

Macro Drawdown Control

A Five-Portfolio Study

Reducing macro-factor-driven drawdowns with weight-only adjustments



Contents

1. The question this answers
 2. The risk measure, in plain English
 3. The guardrail: what it is, and why it is optional
 4. The monthly process
 5. The test
 6. What we found
 7. Two independent referees
 8. The behaviour through time
 9. Where the method has limits — stated plainly
- Appendix A — Backtest methodology
- Appendix B — Implementing this on your own book



1. The question this answers

Portfolio managers who see their macro risk measured for the first time ask the same thing: “Fine — macro is a large share of my risk. Now show me what to do about it, and show me it works.”

This study tests one narrow claim: using MFERM, macro-factor-driven drawdowns can be reduced by shifting weights among the stocks a book already holds — no new names, no cut in gross, no derivatives.

What we claim: lower macro-driven drawdowns, at unchanged gross, net and holdings.

What we do not claim: lower total maximum drawdown, higher total return, or higher Sharpe. Cutting the macro share of risk raises the idiosyncratic share — and what your stock selection then does is your alpha, not our claim. If idiosyncratic risk is strong the book improves overall; if it is weak, it will not. The method controls the macro part, full stop.

2. The risk measure, in plain English

Everything below is built on one number: weekly Expected Shortfall (ES) at 5%. Take every weekly return the book could have had through history, look at the worst 5% of those weeks, and average them. Where Value-at-Risk tells you where the cliff edge is, ES tells you how far you fall once you are over it. Drawdowns live in the tail; ES measures the tail.

We then isolate the macro part. Take each stock’s current macro sensitivities — its MFERM factor betas — run the recorded history of factor moves through them, and ask: if I hold today’s exposures, what would the worst 5% of weeks have done to me? That number, the macro ES, is the control variable for the whole process. It is fully empirical: no normal-distribution assumption anywhere.



3. The guardrail: what it is, and why it is optional

The guardrail is a line in the sand: macro ES must not exceed a set share of the whole book's ES. Throughout this study that share is 20% — “macro should be at most a fifth of my tail risk; the other four-fifths should be my stock picks.”

The 20/80 split is an illustration, nothing more. The level is entirely yours to set — 10/90, 30/70, whatever your mandate implies — and the process is unchanged at any level. You can also run without a guardrail at all.

With a guardrail, the overlay trades only when the line is breached, and then makes the smallest trade that clears it. The advantage is that you do not pay for lower macro risk when there is little macro risk to lower: in quiet regimes there is no trading, no turnover and no deviation from the weights your stock selection chose. It also gives a risk committee one number to approve and audit.

Without a guardrail, the overlay simply takes the largest macro ES reduction available inside the bands, every rebalance, unconditionally. That is maximum protection — and maximum cost: weights sit permanently at their band edges, turnover is higher, and the book carries a full-size tilt even in months when macro is a small share of the tail and the tilt buys little.

Both configurations are covered by this study's results: wherever the guardrail was unreachable, the overlay was effectively running in unconditional maximum-reduction mode, and the drawdown improvements reported include those periods. Our view: keep the guardrail. It makes the overlay a light touch that acts when macro risk is genuinely elevated, rather than a standing tax on the book. But it is a dial, not a doctrine — and it is yours.

4. The monthly process

The process repeats monthly, and every decision uses only information available on the day it is made.

1. At each month-end, measure the book's macro ES using that day's betas and the full factor history to date.
2. If macro ES is at or below the guardrail: do nothing. This is the first and most common lever.



Macro Drawdown Control

3. If it is above: shift weights among the stocks already held to bring it down — subject to four constraints: no name changes; gross unchanged; net unchanged; and no position moved more than $\pm 35\%$ from its original weight (longs stay long, shorts stay short).
4. If the guardrail cannot be reached inside those constraints, take the largest reduction available and record that the month did not clear.
5. Hold for a month, let weights drift naturally, and repeat.

No hindsight enters at any step.

