

SGX Multi-Factor Trading Strategy: Consolidated Research Report

Nicholas Hong · NTU School of Electrical and Electronic Engineering · FYP A1083

Rendered from `REPORT_FINAL.md` (document of record, consolidation 2026-08). Every number is provenance-tagged to `VERIFICATION_LEDGER.md`; excluded and superseded results are catalogued in `EXCLUSIONS.md`.

Repository: <https://github.com/Nicholashsw/sgx-multifactor-fyp>

Systematic Multi-Factor Trading Strategy on the SGX Market

NTU FYP A1083. Final consolidated research report, verified edition. Supervisor: Prof. Patricia Wong, School of EEE, Nanyang Technological University.

This document consolidates the empirical work of the project across all research tracks. Every number in it is VERIFIED in the accompanying verification ledger: it traces to an executed notebook cell output, or to a result file in the archive whose inputs are real-source and whose producing script is identified, on real data only. Bracketed tags cite ledger rows (`VERIFICATION_LEDGER.md`), which in turn cite the exact notebook, cell and input dataset. Claims from earlier drafts that could not be traced, or that traced to fabricated data, are not in this document; they are listed with reasons in `EXCLUSIONS.md`. Null results are reported as findings.

1. Thesis

The Singapore market contains real, correctly signed, replicable cross-sectional return structure, and the frictions of the market remove it before it becomes tradeable at institutional scale. The project's contribution is a measurement of where that boundary lies, together with two methodological findings that generalise beyond Singapore: survivorship bias on SGX runs in the opposite direction to the textbook prediction, and the reference literature's factor premia do not reproduce under honest multiple-testing discipline.

2. Data and universes

Three universes, in strict order of inferential authority.

Broad Compustat universe (canonical for cross-sectional inference). Survivorship-free Compustat Global panel, `fic=SGP`: 704 gvkeys raw, 694 after primary-issue and currency filters, 1985 to 2026, roughly 3.26M daily rows, 338 delisted firms retained, total returns from `prccd/ajexdi x trfd [A1]`. Point-in-

time index eligibility averages 199 names per month (range 37 to 292) over 427 months [C2]. The liquid100 cut used by the momentum and TSMOM work takes the top 100 by lagged market cap with a SGD 0.20 price floor.

PIT STI/MidCap roster (SGX3 track). 54-name point-in-time roster built from 76 quarterly STI constituent snapshots, 2008 to 2026, roughly 30 names per snapshot, with delisting returns taken at each name's true exit [A7]. Membership churn over 75 quarterly reviews: 30 additions, 30 deletions [A8].

Current-constituent 30-name STI panel (diagnostic only). Bloomberg monthly panel, 2001 to 2026, 7,891 security-months. A field audit of the exported panel used by the original diagnostic track established that PX_LAST is its only genuine independent field: TOT_RETURN, PX_VOLUME and market cap in that export are fixed multiples of the price path (for 23 of 30 tickers exactly the fallback formula), and PX_TO_BOOK and DIV_YIELD are constant within ticker [A9]. That is the signature of the synthetic fallback, not of the source workbook: the workbook itself is genuine, with IPO-ragged listing dates and plausible fundamentals, and clean re-extractions from it carry real fundamentals [A10]. The 30-name universe is survivorship-biased by construction and is never used for headline claims.

Cross-vendor validation is executed at three levels: the Bloomberg 30-name equal-weight panel correlates 0.967 with the Compustat point-in-time top-30 (n=300 months) [A3] and 0.947 with the STI index [A4], and 238 per-name return streams match across vendors at median r 0.9998 with median absolute difference 8.1 bps/month [A5].

Supporting data: IBES summary estimates, Markit Securities Finance (1,443,330 daily rows, 572 SG names, 2002-05 to 2026-07, matched to 71.8% of panel name-months) [A12], SDC M&A, JKP Global Factor Data (208 series, 8 markets), CEIC macro merged to a real macro panel [A11], daily Bloomberg market series, Factiva for delisting adjudication. Licensed vendor data does not ship with the repository; see DATA.md.

