

Do Published Equity Anomalies Survive Clean Data?*

Nine US-equity papers replicated on survivorship-free data by an AI research agent

Vladimir Rodeski[†]

Radovan Vojtko[‡]

Abstract

We replicate nine published US-equity anomaly papers on survivorship-free, point-in-time data covering roughly 61,000 active and delisted tickers from 1998 to 2025, using an AI research agent (Claude Opus 4.7 and 4.8, Anthropic). The in-sample window is 2000–22; the out-of-sample window 2023–25 is fixed and pre-registered before any result is read. Each paper runs through a two-pass fidelity review: a code-fidelity pass checks the generated backtest against the extracted paper specification; an execution-fidelity pass reads the trade log against the paper’s stated behaviour. We apply the Harvey, Liu and Zhu (2016) $t > 3$ multiple-testing hurdle on full-sample alphas, value-weighting and a micro-cap screen following Hou, Xue and Zhang (2020), and realistic turnover-aware cost stressing. Three of nine papers show a statistically significant full-sample alpha; only Bali, Cakici and Whitelaw (2011) clears $t > 3$; on paper-faithful builds, none survives out-of-sample. The study’s lone apparent survivor, Zhu, Sun and Stivers (2021) at OOS Sharpe 1.95 on our first build, was a construction artifact the authors caught once the gated source was obtained: built from notes while the paper was paywalled, it diverged from the paper’s design (a corner-sort rather than the specified reversal-within-quintile), and on a faithful rebuild it inverts to -0.95. A parallel look-ahead demonstration shows that sorting on retroactively split-adjusted rather than as-traded prices manufactures approximately 3.5 Sharpe points of illusory out-of-sample edge from a single price field. The fidelity discipline is the methodology result; the null on anomalies on clean data is the empirical result.

1. Executive summary

Each paper carries one verdict: does the documented effect reproduce on clean, survivorship-free data? Reported as the direction it reproduces (or doesn’t), with caveats. This is the scientific question; it is not a trading recommendation.

The headline the benchmark column forces: SPY returned 22.9%/yr at a Sharpe net of cash of 1.43 over 2023–25 (the mega-cap-led AI rally). Among the nine, every long-only strategy posted *negative* excess return versus SPY (−9% to −27%/yr), and none matched its Sharpe (see the scorecard below). A faithful replication of a documented effect is not the same as an effect that pays out-of-sample, and over this window, on faithful builds, none of the nine shows a tradeable out-of-sample edge. A per-paper reproduction summary is in §6.

Consolidated scorecard

All returns net of 5 bps/side. Sharpe is annualized arithmetic: long-only in excess of cash, long/short as a zero-cost spread. α is the annualized CAPM alpha versus SPY (the part of the return that is *not* market exposure), with its t-statistic. “Overlap” is how much of the paper’s original sample our data window covers: a short overlap means our reproduction is a small-sample read on a sub-period the paper didn’t emphasize. The verdict vocabulary is

*AI disclosure: the research pipeline (paper extraction, backtest construction, the two-pass fidelity review, and the first draft) was executed by Claude Opus 4.7 and 4.8 (Anthropic). The paper and source acquisition, research design, and pre-registration of IS/OOS windows and cost assumptions were the authors’; the authors edited and approved the final text from the agent’s draft. The lone apparent survivor (Zhu) was caught once the authors obtained the gated source and the build was graded against the paper.

[†]founder, PolyQuant AI, vlad@polyquantai.com

[‡]CEO & Head of Research, Quantpedia, info@quantpedia.com

fixed at four values (**Reproduces** / **Reproduces (caveated)** / **Does not reproduce** / **Untestable**); per-paper nuance is in §5.

Table A: Performance & significance

Paper	Type	Reproduces?	IS Sharpe	OOS Sharpe	β	OOS α (t)	Full α (t)
RSI trend	LO	Reproduces (caveated)	0.55	0.66	0.99	-7.4% (-1.5)	+2.3% (1.7)
Lottery (MAX)	LO	Reproduces (caveated)	0.76	0.30	0.65	-7.7% (-1.7)	+4.6% (3.4)
Low-beta (BAB)	LO	Reproduces (caveated)	0.73	0.27	0.26	-1.6% (-0.3)	+4.4% (2.5)
Falling knives	LO	Reproduces (caveated)	0.46	-0.23	1.05	-24.2% (-1.7)	+7.0% (0.9)
Market-state mom.	L/S	Reproduces (caveated)	0.33	0.12	0.13	-1.2% (-0.2)	+3.4% (1.4)
52-wk-high reversal	L/S	Does not reproduce	0.13	-0.95	~0	-29% (-1.2)	~0% (-0.2)
Persistent mom.	L/S	Does not reproduce	0.26	0.17	0.09	-0.3% (-0.1)	+5.2% (2.4)
Seasonality	L/S	Untestable	0.52	0.68	0.05	+5.3% (0.9)	+4.4% (1.0)

Table B: Tradeability

Paper	Overlap	Full CAGR	OOS CAGR	OOS MaxDD	Turnover	Break-even
RSI trend	98%	10.5%	13.7%	-17.8%	~100%	~1,300 bps
Lottery (MAX)	16%	10.6%	7.5%	-8.4% ^m	1,200%†	~60 bps†
Low-beta (BAB)	15%	9.7%	7.0%	-6.2% ^m	350%†	~200 bps†
Falling knives	100%	11.5%	-4.2%	-33.8% ^m	1,000%†	n/a‡
Market-state mom.	26%	3.1%	0.6%	-15.6%	~790%	~8 bps
52-wk-high reversal	32%	-1.0%	-29.1%	-55.0%	2,200%†	n/a‡
Persistent mom.	38%	2.3%	1.0%	-11.6%	~1,050%	~10 bps
Seasonality	10%	4.3%	5.9%	-9.4% ^m	1,800%†	~35 bps†

*Full CAGR and Full α are whole-sample; OOS columns are 2023–25. α is the CAPM alpha vs the market, with its t-statistic. **Turnover** = annual one-way traded notional (exact where per-name holdings exist; † = estimated for the reconstructed papers). **Break-even** = the per-side cost that would zero the out-of-sample return (‡ = n/a; the strategy loses money before any costs). SPY out-of-sample: 22.9%/yr, Sharpe net of cash 1.43. Single-factor (CAPM) alphas do not decompose size/value/momentum; at 36 out-of-sample months every t-statistic is small, so read them as effect-direction, not proof. **OOS MaxDD** is from the daily equity curve; ^m marks a monthly-resampled figure that understates the true daily peak-to-trough by ~1.5–2.6× (daily recompute pending; read the ^m values as a lower bound on the true daily drawdown, not as final).*

The price-level look-ahead case (Geertsema & Lu) is excluded from the scorecard above because it is a methodological demonstration, not a tradeable strategy; its before/after numbers are in the box below.