5. The test

Five real portfolio structures, spanning regions, factor models, concentration and net exposure. For each, two books ran side by side for up to 4.5 years (2022–2026): book A, the original weights left alone; and book C, the same book with the monthly adjustment. Same stocks, same size, same direction — the only difference is the overlay, so any difference in outcomes is the overlay's doing.

Book	Structure	Names	Gross / net	Factor model
Japan L/S	Concentrated, asymmetric	215	1.0 / 0.0	APAC, 25 factors
US L/S	Crowded shorts, levered	100	2.0 / 0.0	US, 12 factors
US market-neutral	Diffuse quant book	591	3.9 / 0.0	US, 12 factors
Global net-long	Net long, concentrated	80	2.3 / +0.75	Global, 23 factors
APAC	Concentrated longs	80	1.0 / 0.0	APAC, 25 factors



6. What we found

The claim metric. The maximum drawdown of the macro component of P&L — the direct measure of macro-driven drawdown — fell in all five books, by between a quarter and a half.




Figure 1. Max drawdown of the macro P&L component, unadjusted vs overlay. Reductions of 26–50% across five structures, three regions, three factor models.

The days that matter. On the worst 5% of macro-factor days — the days the overlay exists to protect — cumulative losses were smaller in all five books (for example –17.3% improving to –11.4% on the Japan book, and –24.6% to –14.0% on the APAC book). One note on frequency: the risk estimator is weekly throughout this study, but this test is deliberately daily — it is an outcome measure, not a risk estimate, and daily granularity gives it roughly five times the tail sample.

Real shocks. Across the August 2024 shock, the 2025 tariff episode and the 2026 Iran shock — fifteen book-episode tests — total intra-episode drawdown improved in fourteen and was flat in one (a book with no macro drawdown to save in that window). In 2022, where several books' losses were idiosyncratic rather than macro, no improvement is claimed and none appears: the method's edge and its boundary are the same line.



Book	Aug 2024	Tariffs 2025	Iran 2026
Japan L/S	-3.6% → -2.3%	-3.4% → -2.5%	-1.3% → -1.1%
US L/S	-17.8% → -16.2%	-9.3% → -7.0%	-12.6% → -12.0%
US market-neutral	-6.6% → -6.6%	-3.1% → -2.9%	-2.0% → -2.0%
Global net-long	-9.2% → -8.5%	-11.8% → -8.8%	-13.0% → -10.7%
APAC	-11.3% → -7.3%	-2.9% → -1.7%	-3.4% → -2.8%

Table 1. Intra-episode maximum drawdown of total book P&L, before → after the overlay, across three macro-led shocks.

On a profitable book, the trade is visible — and fair. The APAC book made +19.5% a year unadjusted. With the overlay it made +15.3%, but total max drawdown fell from -12.6% to -9.5% and the August 2024 shock from -11.3% to -7.3% — with Sharpe and Calmar both slightly higher. Some upside was exchanged for shallower drawdowns at a fair rate. The cost is real and should be stated: trimming macro-loaded winners gives up some macro-carry return.

7. Two independent referees

A fair objection: the same model that picks the trades also scores the macro component — if a beta is wrong, both sides share the error. So the result was checked twice, on instruments the model does not touch.

- **Total P&L.** The shock-episode and tail-day results above are scored on total book P&L, which involves no betas at all.
- **Outside market data.** Both books' daily returns were regressed on external proxies — SPX, US 10-year yield changes, KOSPI, high-yield credit, crude, USDJPY and the dollar index. The variance explained by these proxies fell from unadjusted to overlay book in all five cases (for example 9.9% to 2.1% on the Japan book; 37.6% to 22.6% on APAC), and losses on the worst external-macro days were smaller in all five.

Both referees agree with the model's own scoreboard: the overlay books are genuinely less macro-sensitive.



8. The behaviour through time

Two views of the same result: the overlay compresses macro risk month by month, and holds a shallower drawdown path through every episode — not only at the single worst point.




