An Inquiry Into Volatility Regime Based Trading Strategies

Balázs Peisz
balazs.peisz@epfl.ch

Gaël Saugy
gael.saugy@epfl.ch

Axel Turin-Plessia
axel.turin-plessia@epfl.ch

April 2026

Abstract

In this paper we test whether a market volatility regime signal can improve a systematic equity strategy. The signal comes from a two-state Gaussian hidden Markov model, trained on 1993–2009 data and run out of sample as a one-step-ahead filter, which gives a daily probability that the market is in the high-volatility state. The design rests on three hypotheses: that volatility states are identifiable and persistent, that the factor premia we trade are state dependent, and that the conditioning can be implemented within transaction costs. The paper presents two strategies based on this same signal. The first switches discretely between a trend sleeve and a short-horizon reversal sleeve, while the second stays fully invested and slides between a cross-sectional momentum sleeve and a minimum-volatility sleeve in proportion to the regime probability. Both are evaluated on a 2017–2025 holdout set aside before any tuning, against three controls: the same two sleeves blended 50/50 with no signal, the same signal with its timing destroyed, and equal-weighted benchmarks drawn from the same securities. The continuous architecture returns 16.74% annualised against −15.33% for the discrete one, and turns over roughly a third as much, because the two sleeves hold many of the same stocks and their overlapping positions net out as the regime probability moves. Neither architecture separates from its controls, and no variant beats a passive equal weight of the same universe. The explanation is visible in the sleeve returns themselves since over the test period the two sleeves finished 26 basis points apart, leaving a rotation rule almost nothing to exploit. With further elaboration we conclude that the state dependence a regime strategy relies on should be established before a portfolio is built around it.

1 Introduction

Strategies that condition portfolio construction on an estimated market state are a well established response to the non-stationarity of asset returns (Hamilton, 1989; Ang and Bekaert, 2002; Nystrup et al., 2015). The premise is rather simple, if calm and turbulent markets reward different behaviour, and the prevailing state can be identified in close to real time, then an investor should hold a different portfolio in each scenarios.

This paper reports the design, implementation and assessment of such a strategy. Section 2 specifies the two portfolio architectures we constructed. Section 3 develops the theoretical framework underlying the regime model with section 4 stating the economic rationale as three testable hypotheses. Sections 5 and 6 describe the implementation and the experimental protocol whilst section 7 presents the results. Section 8 interprets them, while Section 9 sets out directions for further work.

Our findings separate cleanly along two axes. The architecture question and the signal question have different answers because replacing a discrete switch with a continuous slider improved annualised return by 32.1 percentage points and reduced turnover by a factor of 3.5, with the netting mechanism behaving exactly as derived, while the regime signal itself did not contribute measurably. On the holdout the slider does not separate from a static blend of its own sleeves, nor from a placebo in which the signal's timing has been destroyed. We argue in

Section 8 that the principal reason behind this is attributed neither to the model nor to the architecture, but to the choice of sleeves between which the strategy rotates.

2 Strategy Architecture

2.1 Common structure

Both implemented architectures share a common form. A regime model produces $\pi_t \in [0, 1]$, the probability that the market occupies a high-volatility state on day $t + 1$ conditional on information through day t. Two portfolio sleeves are constructed independently, one intended for calm markets and one for turbulent markets, and an orchestration rule maps π_t onto a combination of the two. (Weights formed on day t are applied to day $t + 1$ returns.)

The architectures differ in the composition of the sleeves and in how the orchestration rule takes π_t.

2.2 Architecture One: discrete directional switching

The first architecture holds exactly one sleeve at a time.

- **Trend sleeve (calm markets).** Take a 20/50-day moving-average crossover on the small-cap panel, taking long position when the fast average exceeds the slow average and taking short otherwise.
- **Reversal sleeve (turbulent markets).** Apply short-horizon mean reversion on the large-cap panel with each security's daily return is standardised against its own five-day rolling mean and standard deviation. Here, the rule enters long below $-\theta$, short above $+\theta$, and exits at zero, with $\theta = 1$.
- **Orchestration.** A hysteresis rule transfers capital to the reversal sleeve when $\pi_t \geq 0.90$ and back when $\pi_t \leq 0.10$. On every confirmed switch the existing book is fully liquidated before the new one is established.