Price-level look-ahead (Geertsema & Lu): a demonstration, not a tradeable strategy

Variant	IS Sharpe	OOS Sharpe	Full CAGR	OOS CAGR	OOS MaxDD	OOS α (t)	Full α (t)
As-traded price (correct)	-0.36	-2.32	-16.6%	-44.3%	-84.7% ^m	-80.7% (-5.1)	-12.5% (-2.0)

Variant	IS Sharpe	OOS Sharpe	Full CAGR	OOS CAGR	OOS MaxDD	OOS α (t)	Full α (t)
Split-adjusted price (look-ahead)	+1.91	+1.24	n/a	n/a	n/a	n/a	n/a

The as-traded row is the PIT-correct build; the split-adjusted row is the illusory look-ahead comparison. The gap between the two Sharpes is the bias, ~ 3.5 Sharpe points out-of-sample. Market exposure of the as-traded spread: β 0.42. This is a methodological warning, not a tradeable strategy; full detail in §5.9.

Sorted by full-sample alpha-t: **Lottery (MAX)** (CAPM $t = 3.4$ full-sample) is the only paper that clears the $t > 3$ multiple-testing hurdle, but its alpha has decayed to negative out-of-sample and lives in capacity-limited micro-caps. **Persistent momentum** and **Low-beta (BAB)** show significant full-sample alpha (t 2.4–2.5) that fades to nothing out-of-sample. **RSI trend**, **Market-state momentum**, **Falling knives** never clear $t > 2$ full-sample: Hill’s return is market beta ($\beta \approx 1.0$, alpha insignificant), Rodon’s mechanism is dormant, and Arendarski’s return is real but statistically indistinguishable from noise. **Seasonality** is untestable at full fidelity: its paper-faithful 20-year-lookback variant is data-limited on our 1998-start history (full depth only from ~ 2015), so its full-sample significance cannot be established. **52-week-high reversal** on a faithful build (see §5.2) carries no full-sample alpha ($t \approx -0.2$) and inverts out-of-sample (OOS Sharpe -0.95). **Price-level look-ahead** is significantly *negative*, which is the paper’s own point. No strategy clears $t > 2$ out-of-sample, let alone $t > 3$.

2. Data, universe, and methodology

Data. A commercial Core US Equities dataset: daily split- and dividend-adjusted prices plus as-traded prices, quarterly point-in-time fundamentals, and a point-in-time S&P 500 membership history. Coverage 1998–2025, $\sim 30,000$ active and $\sim 31,000$ delisted tickers. **Because delisted names are retained, the universe is survivorship-bias-free by construction:** a strategy is only ever allowed to hold names that were actually investable on the formation date.

Universe, per paper. Each paper is run on the universe its mechanism requires: the broad cross-section ($\sim 3,000$ – $3,400$ names/month, with a \$5 price floor and a one-year listing minimum) for breadth-dependent signals; point-in-time S&P 500 members for Hill; rolling top-500 by market cap for the momentum/low-beta papers. Universe construction is annual and point-in-time; data gaps are logged and skipped, never substituted (substitution silently reintroduces survivorship bias, since the replacement name was chosen *because* it survived).

Tradeability screen. Every broad-universe strategy is filtered for tradeable liquidity at formation, a trailing dollar-volume floor plus a stale-price guard, so results reflect names an investor could actually trade rather than data artifacts (see the methodology highlight in §3).

Costs. Paper-comparison results use a flat 5 bps per side. This is conservative for large-caps and optimistic for micro-caps, so every candidate is *additionally* stress-tested at realistic costs keyed to its turnover and universe liquidity (commission + half-spread + borrow for short legs), and we report each strategy’s break-even cost. Because every headline result is a null or a decay, this 5-bps baseline is conservative for the conclusion: more realistic costs can only push a non-surviving strategy further from viability, never rescue it.

Sharpe convention. Annualized arithmetic. Long-only Sharpes are computed in excess of the risk-free rate ($\approx 1.7\%/yr$ in-sample, $\approx 4.7\%/yr$ out-of-sample, material in the current rate regime); long/short spreads are zero-cost and reported without a cash deduction.

Factor decomposition. For each strategy we regress returns on the market excess return (CAPM) and report alpha, its t-statistic, and beta. This separates genuine edge (alpha) from market exposure (beta): a strategy with $\beta \approx 1$ and zero alpha is simply the index. We use single-factor CAPM for the cross-paper scorecard. A multi-factor decomposition of any candidates carried forward is the natural next refinement.

3. Research protocol

Five commitments make a replication study credible. They matter most when the honest answer is “it didn’t hold.”

- **In-sample / out-of-sample lock.** The 2000–22 / 2023–25 split was fixed before any result was read. No variant is reported that was added after seeing the out-of-sample.
- **Paper-window overlap is disclosed.** Several papers’ original samples predate our 1998 data start; where overlap is short (Bali 16%, Frazzini 15%, Heston ~10%), our reproduction is a small-sample read on a sub-period, and we flag it.
- **Multiple-testing discipline.** Across nine papers we evaluated the full set of correlated cells (paper $\times$ bucket $\times$ window). The family-wise-significant threshold at that breadth is far above anything observed out-of-sample, so we report no out-of-sample result as significant and apply the literature’s $t > 3$ hurdle (Harvey, Liu and Zhu 2016) when judging full-sample alphas: the search-adjusted significance bar. Only Bali clears it, and only full-sample; no strategy clears it out-of-sample.
- **Value-weighting and a micro-cap screen** are applied to broad-universe sorts, following the Hou, Xue and Zhang (2020) replication standard: equal-weighted micro-cap sorts are the single largest source of overstated anomalies.
- **Realistic-cost and capacity reporting.** Headline Sharpes are paired with turnover, break-even cost, and capacity notes; the lone candidate is fully cost-stressed.

Methodology highlight: catching data artifacts. Broad-universe replications are acutely vulnerable to single bad data points. In our first Bali (lottery) reconstruction, one halted micro-cap, its price frozen for weeks then a corrupt print, single-handedly produced a +76% basket month and inflated the 27-year curve from $\sim 16\times$ to $\sim 38\times$. The lottery sort actively *selects* such names: a frozen price registers as the “calmest” stock and lands in the long bucket. A point-in-time liquidity-and-stale-price screen removes them, and value-weighting (the paper’s own convention) suppresses what remains. After the fix the artifact month fell to +0.1% and the verdict was unchanged but the magnitude was honest. We apply this screen to every broad-universe strategy and scan all return series for single-month outliers before trusting a curve. This is the concrete form of the survivorship-and-data-integrity discipline that motivated the point-in-time design.