Figure 2. Macro ES at each monthly rebalance. Grey is the unadjusted book's macro ES that day; green is the overlay book's macro ES immediately after that month's adjustment; the dashed line is the guardrail (20% of that month's whole-book ES). The overlay has two modes: where the guardrail is reachable it makes the smallest trade that lands exactly on the line — so green traces the dashed line (US market-neutral, and the calmer years elsewhere); where it is not reachable inside the $\pm 35\%$ bands it takes the maximum reduction available and sits above the line at the feasible floor (global net-long throughout). Where grey is already below the line, no trade occurs and green equals grey.




Figure 3. Drawdown of cumulative macro-factor P&L. The overlay path is a consistently shallower copy of the unadjusted one, through every episode.



9. Where the method has limits — stated plainly

- **Net exposure is irreducible.** A book that is structurally net long carries a directional macro position that no reweighting can remove. On the net-long book the guardrail was never reachable; the overlay still cut macro ES by a third each month. Getting such a book to a 20% macro share is a hedging conversation, not a weights one.
- **Breadth buys control.** The 20% target was reachable every month on the 591-name book, roughly half the time on 100–215-name books, and never on the net-long book. Diffuse books give the constraints the most room.
- **Monthly rebalancing compresses shocks; it cannot dodge them.** A macro event landing mid-month hits both books; the overlay's protection comes from arriving at that day with less macro exposure, not from reacting to it.
- **The price.** Turnover of 1–5% of gross per month one-way; on the profitable book, roughly 4 points of annual return given up, with risk-adjusted returns slightly improved. Transaction costs are not modelled and should be priced per desk from the reported turnover.

The full methodology, every assumption and a step-by-step implementation guide follow in the appendices. The pipeline behind this study runs on a standard five-file data pack; the same test can be run on any book, same-day.

Appendix A — Backtest methodology

This appendix sets out exactly what was done, step by step, with every assumption explicit.

A1. Objective and claim under test

The backtest answers one question: can weight-only adjustments, informed by MFERM factor exposures, reduce macro-factor-driven drawdowns on an equity long/short book? The claim is deliberately narrow. It does not extend to total maximum drawdown, total return, or Sharpe. Reducing the macro share of risk mechanically raises the idiosyncratic share, on which MFERM is silent. Only the macro component is claimed.

Two portfolios are compared over the same period: an unadjusted book (A) holding the original weights, and an overlay book (C) whose weights are adjusted monthly under the



rules below. Both books hold identical names, gross and net at all times.

A2. Data inputs

1. Positions: identifier, signed weight, one snapshot date.
2. Daily total returns per stock, in percent, covering the test window plus two years of prior history for burn-in.
3. Daily MFERM factor exposures (betas) per stock, all model factors, same window and calendar as returns.
4. Daily factor returns for the same factor set, from burn-in start.
5. Daily total/factor/specific risk per stock (optional; used for cross-validation only).

Where factor history begins in 2021, the two-year burn-in moves the live window to January 2023.

A3. Pre-processing

- **Units.** Stock-level files are % daily; returns are converted to decimals for P&L compounding. Where a source mixes conventions, everything is normalised to daily %.
- **Calendar.** The working calendar is the intersection of the stock-return and factor-return calendars; exposures are aligned by exact date.
- **Coverage.** Positions without usable return and exposure data are dropped and disclosed (1–4% of gross across the five books). Remaining sleeve weights are rescaled so long and short sleeves match the book's original sleeve sizes, preserving gross and net.
- **Listings and delistings.** A name is live at a date only if both its return and full exposure vector exist that day. Non-live names are set to zero and sleeves rescaled over live names — identically for books A and C, so universe effects cancel.
- **Exposure outliers.** Betas are winsorised at ± 10 per factor, after one thin-history name showed implausible values in source data. Applied uniformly.



A4. Risk estimator

Horizon and tail. One-week horizon. Weekly observations are overlapping 5-day arithmetic sums of daily returns (not $\sqrt{5}$ -scaled), preserving volatility clustering. Expected Shortfall at 5% is fully empirical: the mean of the worst 5% of weekly observations. No Gaussian assumption anywhere.

