

Common Task Framework: Annual Expanding-Window Ridge

Overview

This project develops a model for predicting next-month stock returns for the Common Task Framework (CTF). These predictions were transformed into portfolio weights.

All 402 characteristics supplied in the CTF feature set were incorporated into the regression to avoid discretionary feature selection and limit the influence of potentially correlated characteristics.

Model development was conducted on observations prior to the CTF evaluation period. Characteristics were cleaned, imputed, and scaled independently within each month. Ridge regression was employed to predict each stock's excess return without overfitting to the data. The model was re-estimated annually using an expanding historical window, and the resulting predictions ranked and converted into portfolio weights to determine which stocks to buy and which to short.

Preprocessing

The dataset contained monthly stock observations of 402 characteristics and next-month excess returns. The prediction target was a stock's excess returns in the following month.

Preprocessing was performed separately for each month to prevent future information from influencing how earlier characteristics were processed. Missing values for each characteristic were replaced using the median across stocks in that same month to limit the effect of extreme observations. Each characteristic was then standardized within the month to have one common scale across the large number of predictors. For characteristic j , stock i , and month t :

$$z_{i,j,t} = \frac{x_{i,j,t} - \bar{x}_{j,t}}{s_{j,t}}$$

Characteristics either entirely missing within a particular month, or with zero or nonfinite standard deviations had their values set to zero for that month. Because zero represents the cross-sectional mean after standardization, assigning a zero provides a neutral value.

Model Selection

Next-month excess returns were predicted using ridge regression. For a given estimation sample, y_i is next month excess return, X_i is the standardized characteristic, and λ controls coefficient shrinkage:

$$\min_{\beta_0, \beta} \left\{ \frac{1}{2N} \sum_{i=1}^N (y_i - \beta_0 - X_i \beta)^2 + \frac{\lambda}{2} \sum_{j=1}^p \beta_j^2 \right\}$$

Ridge regression was chosen because the primary objective was accurate out-of-sample prediction using many potentially correlated predictors. Ridge regression limits the influence of overlapping information contained within the large number of firm characteristics and allows many weak signals of predictive information to contribute to prediction, without dropping those characteristics.

The ridge penalty was chosen by first dividing data into observations from 1951–1969 and 1970–1979. Observations from 1951–1969 were used to fit candidate ridge models which were then used to predict the later 1970–1979 observations. The penalty minimizing mean squared error was $\lambda = 0.04551499$.

A second historical validation exercise used observations from 1980 – 1989 to compare two estimation strategies. The first fixed the fitted model throughout the validation period, while the second re-estimated the ridge regression annually while retaining the selected λ .

Table 1. Historical Validation

Estimation Strategy	Mean Monthly Return	Monthly SD	Annualized Sharpe Ratio	Positive Months
Fixed model	1.54%	1.73%	3.09	84.0%
Annual expanding-window model	1.64%	1.74%	3.27	84.9%

The annual expanding-window strategy was selected because it performed better, producing a higher mean return and higher annualized Sharpe ratio.

Portfolio Construction

At each prediction date, the fitted ridge model produced a predicted next-month excess return for every eligible stock. Cross-sectional ranks were determined for these predictions. For cross-sectional rank $r_{i,t}$ of stock i 's predicted return in month t , the rank is centered around the cross-sectional mean rank for a dollar-neutral portfolio:

$$c_{i,t} = r_{i,t} - \bar{r}_t$$

Portfolio weight was determined by:

$$w_{i,t} = \frac{2c_{i,t}}{\sum_{k=1}^{N_t} |c_{k,t}|}$$

The normalization produced approximately zero net exposure and gross exposure of 2 in each month. Rank-based weighting reduces sensitivity to extreme predicted returns, prioritizing the relative ordering of stocks over their precise magnitudes.

Results

The finalized model was applied to the CTF test observations from December 1989 through November 2023. These observations were not used to modify the model. The mean monthly return was 1.03% (SD 2.83%), and the annualized Sharpe ratio was approximately 1.26. The portfolio produced a positive monthly return in approximately 71.6% of evaluated months.

Figure 1 displays cumulative growth over the evaluation period, with a noticeable decline and recovery around 2008–2010. Cumulative Growth is represented using a logarithmic scale to compare proportional changes over time.




Figure 1. Cumulative Performance of the Long-Short Strategy. Cumulative growth was calculated by compounding monthly portfolio returns over the evaluation period. Portfolio weights are based on cross-sectional ranks of predicted returns and are normalized to approximately zero net exposure and gross exposure of 2.

Table 2. Evaluation Performance by Decade

Period	Mean Monthly Return	Monthly SD	Annualized Sharpe Ratio
1990s	1.59%	1.77%	3.11
2000s	0.97%	4.32%	0.78
2010s	0.53%	1.59%	1.14
2020-2023	1.07%	2.55%	1.45
Overall	1.03%	2.83%	1.26

Performance was not constant over time. The highest Sharpe ratio at 3.11 and strongest performance occurred during the 1990s, while the 2000s had the lowest Sharpe ratio at 0.78 and the most volatility. Despite average monthly returns being lowest during the 2010s, they remained positive across all decades.

Conclusion

The final model uses ridge regression with annual expanding-window estimation and rank-based portfolio construction. Historical validation favored an expanding-window re-estimation over a fixed model. The ridge penalty was determined using data from 1970-1979 and was not adjusted in response to CTF test observations.

Overall, the out-of-sample evaluation produced a positive mean return and annualized Sharpe ratio of 1.26, although performance varied across decades. These results reflect model-based returns on historical data and do not account for all trading costs.