The research agent and its gates

The replication was executed by an AI research agent (Claude Opus 4.7 and 4.8, Anthropic): for each paper the agent extracted a structured construction specification from the source text, generated and ran the backtest, and drafted the per-paper write-up. Agent output is only as trustworthy as the checks it must pass, so the pipeline is built around a fixed set of gates. Each one caught at least one error in this study that an unchecked run would have reported as a result.

Gate	What it checks	What it caught here
Specification extraction	Paper text $\rightarrow$ an explicit construction spec before any code is written	Quality is bounded by source access: Zhu’s spec was built from notes while its PDF was gated, the gap the later review exposed
Code-fidelity review	Generated backtest against the extracted spec	Mis-built sorts and weighting schemes flagged for rebuild
Execution-fidelity review	Realized trade log against the paper’s stated behaviour, once the source is in hand	Zhu’s corner-sort: a clean OOS Sharpe of 1.95 that, rebuilt to the paper’s actual design, inverted to -0.95
Data-integrity scan	Single-month return outliers plus a point-in-time liquidity / stale-price screen	Bali: one halted micro-cap that inflated the 27-year curve from $\sim 16\times$ to $\sim 38\times$
Point-in-time discipline	Level-based signals sorted on as-traded, not retroactively split-adjusted, prices	Quantified at ~ 3.5 out-of-sample Sharpe points of illusory edge (Geertsema-Lu), and applied to every level signal
IS/OOS lock	Out-of-sample window fixed and pre-registered before any result is read	No post-hoc variant reported on the out-of-sample

Gate	What it checks	What it caught here
Independent-engine validation	Four candidates (Hill, Zhu, Chen, Rodon) re-run on a second, independent backtest engine	Agreement within a small Sharpe band treats a result as not an implementation artifact; the other five papers are single-engine

The telling case is what the gates did *not* catch on their own. The Zhu pipeline ran without error and produced a confident, plausible out-of-sample survivor; the divergence from the paper surfaced only once the source PDF (initially inaccessible) was obtained and the build was graded against it. The lesson generalizes: an agent can extract, build, and self-check faithfully and still produce a result that is not the paper’s, because faithfulness to a spec is not faithfulness to a source the agent could not read. Source acquisition, the research design and pre-registration of windows and cost assumptions, the choice to declare rather than absorb each construction deviation, and the final text remained human judgments. The gates make the agent’s output auditable; they do not remove the human’s role in supplying the source and owning the design.

4. Definitions

Term	Plain definition
MAX (lottery) signal	A stock’s single largest daily return over the prior month. “Low-MAX” names have had no recent lottery-like spike.
PTH (price-to-52-week-high)	A stock’s current price divided by its 52-week high. “Low-PTH” names trade far below their 52-week high.
As-traded vs split-adjusted price	The price actually quoted on the day vs a price retroactively rescaled for later stock splits. Using the latter as a <i>signal</i> peeks at the future, the crux of the Geertsema-Lu paper.
Long/short, gross-1.0 sleeve	A market-neutral book holding 50% long + 50% short, so total exposure is $1\times$ capital.

5. Results by paper

Each paper below follows the same template: the paper’s claim, the mechanism in plain terms, our results, the equity curve, and what it means. All nine are treated at equal depth.

5.1 RSI trend-following — Hill (2019)

The paper’s claim. Hill (2019) argues that RSI, usually read as a mean-reversion oscillator, works as a *trend-following* signal on large-cap stocks: names making a fresh RSI high tend to keep rising. The paper reports trade-level statistics (a 2.11 profit-to-loss ratio, ~58% hit rate) over Jul 1998 – Jun 2018 and deliberately builds no portfolio.

Mechanism, plainly. Each month, hold an equal-weight basket of every S&P 500 member whose RSI(14) pushed above 70 within the prior ~6 months (“bull momentum”), point-in-time on index membership, rebalanced monthly. The bet: a recent momentum high predicts continuation. (RSI(14) is the 14-day Wilder momentum oscillator, 0–100; a reading above 70 is conventionally “overbought.”)

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22 (98% overlap)	0.55	10.1%	–60.4%	+3.6% (2.5)

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
Out-of-sample 2023–25	0.66	13.7%	−17.8%	−7.4% (−1.5)

SPY (out-of-sample): 22.9% CAGR · Sharpe 1.43 · MaxDD −18.8%. Full-sample CAPM α : +2.3% (t 1.7), in-sample alpha that reversed out-of-sample.




Figure 1. Growth of \$1, S&P 500 RSI bull-momentum basket vs SPY total-return buy & hold; out-of-sample shaded. The basket tracks SPY closely; its return is market beta, not alpha.

What it means. The signal reproduces the paper faithfully (the trade-level statistics match within ~2%) and it compounds well over the full quarter-century (\$1 → \$13.3 vs SPY’s \$7.8). But it has no demonstrable alpha: its market beta is ≈ 1.0 , its full-sample alpha is insignificant ($t = 1.7$), and out-of-sample it lost to SPY on return and Sharpe. Its 13.7% return is market beta minus a small drag. **Reproduces (caveated):** capacity-robust (large-cap, low-turnover), but its out-of-sample return is market beta, not alpha.

5.2 52-week-high reversal — Zhu, Sun & Stivers (2021)

The paper’s claim. Zhu, Sun & Stivers (2021) show that a one-month reversal signal applied *within* the low-52-week-high quintile earns risk-adjusted returns. The mechanism is conditional: the 52-week-high proximity is a *screen* that selects the universe, and the reversal is the signal within it. Both legs are stocks far from their 52-week highs, long the recent losers among them, short the recent winners. *(The paper PDF was inaccessible when we built this strategy; the construction was from plan notes. We obtained the paper, “Price Anchors and Short-Term Reversals,” SSRN 3092325, after running our initial build.)*

What we built, and what the paper specifies.

Our first build was a blended corner-sort: long stocks *near* their 52-week highs that had pulled back, short stocks *far* from their highs that had popped, equal-weighted across both legs. That produced an OOS Sharpe of 1.95 and was the apparent survivor of this study.

The paper specifies reversal-within-low-PTH: both legs are in the low-52-week-high quintile; the sort is on one-month raw return; the implementation is value-weighted. Our first build’s long leg sat in the *high*-52-week-high names, exactly the names the paper shows have no reversal tendency, compounded by equal-weighting in the micro-cap-heavy spread. This is a neutral mechanism note on why the first build diverged, not a defect we are scoring against the paper.