3. Methodology conventions

Newey-West HAC standard errors at 6 lags. Spearman rank IC at horizon $h=1$ with a 4-month fundamental reporting embargo; the embargo grid shows results are flat beyond a 2-month lag, so 4 months is a convention, not an optimum [F4], and the event-time calibration shows roughly 24% of firm-quarters remain unpublished even at 4 months [F10]. Forward returns on the fundamental and PIT panels use explicit point-in-time lookup rather than bare shift(-1) after the correction in section 12; the price-signal branch uses shift(-1) on gap-free monthly grids where the two coincide, and audits its own alignment with placebo grids (section 9.2). Benchmarks are always the cohort investable at formation date t (section 8). Corwin-Schultz per-name half-spreads with the traded-days correction (the standard recipe inverts the liquidity ordering on illiquid names) [D4]. Markit availability and indicative fees on the short side. Equal weighting by default, with the small-book cap trap disclosed (section 7). Pre-registration discipline where used is reported with its gaps: the 65-test grid's lock date and panel hash were left blank and its amendment log is empty for the 62-to-65 extension, though both multiple-testing corrections are supplied and the conclusion is invariant [G3].

4. Survivorship bias runs backwards on SGX

The central methodological finding, executed on the survivorship-free Compustat panel. Including dead names raises measured universe returns: the bias from dropping them is -0.89 pp/yr on the whole listed universe (8.85% PIT vs 7.96% survivors-only) and -1.01 pp/yr within point-in-time index members [B1]. The mechanism is exit composition: 198 of 338 delistings (59%) are acquisitions or mergers, and acquired names average $+3.44\%/mo$ over their final year against a $+0.96\%/mo$ survivor baseline; exactly one exit is a bankruptcy [B2]. SDC matching confirms 226 of the 338 exits against completed-deal records [B4]. A pre-registered alternative explanation, that index membership flips the sign, was tested and rejected: the sign is identical in both universes [B3].

The -0.89 figure is an upper bound on the bias magnitude, not a point estimate. Compustat freezes suspension prints (one adjudicated wipeout's last price exactly equals its suspension print while the true outcome was -100%), and last prices can predate deal completion by many months [B8]. A terminal-return sensitivity executed in this consolidation brackets the magnitude while leaving the sign untouched: forcing -100% final-month returns on the 37 triaged wipeouts moves the bias to -0.60 pp/yr, forcing the 25 ambiguous exits as well moves it to -0.40 , and forcing all 53 code-10 exits without an SDC deal record moves it to -0.44 . The direction survives the harshest plausible assumption [B6a]. A four-category delisting taxonomy (cash takeout, stock merger, solvent liquidation, suspension-then-insolvency) is established and instantiated on the hand-adjudicated cases; full adjudication of the 115 code-10 exits remains open [B5].

Strategy-level consequences differ sharply by book construction. On the blue-chip 71-name track, the naive survivor backtest earns $17.37\%/yr$ at Sharpe 1.002 with a -44% drawdown; the corrected quarterly-PIT run earns $7.91\%/yr$ at Sharpe 0.447 with a -63% drawdown, and post-GFC the naive Sharpe 1.067 corrects to 0.667 [B9]. On the broad universe the long-short reverses direction: survivor long-short alpha $0.46\%/mo$ ($t=1.01$) against point-in-time $1.32\%/mo$ ($t=3.82$) [B11]. On the 54-name PIT roster the correction is near-symmetric across legs, costing the long leg 5.43 pp/yr and returning 5.26 pp/yr to the short leg [B12]. These reconcile as the same bias with opposite signs by leg: a long-only investor is hurt by the correction, a long-short investor is roughly insulated.

Two implementation traps are documented with executed demonstrations. Applying the membership filter to both eligibility and returns silently reintroduces the bias; the fix (eligibility from the filtered panel, returns from the unfiltered panel) closes the portfolio-return reconciliation to $5.6e-17$, from a documented pre-fix gap of order $1e-2$ [B13]. And Bloomberg's post-delisting price padding must be truncated at the last genuine move; the truncation independently recovers the real delisting dates of CapitaMalls Asia, Keppel Land, NOL, GLP, Noble and CapitaLand Commercial Trust without being told them [B14].

