Forming and Evaluating Portfolios

BUSI 722: Data-Driven Finance II

Kerry Back

From Predictions to Portfolio Weights

The Setup

From Session 4, we have one-step-ahead predicted return ranks for every stock in every month.

The Question

How do we translate these predictions into portfolio weights?

We need a function $w_i = g(\text{rank}_i)$ that maps each stock's predicted rank to a portfolio weight.

- The function should be **monotone**: higher predicted rank $\Rightarrow$ larger weight.
- There are many valid choices, from simple to sophisticated.

Sort-Based Portfolios

The simplest approach: sort stocks by predicted rank and equal-weight within groups.

Long-only (top decile):

- Buy the top 10% of stocks, equal-weighted.
- Weight $= 1/n$ for stocks in the top decile, 0 otherwise.

Long-short (top minus bottom):

- Buy the top 10%, short the bottom 10%, equal-weighted within each.
- Weight $= +1/n$ for the top decile, $-1/n$ for the bottom decile, 0 otherwise.

Sorting Is a Monotone Weight Function

Whether we go long the top decile or long-short top versus bottom, the weights are **monotone** in the predicted rank:

- A stock with a higher predicted rank gets a weight that is at least as large as a stock with a lower predicted rank.

Key Observation

Sorting and equal-weighting within groups is just **one particular** monotone function mapping predictions to weights. There are many others.

Pros and Cons of Sorting

Advantages:

- Simple, transparent, widely used in academic research.
- Robust to outliers in predictions — extreme predictions within a group get the same weight.

Disadvantages:

- **Throws away information:** a stock barely in the top decile gets the same weight as the top-ranked stock.
- Sensitive to the choice of breakpoints (quintiles? deciles?).
- Discontinuous: a tiny change in predicted rank can move a stock from one group to another, causing unnecessary turnover.

Smooth Weight Functions

Weights as a Function of Normalized Rank

Scale predicted ranks to $[0, 1]$:

$$u_i = \frac{\text{rank}_i}{n}$$

Then set portfolio weights as a monotone function of u_i.

- This uses the **entire cross section** of predictions, not just the extremes.
- No arbitrary breakpoints.
- Different function shapes express different views about how aggressively to concentrate in top-ranked stocks.

Linear Weights

The simplest smooth weight function:

$$w_i \;\propto\; u_i - \bar{u}$$

where $\bar{u} = 0.5$ is the median rank.

- Stocks above the median get positive weights; below get negative.
- Dollar-neutral by construction (weights sum to zero).
- Every stock's weight is proportional to its **distance from the median rank**.
- Equal spacing: the difference in weight between any two adjacent-ranked stocks is the same.

Power Functions

Apply a power transformation to the normalized rank:

$$w_i \propto u_i^p - c$$

where c is chosen so that weights sum to zero.

- $p = 1$: linear weights (equal spacing).
- $p > 1$ (convex): **concentrates** weight in the top-ranked stocks. Higher $p =$ more aggressive.
- $0 < p < 1$ (concave): **spreads** weight more evenly. Less concentration at the top.

The parameter p controls the trade-off between **conviction** (concentrating in the best predictions) and **diversification** (spreading weight across many stocks).

Exponential Tilts

Apply an exponential transformation:

$$w_i \; \propto \; e^{\lambda u_i} - c$$

- $\lambda > 0$: tilts toward top-ranked stocks.
- Larger λ: more aggressive concentration.
- $\lambda \to 0$: converges to equal weights.
- $\lambda \to \infty$: concentrates all weight in the top-ranked stock.

The exponential tilt is related to the **softmax** function used in neural networks. With a temperature parameter $T = 1/\lambda$, the softmax assigns:

$$w_i = \frac{e^{u_i/T}}{\sum_j e^{u_j/T}}$$

This produces long-only weights that sum to one.

Comparing Weight Functions

Function	Shape	Character
Step (sort-based)	flat within groups	simple, discontinuous
Linear	straight line	smooth, moderate
Power ($p > 1$)	convex	concentrates at the top
Power ($p < 1$)	concave	more diversified
Exponential	convex	aggressively concentrates
Softmax	convex, long-only	all weights positive

All of these are **monotone**: a higher predicted rank always means a larger weight. They differ in how much they **differentiate** between stocks near the top versus the middle.

Long-Only Considerations

Why Long-Only?

Most investors cannot short stocks. Even those who can face:

- Short-selling costs (borrowing fees, margin requirements).
- Short squeezes and recall risk.
- Regulatory constraints (many funds are long-only by mandate).