Results: faithful rebuild (VW, reversal within low-PTH quintile).

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22	0.13	3.5%	−67.1%	~0% (−0.2)
Out-of-sample 2023–25	−0.95	−29.1%	−55.0%	~29% (−1.2)

First-build corner-sort (for comparison): IS 0.70 / OOS 1.95, a mis-built corner-sort with the wrong quintile and equal-weighting, not the paper’s strategy. Market beta ≈ 0 (low-PTH reversal spread is approximately market-neutral). Equal-weighted variant: IS 0.67 / OOS −2.12.

What it means. The 1.95 OOS Sharpe was never the paper’s strategy. On a faithful build, the strategy carries no detectable alpha in-sample ($t \approx -0.2$) and is significantly negative out-of-sample. The divergence has two sources: (1) equal-weighting inflated the Sharpe in the micro-cap-heavy spread (EW IS 0.67 vs VW 0.13); (2) the first build’s long leg sat in the wrong quintile, stocks near their highs with no reversal tendency. The OOS loss on the faithful build is driven by the short leg: the low-PTH quintile shorts recent winners who kept rising through the 2023–25 rally. One limit: our data starts 1998, so this is a paper-faithful build on a different window than the paper’s 1967–2015 sample. **Does not reproduce:** on a paper-faithful build the effect inverts (OOS Sharpe −0.95); the published sample is a different regime this build cannot test.

5.3 Lottery / MAX — Bali, Cakici & Whitelaw (2011)

The paper’s claim. Bali et al. (2011) show that stocks with extreme recent single-day returns (“lottery” stocks) subsequently underperform; the low-MAX, value-weighted side is the tradeable expression. Original sample Jul 1962 – Dec 2005.

Mechanism, plainly. Each month, sort the broad cross-section by each stock’s largest daily return over the prior month; hold the low-MAX names (value-weighted), rebalanced monthly.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22 (16% overlap)	0.76	11.0%	−34.7% ^m	+6.2% (4.4)
Out-of-sample 2023–25	0.30	7.5%	−8.4% ^m	−7.7% (−1.7)

SPY (out-of-sample): 22.9% CAGR · Sharpe 1.43 · MaxDD −18.8%. Full-sample CAPM α : +4.6% (t 3.4), significant historically, decayed out-of-sample.




Figure 2. Growth of \$1, value-weighted low-MAX basket (liquidity-screened) vs SPY; out-of-sample shaded. A real historical edge (full-sample alpha $t = 3.4$) that decays to negative out-of-sample.

What it means. A real historical effect, full-sample alpha +4.6% clears $t > 3$, that has decayed to negative out-of-sample, and whose remaining edge lives in capacity-limited small-caps where realistic costs (higher spreads, similar high turnover) would erode it. This is also the paper whose first reconstruction exposed the data artifact in §3; the corrected, value-weighted, liquidity-screened result above is right in magnitude. **Reproduces (caveated):** a real historical effect, decayed to negative out-of-sample and concentrated in capacity-limited small-caps.

5.4 Betting-against-beta / low-beta — Frazzini & Pedersen (2014)

The paper’s claim. Frazzini and Pedersen (2014) show that low-beta stocks earn higher risk-adjusted returns than high-beta stocks; the BAB factor (long low-beta levered up, short high-beta levered down) earns a Sharpe of ~ 0.78 in the US. Original sample 1926–2012.

Mechanism, plainly. Sort on market beta; hold the lowest-beta names, top-500 to keep it institutional-grade. We trade the low-beta long leg rather than the paper’s full levered long/short factor out of necessity: on our data the levered construction is degenerate (153 of 276 months contain near-zero-beta names that lever the long leg toward near-infinite exposure), so the long-only low-beta sleeve is the tradeable paper-faithful form.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22 (15% overlap)	0.73	10.0%	−40.3% ^m	+5.6% (3.1)
Out-of-sample 2023–25	0.27	7.0%	−6.2% ^m	−1.6% (−0.3)

SPY (out-of-sample): 22.9% CAGR · Sharpe 1.43 · MaxDD −18.8%. Full-sample CAPM α : +4.4% (t 2.5), significant historically, dead out-of-sample.



Figure 3. Growth of \$1, top-500 low-beta long leg vs SPY; out-of-sample shaded. Significant full-sample alpha ($t = 2.5$) that is dead out-of-sample in a regime that rewarded high-beta mega-caps.

What it means. A significant full-sample alpha ($t = 2.5$) that is dead out-of-sample, unsurprising in a regime that rewarded exactly the high-beta, mega-cap-tech names a low-beta strategy avoids. This is a loose, “BAB-shaped” derivative, not BAB-literal: we use OLS beta (not the paper’s $\sigma \times \rho$ decomposition), no Vasicek shrinkage, a decile sort (not the paper’s median split), equal-weighting (not rank-weighting), and the long leg only. Read the in-sample alpha as directional support for the low-beta effect, not as a reproduction of the BAB factor’s 0.78 Sharpe. **Reproduces (caveated):** a paper-inspired derivative: a significant full-sample alpha that is dead out-of-sample in a regime hostile to low-beta.

5.5 Falling knives — Arendarski (2012)

The paper’s claim. Arendarski (2012) shows that deeply beaten-down stocks (down 50%+ vs the market over the prior ~2 years) that pass a solvency screen subsequently outperform sharply; the best variant (Altman $Z \geq 2.2$) returned +1,722% vs SPY’s +8% over 2001–2011 (the best debt/equity variant returned +901%, not the headline).

Mechanism, plainly. Screen for “falling knives” (down 50%+ vs SPY over 500 trading days), then within that set favour the financially healthiest by an independent solvency overlay; hold equal-weight, rebalanced monthly.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22	0.46	13.9%	−70.3% ^m	+11.1% (1.4)
Out-of-sample 2023–25	−0.23	−4.2%	−33.8% ^m	−24.2% (−1.7)

On the paper’s own 2001–11 window, the strategy reproduces the paper’s order of magnitude ($\sim 9 \times$ SPY). SPY (out-of-sample): 22.9% CAGR · Sharpe 1.43 · MaxDD −18.8%. Full-sample CAPM α : +7.0% but $t = 0.9$, economically large, statistically insignificant given the volatility.



Figure 4. Growth of \$1, falling-knives basket vs SPY; out-of-sample shaded. Faithfully reproduces the paper on its original window; statistically insignificant full-sample alpha and hostile to the 2023-25 regime.