5. Breadth and size are the binding statistical constraints

Down the size ladder, IC t-statistics run 1.06, 4.10, 8.87 (large, mid, small) and long-short alpha t-statistics 0.81, 2.99, 4.13; a liquidity-screened small-cap book reaches alpha $1.92\%/mo$ ($t=5.12$, Sharpe 0.85) gross [C3]. Breadth proper shows the same effect once the benchmark is fixed: under the investable-cohort benchmark the top-100 book sits at $t=1.67$ while the top-150 and top-200 books reach $t=3.30$ and 3.61 , strengthening with breadth [E4]. Nothing estimated on the 30-name cross-section

carries inferential weight; headline claims are estimated on the broad universe or the PIT roster. (An earlier memo's 3-to-25-name book-size sweep pointed the same way but was never archived as an executed run and is excluded.)

6. The cost waterfall: statistically real, economically untradeable

The definitive waterfall on the winning broad composite (2003 to 2026, matched window): gross alpha 1.661%/mo ($t=5.13$) falls to 0.755 ($t=2.25$) after per-name Corwin-Schultz spreads, the dominant cost at 0.906%/mo, then to 0.527 ($t=1.59$) after the Markit availability screen, and to 0.282%/mo ($t=0.85$) after borrow fees [D1].

The short side is constrained by availability, not fees. Markit medians by cap quintile run from 6.00%/yr indicative fee (smallest) to 0.54%/yr (largest), and the bottom two quintiles hold effectively zero lendable inventory [D5]. Real fees alone destroy the naive short book ($-0.02\%/mo$, $t=-0.11$, average 46 bps/month); requiring genuine inventory mechanically forces the book into large caps and rescues it, significant up to S\$500m AUM [D6]. The implementable long-short at S\$50m with real borrow: long $+0.392$ ($t=1.40$), short $+0.372$ ($t=2.72$), combined $+0.764\%/mo$ ($t=2.23$); the short leg stays the more reliable leg even paying observed costs [D7].

On the price-signal book the same structure appears: the borrow-feasibility screen cuts the long-short from 1.661 ($t=4.94$) to 0.712 ($t=2.53$), a 57% reduction, while the long-only top-30% book earns $+0.713\%/mo$ ($t=4.19$, Sharpe 0.85) with no borrow at all; long-only dominates there [D8, D9]. Execution lag is negligible (long-short t 3.74 to 3.59 trading at the next close) [D11].

Capacity: the STI-core top-6 book has a practical capacity of about S\$50M at a 10% participation cap, bound by position feasibility rather than cost (net alpha is still $+0.28\%/mo$ at S\$1bn notional, but 5% of months cannot be traded at S\$50M and 43% at S\$200M) [D12]. The liquid100 long-short book's binding-name capacity without a liquidity constraint at construction is degenerate, of order SGD 0.01 to 0.1m per month [D13].

7. Factor-level anatomy

Value carries the composite, part of it as a REIT effect. Dividend yield and book-to-market are the strongest single signals on the survivor-30 panel, and sector-neutralising roughly halves their ICs (DY t 4.54 to 2.63, BM 3.35 to 2.07, momentum 2.87 to 1.55); on the PIT roster the sector-neutralisation cost is nearer 10% [F1]. On the broad universe the composite IC is 0.049 ($t=5.93$) with DY at $t=12.3$ and B/M dead at $t=0.63$; the two universes disagree about B/M and the broad-universe estimate governs [C1]. Sector-neutral ranking kills the broad long-short outright, $+0.526$ ($t=1.93$) to $+0.169$ ($t=0.69$), re-executed this pass on the real panel; the accompanying Brinson-Fachler decomposition ($+0.51$ allocation, $+1.28$ selection, $+2.98$ interaction, sum $+4.78$ against $+5.12$ actual) refutes the older " $+8.4\%/yr$ selection" figure for this track [C7]. The same re-execution puts the broad long-short's one-way break-even cost at 86 bp against a recomputed real median half-spread of 23.5 bp over 143,978 name-months, and a 12% vol target on the point-in-time book lifts alpha to $+0.663$ ($t=3.08$), a post-hoc specification reported as such [C7, D14].

Quality destroys value added. In the leave-one-group-out attribution, quality contributes -50.2 bps/mo; dropping ROE lifts the top-100 book from $+0.468$ ($t=1.67$) to $+0.97\%/mo$ ($t=4.71$) [E8]. Holdings-based attribution corroborates: the book carries its largest exposure to quality (0.985 z) for zero return ($t=0.26$), while momentum ($+0.427$, $t=2.61$) and value ($+0.257$, $t=3.86$) supply the entire active return and the stock-specific residual is negative ($-0.15\%/mo$): a factor-exposure vehicle, no selection skill [E11]. The ROE-drop improvement is post-hoc on a sample with no holdout and is a hypothesis, not a claim.