Macro (factor-only) ES at date t. Take each live stock's current beta vector at t; apply the full history of weekly factor returns to those betas; the candidate weight vector gives a weekly portfolio factor P&L history; macro ES is the empirical 5% ES of that series. It asks: if I hold today's exposures, what would the recorded history of factor moves have done to me?

Whole-book ES at date t. The empirical 5% ES of the weight vector applied to the history of weekly total stock returns.

Window. Expanding, never rolling, with a two-year burn-in on the factor side. The stock-return side carries a one-year burn-in where its history is shorter — an accepted limitation of data availability.

A5. Guardrail

Target macro/non-macro split: 20%/80%. At each rebalance date the guardrail level equals 20% of the unadjusted book's whole-book ES at that date. It is a fixed level for that month's optimisation, not a ratio re-evaluated on candidate weights. A month clears if macro ES can be brought to or below this level within the constraints; otherwise the overlay takes the maximum feasible reduction and the failure to clear is recorded.

A6. Rebalance schedule and overlay rules

- Rebalance dates: last trading day of each month.
- Conditional action: if unadjusted macro ES is already at or below the guardrail, the minimal-deviation solution is (near-)zero trading — do nothing is the first lever.
- Constraints at every rebalance: signs preserved; gross fixed at the book's original gross; net fixed at the book's original net; every weight within $\pm 35\%$ of its (sleeve-rescaled) original weight.
- Band convention: the $\pm 35\%$ band is applied as a relative bound on each weight. Because it is proportional, it is equivalent to a proportional bound on each name's standalone risk; a genuinely distinct risk-contribution band with cross-name transfer was considered and not implemented.



Optimisation runs in two stages, both exact linear programmes (Rockafellar–Uryasev CVaR formulation, solved with HiGHS):

1. Minimise macro ES subject to the constraints. This establishes the feasible floor for the month.
2. If the floor is at or below the guardrail: re-solve, minimising the L1 deviation from current weights subject to macro ES $\leq$ guardrail — the smallest trade that clears the line. Otherwise take the stage-1 weights and record the month as not cleared.

A7. Simulation of the two books

- **Book A (unadjusted)**. Reset to the (sleeve-rescaled) original weights at every month-end; multiplicative drift within the month. This matches the construction of the source portfolios and makes A its own rebalancing control: any difference between A and C is attributable to the macro overlay, not to rebalancing per se.
- **Book C (overlay)**. Set to the optimiser's weights at each month-end; identical drift within the month.
- **Timing**. Weights are computed on month-end data and applied from the next trading day. No information available after the rebalance date is used in the decision.
- **Costs**. No transaction, borrow or impact costs are modelled. Turnover is reported (1–5% of gross per month one-way) so costs can be layered on externally.

A8. Evaluation metrics

- Ex-ante: macro ES per rebalance date for both books; guardrail clearance rate; distribution of ES reductions.
- Realised macro component: each stock's daily beta $\times$ factor return, aggregated by book weights; reported as factor volatility, factor share of variance, and — the direct claim metric — the maximum drawdown of cumulative macro-component P&L.
- Tail-day test: the worst 5% of days ranked by the unadjusted book's factor P&L; cumulative total P&L of both books on those same dates. The risk estimator is weekly throughout; this test uses daily observations deliberately — it is an outcome measure, not a risk estimate, and daily granularity gives roughly five times the tail sample with no horizon assumption involved.
- Episode windows: the August 2024 shock, the 2025 tariff episode, the February–April 2026 Iran shock, and calendar 2022; total return and intra-window drawdown for both books.



- Whole-book statistics for context (annualised return, volatility, max drawdown), reported but not claimed.

