

Cross-Sectional Return Prediction with Mean-Variance Portfolio Weights

Yohei Ueki · Yale SOM · MGT 924 Statistical Foundations · Problem Set 3

Summary

The model predicts next-month excess returns of US stocks from 402 firm characteristics using an ensemble of ridge and elastic-net regressions, and turns the predictions into monthly long-short weights with a one-factor risk model and mean-variance optimisation. Over the 408 test months (Dec 1989 – Nov 2023) the portfolio earned an annualised excess return of 15.0% at 9.1% volatility, a Sharpe ratio of 1.65 (95% interval about 1.1–2.2) and a maximum drawdown of -22.7%.

Method

- 1. Features and target.** Each month, every characteristic is replaced by its cross-sectional empirical-CDF rank, centred at zero; missing values are set to zero (the cross-sectional median). The target, next-month excess return, is ranked across stocks within the month, which limits the influence of extreme returns.
- 2. Prediction.** Ridge and elastic-net ($\alpha = 0.5$) regressions are fitted on the training months (ctff_test = False; Dec 1951 – Nov 1989, 456 months). Both are linear, so the fit needs only the monthly cross-products $X'X$ and $X'y$. The penalties are chosen by expanding-window validation inside the training sample (five folds; each fold trains on the past and validates on the next 10% of months): ridge penalty $0.3 \times \text{mean}(\text{diag } X'X/n)$; elastic-net penalty $0.05 \times \text{its } \lambda_{\max}$. The elastic net keeps 77–80 of the 402 characteristics, and validation R^2 on ranked returns is about 1%. The final prediction is the average of the two models' standardised predictions.
- 3. Updating.** In the test period both models are re-estimated every 12 months on an expanding window that contains only months strictly before the month being predicted; the penalties stay fixed.
- 4. Portfolio.** Covariance follows the one-factor model $\Sigma = \sigma_m^2 \beta\beta' + \text{diag}(\sigma_u^2)$, with $\beta = \text{betabab_1260d}$ and $\sigma_u = \text{ivol_capm_21d}$ taken from the data. The market variance is inferred each month from the cross-section, as $(\text{total variance} - \text{idiosyncratic variance})/\beta^2$ with total variance from rvol_21d , and averaged over the last 12 months. Expected returns are $\alpha_i = (\text{rvol_21d})_i \times \text{score}_i$; weights are $\bar{\Sigma}^{-1}\alpha$ (computed with the Woodbury identity), capped at four times the mean absolute weight and scaled to gross exposure 2. The portfolio is rebalanced monthly. Missing risk inputs are replaced by the monthly cross-sectional median; volatilities are winsorised at the 5th/95th percentile and betas are limited to the range -0.5 to 3 .
- 5. Choice of weight map.** Four maps (rank, z-score, top/bottom decile, mean-variance) were compared on the last two validation blocks (91 months of the training sample). Mean-variance had the highest Sharpe ratio (3.77 vs. 2.74, 2.61 and 2.51).
- 6. Timing and reproducibility.** Every quantity at month t uses information available at t : ranks use only that month's cross-section, and models use earlier months only. The code has no random components. Removing the last 24 months of data and re-running the model leaves all earlier weights unchanged (maximum difference 10^{-16}).

Results

Table 1. Out-of-sample performance (excess returns, before transaction costs).

Test sample (408 months, Dec 1989 – Nov 2023)	Value
Annualised excess return	15.0%
Annualised volatility	9.1%
Sharpe ratio	1.65 (s.e. 0.26)
Maximum drawdown	-22.7% (trough Jan 2000)
Months with positive return	73%
Worst / best month	-12.3% / +14.5%
Average gross exposure	2.00
Average monthly turnover ($\Sigma \Delta w $)	1.89




Figure 1. Cumulative excess return and drawdown of the portfolio in the test sample.




Figure 2. (a) Average next-month return by decile of predicted return in the validation months: the ordering is monotone. (b) Sharpe ratio by sub-period of the test sample.

Expected performance and limitations

Expected performance. The full-sample Sharpe ratio of 1.65 has a standard error of about 0.26. The sub-period estimates range from 2.71 (1989-99) to 0.69 (2020-23), with 1.92 in 2010-19, so an investor should expect a lower Sharpe ratio than the full-sample figure; the later sub-periods are the more relevant guide.

Costs. Returns are before transaction costs. Average monthly turnover of 1.89 (against gross exposure 2) implies material costs, especially for less liquid stocks.

Risk model. A single market factor ignores other common risks such as size, value and momentum. The implied market volatility (9.6% annualised, median) is below typical levels, so the market-variance estimate is approximate. A multi-factor (Barra-style) covariance matrix is the natural extension.

Fixed choices. The penalties were chosen once, on pre-1990 data, and the weight map was the best of four on the same data (step 5). Test-sample statistics are computed from the provided next-month excess returns; months without a return are ignored.

Disclosure

The cross-sectional rank transformation of the characteristics and the idea of building a long-short portfolio from predictions follow the example on the competition's rules page. The models, the validation design, the updating scheme and the portfolio construction are my own.