2.3 Architecture Two: continuous factor rotation

The second architecture remains fully invested and uses π_t to set the portfolio's risk characteristics rather than its direction:

$$w_t \;=\; \pi_t\, w_t^{\text{def}} \;+\; (1 - \pi_t)\, w_t^{\text{mom}}. \tag{1}$$

- **Risk-on sleeve $w^{\mathbf{mom}}$.** Equal weights on the 20 securities with the highest trailing 126-day price return, ranked across the whole universe.
- **Risk-off sleeve $w^{\mathbf{def}}$.** Equal weights on the 20 large caps with the lowest 63-day realised volatility of daily returns.
- **Orchestration.** The convex combination of Equation (1). Since both sleeves are long-only and fully invested, $\sum_i w_{i,t} = 1$ for every π_t, so the portfolio never holds cash.

3 Theoretical Framework

3.1 Hidden state models of market dynamics

Volatility clustering is one of the most robust regularity in returns which is addressed by two distinct modelling traditions. Generalised autoregressive conditional heteroskedasticity models treat conditional variance as a deterministic function of past shocks, thus yielding a continuously varying volatility estimate. Whereas regime-switching models posit that the market occupies one of a small number of discrete latent states, each with its own return distribution, with transitions between them governed by a Markov chain (Hamilton, 1989).

We adopt the latter for three application specific reasons. It yields a *probability* of occupying each state, which Equation (1) takes as a direct portfolio weight. Its transition matrix parameterises persistence explicitly, so expected state duration is a quantity that can be inspected and validated rather than an emergent property of the fitted model. The hidden Markov formulation also accommodates a multivariate observation vector naturally, allowing realised volatility, implied volatility and volume to be combined within a single filter.

3.2 Model specification

Let $S_t \in \{1, \ldots, K\}$ denote the latent state at time t, evolving as a homogeneous first-order Markov chain with transition matrix

$$A_{ij} = P(S_t = j \mid S_{t-1} = i), \qquad \sum_{j=1}^{K} A_{ij} = 1, \tag{2}$$

and initial distribution $\rho_j = P(S_1 = j)$. Conditional on the state, the observation $X_t \in \mathbb{R}^d$ is drawn from a state-dependent density

$$b_j(X_t) = p(X_t \mid S_t = j) = \mathcal{N}(X_t; \mu_j, \Sigma_j). \tag{3}$$

The complete parameter set is $\Theta = \{\rho, A, \{\mu_j, \Sigma_j\}_{j=1}^{K}\}$.

Two properties follow directly from the specification and are used later. First, the expected sojourn time in state i is geometrically distributed with mean

$$\mathbb{E}[\text{duration in } i] = \frac{1}{1 - A_{ii}}, \tag{4}$$

so persistence is read off the diagonal of A. Second, if A is irreducible and aperiodic, the chain admits a unique stationary distribution φ satisfying $\varphi A = \varphi$, giving the long-run share of time spent in each state.

3.3 Filtering, forecasting and smoothing

Given Θ, inference about the unobserved state proceeds by recursion. We define the filtered distribution as $\alpha_t(j) = P(S_t = j \mid X_{1:t})$. Therefore the forward recursion is

$$\alpha_t(j) \;\propto\; b_j(X_t) \sum_{i=1}^{K} \alpha_{t-1}(i) A_{ij}, \qquad \alpha_1(j) \propto \rho_j \, b_j(X_1), \tag{5}$$

with each step normalised to sum to one. Then the one-step-ahead forecast follows immediately as

$$P(S_{t+1} = j \mid X_{1:t}) = \sum_{i=1}^{K} \alpha_t(i) A_{ij}, \tag{6}$$

and it is this quantity, evaluated at the high-volatility state, that constitutes the signal π_t driving both architectures.

Another important factor in our strategy's inner working is the smoothed distribution $\gamma_t(j) = P(S_t = j \mid X_{1:T})$ which conditions on the entire sample and is obtained by combining Equation (5) with a backward recursion. Smoothing produces a substantially cleaner regime path, because it resolves ambiguity at time t using observations from $t+1$ onward, thus using it to form portfolio weights is therefore a look-ahead error. Among the these α_t and Equation (6) are admissible signals, and we use the latter exclusively.

3.4 Parameter estimation

The likelihood $p(X_{1:T} \mid \Theta)$ cannot be maximised in closed form because the state sequence is unobserved. The Baum–Welch algorithm, an instance of expectation maximisation, addresses this by alternating between computing state responsibilities under current parameters and re-estimating parameters given those responsibilities (Rabiner, 1989). Writing $\xi_t(i,j) = P(S_{t-1} = i, S_t = j \mid X_{1:T})$ for the pairwise smoothed probabilities, the update equations are

$$\hat{A}_{ij} = \frac{\sum_{t=2}^{T} \xi_t(i,j)}{\sum_{t=2}^{T} \gamma_{t-1}(i)}, \qquad \hat{\mu}_j = \frac{\sum_{t=1}^{T} \gamma_t(j) X_t}{\sum_{t=1}^{T} \gamma_t(j)}, \qquad \hat{\Sigma}_j = \frac{\sum_{t=1}^{T} \gamma_t(j)(X_t - \hat{\mu}_j)(X_t - \hat{\mu}_j)^\top}{\sum_{t=1}^{T} \gamma_t(j)}. \quad (7)$$

Each iteration increases the likelihood monotonically, but convergence is only to a local optimum, so initialisation matters. Recursions are carried out in log space to avoid numerical underflow over samples of several thousand observations.