A9. Explicit assumptions and known limitations

1. Static snapshots: books are single-date weight vectors held (with monthly resets) for up to 4.5 years. Any actively managed book decays under this treatment, so total-book statistics are context only; testing on a real weight history remains open.
2. Measurement instrument: the optimiser decides using MFERM betas as of each rebalance date; the realised macro component is measured using MFERM betas as they evolve day by day. Neither uses future information. But decision and score both rest on the same model's beta estimates: if a beta is simply wrong, the optimiser acts on it and the scoreboard credits it. Two beta-free checks address this directly: episode and tail-day results are computed on total P&L, which uses no betas; and the external-proxy regressions (A10) score both books on market data outside the model. Both agree with the model-based measure.
3. Overlapping weekly observations understate effective sample size; no small-sample correction is applied.
4. The 20%/80% split is a design input, not an estimate; guardrail reachability varies with book structure.
5. Arithmetic 5-day sums approximate weekly compounding; the error is second-order at daily magnitudes.
6. Sleeve rescaling after dropped names slightly alters position sizes relative to the delivered book (disclosed per book).
7. No costs, borrow, or capacity constraints; equities only, no derivative hedges — by design of the client question.
8. One model vintage per book; exposures are treated as point-in-time as delivered.



A10. External validation

Both books' daily total returns were regressed on seven external market proxies outside the MFERM model — SPX, changes in the US 10-year yield, KOSPI, HYG, crude, USDJPY and DXY — over each book's live window (contemporaneous OLS, no lags). Variance explained by the proxies fell from unadjusted to overlay in all five books, individually significant loadings shrank in the same direction, and a tail test ranking days by the externally fitted macro component showed smaller losses for the overlay book in all five. The near-unchanged result for the net-long book is consistent with its structural constraint. US-close proxies against Asian books understate fit through time-zone asynchrony, making those reductions conservative.

Appendix B — Implementing this on your own book

This is the operating manual: what to gather, what to compute, and what to do each month. A quant with the data below can reproduce the full process; a PM can run the lighter version at the end.

B1. Data you need

1. Your positions: identifier, signed weight, as of today.
2. Daily total returns for every position, at least three years plus a two-year burn-in (five years total).
3. Daily MFERM factor betas per stock over the same window (available as a standard export).
4. Daily MFERM factor returns over the same window.

One identifier convention across all files; blank (not zero) before listing; units stated once. These format rules caused every delay in the study — fix them at export.

B2. Build the two ES numbers (repeat monthly)

1. Form weekly observations: rolling 5-day sums of daily returns, for both stock returns and factor returns. Use the full history to date — the window expands, it never rolls.
2. Whole-book ES: apply today's weights to the history of weekly stock returns; take the worst 5% of those weekly outcomes; average them. That average loss is your whole-book ES.



3. Macro ES: take each stock's betas as of today; multiply by each historical week's factor returns to get that week's macro-driven return per stock; apply today's weights; take the worst 5% of the resulting weekly series; average them. That is your macro ES.
4. Sanity checks before trusting either number: at least two years of factor history including one real stress; betas winsorised at ± 10 ; every position covered or its exclusion disclosed.

B3. Set the guardrail — or choose not to

Guardrail = your chosen macro share $\times$ whole-book ES, refreshed each rebalance. 20% is this study's illustration; substitute whatever your mandate implies — the process is unchanged. Optionally add an amber line at 75% of the guardrail: no action required, but no new macro-additive positions until it clears. Running without a guardrail is also legitimate: every rebalance then takes the maximum reduction available. Expect permanently tilted weights and higher turnover in exchange — you are paying for macro reduction in months where there is little macro risk to reduce, which is precisely what the guardrail exists to avoid.

B4. The monthly decision

1. Macro ES at or below the guardrail $\rightarrow$ do nothing. Log the number and stop. This will be the most common outcome on diffuse, well-hedged books.
2. Above the guardrail $\rightarrow$ reallocate. Hold gross, net and the name list fixed; move each weight at most $\pm 35\%$ from its base; keep every sign.
3. Screen every candidate reallocation on two numbers: macro ES must fall (the binding test), and macro share of risk must not rise. A trade that cuts total risk without cutting macro risk — trimming a high-idio name is the classic — passes the first test and fails the second, leaving the book more macro-exposed. An unconstrained search returns exactly these trades by default.
4. If the guardrail is unreachable inside the bands, take the biggest reduction available and record it. Do not widen the bands mid-process to force a clear — that decision belongs to the mandate, not the month.



