Frictions All the Way Down

A Fully-Reproducible Audit of Time-Series Momentum from Frozen Specification to Integer Contracts

Vincent W.

June 2026 | working paper

Without giving up the index position, can an individual investor systematically improve on simply holding it?

Representative scale $500,000 (not a disclosure of actual funds or positions)

| Beating the index is hard: three barriers

01 | Overfitting

Parameters

Pretty parameters picked in a backtest collapse out of sample.

02 | Wrapper leakage

Wrapper

Switching from ETFs to tradable instruments—does the strategy survive?

03 | Account too small

Granularity

Integer contracts and instrument life events eat the paper return.

Our contribution: take one strategy across all three barriers, end to end.

[Schematic] Trend following = crisis convexity

Follow the trend, don't predict

The signal is the sign of the trailing return, which sets the position direction.

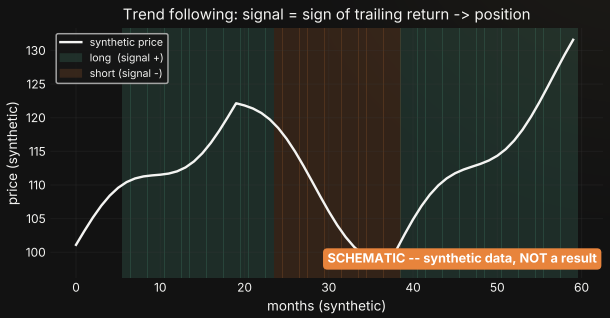

! The only synthetic-schematic figure in the deck—not a measured result.

A fifty-year proof of concept

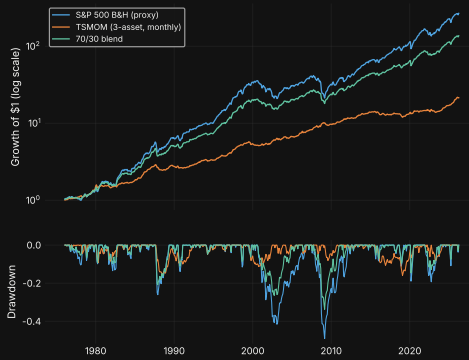

1.09
Sharpe

6.3%
CAGR

-15.9%
Max drawdown

0.21
Corr. to S&P

1976–2026 (603 months), three assets. Blue = S&P proxy, orange = TSMOM, green = 70/30; top = log growth, bottom = drawdown. Declared approximations: bond-duration proxy; no risk-free deduction (Sharpe overstated ≈0.2–0.3); monthly; non-investable indices.

| Credibility: the watershed vs a backtest

Frozen specification

Parameters set from literature priors and frozen **before** estimation.

Look-back 3/6/12 mo
Vol window 60 d
Target vol 10%
Leverage cap 1.50

Unit tests vs look-ahead

An independent shift test and a “double the future return” test enforce no look-ahead.

Sample split: 2015-01-01

Full presentation

108-cell robustness, 6 crisis windows, pre/post in full—**no cherry-picking**.

NW (lag 6)
Circular block bootstrap (10000×6 mo)

Crises are the main battlefield

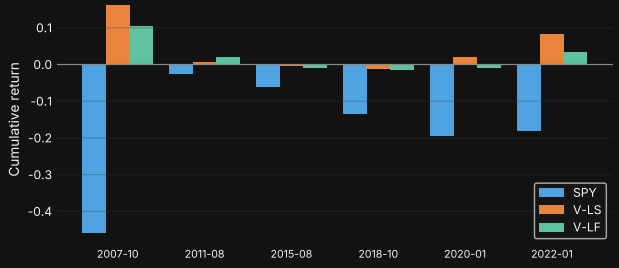

6/6

crisis windows all won

GFC window

Overlay sleeve
16.1%

SPY, same window
-46.0%

Crisis convexity is the core source of economic value.