What it means. A faithful replication: it reproduces the paper’s order-of-magnitude on the original window, yet its full-sample alpha is statistically insignificant ($t = 0.9$): the returns are real but so volatile and micro-cap-concentrated that they are indistinguishable from noise, and out-of-sample the strategy lost 4%/yr with a −34% drawdown. Deep-value falling-knives are hostile to a mega-cap-led regime. **Reproduces (caveated):** faithful to the paper on its own window, but the full-sample alpha is statistically insignificant and the effect is hostile to the out-of-sample regime.

5.6 Market-state momentum — Rodon Comas (2025)

The paper’s claim. Rodon Comas (2025) proposes a cross-sectional momentum signal that *reverses* in down-states of the market (long past losers, short past winners during stress), improving robustness. Reported gross Sharpe ~ 0.79 , net ~ 0.4 – 0.5 .

Mechanism, plainly. Rank on intermediate-horizon momentum; in market up-states hold winners-minus-losers, in down-states flip to losers-minus-winners, using overlapping cohorts. Market-neutral.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22	0.33	3.4%	−42.9%	+3.9% (1.5)
Out-of-sample 2023–25	0.12	0.6%	−15.6%	−1.2% (−0.2)

Market-neutral ($\beta \approx 0.1$). Full-sample CAPM α : +3.4% (t 1.4), not significant; the down-state mechanism never fired out-of-sample.




Figure 5. Growth of \$1, deployable top-500 market-state momentum vs SPY; out-of-sample shaded. The down-state reversal mechanism never fired out-of-sample; the strategy is effectively dormant momentum.

What it means. This is the most paper-faithful spike here (every deviation is explicitly declared), and it reproduces the paper’s direction in-sample. But its load-bearing mechanism, the down-state reversal, never fires out-of-sample (2023–25 had essentially no sustained down-state), so the out-of-sample is just dormant momentum at a near-zero, insignificant alpha. **Reproduces (caveated):** reproduces in-sample with declared deviations, but its load-bearing down-state mechanism never fired out-of-sample, so the window cannot test it.

5.7 Persistent momentum — Chen (2016)

The paper’s claim. Chen (2016) argues that “persistent” winners (top-decile in consecutive formation windows) continue to outperform persistent losers; the long/short spread earns ~1.26%/month ($t = 3.35$), though the paper’s own four-factor alpha is marginal.

Mechanism, plainly. Identify names that were top-decile in *both* of two consecutive formation windows (persistent winners) and bottom-decile in both (persistent losers); hold winners-minus-losers, market-neutral, top-500.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22	0.26	2.5%	−36.9%	+5.3% (2.2)
Out-of-sample 2023–25	0.17	1.0%	−11.6%	−0.3% (−0.1)

Market-neutral ($\beta \approx 0.1$). Full-sample CAPM α : +5.2% (t 2.4), significant historically, exactly zero out-of-sample.




Figure 6. Growth of \$1, persistent winners-minus-losers (top-500) vs SPY; out-of-sample shaded. Significant full-sample alpha ($t = 2.4$); exactly zero out-of-sample as the short leg was crushed by the mega-cap rally.

What it means. The reading is “reproduced historically, gone out-of-sample.” Full-sample the long/short spread carries a significant alpha ($t = 2.4$); out-of-sample it is exactly zero ($t = -0.1$): the short leg, in particular, was crushed by the mega-cap rally. The paper’s central conditioning variable, analyst-coverage dispersion (IBES), is not in our data, so a fuller test remains open; the concrete path is to substitute standardized unexpected volume (SUV), computable from price and volume alone, which the paper’s own Table 8 shows interacts with the effect (persistent-winners $\times$ SUV $+0.34\%/mo$, $t = 4.5$; persistent-losers $\times$ SUV $-0.76\%/mo$, $t = -5.9$). Two caveats on our construction: we proxy the paper’s “top-decile in two *consecutive* windows” with a count rule (≥ 4 of N months) and omit the paper’s one-month skip between formation and holding, neither of which rescues the unconditional spread, which does not survive 2023–25. **Does not reproduce:** the unconditional long/short claim is exactly zero out-of-sample; the conditional claim (on analyst-coverage dispersion, absent from our data) remains untested, with an SUV substitution as the open path.

5.8 Same-month seasonality — Heston & Sadka (2008)

The paper’s claim. Heston-Sadka (2008) document that a stock’s return in a given calendar month persists year over year, same-month winners keep winning that month, strongest at long (multi-year) lags. NYSE/AMEX only; original sample 1965–2002.

Mechanism, plainly. Each month, rank stocks by their average historical return *in that same calendar month* over a long lookback; hold a long/short spread on that ranking.

Results.

Window	Sharpe	CAGR	MaxDD	CAPM α (t)
In-sample 2000–22 (10% overlap)	0.52	4.1%	$-28.5\%^m$	$+4.2\%$ (2.4)
Out-of-sample 2023–25	0.68	5.9%	$-9.4\%^m$	$+5.3\%$ (0.9)

Market-neutral ($\beta \approx 0.05$). The paper-faithful variant is the long (20-year) lookback, but it is data-limited on our 1998-start history (full depth only from ~ 2015), so Heston’s full-sample significance is untestable rather than established. (A shorter 5-year-lookback variant gives full-sample $t \approx 1.0$ but is not the paper’s signal.)




Figure 7. Growth of \$1, broad long/short seasonality spread vs SPY; out-of-sample shaded. Qualitatively positive out-of-sample, but the paper-faithful 20-year lookback is data-limited on a 1998 start.

What it means. Qualitatively the seasonality persists. The paper’s signal requires multi-decade lags (up to ~20 years), and our 1998 data start leaves too little history for that paper-faithful long-lookback variant (full depth only from ~2015), so its full-sample significance cannot be established. Out-of-sample the spread is positive (though insignificant). It is best read as **Untestable** at full fidelity on this history rather than confirmed or refuted: the paper-faithful 20-year lookback is paper-faithful but data-limited on our 1998 start, and the shorter 5-year variant is not the paper’s signal.

5.9 Price-level look-ahead — Geertsema & Lu (2019)

The paper’s claim. The well-known “low-price anomaly” (cheap-nominal-price stocks outperform) is, the paper argues, largely a **look-ahead artifact**: it appears only when you sort on *retroactively split-adjusted* prices, which encode future splits. On the as-traded price an investor actually saw, the effect should vanish.

Mechanism, plainly. Run the same low-price-minus-high-price long/short sort two ways, once on as-traded prices (correct), once on split-adjusted prices (look-ahead), and compare.

Results.