B5. Finding the weights: two routes

- **Full version (quant).** Solve it as a linear programme using the Rockafellar–Uryasev CVaR formulation: minimising empirical ES, and constraining it, are both exactly linear once one auxiliary variable and one slack per historical week are added. Stage 1 minimises macro ES subject to the constraints; if the floor clears the guardrail, stage 2 finds the smallest trade (minimum total weight moved) that just clears it. Any LP solver handles this at realistic book sizes in seconds.
- **Light version (PM, no optimiser).** Rank positions by their contribution to macro ES (contribution, not each stock's own macro share — a small position with a huge macro share may contribute little, and a short may be hedging you). Trim the largest positive contributors toward their -35% bound and add the freed weight to the largest negative or smallest contributors, sleeve by sleeve so gross and net hold. Re-measure. Two or three passes typically capture most of what the optimiser finds.

B6. After each rebalance and each month

1. Re-measure macro ES on the post-trade book before execution, not after — confirm the trade does what it claims.
2. Log four numbers monthly: macro ES, whole-book ES, the share, and turnover. The time series is the evidence your risk committee will ask for.
3. Quarterly, compare realised macro-driven P&L against what the model implied. Persistent misses mean the beta estimates need attention, not the book.

B7. What to expect, calibrated by book shape

- Diffuse market-neutral books (hundreds of names): guardrail typically reachable; turnover around 1% of gross per month; act rarely.
- Concentrated books (≤ 80 –200 names): expect roughly a 35–50% cut in macro ES per action month within the bands; the guardrail will clear only sometimes.
- Net-long books: the net position is itself a macro exposure no reweighting removes. Use the overlay for what it can do (compress the residual), and take the rest to a hedging discussion.
- Everywhere: protection arrives by holding less macro into the shock, not by reacting to it. A shock landing mid-month hits you at whatever exposure you brought to it — which is the point of keeping it low.



Macro Drawdown Control

Important notes. Results are historical simulations, not live performance, and do not represent any client account. Simulated and past performance are not reliable indicators of future results. This document is for information purposes only and does not constitute investment advice or an offer to buy or sell any security.



Disclaimer

This document is being sent only to investment professionals (as that term is defined in article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 ("FPO")) or to persons to whom it would otherwise be lawful to distribute it. Accordingly, persons who do not have professional experience in matters relating to investments should not rely on this document. The information contained herein is for general guidance and information only and is subject to amendment or correction.

This document is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation. This document is provided for information purposes only, is intended for your use only, and does not constitute an invitation or offer to subscribe for or purchase any securities, any product or any service and neither this document nor anything contained herein shall form the basis of any contract or commitment whatsoever.

This document does not constitute any recommendation regarding any securities, futures, derivatives or other investment products. The information contained herein is provided for informational and discussion purposes only and is not and may not be relied on in any manner as accounting, legal, tax, investment, regulatory or other advice. Information and opinions presented in this document have been obtained or derived from sources believed to be reliable, but Quant Insight Limited (QI) makes no representation as to their accuracy or completeness or reliability and expressly disclaims any liability, including incidental or consequential damages arising from errors in this publication.

No reliance may be placed for any purpose on the information and opinions contained in this document. No representation, warranty or undertaking, express or implied, is given as to the accuracy or completeness of the information or opinions contained in this document by any of QI, its employees or affiliates and no liability is accepted by such persons for the accuracy or completeness of any such information or opinions.

Any data provided in this document indicating past performance is not a reliable indicator of future returns/performance. Nothing contained herein shall be relied upon as a promise or representation whether as to past or future performance. This presentation is strictly confidential and may not be reproduced or redistributed in whole or in part nor may its contents be disclosed to any other person under any circumstances without the express permission of Quant Insight Limited.