Machine Learning*, pages 337–366. Springer International Publishing, Cham, 2022. ISBN 978-3-031-15149-1. doi: 10.1007/978-3-031-15149-1.