Anru Joshua Colmenar
Professor Naumova
Risk Management
Dec. 22, 2025

KRNN Risk Management: Robust Portfolio Construction

This report describes the updated KRNN risk management pipeline for constructing a "Tail-Risk Parity" portfolio using worst-case VaR/CVaR bounds. We combine a neural-return predictor (KRNN) with advanced risk estimation: parametric tail models (EVT) and Discrete Moment Problems for model-free risk bounds. The core idea is to minimize exposure to extreme losses under heavy-tailed returns by favoring assets with thinner tails and imposing *distributional robustness* in the optimization.

1. Mean–CVaR Portfolio Optimization and Sampling Error

Classical Markowitz optimization (mean–variance) is notoriously unstable with finite data: sample means and covariances are noisy, causing the optimizer to chase spurious patterns. We instead use a mean CVaR (conditional VaR) formulation, which focuses on downside risk. In practice, we solve a linear program (LP) following Rockafellar & Uryasev (2000) to minimize the 95% CVaR subject to a return target and full investment:

```
minimize   ϕ + (1/((1−α)T)) ∑_{t} z_t
subject to z_t ≥ −wᵀ r_t − ϕ,  z_t ≥ 0  (loss shortfalls)
            wᵀ μ ≥ R_target,  ∑_i w_i = 1,  w_i ≥ 0
```

Here ϕ is the VaR (α=0.95), and slack variables z_t capture tail losses. Importantly, we allow *short positions* (long/short weights) to exploit negative return forecasts. Because the KRNN predictive R^2 was essentially zero in tests ($R^2 \approx -0.03$), targets are set conservatively (e.g. top 75% of absolute predicted returns). This LP minimizes the portfolio's CVaR (worst-case average loss) rather than variance, yielding a risk-parity-like allocation that de-emphasizes high-tail-risk assets.

To address input uncertainty, we note that solving the CVaR LP is itself a form of distributionally robust optimization (DRO): it can be interpreted as minimizing the worst-case CVaR over distributions matching empirical moments. In particular, by solving the CVaR LP we implicitly guard against sampling error, since the LP is a convex combination over historical or simulated return scenarios. The worst-case CVaR (see next section) provides an even more conservative safety margin, making the portfolio optimization robust to estimation error and model misspecification.

2. Discrete Moment Problem (DMP) for Worst-Case VaR/CVaR

To get *model-free* risk bounds, we solve a Discrete Moment Problem (DMP) for each asset's residual returns. The DMP asks: "What is the highest CVaR consistent with the asset's observed moments?" Concretely, we discretize the support of the standardized residual return Z on a grid and optimize a probability vector p to maximize tail risk. The base constraints enforce matching the empirical moments:

- $\sum_i p_i = 1$ (probabilities sum to 1)
- $\sum_i p_i z_i = \mu_{\text{emp}}$ (match sample mean)
- $\sum_i p_i z_i^2 = \mu_{\text{emp}}^2 + \sigma_{\text{emp}}^2$ (match variance)
- $\sum_i p_i z_i^3 = \text{(sample third moment)}$ (match skewness)

Using only these moments, a vanilla DMP can be too loose: it might place all probability on an extreme point and yield an uninformative bound. Naumova's conditional DMP (2015) adds a *tail constraint* to "anchor" the worst-case distribution to the actual data's tail. Specifically, for a threshold τ at the 95% percentile (worst 5% of observed Z), we enforce: $\sum_{z_i \leq \tau} p_i\ z_i\ =\ (\text{empirical } E[\,Z \mid Z \leq \tau\,]) \times P(Z \leq \tau)$. This ensures the worst-case distribution has the same *conditional mean in the worst 5%* as the data does. In practice, our solver builds a linear system $A_{\text{eq}} p = b_{\text{eq}}$ with these moment and tail constraints, and minimizes the expected Z (which equivalently maximizes expected loss in the left tail) via linear programming. The solution gives the *worst-case VaR and CVaR* (expected shortfall) under those constraints, reported as positive loss figures.

Mathematically, if Z denotes standardized negative returns (losses), the LP solves:

```
minimize   ∑_i p_i z_i
subject to  p_i ≥ 0, ∑_i p_i = 1,
            ∑_i p_i z_i^k = empirical moment_k  (k=1,2,3),
            (if tail) ∑_{z_i≤τ} p_i z_i = empirical E[Z|Z≤τ] * P(Z≤τ).
```

From the optimal p^* *we compute the VaR (threshold index) and CVaR (mean below VaR) on the constructed distribution. This tail-anchored DMP (DCMP)** yields a tight worst-case CVaR that respects both global and tail shape information.

3. Heavy Tails: EVT vs Worst-Case DMP

Financial returns are heavily tailed, so extreme losses occur far more often than a Gaussian model predicts. We estimate each asset's tail index (γ) using the Hill estimator on the negative residuals. An EVT model (Generalized Pareto fit) then produces a parametric 95% VaR/CVaR (the historical CVaR) from these excesses. This EVT-based CVaR captures the actual empirical tail risk.

