

Momentum Fusion, *Through the STORM*

A structured distillation of J.P. Morgan's European Momentum Fusion framework (Hanif, Summer & Chaudhry, 29 June 2026) — mapped onto the STORM protocol, stress-tested against the AQI research stack, and mined for what we should adopt, what we should challenge, and where the two houses genuinely part ways.

DATE 11 JUL 2026 SOURCE JPM GLOBAL QUANT — EUROPE UNIVERSE MSCI EUROPE, 1996–2026

CLASSIFICATION INTERNAL — AQI

EXECUTIVE SUMMARY

One premium, many labels, and a very polite civil war

JPMorgan's European quant team tested 37 momentum signals across six families over thirty years of MSCI Europe data, blended the survivors into a new Momentum Composite, ran it head-to-head against their own incumbent Momentum Style, and concluded — after producing a 5.29 t-stat of incremental predictive power — that no framework change is warranted. That sentence is the whole note in miniature: rigorous signal engineering, genuinely strong evidence, and a recommendation shaped as much by institutional governance as by statistics. It is excellent work, and it is work built on a different chassis than ours.

The headline machinery: a **Momentum Composite** (equal-weight blend of vol-adjusted 12M–1M price momentum, high-to-price momentum, 6M earnings momentum, and 3M net cashflow revisions) delivering 14.3% annualized long/short at a 0.88 Sharpe against the incumbent Style's 10.6% at 0.75. A crash-risk toolkit horse race where regime screening (QMI) wins on Sharpe, volatility thresholds win on drawdown, and the humble stoploss is sent home without dinner. And a drawdown anatomy showing crowding rising +247.6% and volatility *falling* 10.5% into the six months before momentum crashes — the volatility paradox, quantified.

Our read, developed in full below: the framework converges with AQI on the destination (drawdown-first evaluation, composite construction, regime awareness) while diverging on the vehicle. JPM harvests momentum *unconditionally* and bolts risk management on afterward; AQI conditions entry on events and sizes with Kelly-Thorp, which makes several of their overlays redundant and one of their conclusions — the stoploss verdict — vulnerable to exactly the counterfactual-asymmetry critique we formalized in AQI-WP-2026-001.

THE BOTTOM LINE

Read it as a parts catalog, not a blueprint. Steal the drawdown-first scorecard, the cashflow-revisions confirm signal, the transcript-over-news NLP finding, and the volatility-paradox early warning. Decline the fixed vol targets, the return-based crowding proxy, and the realized-path tool grading. The chassis stays Kelly-Thorp.

PART I — THE DISTILLATION

The JPM framework, refactored through STORM

JPM's note is organized as a signal taxonomy followed by a risk-tool bake-off. Refactored into our protocol, its actual argument structure looks like this.

S

SITUATION

Momentum works everywhere and crashes everywhere

The premium is the most replicated result in cross-sectional finance — Jegadeesh & Titman (1993) forward. JPM's contribution to the situation-setting is scope: 37 signals across six families (price momentum, technical, linear trend, fundamental revisions, analyst sentiment, NLP), backtested monthly on MSCI Europe from January 1996 through March 2026, long-only and long/short. Two families emerge as the load-bearing walls:

- **Technical momentum** — what the market is doing. Best expression: 12M–1M price momentum scaled by 3-month volatility (12.2% L/S, 0.75 Sharpe).
- **Fundamental momentum** — whether the market has a reason to keep doing it. Best expression, and the quiet star of the entire note: **Net Cashflow Revisions 3M** (12.2% L/S, 1.53 Sharpe, 6.76 t-stat, max drawdown of just –14%).

The two families capture different phases of information diffusion — market action versus analyst reaction — which is what makes them blendable rather than redundant.

T

TENSION

Three tensions, one of them political

Return versus ruin. Every one of the 37 factors suffered at least one >9.7% drawdown event over 2005–2025; the mean worst drawdown across the universe was –40.9%, the worst single episode –82.3%. Momentum's crash risk is structural, not incidental. Worse, the pre-crash signature is perverse: returns elevated and *accelerating*, crowding nearly tripling, and volatility *compressing* into the event. Every standard risk gauge reads calm precisely when the strategy is most fragile.