| Robustness: 108/108 cells positive, no cliff

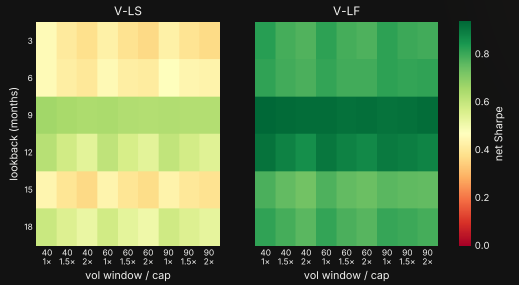

0.36
minimum net Sharpe

Per column: upper row = vol window, lower = leverage cap; rows = look-back months; left V-LS, right V-LF; bar = net Sharpe.

No cliff anywhere on the parameter surface.

All strict pairs clear the 0.85 correlation gate

strict ETF / futures pair	V-LS corr.	V-LF corr.
ES / SPY	0.95	0.97
ZN / IEF	0.90	0.95
ZB / TLT	0.94	0.94
GC / GLD	0.94	0.94

4/4
all clear 0.85

0.90–0.97
observed range

0.81
portfolio V-LS

Swapping the instrument wrapper is cheap; the strategy itself ports.

Portfolio correlation is a separate aggregate metric. CL/DBC, 6E/UUP, 6J/UUP are context pairs, not strict-pair evidence; full diagnostic in backup.

Decision geometry: the 80/20 + 1.5× corner

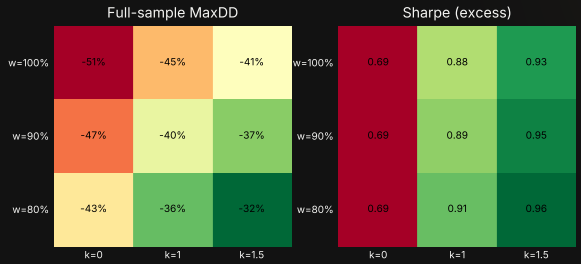

Corner = base 80% + return-stacking overlay 1.5×: CAGR 14.7%, drawdown -31.8%, Sharpe 0.96; pure SPY gives CAGR 11.3%, drawdown -50.8%.

Only this direction simultaneously raises return, lowers vol, and narrows drawdown.

Smaller accounts incur more tracking error

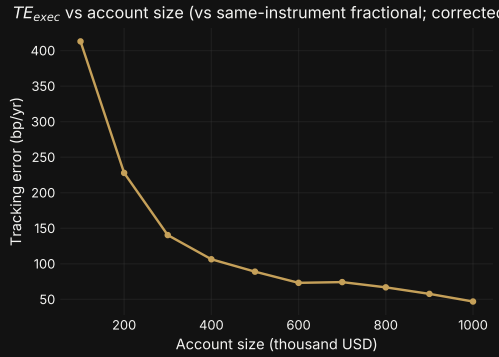

413 bp
$100K account

228 bp
$200K account

327 bp
final vs frozen model

Two error components co-vary: rounding+commission 110 bp; instrument basis 355 bp; not additive in quadrature. Decomposition in backup.

Executable form: eight tradable legs

instrument	contract	signed contracts	status
MES	MES-2024-06	3	OK
MNQ	MNQ-2024-06	2	OK
MGC	MGC-2024-06	3	OK
MCL	MCL-2024-07	7	OK
M6E	M6E-2024-06	-3	OK
6J	6J-2024-06	-1	OK
ZN	ZN-2024-09	-1	OK
ZB	ZB-2024-09	-1	OK

! Representative scale $500,000; a process record only—not an order or evidence of a completed trade.

Execution boundary: sizing and pre-committed triggers

Sizing boundary

$500,000

representative scale

Per-leg ratio of target notional to contract value capped at 1.22.

Pre-committed triggers

Data staleness >5 business days
⟶ suspend that leg.

Instrument lifecycle event
⟶ follow the spec's fallback chain.

Evidence boundary

The full-precision audit table is in backup.