Variant	IS Sharpe	OOS Sharpe	Reading
As-traded price (correct)	−0.36	−2.32	no edge, confirms the paper
Split-adjusted price (look-ahead)	+1.91	+1.24	the illusory anomaly




Figure 8. Growth of \$1 (log scale): the look-ahead-biased sort (split-adjusted prices) compounds to ~500x while the correct sort (as-traded prices) decays to ~0. The gap between the lines is approximately 3.5 Sharpe points of illusory out-of-sample edge from a single price field.

What it means. The clearest chart here: sorting on split-adjusted prices manufactures a ~3.5-Sharpe-point edge out-of-sample that does not exist on the prices an investor could trade, because a low split-adjusted price today silently flags a stock that *will* split (typically a past winner). **Does not reproduce** as a tradeable anomaly: the look-ahead thesis is confirmed, a methodological warning, not a strategy. It is also the reason this entire study sorts level-based signals on as-traded prices throughout.

6. Reproduction summary

By verdict (full per-paper detail in §5, scorecard in §1):

- **Reproduces (caveated), 5:** RSI trend, Lottery (MAX), Low-beta (BAB), Falling knives, Market-state momentum. Each reproduces the documented effect in-sample but is undeployable out-of-sample: decayed alpha, capacity-limited micro-caps, market beta without alpha, or a mechanism that never fired.
- **Does not reproduce, 3:** 52-week-high reversal (faithful rebuild inverts to OOS -0.95), Persistent momentum (unconditional spread is zero out-of-sample), Price-level look-ahead (thesis confirmed, a negative result by design).
- **Untestable, 1:** Seasonality (paper-faithful 20-year lookback is data-limited on a 1998 start).

7. Cross-paper observations

- **Out-of-sample decay is the rule.** Three of nine effects were statistically real full-sample; all three faded to insignificance over the fixed 2023–25 out-of-sample window. This mirrors McLean and Pontiff (2016), though we measure decay against a common out-of-sample date, not from each paper’s publication year, and is the central calibration result: a published equity anomaly is more likely than not to be decayed or undeployable by the time you trade it.
- **Nothing beat passive in 2023–25** (see §1). The out-of-sample regime, a narrow mega-cap rally, was hostile to every value, low-beta, lottery, and short-leg strategy and kind to passive cap-weighted beta. The benchmark, not the signals, dominated.

- **Equal-weighting and micro-caps inflate anomalies.** The data-artifact in §3 and the micro-cap concentration of the “best” raw returns (Bali, Arendarski) both point the same way, consistent with Hou, Xue and Zhang (2020): value-weight and screen for tradeability or you are measuring frictions, not edge.
 - **Verifying the build against the source paper is a gate, not an option.** Zhu is the sharpest case: the pipeline produced an OOS Sharpe of 1.95 on a build that diverged from what the paper described; only checking the build against the source paper (once the gated PDF was obtained) caught it. A pipeline that cannot check its build against the source text will surface results that are not the paper’s.
 - **Look-ahead bias is quantifiable and large.** Geertsema-Lu puts a number on it: ~3.5 Sharpe points out-of-sample of pure illusion from one price field. Point-in-time as-traded data flips the sign of the result, not just its size.
-

8. Conclusions

Nine published US-equity anomalies, replicated on survivorship-free, point-in-time data with factor decomposition, realistic-cost stressing, and multiple-testing discipline. Two results stand.

- **The base rate.** 3 of 9 show a statistically significant full-sample alpha; only Bali clears the $t > 3$ hurdle; on paper-faithful builds, none survives out-of-sample. The lone apparent survivor (Zhu, OOS Sharpe 1.95) was a construction error the fidelity discipline caught; rebuilt faithfully it inverts to -0.95 . On clean data, with realistic costs and a real multiple-testing hurdle, published equity anomalies mostly do not survive to be tradeable, and the rare one that does is gated by microstructure, not by signal.
 - **The regime caveat (a business-cycle caveat).** The 2023–25 out-of-sample is a single, passive-favouring regime, a ~2-year AI-led mega-cap bull, not a full business cycle. Equity factor strategies need a complete cycle, a bear *and* a bull, to judge fairly, so reading a fade over this window as a strategy being “broken” overclaims: what the evidence supports is that the effect did not pay in *this* regime, not that it is dead. Several of the decayed effects could read differently after a regime shift, which is why a single-window fade is regime-specific evidence rather than proof an effect is dead. The natural robustness extension is to move the out-of-sample start back to Jan 2022, so the window spans the 2022 bear market as well as the 2023–25 bull.
-

9. Appendix

9a. Per-paper notes: what worked, what didn’t

A qualitative companion to §5: where each reproduction got traction, where it broke, and what surprised us.

RSI trend — Hill (2019).

- *Worked:* the signal reproduces the paper’s trade-level statistics within ~2%, and in-sample it carried a real alpha ($t = 2.5$).
- *Didn’t:* out-of-sample the alpha reversed ($\beta \approx 1.0$, $\alpha -7.4\%$), what’s left is market beta.
- *Surprising:* the paper’s *per-trade*-preferred stricter variant collapses the portfolio to ~12 names; the “second-best” trade-level signal is the better *portfolio* (per-trade quality and portfolio diversification pull in opposite directions).

52-week-high reversal — Zhu, Sun & Stivers (2021).

- *Worked:* the first build ran without error and produced plausible numbers (IS 0.70, OOS 1.95 on the mis-built corner-sort).
- *Didn’t:* the paper PDF was inaccessible when we built it; the first build diverged structurally from the paper (corner-sort vs reversal-within-low-PTH; equal-weight vs value-weight). When the paper was obtained and

the build graded against it, the paper-faithful rebuild (reversal within the low-PTH quintile, VW) collapses to IS 0.13 / OOS -0.95.

- *Surprising*: the 1.95 came entirely from the divergent construction, the wrong quintile for the long leg plus equal-weighting in a micro-cap-heavy spread each contributed. A confident result from a working pipeline, caught only by checking the build against the source text.

Lottery (MAX) — Bali, Cakici & Whitelaw (2011).

- *Worked*: direction reproduces; full-sample alpha clears the $t > 3$ bar.
- *Didn't*: decayed to negative out-of-sample, and the edge lives in capacity-limited micro-caps.
- *Surprising*: the data artifact (see §3), one halted micro-cap inflated the curve from $\sim 16\times$ to $\sim 38\times$ until the liquidity screen and value-weighting removed it. A cautionary tale for any broad equal-weighted sort.

Low-beta (BAB) — Frazzini & Pedersen (2014).

- *Worked*: direction reproduces in-sample.
- *Didn't*: our build is a simplified “BAB-shaped” long-only low-beta sleeve, not the paper’s literal levered beta-neutral factor, a loose replication; out-of-sample the signal was dead in a regime that rewarded exactly the high-beta names it avoids.
- *Surprising*: its low turnover makes it the most cost-robust strategy here; the regime, not costs, is what killed it.