Long-only weight functions:

- Softmax: $w_i = e^{\lambda u_i} / \sum_j e^{\lambda u_j}$
- Truncated: set $w_i = 0$ for stocks below a rank threshold, then equal-weight or rank-weight the rest.
- Score-tilted market cap: $w_i \propto \text{mcap}_i \cdot g(u_i)$

Score-Tilted Market-Cap Weights

Start with market-cap weights and tilt toward higher-ranked stocks:

$$w_i \;\propto\; \text{mcap}_i \times g(u_i)$$

where g is a monotone increasing function (e.g., $g(u) = u^p$ or $g(u) = e^{\lambda u}$).

- Stays close to the investable market-cap benchmark.
- Large stocks retain significant weight even with low ranks.
- More realistic for large portfolios — avoids overweighting micro-caps.
- The tilt function g controls how much the portfolio deviates from the benchmark.

Evaluating Portfolio Returns

Portfolio Returns

Once we choose a weight function, we compute the portfolio return each month:

$$r_{p,t} = \sum_{i=1}^{n_t} w_{i,t} \, r_{i,t}$$

- $w_{i,t}$ = weight of stock i in month t (from the weight function applied to predicted ranks).
- $r_{i,t}$ = realized return of stock i in month t.
- This gives us a time series of monthly portfolio returns: $r_{p,1}, r_{p,2}, \ldots, r_{p,T}$.

Sharpe Ratio

The **Sharpe ratio** measures risk-adjusted excess return:

$$\text{Sharpe} = \frac{\bar{r}_p - r_f}{\sigma_p}$$

- $\bar{r}_p$ = mean monthly portfolio return.
- r_f = risk-free rate.
- σ_p = standard deviation of monthly portfolio returns.
- Annualize by multiplying by $\sqrt{12}$.

A higher Sharpe ratio means more return per unit of risk. For a long-short portfolio, r_f is often set to zero (the portfolio is self-financing).

CAPM Alpha

Regress the portfolio's excess return on the market's excess return:

$$r_{p,t} - r_{f,t} = \alpha + \beta(r_{m,t} - r_{f,t}) + \varepsilon_t$$

- $\alpha =$ the portfolio's **average return not explained by market exposure**.
- $\beta =$ the portfolio's sensitivity to the market.
- $\varepsilon_t =$ residual (idiosyncratic) return.

A positive α means the portfolio earns more than its market beta would predict. This is what active management aims to deliver.

Is Alpha Statistically Significant?

The regression also gives us a standard error for α, and hence a t-statistic:

$$t = \frac{\hat{\alpha}}{\mathrm{SE}(\hat{\alpha})}$$

- $|t| > 2$ is a rough rule for significance at the 5% level.
- With monthly data over 10 years (120 observations), we need economically large alphas to achieve significance.
- **Always report the t-statistic**, not just the alpha estimate.

CAPM Information Ratio

The **information ratio** measures alpha relative to the volatility of the residuals:

$$\text{IR} = \frac{\hat{\alpha}}{\sigma(\varepsilon)}$$

- $\hat{\alpha}$ = estimated CAPM alpha.
- $\sigma(\varepsilon)$ = standard deviation of the regression residuals ("tracking error" relative to the CAPM benchmark).
- The IR measures how **consistently** the portfolio delivers alpha.

Interpretation:

- The Sharpe ratio measures total risk-adjusted return.
- The information ratio measures **skill** — alpha per unit of idiosyncratic risk.
- Annualize by multiplying by $\sqrt{12}$.

Comparing the Three Metrics

Metric	**Numerator**	**Denominator**
Sharpe ratio	excess return $\bar{r}_p - r_f$	total volatility σ_p
CAPM alpha	$\hat{\alpha}$	(reported with t-stat)
Information ratio	$\hat{\alpha}$	residual volatility $\sigma(\varepsilon)$

- The Sharpe ratio evaluates the portfolio as a **standalone investment**.
- Alpha and the IR evaluate it as an **addition to the market portfolio**.
- A portfolio can have a high Sharpe ratio but zero alpha (if it just takes on market risk).
- A portfolio can have a high IR but a low Sharpe ratio (if alpha is consistent but small).

Putting It Together

1. Generate one-step-ahead predicted ranks (Session 4).
2. Choose a weight function: sort-based, linear, power, exponential, softmax, or score-tilted.
3. Compute monthly portfolio returns.
4. Evaluate: Sharpe ratio, CAPM alpha (with t-stat), and information ratio.
5. In Session 6: extend evaluation to Fama-French factors and other considerations.