Momentum's IC does not convert to a tradeable spread, and the reason is decile shape.

Momentum's quintile payoff is U-shaped and ROE's is reverse-J, so rank IC overstates what a top-minus-bottom book harvests; with alignment matched, IC t-statistics and portfolio t-statistics fall into a 0.9 to 1.4 ratio across all fourteen tested cells: the gap was construction, not alignment [F8].

Alignment. Contemporaneous IC is structurally wrong-signed for price-denominated ratios (B/M contemporaneous $t=-1.72$ vs predictive $+3.35$; DY -1.05 vs $+4.54$) because price sits in the denominator; predictive $h=1$ with the embargo is the honest measurement, and the notebooks audit themselves with placebo and contemporaneous grids, including a reversal signal whose contemporaneous IC is exactly -1.0 by construction [F3, I6].

Concentration and regime dependence. The top-100 book's active return is concentrated: top name 27.9%, top five 69.9%, top ten 96.7% of cumulative active return; leave-out-top-5 drops it to 0.210 ($t=0.94$). It exists post-2020 only (2008-13 $t=0.13$; 2014-19 $t=0.43$; 2020-26 $t=3.52$) and in calm or bull months only (calm $+1.205$, $t=3.91$; turbulent -0.12) [E9].

IBES analyst revisions add nothing. Pre-registered orthogonality test: all four revision signals subsumed by momentum (REV_NET IC 0.022, $t=1.17$; marginal Fama-MacBeth $-0.179\%/mo$, $t=-0.90$; none clears $|t|=1.96$ standalone), with the mechanism confirmed in the correlation structure: revisions load on the momentum horizons and on nothing else the composite does not already own [E10].

Macro interactions are risk, not alpha. The 10Y-rate interaction explains contemporaneous returns (IC $t=5.02$) and predicts nothing ($t=-0.70$); dSORA is the only persistent macro shock ($AR(1)=0.76$) and the only interaction with a positive predictive t-statistic, and even that is insignificant ($t=1.12$). Macro terms belong in the risk model [F2].

Weighting is load-bearing. Cap-weighting kills the alpha everywhere it is verifiably tested: broad-universe equal-weight long-short $+0.53$ ($t=1.93$) against value-weight -0.08 ($t=-0.27$) [C4]; PIT roster value-weight 0.339 ($t=1.21$) against equal-weight 0.920 ($t=3.11$) [F6, F5]. Two cap traps are documented with executed evidence: an absolute 15% cap on a 6-name book is infeasible ($6 \times 15\% < 100\%$) and silently reverts to equal weight, manufacturing false robustness [F7], and a 5% cap on a 100-name book is equal weighting in disguise (max weight 5.4% vs EW 5.2%) [C4].

8. The reference results do not reproduce, and the benchmark matters

Table I of the reference paper fails to reproduce: zero of nine factors survive Bonferroni at $t=2.77$, best univariate $|t|$ 2.767 (book-to-market) [E2]. The published headline $t=2.82$ against the reproduced 2.26 traces to implied book breadth (a 7-to-8-name book against the stated top-6); all consolidated numbers use the top-6 construction at $t=2.26$ [E1].

A second inferential defect was found and corrected in this project's own work: a benchmark built from members present at $t+1$ credits the benchmark with entrant appreciation that no portfolio formed at t can earn. The measured contamination is $+0.395\%/mo$. The STI-core headline is unaffected (2.56 vs 2.52), but the broad-universe breadth-decay null of the earlier expansion drafts reverses under the investable benchmark: top-150 and top-200 books become significant ($t=3.30, 3.61$) and strengthen with breadth [E4]. Both the earlier null and its reversal are then subject to the grid discipline below.

The pre-registered specification grid formalises multiple testing project-wide: 65 tests on the development window (1995-01 to 2019-12), of which 24 clear $|NW t| > 1.96$ unadjusted, exactly one clears $|t| > 3.0$ (an index-universe equal-weight long-short at $t=3.674$, which also uniquely passes the bootstrap reality check), and zero survive BHY false-discovery control (minimum adjusted p 0.074) [G1, G2]. The parallel $M=55$ family on the price-signal cells reaches the same verdict net of costs: nothing survives any correction (0 of 22 net tests; DSR $p=0.001$) [G4]. The 2020+ holdout has no result file in the archive, consistent with it being untouched.