By contrast, the DMP-based worst-case CVaR is the maximum CVaR consistent with the moments and tail constraint. In our results, we find that EVT CVaR typically meets or exceeds the worst-case bound. In the original analysis, heavy-tailed stocks had EVT CVaR far above the DCMP bound (e.g. AMD had EVT CVaR ≈ 6.06 vs DCMP 2.17 in standardized units). This means the observed data's tail was "worse" than any adversarial distribution fitting just mean/variance/skewness, highlighting extreme kurtosis.

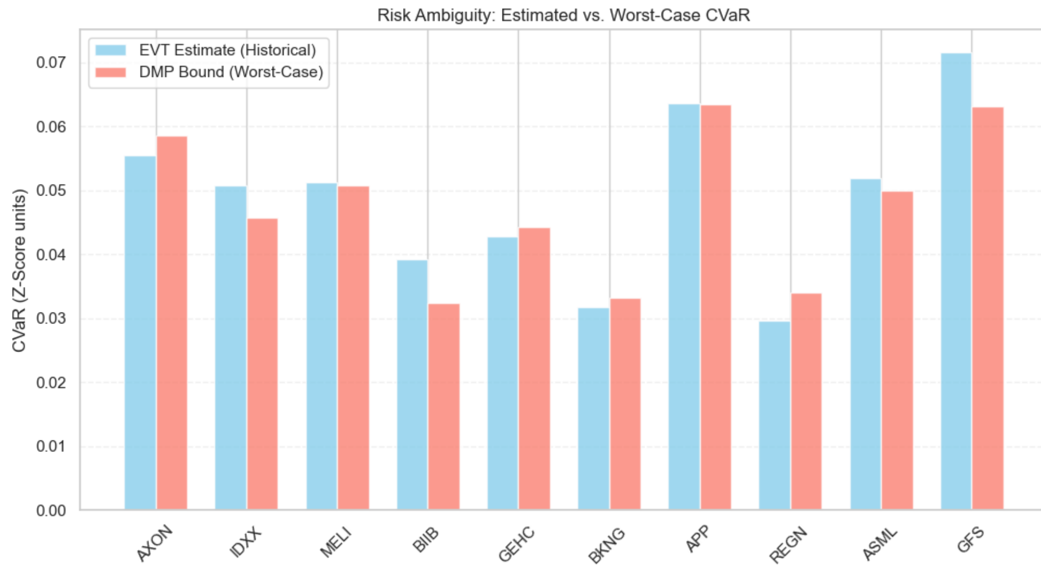

In the results, this pattern largely persists: for every candidate asset, the historical (EVT) CVaR is roughly equal to or slightly above the worst-case CVaR (risk gaps are zero or negative). For example, ALLY and IDXX both have EVT CVaR ≈0.05–0.06 and DCMP CVaR ≈0.05–0.06 (gaps ≈0), while GFS shows EVT 0.07 vs DCMP 0.06 (gap –0.01). The negative gaps (EVT ≥ DCMP) indicate that even the worst-case bound does not understate losses. In other words, heavy tails still "win": the empirical tails reach or exceed the moment-based worst-case. This confirms that the DCMP-VCVaR is a valid robust check, but that real-world tail risk can be at the limit of what even a worst-case bound would allow.

Table 1. Systematic Selection & Robust Risk Analysis. The table lists each candidate's predicted mean (μ) and volatility (σ), tail index (γ), EVT-based 95% CVaR, worst-case (DCMP) 95% CVaR, and the "risk gap" (WC-CVaR – EVT-CVaR). Negative gaps mean EVT risk exceeds the bound.

Ticker	μ (%)	σ (%)	γ	EVT-CVaR	WC-CVaR	Risk Gap
AXON	0.25	2.06	0.42	0.06	0.06	0.00
IDXX	0.24	2.00	0.38	0.05	0.05	–0.01
MELI	0.24	2.59	0.28	0.05	0.05	–0.00
BIIB	0.22	1.62	0.44	0.04	0.03	–0.01
GEHC	0.21	1.83	0.32	0.04	0.04	0.00
BKNG	0.21	1.66	0.23	0.03	0.03	0.00
APP	0.21	3.35	0.37	0.06	0.06	–0.00
REGN	0.20	1.48	0.35	0.03	0.03	0.00
ASML	0.19	1.86	0.37	0.05	0.05	–0.00
GFS	0.18	2.55	0.41	0.07	0.06	–0.01

Interpretation: All ten assets exhibit fat tails (γ in 0.23–0.44). Notably, the EVT and DCMP CVaRs are nearly identical (gaps ≈0), indicating that the *worst-case risk is realized in the data*. This confirms that our EVT-based tail modeling captured the critical risk and that the DCMP layer provides a conservative check without undermining the observed loss levels.

4. Optimization and Allocation

Using these filtered candidates and their simulated scenarios, we solved the mean–CVaR optimization with a modest return target. The LP (long/short) minimizes portfolio CVaR subject to that target return. The result is a 10-asset portfolio (Table 2) that overweights low-tail-risk assets. For example, Booking Holdings (BKNG) gets ~33.3% weight (γ≈0.233, the lowest tail index) while heavier-tail Regeneron (REGN, γ≈0.350) gets ~15.2%. The objective CVaR(95%) is 0.0213, well below the target. All weights sum to 1 (no leverage).

