

EWMA_63_Strat — validation report

Public strategy from github.com/Filip303/EWMA_63_Strat · 197 monthly decisions, January 2010 – May 2026 · author's declared verdict: NO-GO

VERDICT

The EWMA63 layer adds no measurable value. The author's NO-GO is confirmed, and understated.

The published figures reproduce exactly and the comparison is correctly built. What the original report does not say is that the layer's small gross edge is entirely cancelled by a trading cost of 38.7 basis points — inside the normal range for the small-cap names this portfolio holds.

1. What the strategy does

The portfolio starts as 200 US stocks, bought and held. On top of that sit two layers that act together whenever a stress indicator fires: **first it cuts total exposure** (from 100% to 70% or 50%), and **while cutting, it reduces high-beta names more** than low-beta ones. That second layer is EWMA63.

The first layer is simple and obvious: hold less when the market falls and you fall less. **The question that settles everything is whether the second layer adds anything on top of the first**, or whether choosing *which* stocks to cut makes no difference as long as you cut the same amount.

2. What it is compared against

✗ Buy-and-hold

Does not answer the question. Against staying fully invested the strategy looks better on Sharpe and on drawdown — but the first layer produces that, by holding less. Comparing this way makes EWMA63 look valuable when it has not been measured at all.

✓ Exposure-matched control

A portfolio that cuts exactly the same exposure at every decision, but spreads the cut without looking at beta. It is the only comparison in which a difference can only come from EWMA63.

That control carries the whole conclusion, so the first thing checked was whether it really is matched. It is: across all 197 decisions the two portfolios never differ in exposure by any meaningful amount. The comparison is clean.


3. The result

PORTFOLIO	CAGR	SHARPE	MAX DRAWDOWN	ANNUAL TURNOVER
Strategy with EWMA63 layer <i>the strategy under test</i>	22.77%	1.343	-21.82%	0.873x
Exposure-matched control <i>the correct control: same exposure, beta-blind</i>	22.70%	1.324	-23.46%	0.802x
Difference attributable to EWMA63 <i>a positive sign favours the EWMA63 layer</i>	+0.07 pp	+0.018	+1.64 pp	+0.071x

At first glance the strategy wins slightly: 0.018 more Sharpe and 1.64 points less drawdown. But that is **a single historical realisation**, not a measured effect. The return difference — the only thing that can be tested statistically — is +0.02% per year with a t-statistic of 0.12. Anything below 2 is treated as indistinguishable from chance; 0.12 is nowhere near.

And there is something the original report does not break down: **where that +0.02% comes from**. The layer does add value before costs, but it requires more trading, and the two cancel:

+0.0274%	0.0707x	38.7 bps
what the layer adds before costs, per year	extra turnover it requires per year (0.873x versus 0.802x)	trading cost at which the two cancel out



Net incremental annual return of EWMA63 against the exposure-matched control, by trading cost. It is a straight line because the two portfolios differ only in which positions they cut and how much they trade to do it.

ASSUMED COST	NET ANNUAL CONTRIBUTION	T-STATISTIC
0 bps	+0.0274%	+0.16
10 bps	+0.0203%	+0.12
25 bps	+0.0097%	+0.06

38.7 bps	+0.0000%	+0.00
50 bps	-0.0080%	-0.05
100 bps	-0.0433%	-0.26

Above 38.7 bps the layer destroys value. And because the effect is not distinguishable from zero at any cost level, the honest reading is not that it is profitable below that threshold: it is that **there is no measurable effect to price**. As confirmation: the layer changes the allocation in 36 of 197 months, but changes turnover in 43. In 11 months it pays a cost without changing the outcome.

4. Context

PORTFOLIO	CAGR	SHARPE	MAX DRAWDOWN	ANNUAL TURNOVER
200 stocks, buy-and-hold <i>always invested, no layers</i>	24.57%	1.202	-33.61%	0.123x
Exposure gate only <i>first layer only: cuts exposure, beta-blind</i>	22.58%	1.314	-24.12%	0.803x
SPY, buy-and-hold <i>the market</i>	14.49%	0.985	-23.31%	0.061x

These three settle nothing about EWMA63, but they explain why it is easy to fool yourself. Against staying fully invested, the strategy lifts Sharpe from 1.202 to 1.343 and cuts drawdown from -33.61% to -21.82%. That is a real and large improvement. **All of it comes from the first layer**, which is far simpler: the exposure gate alone already reaches 1.314 and -24.12%.

5. How it was checked

The analysis was re-run on the same frozen price data, confirmed to be the exact files the original used, and 180 published figures were compared one by one: 0 unexplained differences. Return, Sharpe, drawdown, turnover, exposure and the statistical tests were then recalculated from the monthly result series alone — 37 values — without using the author's own calculations, because an error inside a calculation reproduces itself faithfully every single time.

That was the scope and nothing more. Three limitations the author discloses remain open and no re-run can close them: the parameters were chosen after seeing the history they are tested on, so there is no untouched forward period; the 200-name universe is companies that exist today, which inflates the buy-and-hold reference by construction; and the price history contains corporate-action distortions the author documents but does not trace to source.

Prepared by Strateva.ai — independent quant research validation. A sample of the deliverable format, produced on a public repository so every figure can be checked by the reader. Not investment advice; no representation is made about future performance.
stan@strateva.ai · strateva.ai