9. Distinct strategy modules

9.1 Time-series momentum (TSMOM)

Moskowitz-Ooi-Pedersen construction, sign of trailing k -month excess return scaled to 15% ex-ante vol, $|w|$ capped at 2, vol estimated on lagged data. On the survivorship-clean liquid100 universe at $k=3$ the diversified strategy earns Sharpe 0.974 ($t=4.65$, 2001-2026); the 30-name Bloomberg version's 0.876 is survivorship-inflated (its PIT counterpart is 0.565), and the pre-2001 window is a null, so the result is one strong window out of three [H1, H2]. Horizon structure is shifted forward relative to the US benchmark: predictability peaks at $h=3$ to 5 ($t=2.55$ at $h=3$) and fades by $h=12$ ($t=1.12$); $h=3$ is a searched cell and is priced as such (block-bootstrap $p=0.001$ across the searched grid) [H3]. Adding a one-month implementation lag halves the wide-universe Sharpe (0.727 to 0.510); the lagged number is the realistic one [H5].

A disclosure applies to this module's combination and attribution results: they execute from in-branch cached panels (the branch's own upstream builds), so their lineage is real but re-execution requires those caches [H8].

The module's finding is the combination reversal. Pre-borrow, the fundamental sleeve plus TSMOM improved Sharpe on a -0.14 sleeve correlation. On borrow-constrained, fully costed sleeves over the matched 2003+ window the correlation rises to $+0.333$ and the combination Sharpe (0.460, $t=1.71$) falls below TSMOM alone (0.543, $t=2.11$). The diversification premise dissolves under realistic constraints [H4]. Attribution explains the asymmetry: the fundamental sleeve's leakage waterfall runs 0.810 to $-0.013\%/mo$ and its P&L is concentrated (top ten names 83%), while TSMOM survives its own waterfall (0.773 to 0.576) and is breadth-robust [H5].

9.2 Price-level signals (52-week high, short-term reversal)

The 52-week-high and momentum signals subsume each other conditionally on share turnover, a mechanism identified after two corrected framings (breadth, then liquidity, then turnover; a placebo on 12 random same-size subsets confirms the screen rather than sample size drives it) [I2, F9]. In high-turnover, liquid names the 52WH spread is $+1.45\%/mo$ gross ($t=3.74$) and clears Bonferroni and the

bootstrap reality check; the subsumption residual over momentum clears only BHY/HLZ. In low-turnover names momentum subsumes the 52-week high with a residual $t=5.67$ that survives every gross correction and breakpoint variation. Net of Corwin-Schultz spreads nothing in the family survives any correction (0 of 22) [I3, I4, I5]. Reported as a statistical finding about signal structure, not tradeable alpha. The jump-momentum conditioning result is not claimed: its K-sensitivity evidence is incomplete and what exists says it is fragile [I7].

9.3 Calendar seasonality

Month-of-year, Halloween and Heston-Sadka effects were tested on SGX and validated against long US series. No tradable seasonality exists on SGX: Halloween is economically large but insignificant ($t=1.87$), Heston-Sadka is statistically present on the wide universe and tradably absent, and turn-of-month was only establishable on US data. One statistical survivor is reported for completeness: the August effect ($-2.45\%/mo$, $t=-3.80$) clears Bonferroni but does not constitute a strategy [J1].

9.4 Volatility management and tail shape

The Moreira-Muir overlay's benefit is skew-conditional: real ex post across 208 JKP series in 8 markets (slope of Sharpe improvement on skew -0.0186 , $t=-3.95$, with country and theme fixed effects; improvement $+0.127$ for skew below -2 vs -0.028 above $+1$) [K1], not a volatility confound (horse race skew $t=-3.77$ vs vol $t=+0.25$) [K2], and unexploitable ex ante because skewness is near-unforecastable (predictive R-squared 0.003 at 12 months against 0.247 for volatility, $n=98,932$) [K3]. Searched ex-ante gates fail White's reality check ($p=0.576$) [K4].