Two error classes, both recorded

Research erratum

$200K
tracking error

353 → 228 bp

A 100× contract-value defect, detected, corrected, and restated.

Production self-audit

108 rows
appendix

dropped → all restored

The overflowing `tabular` became a `longtable`; every row re-checked against the source CSV.

Sources: paper §10, `erratum_record.csv`, P1 build note; research-number errors and production errors disclosed separately. Second erratum in backup.

| The audit chain is itself the product

01 | FOUNDATION

Frozen spec

Core parameters and arithmetic are frozen before any operation, so the logic cannot be rewritten.

02 | EXECUTION

Agent build

Pre-registered thresholds bind roll, substitution, and sizing rules, blocking in-flight rationalization.

03 | GOVERNANCE

Human call

Capital allocation and the spec-change pen stay in human hands for final control.

| Decoupled authority: the core landing

MACHINE / ARITHMETIC

Arithmetic

machine-executed

HUMAN / CAPITAL

Capital

human-adjudicated

Arithmetic can be frozen; capital must be adjudicated by a human—the transferable result.

Human-adjudication example: the 6J rule tie

2.0%

MJY impact vs zero full-size

Rule outcome

The 6J rule produced a statistical tie within the noise band.

Liquidity-depth adjudication

On liquidity depth a human chose the full-size contract: MJY impact 2.0% versus zero for full-size.

Rule outcome and human override are **both recorded**.

Capital allocation must be human: a threshold once overreached to bind account size (1M); a human decoupled it and set $500,000. This is the transferable result—the landing of "the audit chain is the product."

| Limitations and the evidence boundary

Evidence boundary

Frozen priors $\neq$ out-of-sample validation.

All results are **in-sample**; spliced series are a research proxy, **not directly reproducible live**.

Observed decay

Long/short Sharpe: full 0.57; pre 0.73 $\rightarrow$ post 0.40.

Long/flat: pre 0.96 $\rightarrow$ post 0.73.

Forward anchor

Forward expectations should anchor to the post-split band of about 0.8 (paper §9.5), not the full-sample corner 0.96.

Execution convention, validation, and readiness

Same-close convention

Fill at the signal-day close (standard MOP)—an **idealized convention**.

No look-ahead in accounting, but not executable timing.

Delayed-execution sensitivity

One-day delay proxy: excess Sharpe 0.32/0.58.

Sensitive to the entry gap; validation requires post-signal fills.

Current readiness

Internal research archiving and controlled paper-trade validation: **conditionally usable**.

No live-deployment evidence; live-trading readiness is not established.

Not investment advice. Historical in-sample results do not predict future returns.

| Backup | Significance: bootstrap CI and Newey–West t

Long/short V-LS

0.57

excess Sharpe

Circular block bootstrap 95% CI

[0.19, 0.96]

NW $t = 2.86$

Long/flat V-LF

0.85

excess Sharpe

Circular block bootstrap 95% CI

[0.49, 1.24]

NW $t = 4.45$

Source: stats_summary.csv; NW (lag 6), circular block bootstrap (10000×6 mo block), seed 42.

| Backup | Tracking-error components co-vary

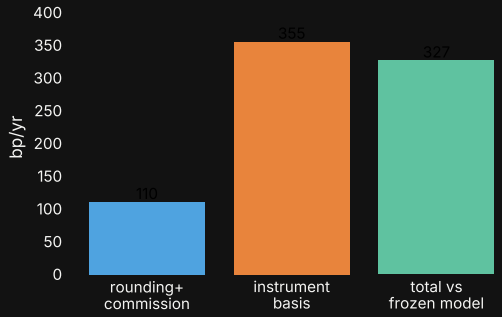

110 bp
rounding + commission

355 bp
instrument basis

327 bp
final vs frozen model

Components co-vary, not additive in quadrature; acceptance channel [261, 653] bp.

| Backup | Instrument lifecycle events require fallbacks

01 | M6J delisted

2024-03

event date

Micro-yen contract no longer available.