Falling knives — Arendarski (2012).

- *Worked*: a faithful replication that reproduces the paper’s order of magnitude ($\sim 9\times$ SPY) on the original window.
- *Didn't*: the full-sample alpha is statistically insignificant ($t = 0.9$) despite a large point estimate, too volatile and micro-cap-concentrated; out-of-sample it lost 4%/yr with a -34% drawdown.
- *Surprising*: the -50% price pre-screen does the work; the solvency overlays add little.

Market-state momentum — Rodon Comas (2025).

- *Worked*: the most paper-faithful reconstruction (every deviation declared), reproducing the paper’s direction in-sample, with the down-state reversal adding value where down-states occurred.
- *Didn't*: the load-bearing down-state mechanism never fired out-of-sample (2023–25 had no sustained down-regime), so the out-of-sample is dormant momentum at a near-zero alpha.
- *Surprising*: the broad-universe version is commission-impractical at realistic size, the deployable narrow version is where the edge disappears.

Persistent momentum — Chen (2016).

- *Worked*: a significant full-sample alpha existed historically ($t = 2.4$).
- *Didn't*: the published long/short claim is exactly zero out-of-sample, the short leg was crushed by the mega-cap rally.
- *Surprising*: the verdict is “reproduced historically, gone now,” not a flat failure; the paper’s central conditioning variable (analyst-coverage dispersion) isn’t in our data, leaving a conditional version untested.

Seasonality — Heston & Sadka (2008).

- *Worked*: the same-calendar-month effect qualitatively persists (in-sample $t = 2.4$ on the paper-faithful 20-year lookback), and out-of-sample the spread is positive.
- *Didn't*: the paper’s signal needs multi-decade lags that our 1998 data start can’t fully supply, so the paper-faithful long-lookback variant is sample-starved and its full-sample significance can’t be established; the in-sample reading does not carry out-of-sample ($t = 0.9$).
- *Surprising*: a shorter 5-year-lookback variant (not the paper’s signal) gives a lower full-sample $t \approx 1.0$; the faithful variant is simply too data-limited to call, hence the “Untestable” label.

Price-level look-ahead — Geertsema & Lu (2019).

- *Worked*: the thesis is confirmed and the look-ahead bias is quantified at ~ 3.5 Sharpe points out-of-sample (split-adjusted compounds to $\sim 515\times$, as-traded decays to $\sim 0.01\times$).
- *Didn't*: n/a, it's a confirmation, not a strategy.
- *Surprising*: the bias also surfaced in the *same* 2007 month as the Bali artifact, a shared underlying data glitch that two independent reconstructions both exposed, reinforcing why single-month outlier scans are part of the protocol.

9b. Methodology notes

- **Universe.** Per-paper, point-in-time, refreshed annually. Delisted names are retained, so the universe is survivorship-bias-free by construction; data gaps are logged and skipped, never substituted. Broad-cross-section runs apply a \$5 price floor and a one-year listing minimum.
- **Membership-PIT verification status.** The S&P 500 membership universe (Hill) is verified point-in-time: membership is reconstructed by replaying the index add/remove event-log into [start, end) intervals and snapshotting annually, not read from a current roster. Spot-checks confirm correct entry dates (Tesla first eligible 2021, Meta 2014, Nvidia 2002) and that delisted names are retained in their membership years (Enron through 2001, Lehman through 2008, AMR through 2003, names a survivorship-biased roster would omit). For the **broad cross-section** (the reconstructed papers), survivorship-safety rests on the same delisted-retention mechanism, and a per-paper audit of delisted-name coverage has now been run: eventually-delisted names make up 43–88% of held positions in the older years, confirming the universe retains them rather than dropping them. A further check booked the loss on every name that delisted while held (rather than freezing it at its last price) and re-priced the survivorship worst case: Rodon's broad paper-faithful variant, whose micro-cap-heavy long-loser leg is the most exposed. Booking those losses moved its out-of-sample Sharpe only $\sim 1.87 \rightarrow \sim 1.84$, confirming the broad-universe results are not survivorship-inflated. That ~ 1.84 is a flat-5-bps, capacity-unconstrained figure, not a tradeable survivor: the broad variant has near-zero realistic capacity (its smallest-ADV names bind, and $\sim 53\%$ of its short leg is un-shortable), so liquidity-tiered costs cut it to ~ 0.67 (negative under a harsh borrow assumption). This is why the scorecard reports the deployable top-500 build, at out-of-sample Sharpe 0.12, as Rodon's tradeable form.
- **Tradeability screen.** Broad-universe sorts are filtered at formation by a trailing dollar-volume floor plus a stale-price guard (no run of identical prices), removing untradeable names before ranking. This follows the value-weighting + micro-cap-screen standard of the anomaly-replication literature (Hou, Xue and Zhang 2020).
- **Weighting.** Value-weighted where the paper's primary specification is value-weighted (e.g. Bali); equal-weighted where the paper specifies it. Equal-weighting an unscreened micro-cap universe is the single largest source of overstated anomalies and is avoided.
- **Risk-free rate** is period-appropriate ($\approx 1.7\%$ in-sample, $\approx 4.7\%$ out-of-sample); a flat low rate would overstate every Sharpe in the current regime by $\sim 0.3\text{--}0.5$.
- **Sharpe conventions.** Long-only Sharpes are in excess of cash; long/short Sharpes are the zero-cost spread (no cash deduction). All Sharpes are annualized arithmetic.
- **Drawdown convention.** MaxDD is the peak-to-trough of the **daily** equity curve. An earlier draft reported MaxDD from monthly-resampled returns, which masks intra-month troughs and understates the true draw-down $\sim 1.5\text{--}2.6\times$ (e.g. SPY out-of-sample is -18.8% daily vs -8.3% monthly over 2023–25, the April-2025 selloff being the binding trough). Figures recomputed from a persisted daily curve are unmarked; the five reconstructed papers whose spike persisted only monthly returns retain a monthly MaxDD, marked ^m, pending a daily rerun. Read every ^m value as a lower bound on the true daily trough, not as a final number. Sharpe and CAGR are resolution-robust and unaffected.
- **Factor decomposition.** CAPM (single-factor): returns are regressed on the market excess return; long-only in excess of cash, long/short as the raw zero-cost spread. Alpha and its t-statistic separate edge from market exposure. The cross-paper scorecard uses CAPM for all nine papers. A multi-factor decomposition of the surviving candidates is the natural next refinement.
- **Paper-window overlap** is reported because several papers' samples predate the 1998 data start; short-overlap reproductions are small-sample reads, flagged as such.
- **Seasonality lookback (Heston).** Heston-Sadka's signal is the same-calendar-month historical average re-

turn, built from up to ~20 years of prior history; the 20-year variant shown here is therefore the paper-faithful one (a 5-year variant was also run but is a shorter-horizon deviation, not the paper's signal). Our reproduction is bounded by the 1998 data start: full 20-observation depth is only available from ~2015, with earlier formations using ≥ 10 observations, so treat this as a data-limited reproduction of the paper's longer-horizon effect.