9.5 Regime and jump-risk measurement: a defensible null for alpha, real for risk

On the STI-core daily return series, Merton jump-diffusion fits decisively: 57.9% of return variance is jump risk, roughly 58 implied jumps per year at 1.86% jump volatility, KS statistic 0.016 against 0.072 Gaussian. At the 99% level the Gaussian benchmark understates Expected Shortfall by 1.60 pp (-3.13% against the empirical -4.73% , a 34% relative understatement); the jump model is off by 0.30 pp. Gaussian-calibrated limits therefore permit roughly 1.5x the intended tail exposure [L1]. As a conditional forecaster, however, Merton is superseded: GJR-GARCH with skew-t innovations beats it by 713 BIC points and is the only model passing conditional VaR backtests at both the 95% and 99% levels; volatility clustering matters more than jump structure for risk forecasting [L2]. These results license claims about position sizing and risk limits only; Merton jump arrivals are homogeneous Poisson and carry no timing information.

Overlays do not convert any of this into return. Eight overlays were tested and eight failed to add risk-adjusted return: Barra-style risk model, HMM, gradient-boosted trees, the switch-penalised statistical jump model (its best configuration never switches), IBES revisions, jump-aware ES sizing (timing worth Sharpe $+0.05$ with a CI excluding zero, below the pre-registered $+0.10$ bar), GARCH vol targeting, and Markov switching (filtered probabilities actively hurt) [E21, L5]. Mean-variance optimisation is the only overlay with a genuine benefit, halved drawdown at equal Sharpe and no alpha. Two estimation traps were caught and quantified: a GARCH look-ahead (month-end forecast applied to the same month) inflated Sharpe from 0.73 to 1.05 before correction, and Markov smoothed probabilities manufacture a $+0.525$ Sharpe illusion over filtered ones (correlation between the two signals only 0.21) [L3, L4]. Limitation carried forward: the overlay family was never benchmarked against naive inverse-realised-volatility scaling, so the null is a null against buy-and-hold-style baselines only [L6].

10. The STI-core strategy and its risk characterisation

The STI-core primary (top-6 by composite on the PIT roster, exit-classified delisting, 4-month lag) earns +0.396%/mo, NW-t 2.26, 10.1% annualised [E1]. It reproduces on a fully Bloomberg-independent Compustat stack at +0.36%/mo ($t=2.26$) with quarterly-hybrid fundamentals (+0.27, $t=1.82$ with annual-only, the attenuation being a data-frequency artifact) [E3]. Distribution-free inference corroborates the parametric result (block-bootstrap $P(\alpha>0)=98.7\%$, stream-verified) and the -50.3% max drawdown is a typical path for the strategy, not bad luck (bootstrap median -46%) [E18]. The bootstrap confidence interval and the data-snooping probability quoted in earlier drafts survive only in figure legends whose .npy inputs are not in the archive, and are excluded pending recomputation [E18, G5 notes].

The corrected cross-section behind it is monotone: deciles run from $-0.518\%/mo$ ($t=-3.08$) at the bottom to $+0.476$ at the top, D10–D1 $+0.995\%/mo$ ($t=2.74$); the earlier U-shape was a forward-return bug [E5]. The spread splits 47/53 between legs with the short leg far more reliable (long $+0.492$, $t=1.75$; short $+0.547$, $t=3.30$) [E6], which is why borrow is the binding constraint and why the Markit-screened long-short (section 6) is the preferred implementable form. The alpha is not market beta (CAPM alpha $+0.45\%/mo$, $t=2.59$, 25% idiosyncratic variance) and is fully explained by its designed style exposures (multi-factor alpha $+0.16$, $t=1.17$), compensation for value and momentum style risk rather than an anomaly [E15]. It survives institutional constraints, modestly improving under sector caps and beta bands (33% sector cap: $+0.489$, $t=2.51$) [E16], which matters because the unconstrained book is majority Financials/REITs in 61% of months, peaking at 83% [E17]. Full equity tail risk applies: daily 99% VaR/ES $-3.14\%/-4.73\%$, no crash protection, mildly defensive only in rotational selloffs [E14].