3.5 Specification choices and their consequences

Three choices constrain what the fitted model can represent which are the following:

Number of states. We fix $K = 2$, because this is the smallest specification consistent with the calm-versus-turbulent distinction the strategy requires, and Equation (1) consumes a single scalar probability, which a two-state model supplies directly. Larger K would require an additional mapping from the state simplex onto sleeve weights.

Covariance structure. We also restrict Σ_j to be diagonal, which is a genuine simplification, since realised and implied volatility are strongly correlated, but it reduces the free parameter count to $K(2d) + K(K-1) + (K-1) = 23$ for $K = 2$, $d = 5$, which materially improves estimation stability on a sample of this size.

Label identifiability. Expectation maximisation is invariant to permutation of the state labels, so the correspondence between fitted states and economic interpretation is not determined by the algorithm. We therefore assign labels *ex post* by comparing responsibility-weighted feature means across states, designating as high-volatility whichever state exhibits the larger average across the volatility-related features. This renders the labelling reproducible and independent of initialisation.

4 Economic Rationale and Research Hypotheses

The strategy rests on three hypotheses, stated here so that the results of Section 7 can be read against them afterwards.

4.1 H1: volatility states are identifiable and persistent

The framework of Section 3 is only useful if the estimated states are economically distinct and sufficiently persistent so a portfolio can be repositioned before the state changes again. Nystrup et al. (2015) apply a closely related two-state Gaussian hidden Markov model to asset allocation and report that it adds value over static weights after transaction costs, which establishes the specification as a reasonable starting point. Equation (4) makes persistence directly testable: the diagonal of the estimated transition matrix should imply state durations long relative to the rebalancing interval.

4.2 H2: the relevant factor premia are state dependent

This is the load-bearing hypothesis of our approach and to our best knowledge the empirical literature supports each of its components as detailed in the following.

Momentum weakens, and crashes, in turbulent states. Jegadeesh and Titman (1993) document cross-sectional momentum with the leading explanations being behavioural, involving gradual information diffusion and investor underreaction (Hong and Stein, 1999). Critically for our design, Daniel and Moskowitz (2016) show that momentum crashes are partially forecastable and occur specifically in panic states — following market declines, when market volatility is elevated, simultaneously with market rebounds. This means that a momentum sleeve is appropriate for the calm state and inappropriate for the turbulent state.

Returns to liquidity provision increase with volatility. At daily-to-weekly horizons return autocorrelation turns negative (De Bondt and Thaler, 1985; Lehmann, 1990). Nagel (2012) interprets short-horizon reversal returns as a proxy for the returns to liquidity provision and demonstrates that these are highly predictable using the VIX, with expected returns and conditional Sharpe ratios rising sharply during periods of market turmoil. This constitutes near-direct support for Architecture One's reversal sleeve since the premium the design seeks to harvest is documented to be strongest precisely when the regime model indicates a switch into it.

Low-volatility equities provide defensive exposure without exiting the asset class. Low-volatility portfolios have historically delivered risk-adjusted returns at least comparable to the market, contradicting the prediction of the capital asset pricing model (Blitz and van Vliet, 2007; Baker et al., 2011); Frazzini and Pedersen (2014) attribute this to leverage constraints. For Architecture Two the relevant property is that a minimum-volatility book is a low-beta exposure that remains an equity exposure, and therefore participates in recoveries in a way cash does not.

Based on these findings, H2 predicts that a portfolio tilting between momentum and defensive equity as a function of the estimated state should dominate a fixed allocation to either.

4.3 H3: the conditioning is implementable within transaction costs

Anomaly returns measured gross of costs frequently fail to survive net of them (Korajczyk and Sadka, 2004; Novy-Marx and Velikov, 2016). For regime-conditioned strategies this is a first-order design constraint rather than an accounting detail, because a rule reallocating on every state change incurs turnover proportional to the number of state *crossings* — a quantity sensitive to the classification threshold and potentially large even when the underlying state is persistent.