02 | MJY listing gap

22 mo

interruption

Alternate yen not continuously listed.

03 | 30Y liquidity death

26 d

lag

Median volume 1 lot.

Instrument lifecycle is an execution constraint, not a footnote.

Backup | Full availability and substitution timeline

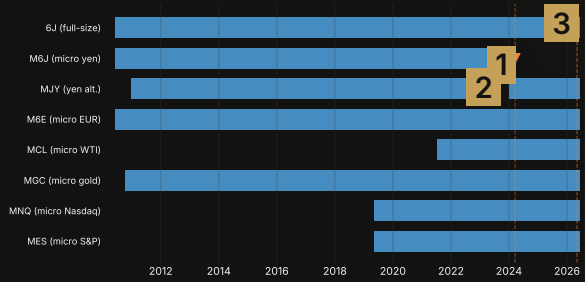

1 | M6J delisted

2024-03

2 | MJY listing gap

22 months

3 | 30Y liquidity death

26-day lag; median 1 lot

| Backup | Rolling 36-month Sharpe & correlation

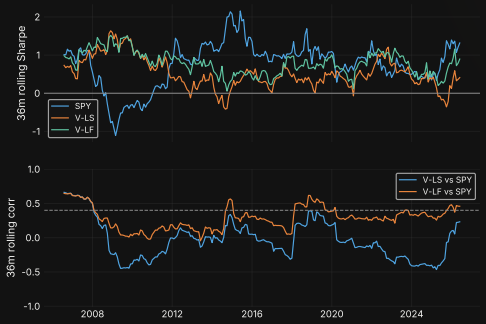

Rolling 36-month Sharpe and correlation to SPY; **all in-sample**.

| Backup | Full-precision audit table

leg	instrument	contract	close_used	multiplier	cv_usd	n_contracts	status
ES	MES	MES-2024-06	5299.500000	5.000000	26498.000000	3	OK
NQ	MNQ	MNQ-2024-06	18571.750000	2.000000	37144.000000	2	OK
GC	MGC	MGC-2024-06	2325.500000	10.000000	23255.000000	3	OK
CL	MCL	MCL-2024-07	77.160000	100.000000	7716.000000	7	OK
6E	M6E	M6E-2024-06	1.085600	12500.000000	13570.000000	-3	OK
6J	6J	6J-2024-06	0.006371	12500000.000000	79631.000000	-1	OK
ZN	ZN	ZN-2024-09	108.890625	100000.000000	108891.000000	-1	OK
ZB	ZB	ZB-2024-09	116.281250	100000.000000	116281.000000	-1	OK

! Representative $500,000 paper-trade target record; not an order or evidence of a completed trade. Exact values preserved from the generated audit table.

| Backup | Migration: strict vs context pairs

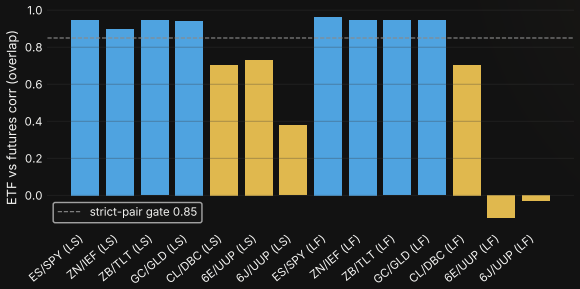

Blue | strict pairs

Overlap-qualified; support the strict-pair claim.

Orange | context pairs

Completeness only; not evidence. Dashed = 0.85 gate.

| Backup | Research and production audit details preserved

Research errata

- 100× contract-value inflation: $200K tracking error 353 → 228 bp.
- Folding ZB duration into ZN made it *worse*: 327 to 377 bp.

Production self-audit

- A 108-row `tabular` overflowed and dropped rows.
- Replaced with `longtable`; every row re-checked against the source CSV.

Sources: paper §10, `erratum_record.csv`, P1 build note; production notes are not paper research results.