Signal proliferation versus independence. Thirty-seven signals sounds like diversification; JPM's own footnote concedes cross-factor correlations remain "exceptionally high for what should be random variables" — i.e., one underlying premium wearing many costumes. How many independent bets actually exist here is the note's least-examined question.

New Composite versus incumbent Style. The political tension. The European Momentum Style is embedded in JPM's live process with a long research history. Any challenger must clear not a statistical bar but an institutional one: demonstrate incremental efficacy against a model the franchise already ships.



The menu, as tested

SIGNAL CONSTRUCTION OPTIONS – MSCI EUROPE, L/S ANNUALIZED

Construction	Ann Ret	Sharpe	Max DD	Turnover
Incumbent Momentum Style	10.6%	0.75	–46%	42%
Momentum Composite (new)	14.3%	0.88	–45%	30%
50:50 Style/Composite blend	9.8%	0.70	–47%	44%
Factor-level blend	13.0%	0.93	–40%	30%
Net Cashflow Revisions 3M (standalone)	12.2%	1.53	–14%	31%

CRASH-MITIGATION OPTIONS – THE HORSE RACE, EUROPE

Tool	Sharpe	Max DD	Ann Ret	Verdict
Stoploss (10% SL)	0.44	–47%	8.4%	Wins no category
Vol threshold (30%)	0.54	–31%	9.6%	Best drawdown insulation
QMI screen (skip Recovery)	0.63	–31%	8.9%	Best risk-adjusted; wins 5 of 6 universes

Plus the sizing layer: vol scaling at 5% or 10% targets versus per-leg beta scaling. Vol scaling at 10% lifts the Composite's L/S Sharpe to 1.27 with a –25% drawdown and near-normal tails; beta scaling buys more raw return (19.4%) at the price of 16.6 kurtosis and a –40% drawdown — conditional leverage that levers into precisely the regimes where momentum fails.



Keep the incumbent; upgrade at the margins

JPM's stated recommendations, faithfully rendered:

1. **Keep the European Style as-is** — "no framework change warranted." The opportunity sits in a factor-level blend, explicitly not a 50:50 mix (which merely dilutes Style).

2. **Vol scaling as the default sizing rule:** 10% target for long/short mandates, 5% where budgets are tighter, and on the long side for any long-only deployment. Beta scaling only where market-neutrality is a hard constraint.
3. **QMI screening for return efficiency, volatility thresholds for downside insulation, and stoploss rules for neither.** Skipping the Recovery regime is the single most powerful configuration — in Europe, skipping Expansion helps too, while skipping Slowdown or Contraction is destructive (total returns collapse to 82% and 66% versus a 301% baseline).

THE TELL

The same note that calls the Composite "the superior signal for portfolio construction" and reports its incremental Fama-MacBeth gamma at a Newey-West t-stat of 5.01 concludes that the framework should not change. Both statements can be true — governance is a real constraint — but the asymmetry in where the burden of proof lands is doing unacknowledged work.



Where the alpha survives, or doesn't

- **Turnover discipline.** The factor blend runs ~30% L/S turnover versus 42–44% for Style and the 50:50 mix — the blend is cheaper to run than the thing it improves on.
- **Statistical hygiene.** Bivariate Fama-MacBeth with Newey-West errors; VIF diagnostics (median 1.04, max 1.30) confirming the two signals can coexist in one model. The signals' cross-sectional correlation is persistently *negative* in the monthly regression frame — genuine diversification, not relabeling.
- **Capacity caveat.** Vol scaling earns its Sharpe by varying exposure through time, making financing terms and capacity first-order for 5–10% vol sleeves running 21–42% turnover.
- **Re-entry rules dominate.** Their stoploss discussion concedes that implementation details "often dominate outcomes once you move from backtest to live trading" — the most Kelly-adjacent sentence in the document.
- **Equal-weight over optimized.** A footnote notes the FMB gammas imply overweighting the new Composite, then declines to do so on information-leakage grounds. Correct call, quietly made.