Architecture Two is our response to this constraint. Let us consider a security j held by both sleeves. Its blended weight under Equation (1) is

$$w_{j,t} \;=\; \pi_t \cdot \frac{1}{N} \;+\; (1-\pi_t) \cdot \frac{1}{N} \;=\; \frac{1}{N}, \tag{8}$$

independent of π_t. A security in the intersection of the two sleeves generates no turnover however sharply the regime probability moves. Only securities which are held by exactly one sleeve trade, and they trade in proportion to $\Delta\pi_t$ rather than discretely at a threshold crossing. This means that turnover scales with a smooth quantity rather than a more jagged one, and declines further in proportion to sleeve overlap.

5 Data and Implementation

5.1 Investment universe

The traded universe comprises of 41 US equities of which 29 are large caps and 12 are small caps loaded as per-security daily files. Prices are split-adjusted close prices without dividend reinvestment, so all returns reported here are **price returns**. We apply this convention to all the strategies and benchmarks as well.

There are two properties of the universe which are quite apparent. First, it is rather **narrow** because both sleeves in Architecture Two hold 20 securities, selecting 20 from 41 and 20 from 29 respectively, which limits cross-sectional ranking. Second, it is **survivorship-selected** since both ticker lists were specified manually and filtered on market capitalisation and average volume as per construction time. This means that securities which appear in the large-cap panel are considered as such because they were large caps at the time of construction rather than at the start of the sample. This inflates return levels for all series, active and passive alike, and is the reason we rely on the ordering of results rather than their levels.

5.2 Regime model inputs and return conventions

The market return series driving the regime model is cached in `portfolio_returns_equal_weight.csv`, which holds separate large-cap and small-cap columns. It extends to 1980 and therefore supports a long training window, whereas the traded panels begin in 2010. Apart from that, VIX closes and SPY volume complete the feature set.

It is important to mention about this data is that its columns are cross-sectional means of *logarithmic* returns, $\frac{1}{N}\sum_i \log(P_{i,t}/P_{i,t-1})$. This is a well-behaved input to a volatility filter – symmetric, additive in time, and with dispersion equal to the quantity the model characterises – but it is not a portfolio return. An equal-weighted portfolio earns the cross-sectional mean of *simple* returns, and the two differ by approximately half the average idiosyncratic variance: on our panels, 5.4 percentage points per annum for the large caps and 7.4 for the small caps. We therefore use this series for regime estimation only and construct benchmarks directly from the traded panels.

5.3 Benchmark construction

Table 1: Annualised price return of candidate benchmark series, 2010–2025.

Series	Annualised return
Equal weight of the traded universe (41 securities)	18.52%
Equal weight of the large-cap panel (29 securities)	20.36%
Equal weight of the small-cap panel (12 securities)	13.07%
Cached market series used for regime estimation	10.61%

All four rows cover the same 41 securities; the dispersion is entirely attributable to construction. The cached series trails an equal weight of the traded universe by 7.9 percentage points per annum for two compounding reasons: it is a mean of logarithmic rather than simple returns, and it weights the two legs 50/50 where the traded universe is 71/29 by security count. The cached columns reproduce the panel log-return means to six decimal places with a correlation of 1.0000, confirming that no constituent difference is involved.

We report against the first two rows of Table 1. The large-cap benchmark is required because the minimum-volatility sleeve draws exclusively from the large-cap panel and a comparison against the 41-security universe would confound a low-volatility effect with a size tilt. Sharpe ratios use daily returns on a 1–3 month Treasury bill exchange-traded fund, averaging 1.24% annualised over the sample.

5.4 Estimation procedure

We estimate the model of Section 3 with $K = 2$ on the observation vector

$$X_t = \big(r_t,\ |r_t|,\ \widehat{\sigma}_t^{(21)},\ \mathrm{VIX}_t,\ \Delta \log V_t^{\mathrm{SPY}}\big)^\top, \tag{9}$$

where r_t is the market return and $\widehat{\sigma}_t^{(21)}$ its 21-day realised volatility annualised by $\sqrt{252}$. Including both r_t and $|r_t|$ permits the model to separate the sign of the drift from the scale of the shock.

Features are standardised using training-sample moments exclusively. Initialisation sorts observations on the mean of the standardised volatility features, assigns the upper half to one state, and sets the initial transition matrix to 0.95 on the diagonal while variances are floored at 10^{-4}. During the execution we iterate until the log-likelihood improvement falls below 10^{-4} or after 200 iterations.

The model is fitted once on **1993-02-01 to 2009-12-31** (4,263 observations), which can be a bit restrictive but the sample cannot begin earlier due to the fact that SPY volume is unavailable before 1993. Over the evaluation window it operates purely as a filter via Equations (5) and (6). Additionally, we audited signal timing throughout the pipeline and identified no look-ahead.