11. The 30-name diagnostic track, corrected

The original three-notebook pipeline contained a silent synthetic fallback: on a missing workbook or column, a seeded generator fabricated the panel (and a ticker-hash fallback fabricated every fundamental column of the surviving export), with only a console print. Every downstream output of that chain is excluded, and the shipped code fails loudly on missing data [M2]. The corrected re-run on the certified real fields (prices only): no signal reaches $|t|>1.96$ (best momentum $t=1.38$); the long-only book trails buy-and-hold even at zero cost, a signal problem, not an execution problem; the long-short leg is a genuine diversifier (correlation -0.11 , beta -0.13 , $+1.10\%/mo$ in benchmark-down months); and a 20% long-short overlay is the efficient combination point (Sharpe 0.326 vs 0.316, drawdown -47% vs -56%). The overlay is in-sample and inside noise, and is carried only as a hypothesis to retest on the PIT universe with borrow costs [M1].

For the record, the survivorship figures previously published on the project site (Sharpe 0.98 to 0.61 long-only, 0.64 to 0.25 long-short, composite IC 0.075 at $t=5.05$) trace to a superseded first-pass survivorship run whose book-to-market factor was backfilled from a contaminated file, and to the survivor-30 IC re-run. The verified replacements are the canonical blue-chip pair (Sharpe 1.002 to 0.447; long-short 0.586 to 0.313) and the 71-name PIT IC table (0.0697, $t=5.90$) [O1, O2].

12. Corrections log

Every material error found and fixed, in the open, because the corrections are part of the contribution. Each entry is traced in the ledger [N1-N13].

1. Silent synthetic-data fallback in the original pipeline; all downstream results rejected and rebuilt; the generator was reproduced exactly (5 of 5 tickers) to prove the contamination.
2. Forward-return alignment bug (groupby shift(-1) instead of explicit PIT lookup) manufactured a false U-shaped payoff and a flat long-short spread; corrected to a monotone payoff and +0.995%/mo (t=2.74).
3. GARCH look-ahead (same-month forecast) inflated Sharpe 0.73 to 1.05; corrected to a null.
4. Markov smoothed-probability illusion, +0.525 Sharpe over filtered; filtered used.
5. Amihud rolling window included month t; lagging it moved the fundamental sleeve Sharpe 0.606 to 0.563.
6. Turnover measured on unlevered weights against levered returns produced a spurious 1,400%/month figure; corrected.
7. Corwin-Schultz standard recipe inverts the liquidity ordering on illiquid names; traded-days correction applied before use.
8. Double application of the membership filter reintroduced survivorship; fixed and verified by reconciliation to 5.6e-17.
9. Returns and momentum originally computed from PX_LAST, discarding 4.55 pp/yr in a high-yield market; total returns adopted.
10. Published t=2.82 reconciled to 2.26 via implied book breadth; manuscripts corrected (residual stale copies listed in EXCLUSIONS.md).
11. Benchmark contamination: t+1-membership benchmark inflated the passive alternative by +0.395%/mo; investable-cohort benchmark adopted, reversing the earlier breadth-decay null.
12. Absolute position caps on small books are infeasible and silently revert to equal weight (the 15% cap on a 6-name book); per-name cap corrected to cap/n.
13. STI churn misquote (51/51) corrected to the executed 30 additions, 30 deletions.

13. Limitations and data-gated work

The blue-chip survivorship alpha regressions (the t=4.60 to 1.19 pair of earlier drafts) could not be re-executed because the final-spec workbook and monthly series are not in the archive; the finding is carried at Sharpe/return level only. The delisting-bias magnitude awaits terminal-return adjudication of 115 code-10 exits (worksheets and taxonomy are ready; SDC covers 61 of them). Value, quality, size and liquidity on the 30-name panel await a clean Bloomberg re-pull. Registered but data-gated: SGX/HKEX cross-listing arbitrage, the STI futures-cash basis study, the true STI index-effect study, and the REALIS ownership crosswalk (the export lacks a developer field; the statistical alternative found one suggestive differential, property-volume growth t=-2.61, which fails the project's own Bonferroni bar and is reported as suggestive only). The WRDS S&P-BMI track's results exist only as external-session write-ups and are excluded pending artifact recovery. The regime-overlay null still needs the inverse-realised-vol benchmark. Sample end is 2026-06/07 across panels.

14. Reproducibility

The repository ships executable code, schema documentation, and labelled synthetic replicas so the pipeline runs end to end for cloners; no licensed vendor data ships. Every licensed input is documented in DATA.md with fields, windows and entitlement route. The fabricated legacy CSVs still resident in the

project workspace share filenames with clean re-runs; consumers must check provenance by content, and the quarantine action (deleting the legacy blobs) is flagged for the owner.