Where the houses agree

Before the divergences, the convergences — because they're substantial, and they validate several standing AQI positions:

- **Drawdown is the metric that tells you whether a backtest is livable.** JPM's scorecard scores signals on Sharpe *and* drawdown jointly, arguing the best signal is the one an investor can actually hold through adversity. This is the philosophical core of our Kelly-Thorp system — drawdown is the binding constraint on the leverage a real edge can support.
- **Composites over hero signals.** Orthogonal technical and fundamental confirmation beats maximizing exposure to the single strongest variant. Our event studies pair price structure with breadth and fundamental context for the same reason.
- **Regime conditioning is not optional.** Their QMI finding — identical pre-crash signatures mean opposite things in Contraction versus Recovery — is the same lesson our DXY regime study delivered across the ETF universe: the signal doesn't change, the regime changes what the signal means.
- **Backward-looking triggers arrive late.** By T-1 the deterioration (acceleration, crowding, vol compression) is already well advanced. Forward-looking conditioning beats reactive exits. Agreed, with an asterisk we'll get to.

Unconditional harvesting vs. conditional event studies

JPM

Hold the top decile always; manage the crashes with overlays bolted on after construction.

AQI

The condition is the signal. Enter after a defined event (BORD, failed breakdown, reclaim-breakout) with a mapped forward-return distribution; no event, no position.

This is the deepest fork. JPM answers "when does momentum pay?" with a regime label — Contraction, Slowdown. Our event-study program answers it with a price event and a conditional forward-return distribution. Both are conditioning; theirs is macro and slow (QMI phases persist for quarters), ours is micro and episodic. The practical consequence: their framework must carry momentum exposure through the pre-crash fragility window and hope the overlay catches it, while an event-conditioned book is frequently flat by construction when no setup exists. The frameworks are complementary — QMI-style phase attribution is a legitimate macro filter on top of event entries — but they are not the same discipline, and the note's overlay complexity is partly the cost of refusing to condition at entry.

Fixed vol targets vs. Kelly-Thorp dynamic sizing

JPM

Constant risk budget: scale positions to hold annualized vol at 5% or 10%, monthly, using trailing realized volatility.

AQI

Fractional Kelly: exposure proportional to edge over variance. Both the numerator and denominator move.

Vol targeting is Kelly with the numerator frozen. It works — their own tables show the 10% target lifting the Composite Sharpe from **0.89** to **1.27** while normalizing the tails — but it contains a contradiction the note itself documents and doesn't resolve. Their drawdown anatomy shows volatility *falling* 10.5% into the six months before crashes. A trailing-vol denominator therefore instructs the strategy to *lever up* into exactly the fragility window; the note even concedes a volatility-based exit "would miss the key signal." Kelly sizing routes around this: when crowding triples and acceleration turns convex, the forward edge estimate decays, and the position shrinks *before* realized vol confirms anything. JPM's vol scaling manages the crash; edge-conditioned sizing declines the invitation. Our tranche state machine already encodes this — the note is, unintentionally, a 51-page argument for it.

The stoploss verdict and the Flummox Problem

JPM

"Stoploss rules for neither." The tool wins no category in the realized horse race and is dismissed framework-wide.

AQI

Counterfactual Symmetry (AQI-WP-2026-001): grade the decision rule across the distribution of paths, not the one path history happened to deal.

Here's the uncomfortable table entry: the 10% stoploss produced the *highest total return in its entire section* — 1043% versus 417% unhedged — and a better Sharpe than baseline, and it still gets banished because in the three-way category race it never finishes first. That is outcome-graded tool evaluation on a single realized path, which is precisely the Seahawks play-call error in institutional dress. The relevant question is not "which tool won 1996–2026" but "which tool's conditional performance distribution do you want to own across the paths that could have occurred?" A stoploss's value is concentrated in left-tail paths that are, by definition, underrepresented in any single sample. The same critique cuts at the QMI verdict from the other side: "skip Recovery" is graded on regime labels assigned with full-sample information — the note never establishes that QMI phase classification was knowable in real time without revision. We'd want vintage-data regime calls before promoting that from attribution exercise to trading rule.

Return-based crowding vs. breadth-based crowding

JPM

Crowding proxy = rolling 6M cumulative factor return relative to benchmark. It rises +247.6% into drawdowns.

AQI

Measure crowding from participation, not from the factor's own P&L. Cap-weighted breadth (CWB) reads narrowing leadership directly.