6 Experimental Design

As we know, searching over a large space of trading rules generates significant in-sample performance from noise alone (Sullivan et al., 1999), and conventional significance thresholds are insufficiently restricting given the number of strategies collectively tested (Harvey et al., 2016). This is the reason why we decided to fix the following protocol prior to examining any results:

Sample partition. We use two windows, the first one is from 2010-01-04 to 2016-12-30 (1,762 trading days) which is used for parameter selection while the window from 2017-01-03 to 2025-12-31 (2,262 trading days) is held out and was inspected only after parameters were fixed. Holdout figures constitute the reported results and tuning figures are presented alongside so that in-sample to out-of-sample degradation is observable.

Parameter grid. Four friction parameters were fitted on the tuning window: rebalancing frequency (daily, weekly, monthly), a per-security no-trade band, an exponentially weighted moving-average half-life applied to π_t, and a tempering coefficient λ in $\pi_t^{\mathrm{eff}} = 0.5 + \lambda(\pi_t - 0.5)$. Throughout the process 36 combinations were evaluated, spanning tuning-window Sharpe ratios from 0.82 to 1.05. The selection was weekly rebalancing, a 50 basis point band, a 10-day half-life and $\lambda = 1$. Sleeve parameters like the 126-day lookback, the 63-day volatility window and the 20-security breadth were fixed from the literature.

The coefficient λ interpolates between a static blend ($\lambda = 0$, no regime information) and the unmodified signal ($\lambda = 1$), reducing the central question to a single estimable parameter. The no-trade band implements, in simplified form, the buy/hold spread that Novy-Marx and Velikov (2016) identify as the most effective single cost-mitigation technique.

Transaction cost accounting. Between rebalancing dates the realised returns move the book away from its target, but the trades required to restore this are achieved through genuine trades. We therefore measure turnover against drifted rather than target weights. One-way costs are set at 15 basis points for large caps and 25 basis points for small caps.

Control specifications. A **static** 50/50 **blend** of the same two sleeves isolates the contribution of the signal from the contribution of holding the sleeves. While a **placebo signal**, constructed by circularly shifting π_t by 400 trading days, preserves the marginal distribution and persistence of the signal whilst destroying its timing. We use **size-matched passive benchmarks** in order to prevent a size tilt from being interpreted as strategy performance.

7 Results

7.1 Estimated regime model

Table 2: Two-state Gaussian hidden Markov model estimated on 1993-02-01 to 2009-12-31. The state-conditional feature means are reported in original units.

	`low_vol`	`high_vol`
Market return r_t (daily)	0.000632	−0.000214
Absolute return $\lvert r_t \rvert$ (daily)	0.006659	0.016716
Realised volatility (annualised)	0.1334	0.3162
VIX close	16.96	30.77
SPY log volume change	0.001289	0.000398
A_{ii}	0.99075	0.97376
Expected duration (trading days)	108.1	38.1
Stationary distribution	73.9%	26.1%
Training log-likelihood	−21,901.24	

We can see in Table 2 that the two states separate in the anticipated direction. The high-volatility state exhibits 2.5 times the absolute return, 2.4 times the realised volatility and 1.8 times the VIX level of the low-volatility state, together with a negative mean return. Applying Equation (4) to the estimated diagonal yields expected durations of 108.1 and 38.1 trading days, both long relative to the weekly rebalancing interval selected in Section 6.

7.2 Regime paths out of sample

Over the evaluation window the mean forecast probability of the high-volatility state is 23.5%, close to the 21.2% of days classified as high volatility by a simple VIX rule; binary agreement between the two classifications is 69.9%.

Table 3: Regime path statistics under three classification rules, 2010–2025.

	Argmax of forecast	Switch only if $\pi_t > 90\%$	90% threshold + 2 consecutive days
Regime changes	140	82	78
Changes per year	8.76	5.13	4.88
High-vol share of days	23.4%	24.1%	24.1%
Median high-vol run (days)	2	10	12
Median low-vol run (days)	19	46.5	42

As presented in Table 3, the unfiltered classification changes state 8.76 times per year with a median high-volatility run of two days, notwithstanding the chain's estimated high-volatility duration of 38 days. This divergence between model persistence and classification persistence is the phenomenon anticipated by H3 and motivates the hysteresis rule employed in Architecture One.

7.3 Portfolio performance

Table 4 reports the holdout performance of both architectures alongside the two controls and the passive benchmarks, Figure 1 presents the same results graphically, while Figure 2 shows the resulting equity curves over the full sample.

Table 4: Holdout window, 2017-01-03 to 2025-12-31, net of transaction costs.