- **Independent-engine validation.** Four candidates (Hill, Zhu, Chen, Rodon) were re-run on a second, independent backtest engine; agreement within a small Sharpe band is treated as evidence a result is not an implementation artifact. The other five papers, including Bali (the only full-sample $t > 3$) and Frazzini, are single-engine, and the Zhu cross-check covered the first, divergent build rather than the faithful rebuild reported here.
- **Turnover and break-even.** Turnover is annual one-way traded notional, computed exactly from per-name holdings where available (Zhu, Chen, Rodon) or derived from trade counts (Hill); for the five reconstructed papers only basket sizes are available, so turnover/break-even are basket-churn estimates, flagged †. Break-even is the per-side cost that drives the out-of-sample return to zero; ‡ marks strategies that lose money before any cost.
- **Cost model.** Paper-comparison runs use a flat 5 bps/side. The candidate (Zhu) is additionally stressed at commission + half-spread keyed to universe liquidity + borrow on the short leg, applied to measured turnover (see §5.2). Multiple testing is handled by the $t > 3$ hurdle (Harvey, Liu and Zhu 2016) on full-sample alphas.

9c. Full metric table (in-sample and out-of-sample, all nine)

Sharpe, CAGR and α are monthly-based; **MaxDD is computed from the daily equity curve** where one was persisted (SPY, Hill, Zhu, Chen, Rodon) and is ^m-marked where only a monthly series exists (Bali, Frazzini, Arendarski, Heston-Sadka, Geertsema-Lu; daily recompute pending; see the drawdown convention in §9b). Sharpe annualized arithmetic (long-only excess-of-cash, long/short zero-cost); α is CAPM annualized alpha vs the market with its t-statistic. Turnover/break-even per the scorecard in §1 († estimated, ‡ n/a).

Paper	Window	Sharpe	CAGR	MaxDD	CAPM α (t)
RSI trend	In-sample	0.55	10.1%	-60.4%	+3.6% (2.5)
RSI trend	Out-of-sample	0.66	13.7%	-17.8%	-7.4% (-1.5)
52-wk-high reversal (faithful rebuild)	In-sample	0.13	3.5%	-67.1%	~0% (-0.2)
52-wk-high reversal (faithful rebuild)	Out-of-sample	-0.95	-29.1%	-55.0%	~-29% (-1.2)
Lottery (MAX)	In-sample	0.76	11.0%	-34.7% ^m	+6.2% (4.4)
Lottery (MAX)	Out-of-sample	0.30	7.5%	-8.4% ^m	-7.7% (-1.7)
Low-beta (BAB)	In-sample	0.73	10.0%	-40.3% ^m	+5.6% (3.1)
Low-beta (BAB)	Out-of-sample	0.27	7.0%	-6.2% ^m	-1.6% (-0.3)
Falling knives	In-sample	0.46	13.9%	-70.3% ^m	+11.1% (1.4)
Falling knives	Out-of-sample	-0.23	-4.2%	-33.8% ^m	-24.2% (-1.7)
Market-state mom.	In-sample	0.33	3.4%	-42.9%	+3.9% (1.5)
Market-state mom.	Out-of-sample	0.12	0.6%	-15.6%	-1.2% (-0.2)
Persistent mom.	In-sample	0.26	2.5%	-36.9%	+5.3% (2.2)
Persistent mom.	Out-of-sample	0.17	1.0%	-11.6%	-0.3% (-0.1)
Seasonality	In-sample	0.52	4.1%	-28.5% ^m	+4.2% (2.4)
Seasonality	Out-of-sample	0.68	5.9%	-9.4% ^m	+5.3% (0.9)
Price-level look-ahead (as-traded)	In-sample	-0.36	-16.6%	-97.0% ^m	-3.7% (-0.6)
Price-level look-ahead (as-traded)	Out-of-sample	-2.32	-44.3%	-84.7% ^m	-80.7% (-5.1)

Geertsema-Lu shown on the PIT-correct as-traded price; the look-ahead split-adjusted variant (IS Sharpe +1.91 / OOS +1.24) is the illusory comparison in §5.9. Zhu numbers are the paper-faithful rebuild (VW reversal-within-low-PTH); the first-build corner-sort (IS 0.70 / OOS 1.95) is the divergent build documented in §5.2.

References

1. Hill, "Finding Consistent Trends with Strong Momentum: RSI for Trend-Following and Momentum Strategies" (2019). SSRN 3412429. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3412429
2. Zhu, Sun & Stivers, "Price Anchors and Short-Term Reversals." *Financial Management* 50(2) (2021). SSRN 3092325. <https://ssrn.com/abstract=3092325>
3. Bali, Cakici & Whitelaw, "Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns." *Journal of Financial Economics* (2011). NBER w14804. <https://www.nber.org/papers/w14804>
4. Frazzini & Pedersen, "Betting Against Beta." *Journal of Financial Economics* (2014). <http://pages.stern.nyu.edu/~lpederse/papers/BettingAgainstBeta.pdf>
5. Arendarski, "Tactical Allocation in Falling Stocks: Combining Momentum and Solvency Ratio Signals." WNE Working Paper No. 1/2012 (67), University of Warsaw (2012). http://www.wne.uw.edu.pl/inf/wyd/WP/WNE_WP67.pdf
6. Rodon Comas, "Winners & Losers in Motion: A Market-State Momentum Signal" (2025). SSRN 5130289. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5130289
7. Chen, "Persistency of the Momentum Effect" (2016). SSRN 2652592. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2652592
8. Heston & Sadka, "Seasonality in the Cross-Section of Expected Stock Returns." *Journal of Financial Economics* (2008). SSRN 687022. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=687022
9. Geertsema & Lu, "Revisiting the Price Effect in US Stocks" (2019). SSRN 4013958. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4013958
10. Hou, Xue & Zhang, "Replicating Anomalies." *Review of Financial Studies* (2020).
11. Harvey, Liu & Zhu, "... and the Cross-Section of Expected Returns." *Review of Financial Studies* (2016).
12. McLean & Pontiff, "Does Academic Research Destroy Stock Return Predictability?" *Journal of Finance* (2016).