

# Symmetric Löwdin Orthogonalisation of Characteristics

Haoyu Wang  
University of Auckland

September 20, 2026

**TL;DR:** We orthogonalise the 402 supplied characteristics within each formation month before using them, via the symmetric (Löwdin) transform. This tends to raise the Sharpe ratio of the resulting portfolio, mainly by lowering its volatility rather than by raising its average return.

## 1 The method

For each formation month  $t$ , the characteristics are first cleaned within that month’s cross-section: missing values are filled with the cross-sectional median, values are winsorised at the 1st and 99th percentiles, and each column is demeaned and scaled to unit variance. Call the cleaned matrix  $Z_t$ .

The symmetric Löwdin transform is

$$\tilde{Z}_t = Z_t G_t^{-1/2}, \quad G_t = \frac{1}{N_t} Z_t^\top Z_t,$$

where  $G_t$  is the cross-sectional Gram (correlation) matrix of the cleaned characteristics, and its eigenvalues are floored at  $10^{-6}$  before the inverse square root is taken to keep the transform well defined. The result  $\tilde{Z}_t$  is exactly orthogonal,  $\tilde{Z}_t^\top \tilde{Z}_t / N_t = I$ .

What makes the Löwdin transform particular, relative to other ways of orthogonalising  $Z_t$ , is that it is symmetric: among every orthogonal transform of  $Z_t$ , it is the one closest to  $Z_t$  itself in a least-squares (Frobenius) sense. An asymmetric orthogonalisation such as Gram–Schmidt would instead privilege whichever characteristic happens to be ordered first, and a PCA-based orthogonalisation would require choosing how many components to keep. The symmetric map does neither: every characteristic is treated the same way, and all 402 orthogonalised characteristics are retained.

## 2 Why it helps

The 402 raw characteristics are highly correlated with one another. Feeding correlated regressors into a linear predictive model makes the fitted coefficients unstable: small changes in the training sample can produce large swings in individual coefficients, even when the overall fit barely changes. Those unstable coefficients translate into positions that move around more from one formation month to the next than the underlying signal warrants.

Orthogonalising first removes that instability. Because  $\tilde{Z}_t$ ’s columns are uncorrelated by construction, the model’s coefficients are identified independently of one another, and the resulting predicted returns, and the portfolio built from them, are steadier through time.

In our tests, that steadiness shows up mainly as a reduction in portfolio volatility rather than as an increase in average return: the orthogonalised version of the model earns a similar average return to the un-orthogonalised one, but with a smoother return path, so the Sharpe ratio, return divided by volatility, improves primarily through the denominator.

### 3 Controlled ablation

The comparison below changes exactly one component: the baseline feeds the cleaned and standardised characteristics directly into PLS(5), while the treatment applies the symmetric Löwdin transform first. Everything else (the expanding-window training schedule, five PLS components, the 756-day regularised covariance estimator, and the zero-net, unit-gross MVO operator) is identical between the two.

| Metric                        | PLS5 without Löwdin | Löwdin PLS5    | Change      |
|-------------------------------|---------------------|----------------|-------------|
| Annualised mean excess return | 5.23%               | <b>6.08%</b>   | +0.85 pp    |
| Annualised volatility         | 1.95%               | <b>1.94%</b>   | −0.01 pp    |
| Sharpe ratio                  | 2.687               | <b>3.132</b>   | +0.445      |
| Maximum drawdown              | −3.84%              | − <b>2.59%</b> | +1.25 pp    |
| Positive months               | 77.70%              | <b>82.84%</b>  | +5.15 pp    |
| Worst month                   | −0.81%              | − <b>0.78%</b> | +0.03 pp    |
| Best month                    | 2.63%               | <b>2.78%</b>   | +0.15 pp    |
| Evaluation months             | 408                 | 408            | same window |

Table 1: Controlled comparison over the common evaluation window, December 1989 through November 2023. The annualised mean and volatility use factors of 12 and  $\sqrt{12}$  respectively; the Sharpe ratio uses a zero benchmark since the target is already an excess return. Results are gross of transaction costs.




Figure 1: Cumulative value of one unit of capital, log scale, over the evaluation window.

The raw factors show pronounced blocks of positive and negative dependence. After the symmetric Löwdin transform, off-diagonal correlation is numerically negligible, confirming a full whitening rather than a partial rotation or rescaling.




Figure 2: Pairwise Pearson correlations across the same 402 active factors, before and after the symmetric Löwdin transform, with an identical hierarchical-clustering order in both panels.

| Absolute off-diagonal correlation | Before | After                  |
|-----------------------------------|--------|------------------------|
| Mean                              | 0.1025 | $3.18 \times 10^{-15}$ |
| Median                            | 0.0619 | $8.37 \times 10^{-16}$ |
| 95th percentile                   | 0.3412 | $1.11 \times 10^{-14}$ |
| Maximum                           | 0.9993 | $2.54 \times 10^{-12}$ |

Table 2: Summary of absolute pairwise correlation outside the diagonal, same 402-factor cross-section, November 2023 ( $N = 2,257$  securities).