	Terminal wealth	Ann. return	Ann. vol.	Sharpe	Max drawdown
Passive benchmarks					
Buy & hold, large caps (29)	5.81	21.65%	18.78%	**1.02**	−29.9%
Buy & hold, traded universe (41)	4.58	18.48%	19.35%	0.86	−35.3%
Individual sleeves					
Minimum volatility only	4.11	17.05%	16.15%	0.92	−27.6%
Momentum only	4.03	16.79%	19.47%	0.78	−31.9%
Controls					
Static 50/50 blend, no signal	4.09	16.98%	17.24%	0.87	−29.9%
Placebo signal	4.07	16.93%	18.48%	0.82	−31.6%
Regime-conditioned strategies					
Architecture Two (slider)	4.01	16.74%	18.65%	0.80	−30.9%
Architecture One (switch)	0.22	−15.33%	18.42%	−0.93	−77.6%

Average annual turnover, two-sided notional: passive 0%; minimum volatility 357%; static blend 641%; Architecture Two 965%; Architecture One 3,363%.

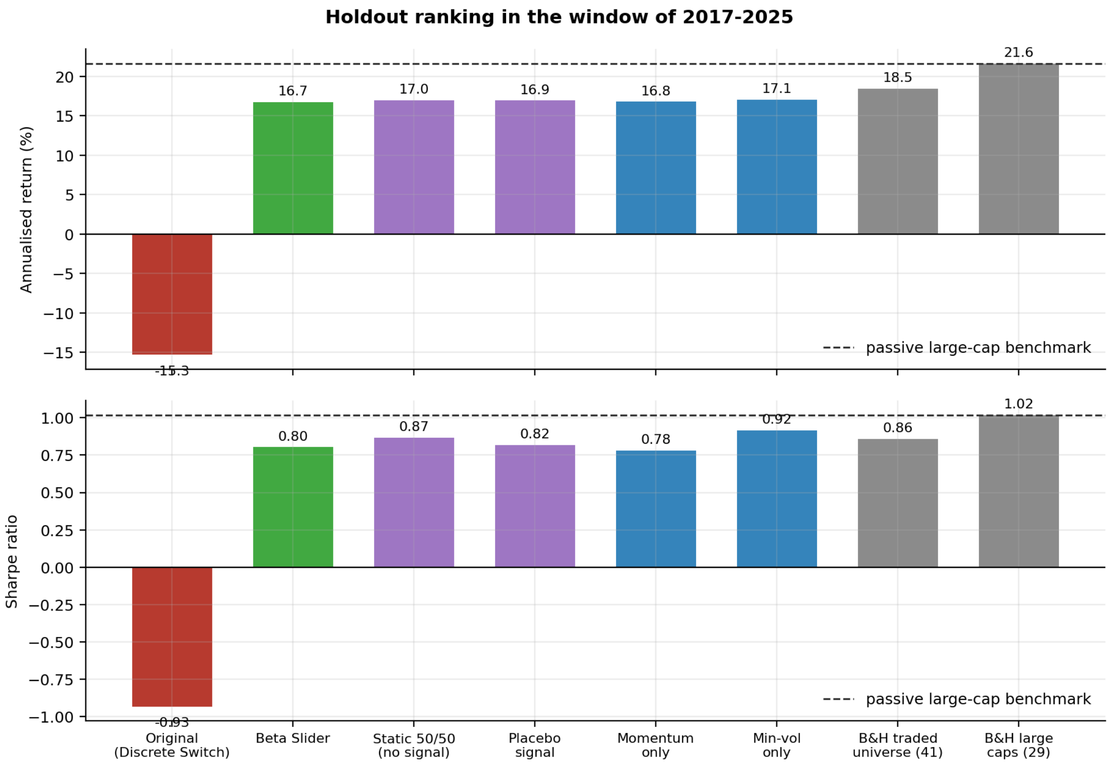

Figure 1: Holdout performance across all variants and controls. The dashed line marks a passive equal weight of the large-cap panel.

