Conditional risk forecasting

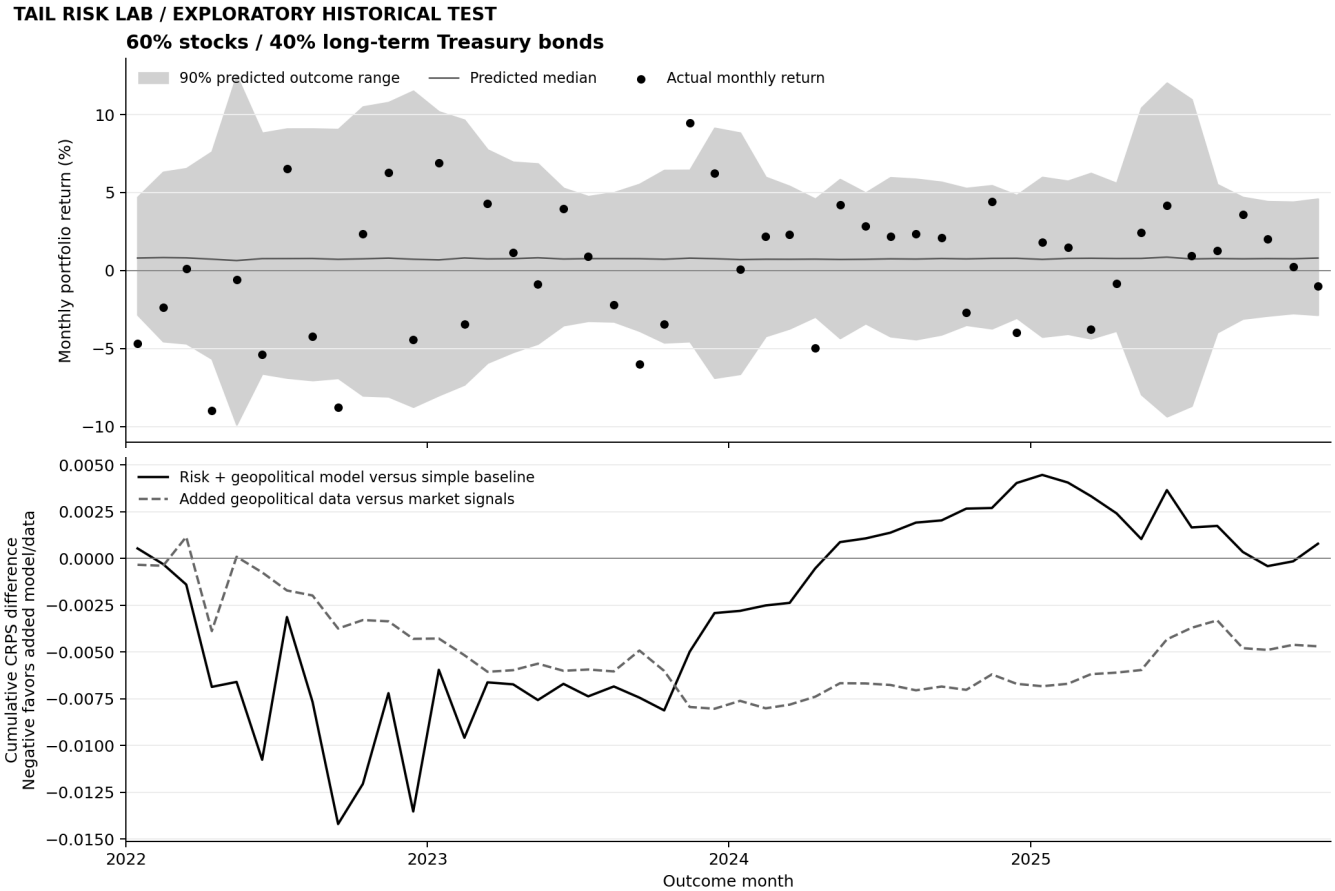
We tested whether geopolitical information helps estimate the size and downside of portfolio moves, with conservative expected growth. Conditions affect volatility rather than directly predicting return direction.

Study: 48 monthly origins in 2022-2025, three related portfolios, five models and 47 pinned geopolitical data vintages. Twelve implementation tests passed. Of 720 forecasts, 715 passed numerical checks; five failed fits remain explicit.

60% SPY / 40% TLT: all 48 months

Model	CRPS (lower better)	90% coverage
Simple volatility baseline	0.02343	85.4%
Conservative drift / constant risk	0.02317	85.4%
Market signals affect risk	0.02342	85.4%
Market + geopolitical risk	0.02332	85.4%
Context + Student-t tails	0.02345	85.4%

Finding: adding geopolitical data to the market-risk model improved average distribution scores by about 0.42% for 60/40 and 0.46% for the stock/bond/gold mix, but worsened them by 0.80% for stocks alone. All paired block uncertainty intervals include no improvement. A reliable predictive advantage has not been established.