The +247.6% crowding build-up is the note's most dramatic exhibit and its most circular. Defining crowding as recent factor outperformance guarantees that crowding "predicts" crashes, because crashes are definitionally preceded by the run-up that defines the proxy. It's not wrong, but it's partially tautological. The CWB oscillator attacks the same phenomenon from an independent axis: when a momentum advance is carried by fewer and fewer names while the cap-weighted index grinds on, participation narrows before the factor return series shows anything. Pairing JPM's return-path diagnostics (acceleration, vol compression) with a holdings- and breadth-based crowding read gives a two-factor confirmation where they currently have one factor measured twice.

Scorecard selection vs. multiplicity discipline

JPM

Rank 37 signals on a z-scored Sharpe-and-drawdown scorecard; blend the family winners; validate in-sample with FMB.

AQI

Block-bootstrap Monte Carlo on the selected composite; report the effective number of independent bets, not the nominal signal count.

Selecting the best of 37 correlated signals and then measuring the winner on the same 1996–2026 window is a selection-bias machine, however honest the scorecard. JPM's own footnote flags the smoking gun — factor correlations "exceptionally high," the mispricing source "fundamentally the same across factors." Translation: the effective bet count is closer to two or three (a price-persistence bet, a revisions bet, arguably an anchoring bet) than to thirty-seven. The STORM protocol's block-bootstrap layer exists exactly for this: resample the selection *and* the evaluation jointly and see whether the Composite's edge over Style survives. Our prior, given the negative cross-signal correlation in their FMB frame, is that it does survive — but at a haircut to those 5-handle t-stats, and the haircut is the number a position-sizer actually needs.

The incumbent's thumb on the scale

JPM

Composite wins on return, Sharpe, reliability, IC stability, and marginal FMB contribution — conclusion: keep Style, no framework change warranted.

AQI

Unburdened by a legacy franchise: run the factor-level blend (0.93 Sharpe, -40% DD, 30% turnover) as the default momentum expression.

To be fair to the authors, "keep the framework" is defensible: live research history has value, model governance has costs, and the factor-level blend they endorse quietly imports most of the Composite's content anyway. But notice the shape of the reasoning — the challenger needed a 5.29 t-stat, superior reliability through the 2018–2021 inversion, *and* better drawdown behavior merely to earn a seat as a "complement." That's escalation-of-commitment architecture: the burden of proof scales with the incumbency of the incumbent, not with the evidence. For a shop like ours with no legacy Style to defend, the correct reading of their own tables is simpler: the factor blend is the best object in the document and should be the starting point, not the diplomatic compromise.

Regime maps don't travel

JPM

Global prescription: momentum overweight in Expansion and Slowdown. Europe inverts it — Expansion momentum is muted (12M PMOM: 0.16%), Contraction is the strongest phase (21.96%).

AQI

Regime conditioning is universe-specific and must be re-estimated per geography, per asset class — the DXY regime study lesson.

The Europe QMI table is the note's best supporting evidence for a position we already hold: you cannot port a regime map. Europe's defensive sector composition lets quality names sustain relative outperformance through downturns, so trend-following *thrives under macro stress* there — the opposite of the global pattern, and the reason skipping Slowdown or Contraction destroys the European strategy while skipping Recovery rescues it. For the Query Breadth engine this is a design requirement, not a footnote: regime layers get estimated within-universe, with the phase-attribution harness (switch off one regime at a time) as the standard diagnostic.

What AQI should steal, starting this quarter

Divergence isn't dismissal. Five components of this note earn immediate adoption into the stack — several ahead of the August 1 beta sprint.

01 The drawdown-first scorecard, ported to the Query Breadth engine

Z-score candidate signals jointly on risk-adjusted return and max drawdown, long and L/S dimensions aggregated. Cheap to implement, philosophically aligned with Kelly (drawdown bounds sustainable leverage), and a better factory gate than Sharpe-only ranking for anything feeding the paper-trading loop.

02 Net Cashflow Revisions 3M as a fundamental confirm layer

The cleanest object in the entire document: 1.53 L/S Sharpe, 6.76 t-stat, -14% max drawdown, ~31% turnover. Breadth-of-revisions construction (count of upward minus downward CFPS estimate changes over total) makes it robust to single-analyst noise. Natural pairing: require positive cashflow-revision breadth as a filter on technical event-study entries. FactSet estimates data covers this today.