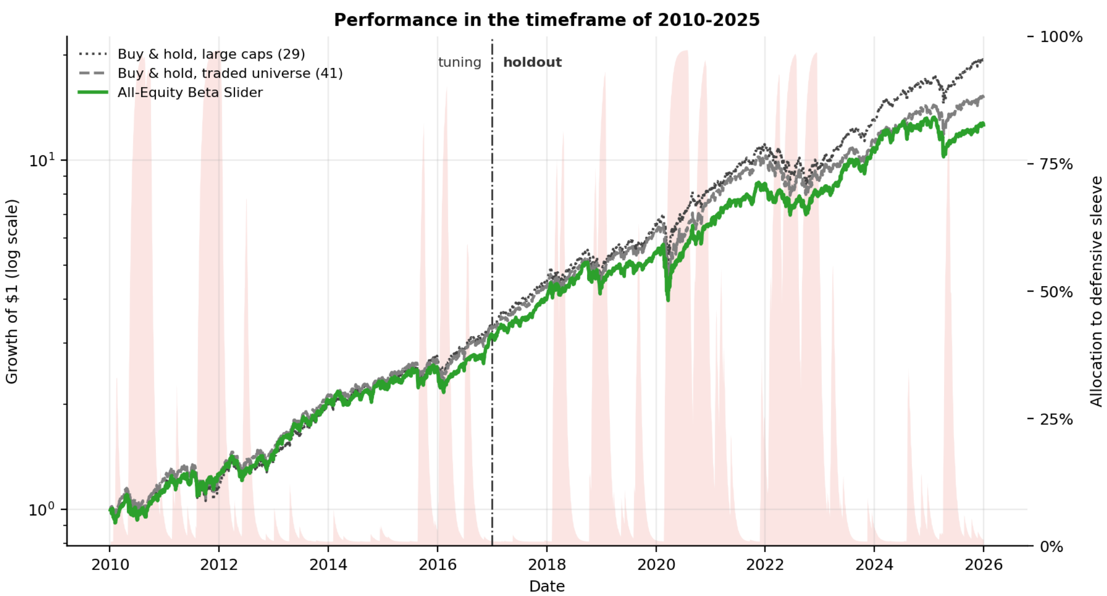

Figure 2: Growth of $1 in Architecture Two against both size-matched passive benchmarks, 2010–2025, net of transaction costs and on a logarithmic scale. The shaded area is the regime probability π_t, which is the share of capital allocated to the defensive sleeve. Only the segment to the right of the boundary is out of sample, the parameters of Section 6 were selected on the segment to its left.

7.4 Comparison of architectures

We found that Architecture Two exceeds Architecture One by 32.1 percentage points of annualised return, raises the Sharpe ratio from -0.93 to 0.80, reduces maximum drawdown from

77.6% to 30.9%, and lowers turnover by a factor of 3.5. As expected, mean overlap between the two sleeves of Architecture Two is measured at 9.8 of 20 securities, which means that approximately half the book is structurally insensitive to the signal, consistently with Equation (8). Furthermore, against the static 50/50 blend, which holds the same two sleeves without regime information, Architecture Two returns 16.74% against 16.98%, a difference of −0.25 percentage points with a Sharpe ratio of 0.80 against 0.87. Against the placebo the corresponding differences are −0.19 percentage points and −0.01 of Sharpe.

7.5 Comparison with passive benchmarks

A passive equal weight of the large-cap panel returns 21.65% at a Sharpe ratio of 1.02, which exceeds every active variant we tested on both measures. The same ordering holds in the tuning window as well, where the passive large-cap benchmark returns 18.73% at a Sharpe ratio of 1.14 against Architecture Two's 17.66% and 1.05. In other words, there is no window in this study in which any of our strategies beat simply holding the securities they trade.

It is worth looking at the minimum-volatility sleeve separately, since it records the highest active Sharpe ratio (0.92) and the smallest drawdown (−27.6%) of anything we built. Measured against the 41-security universe it appears to add value, but this is misleading because the sleeve draws exclusively from the large caps. Measured against its own opportunity set, which is the large-cap panel, it forgoes 4.6 percentage points of return and 0.10 of Sharpe in exchange for 2.2 percentage points of drawdown reduction. What looks like a low-volatility effect is therefore largely a size tilt. We also note that its realised full-sample market beta is 0.77, which is materially below unity as the low-volatility literature predicts, but above the 0.4 to 0.5 typical of minimum-variance portfolios constructed on broad universes. Selecting 20 securities out of 29 simply leaves little room to diversify away market exposure.

7.6 In-sample and out-of-sample comparison

In the tuning window Architecture Two exceeds the static blend by 0.11 of Sharpe and the placebo by 0.06, whereas on the holdout it trails them by 0.06 and 0.01 respectively. It is therefore the variant exhibiting the largest degradation among those plotted in Figure 3, while the minimum-volatility sleeve is the only one to improve out of sample and the static blend degrades least of those that do not. This is exactly what Sullivan et al. (1999) would predict for a strategy whose apparent edge is a fitted artefact rather than a real one.