The figure uses the Student-t context model. Gray: 90% outcome range. Black dots: realized monthly returns. Lower panel: cumulative score differences, with negative values favoring the added model/data. Scores concern log-growth distributions, not investment profits.

Model and posterior uncertainty

These are parametric Bayesian heteroskedastic models. They are not a new statistical invention or a nonparametric DP/GP model. Detailed equations are supplied in mathematical_model.tex.

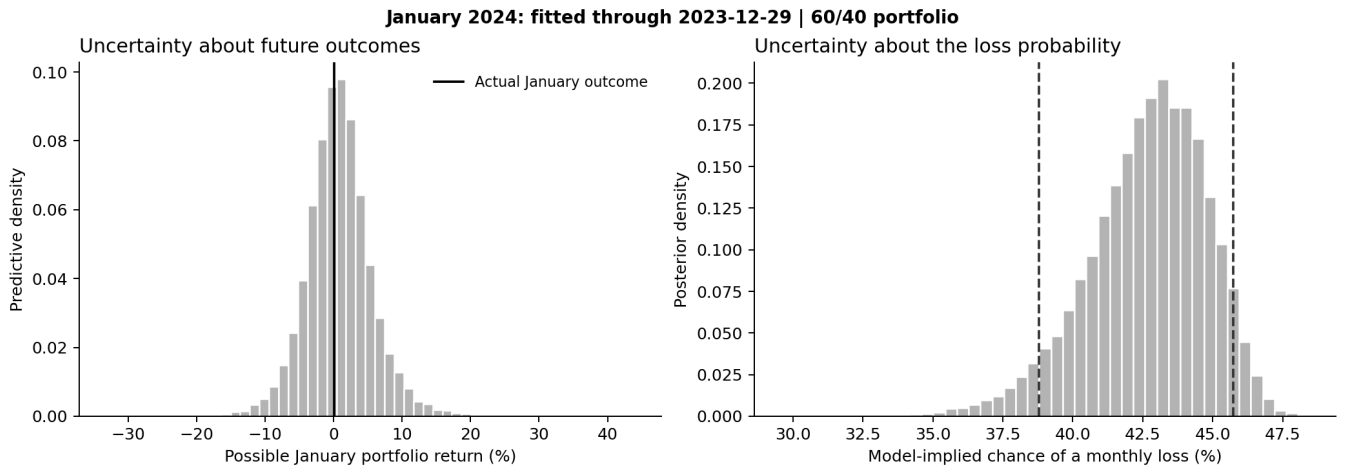
$$r_i = \mu + s_i \exp\left(\frac{x_i^\top \gamma}{2}\right)\varepsilon_i$$

The known scale is preceding 63-session volatility times sqrt(21). Predictors contain only prior market signals and the selected archived geopolitical threats/acts indexes. Drift has prior N(0, 0.005 squared); log-variance slopes have prior N(0, 0.15 squared). Errors are Gaussian or Student-t with 7 degrees of freedom.

Four chains jointly infer drift, log-variance responses and Student-t precision multipliers using Gibbs and elliptical slice sampling. A failed numerical check triggers one longer retry without consulting the outcome.

Historical illustration: January 2024

Fitted through December 29, 2023: for a hypothetical $1,000 in 60/40, the model gave a 90% future-value range of about **$934-$1,088**, a median near $1,007, and a **42.7%** loss probability. The 90% posterior credible interval for that probability was **38.8%-45.7%**. Actual end value in the saved data: about $1,000.66.



Two different uncertainties: the left distribution concerns possible future returns; the right concerns parameter uncertainty about a model-implied loss probability. One contained outcome is not evidence of calibration.

Limits and data attribution

Developed after inspecting earlier results on these years; not untouched or prospective validation. GPR availability follows conservative archive labels, not independently verified publication timestamps. Older training covariates are reconstructed from each outer-origin snapshot. ETF prices are revised vendor history. GPR is newspaper coverage, not war probability or causal identification. Oil/yield inputs are excluded from this follow-up because ALFRED was unavailable.

No fees, taxes, inflation, trading strategy or profit claim. Unbounded scale priors do not generally yield finite mean-dollar wealth; report medians, quantiles and probabilities. The public website model has not been replaced by these research models.

GPR: Caldara and Iacoviello (2022), Measuring Geopolitical Risk, AER 112(4), 1194-1225; matteoiacoviello.com/gpr.htm, downloaded November 28, 2026, CC BY. Sampler: Murray, Adams and MacKay (2010), proceedings.mlr.press/v9/murray10a.html.