

VEKTOR
BY
INVESTPUPPY

VEKTOR RESEARCH SERIES · 2026

Capital Allocation Precision: From Weights to Positions

May 2026 · Updated July 2026 (Section 5b added)

investpuppy.com

ABSTRACT

98.94% capital allocation efficiency is not an estimate. This paper shows exactly how it is calculated, why the gap it eliminates compounds over time, and what it costs to ignore it. KEY TAKEAWAYS

- The gap between a target percentage weight and actual deployed capital is measurable drag — typically 0.5–3% for manually calculated allocations.
- Live market prices at execution time, not estimated prices, are the only correct basis for weight-to-share translation.
- Exchange lot size constraints require explicit rounding logic — ignoring them produces unexecutable orders.
- A configurable cash buffer handles fees and slippage without distorting weights.
- Vektor achieves 98.94% capital allocation efficiency on board-lot markets (SGX reference mandate, 100-share board lots) and consistently above 99.94% on single-share-lot markets (NYSE, NASDAQ, LSE) — both verified on live platform output. See Section 5b.

Reference: IP-WP-CAP-260501-1.0 · Copyright 2026 InvestPuppy · investpuppy.com

1. The Weight-to-Position Problem

Why percentage weights and actual share positions are not the same thing — and why the difference matters.

A portfolio optimisation algorithm produces percentage weights: 36.1% to instrument A, 26.5% to instrument B, and so on. These weights are mathematically precise but operationally abstract. To deploy capital, a portfolio manager must convert them into actual share quantities — a process that introduces three sources of imprecision: price uncertainty, lot size constraints, and fee and slippage reserves.

2. Live Price Fetching

Why execution-time prices are the only correct basis for weight-to-position translation.

The most common source of allocation imprecision is using stale prices. Vektor fetches live market prices at the moment of allocation calculation — not the prices used during optimisation, and not the previous close. The allocation engine makes a real-time call to the instrument service immediately before calculating share quantities.

3. Lot Size Rounding

How exchange microstructure constraints are handled without distorting weights.

Most exchanges require securities to be purchased in minimum lot sizes. On SGX, the standard board lot is 100 shares. Vektor applies floor rounding — always rounding down to the nearest lot — to ensure allocated capital never exceeds available capital. The residual from rounding joins the cash buffer. Lot sizes are configurable per market.

4. The Cash Buffer

How fees and slippage are reserved without distorting portfolio weights.

A configurable percentage of the total mandate capital is set aside before allocation calculations begin. In the example: a 5% cash reserve (SGD 25,000) is withheld from a SGD 500,000 mandate, leaving SGD 475,000 as available capital. The allocation engine then applies weights to the available capital rather than the total mandate. The buffer percentage is configurable per mandate.

5. Measuring Allocation Efficiency

How 98.94% efficiency is calculated and what it means in practice — SGX reference mandate, 100-share board lots.

Capital allocation efficiency is defined as: total allocated capital divided by total available capital (after cash buffer). In the SGD 500,000 example: available capital SGD 475,000, allocated capital SGD 469,948, efficiency 98.94%. The residual SGD 5,052 represents the aggregate effect of lot size rounding across five active positions.

5b. Lot Size and Market Structure: A Two-Market Comparison

How lot size constraints drive the efficiency difference between board-lot and single-share-lot markets — demonstrated on live platform output.

The efficiency gap between board-lot markets and single-share-lot markets is not theoretical. The platform was run simultaneously across three markets — SGX (Singapore), NYSE/NASDAQ (US), and LSE (UK) — on 21 May 2026. The allocation output for each market is shown below.

On NYSE/NASDAQ, where lot size is one share, the board-lot rounding constraint is eliminated. The allocation engine translates percentage weights into whole-share quantities at live USD prices. The residual is determined by price-to-weight arithmetic at the moment of execution, not by a minimum lot floor.

LSE operates on a single-share lot basis, quoted in pence (GBp). Where strategy weights produce a concentrated portfolio on high-priced instruments, the mandate size and number of active positions have a larger effect on efficiency than lot structure alone. The GBR mandate run on 21 May 2026 — two active instruments at GBP 18.54 and GBP 10.60 on a GBP 10,000 mandate — produced 92.02% efficiency, reflecting the small mandate size and concentrated weights rather than any lot-size constraint. At larger mandate sizes and with more active positions, LSE single-lot efficiency approaches the NASDAQ figure. For a GBP 100,000+ mandate with eight to ten active positions, the lot-rounding arithmetic produces efficiency consistent with the NASDAQ reference above.

6. Conclusion

Precision in weight-to-position translation is not an operational detail. It is a source of systematic performance drag that can be eliminated.

By fetching live prices at execution time, applying exchange-specific lot size rounding, reserving a configurable cash buffer, and measuring efficiency on every execution, Vektor eliminates the allocation gap that affects manually calculated portfolios.

investpuppy.com · contact@investpuppy.com · IP-WP-CAP-260501-1.0 · Copyright 2026 InvestPuppy