02 / AI RISK

Different assets. The same bet?

A varied portfolio can still depend on one economic outcome. Hokma lets learned market relationships shape position sizing, without replacing the allocator that sizes them.



Different holdings can share more beneath the surface than their labels suggest.

OPERATING
construction route with learned similarity · pre-trade gate

UNDER VALIDATION
value over a correlation route (R2), after a configuration check (R1)

PROPOSED
realised-outcome review loop

A portfolio statement can look reassuringly varied: equities in one line, bonds in another, currencies further down. Yet several positions may be exposed to the same change in growth, inflation or funding conditions. Different labels can conceal a shared bet.

Portfolio managers already address this through correlation, covariance and risk limits. Across Hokma's own instruments, mean trailing 63-day pairwise correlation has been near 0.40 in 2026, not the correlation of independent bets. A risk budget that counts positions rather than risks sizes the wrong thing.

Hokma's answer changes the allocator's input, not the allocator. A convolutional autoencoder reads windows of each instrument's returns, refitted at every run, and learns a compact code for how it has behaved. Similarity between codes is the one point at which learning enters.

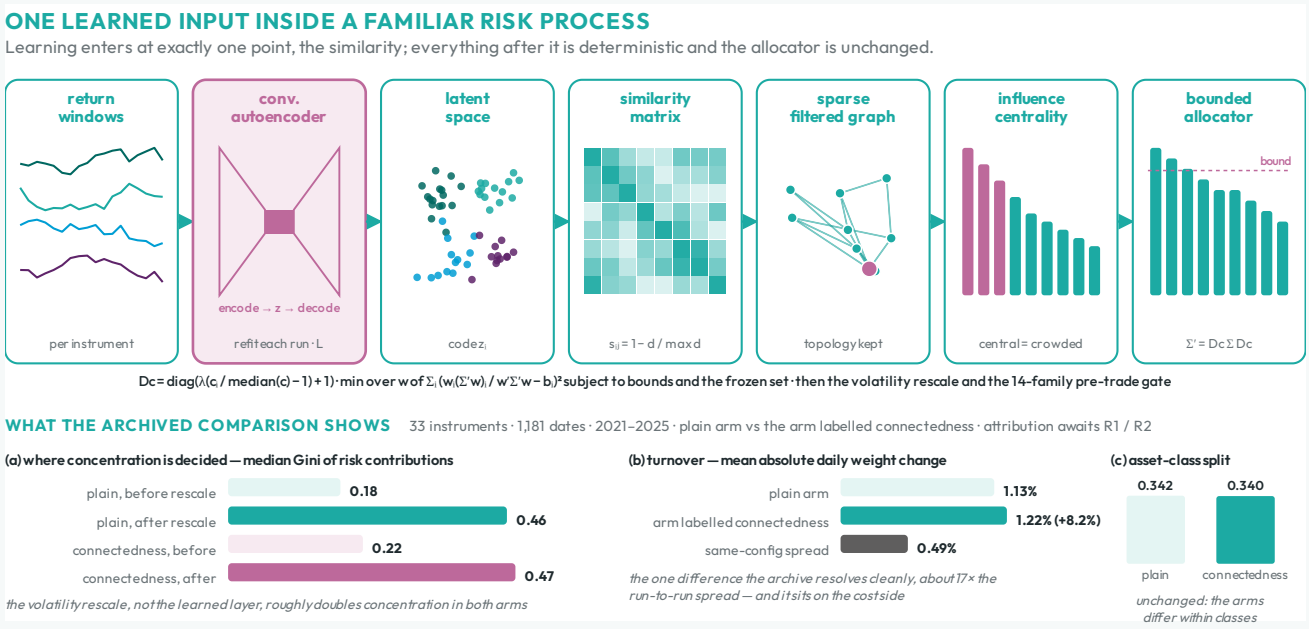
Everything after it is deterministic. The similarity is passed through a sparse topology-preserving filter, each instrument is scored by an influence centrality on the filtered graph, and the score rescales the covariance on which an established bounded risk-budgeting allocator solves. A central, crowded instrument looks riskier than its own history says, and the allocator holds less of it within existing limits.

The decision that changes is each instrument's entitlement to the risk budget, set before the bounds and the volatility target act rather than corrected afterwards. The output is constrained portfolio weights, checked by a 14-family pre-trade gate before release.

Diversification is about the risks you share, as well as the assets you own.

From relationships to position sizes

An appealing network picture is only a starting point. The investment question is whether the learned input changes the portfolio in a useful way, net of the turnover it adds.



Judge the approach by the portfolio it produces. Does it reduce unwanted common exposure? Does any improvement survive trading costs? Do the limits leave room for the new information to reach the final weights?

The archive answers part of this. On a 33-instrument backtest panel over 1,181 dates, the arm labelled connectedness differed from the plain arm within asset classes by about the run-to-run variation, left the asset-class split unchanged and had 8.2% higher turnover. In both arms the volatility rescale applied after the solve roughly doubled concentration: most of the released book's concentration is set downstream of the learned layer.

Attribution is still open. Both arms compute the same learned similarity and their labels differ in two configuration tokens, so the paper writes "the arm labelled connectedness", not "the layer", until a configuration check or the controlled comparison against a correlation-based similarity settles it. Turnover is the hurdle any benefit must clear.

The construction route is in daily operation on the macro book, and turnover is the limit the pre-trade gate warns on most consistently. Whether the learned representation improves construction at common risk, net of turnover, is the experiment the paper defines.

33 × 1,181

instruments and dates in the archived comparison

2021 to 2025; plain arm against the arm labelled connectedness

+8.2%

turnover in the arm labelled connectedness

about 17× the spread between two runs of one configuration

0.18 → 0.46

median Gini of risk contributions, before and after the volatility rescale

plain arm; the rescale, not the learned layer, decides most of the concentration

EXPLORE THE FULL WHITE PAPER

Learning Market Connectedness for Portfolio Construction

Jonathan Lee · Valeriia Pervushyna · Hokma IQ Research · AI Quant series

Request a copy: inquiries@hokma.ae

www.hokma.ae

Source: accompanying white paper, version 1.0. Risk contributions are model-estimated, not realised. Illustrations are conceptual. For information only.