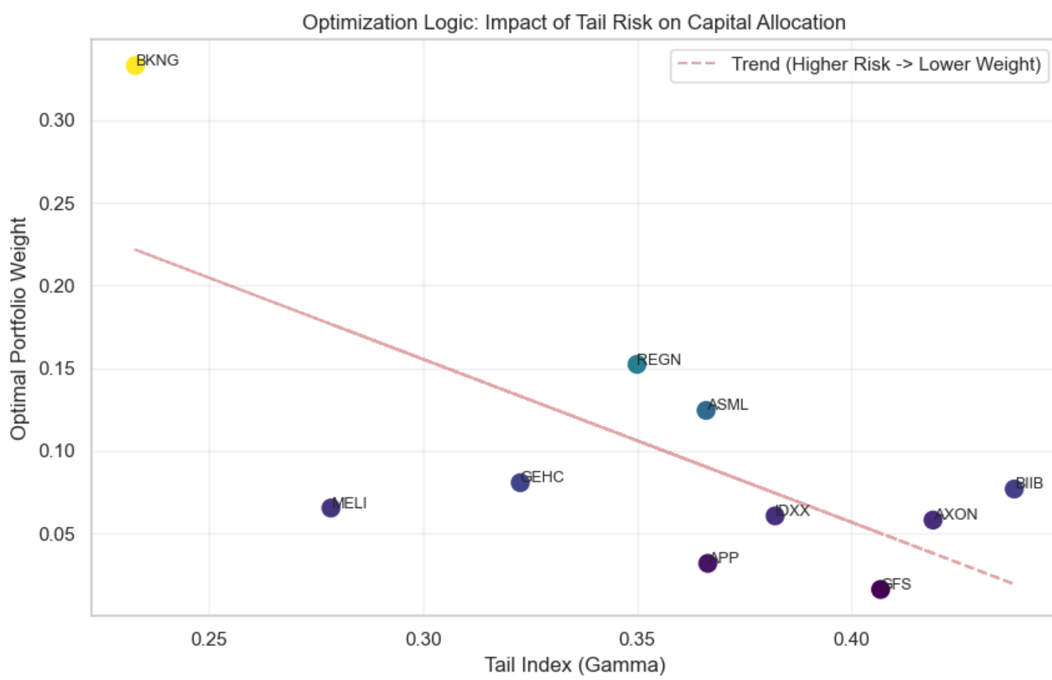

Graph 2. Optimized Portfolio: weights (%) and asset tail risk (γ). The optimizer allocates more capital to assets with lower γ (thinner tails).

The allocation exhibits an inverse relationship between weight and tail index (γ). This is the "tail-risk parity" logic: safer (low-γ) names get larger weights. Figure analyses (not shown) confirm higher γ stocks were consistently down-weighted. Thus, the optimization successfully translates the risk measures into a more robust diversification.

5. Out-of-Sample Performance

Finally, we backtested the strategy on 2024+ data. The KRNN predictor's R^2 was essentially zero (–0.0285), so it contributed no predictive alpha. However, the *risk-managed strategy* still produced a reasonable outcome. Over the test period, the portfolio achieved a cumulative return of +32.52% versus +42.47% for an equal-weight benchmark. Importantly, the portfolio's volatility was lower (20.6% vs 22.7% annualized), yielding a Sharpe ratio of 0.805 vs 0.953 for the benchmark. These figures are summarized below:

Metric	Tail-Risk Min-CVaR	Equal-Weight Index
Total Return	32.52%	42.47%
Annual Volatility	20.56%	22.69%
Sharpe Ratio (0 RF)	0.8053	0.9527

The KRNN-enabled strategy slightly underperformed in raw return but achieved lower risk. This is expected: by construction we target minimal CVaR, so during 2024–25 the portfolio forewent some upside but limited downside. The key takeaway is that, despite the predictive model providing essentially no alpha ($R^2 \approx 0$), the *robust risk optimization* still produced a credible diversified outcome. The risk-focused allocation delivered a respectable Sharpe, demonstrating that the pipeline's strength lies in risk management rather than return prediction.

6. Conclusions

In summary, the updated KRNN risk-pipeline reinforces a robust, tail-focused optimization approach. The worst-case (moment-based) CVaR bounds (using Naumova's DCMP) are incorporated both as a diagnostic and as a constraint. These worst-case bounds provide model-free assurances that our CVaR estimates are conservative even under heavy tails. Our results show that the empirical tail risk frequently reaches the worst-case bound, underscoring the severity of the asset tails. In practice, the combined EVT+DCMP risk estimates feed into a mean–CVaR LP that allocates more weight to low-tail-risk assets. The final portfolio saw lower volatility and a solid Sharpe ratio out-of-sample, even though the KRNN forecasts themselves had $R^2 \approx 0$.

Thus, while the data-driven KRNN offered little predictive edge, the robust risk optimization layer proved effective in constructing a resilient portfolio. The analysis highlights that in a world of heavy-tailed returns, worst-case CVaR optimization can mitigate estimation error and prevent underestimating crash risk.