03 Transcript NLP over news NLP — a pipeline decision

Their finding is unambiguous: earnings-call transcript sentiment carries real alpha (0.70 L/S Sharpe, -12% DD) while news-flow sentiment is noise (-0.13 Sharpe L/S). For our MT Newswires / FactSet ingestion priorities, this reorders the queue: transcript-derived features first, headline sentiment demoted to context.

04 The volatility-paradox fragility flag

Encode the pre-crash signature directly: *low and falling* realized vol + elevated and *accelerating* returns + narrowing breadth (our CWB substitution for their crowding proxy) = fragility state. Wire it as a de-risking input to the Kelly tranche state machine — an edge-decay signal, not an exit trigger.

The switch-off-one-phase exercise is elegant and generalizes: for any strategy in paper trading, systematically zero exposure during one regime at a time (theirs: QMI phases; ours: DXY regimes, El Niño-linked commodity seasons, breadth states) and observe where the drawdown risk concentrates. Attribution first, trading rule only after vintage-data validation — see Divergence 03.

PART V – SIDE BY SIDE

The STORM verdict, letter by letter

TWO HOUSES, ONE FACTOR

Letter	JPM's answer	AQI's answer
Situation	Momentum premium is universal; crash risk is structural; 37 signals, 6 families, 30 years.	Agreed on the physics. Add: the premium is episodic at the security level – which is where we hunt.
Tension	Return vs. ruin; signal redundancy; challenger vs. incumbent franchise.	Same three, plus a fourth they under-weigh: realized-path evaluation vs. counterfactual robustness.
Options	Composites, blends, and a three-tool overlay race graded on one historical path.	Event-conditioned entries + Kelly sizing collapse much of the overlay menu; overlays audition as edge inputs, not exits.
Recommendation	Keep Style; factor-level blend at the margin; vol-scale; QMI for efficiency, vol thresholds for insulation, stoploss for neither.	Run the factor blend as default; size with fractional Kelly; adopt the fragility flag; hold the stoploss verdict for counterfactual retrieval.
Mechanics	Turnover discipline, Newey-West rigor, VIF hygiene, capacity caveats – genuinely first-rate.	Import wholesale, add block-bootstrap multiplicity control and vintage-data regime validation.

Craft note, offered with affection: the document's prose and tables occasionally disagree with each other — the executive summary quotes the Style at 10.3%/0.72 while Table 15 says 10.6%/0.75; the Composite

scaling section narrates an unscaled 0.72 Sharpe that Table 21 reports as 0.89. Different sample starts (1994 vs. 1996) likely explain some of it, but a note this good deserves reconciled exhibits. We hold our own STORM reports to single-source-of-truth numbers for exactly this reason.

APPENDIX — WORKING QUEUE

Proposed follow-ons

1. **Replication sketch (US universe):** rebuild the four-factor Composite (vol-adj 12M–1M, high-to-price, 6M earnings momentum, 3M net cashflow revisions) on a US large-cap universe via FactSet, score it on our drawdown-first card, and compare against SPMO as the live-momentum benchmark we already track.
2. **Counterfactual stoploss study:** block-bootstrap the 12M momentum series and re-run the stoploss/vol-threshold/regime-screen horse race across resampled paths. Deliverable: distributional grades per tool — the Flummox correction to Table 27–29.
3. **CWB × crowding cross-validation:** test whether cap-weighted breadth deterioration leads their return-based crowding proxy at momentum-factor turning points. If it leads, the fragility flag gets a genuinely independent second input.
4. **Fragility-state spec for the Kelly engine:** formalize the low-vol / accelerating-return / narrowing-breadth state as an edge-decay multiplier in the tranche state machine ahead of IBKR paper trading.

ALPHAQUERY INTELLIGENCE

AQI-RN-2026-071 · Momentum Fusion Through the STORM Lens · 11 July 2026
Prepared for internal research use. Source under review: J.P. Morgan Global Markets Strategy, "Momentum Fusion: European Equity Trend-Following," 29 June 2026 (Hanif, Summer, Chaudhry). All source metrics paraphrased and attributed; interpretive claims are AQI's own.