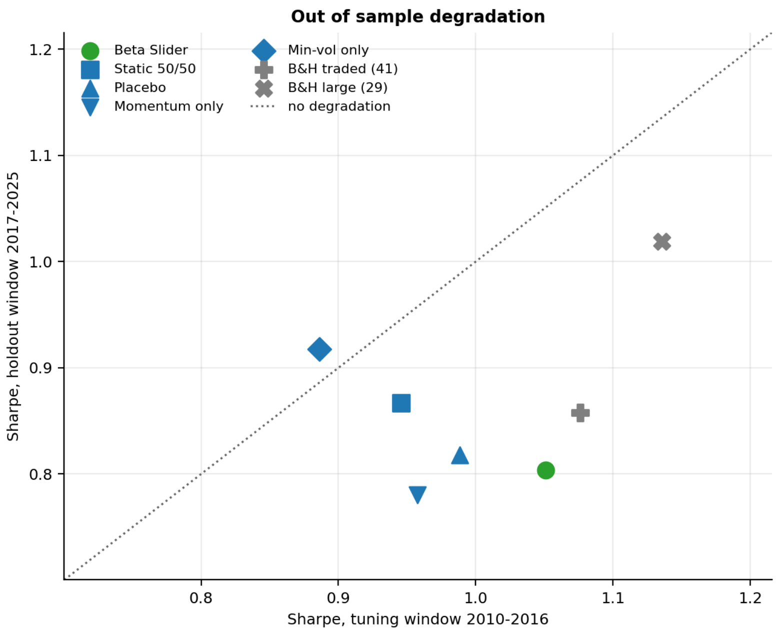

Figure 3: Tuning-window Sharpe ratio against holdout Sharpe ratio. Points below the diagonal degraded out of sample.

7.7 Decomposition of the design change

In order to localise the source of the difference between the two architectures we varied one design element at a time while holding the remainder fixed. The results are presented in Table 5.

Table 5: Attribution of the difference between architectures, holdout window.

Configuration	Annualised return	Change
(a) Discrete switch with liquidation	−15.33%	—
(b) Identical sleeves, liquidation removed	−15.80%	−0.47
(c) Identical sleeves, continuous blend	−13.20%	+2.60
(d) Revised sleeves, continuous blend	+16.74%	+29.94

Removing the liquidation policy contributes −0.5 percentage points, substituting a continuous blend with the sleeves held fixed contributes +2.6, and revising the sleeves with the orchestration held fixed contributes +29.9. This means that the orchestration rule accounts for roughly 7% of the total improvement while the choice of sleeves accounts for the remainder. It is important to note that this was not what we expected when we designed the second architecture, as we had assumed that staying invested and netting the turnover would be the dominant effect.

8 Discussion

8.1 Assessment of the hypotheses

When it comes to assess our hypotheses we can conclude that two of our three hypotheses are supported by the results while the third is not.

H1 holds since the estimated states are economically distinct and both are highly persistent, with expected durations of 108 and 38 trading days against a weekly rebalancing interval. H3 holds as well, given that the netting property derived in Equation (8) is confirmed by a measured sleeve overlap of 9.8 out of 20 securities and a turnover reduction of a factor of 3.5 relative to the discrete architecture.

However, H2 is not supported because Architecture Two does not separate from a static blend of its own sleeves or from a placebo signal, and Table 5 attributes only 7% of the improvement over Architecture One to the orchestration rule. Since this is the hypothesis the whole strategy depends on, the following subsection considers why it failed.

8.2 Why the regime signal did not contribute

The explanation we find most compelling is visible directly in Table 4. Over the holdout the momentum sleeve returned 16.79% and the minimum-volatility sleeve 17.05%, which is a differential of only 26 basis points per annum.

A rotation rule can add value only if the assets between which it rotates earn different returns and the rule allocates towards the better one at the right times. Where two sleeves earn materially identical returns over a sample, no allocation rule between them can add much irrespective of how good the timing signal is. The attainable range for any rotation strategy is then bounded close to the sleeve returns themselves, which is precisely what we observe since every blended variant falls between 16.74% and 16.98%. This reframes the interpretation of our results, because the experiment was not a powerful test of whether regime conditioning adds value. It was a test conducted on a sleeve pair that offered very little to condition upon in the first place.

Three further factors most likely contributed to this outcome. First, the composition of the sleeves themselves works against differentiation. The netting property that makes Architecture Two inexpensive has a counterpart, namely that the overlap suppressing turnover equally suppresses the strategy's capacity to distinguish between states. With half the book common to both sleeves, the maximum attainable dispersion between the two regimes is halved before the signal is even consulted, and the proximate cause of this is universe breadth. Selecting 20 securities from 41, and 20 from 29, constitutes near-total inclusion rather than selection.

Second, the factor structure of the sample period does not match what our signal measures. Over 2010–2025 the large-cap panel returned 20.36% against 13.07% for the small-cap panel (Table 1), which means the dominant source of cross-sectional return dispersion in our sample was size and mega-capitalisation concentration rather than volatility state. A signal measuring market-level volatility carries no information along that dimension, so conditioning upon it could not help, and any defensive tilt away from the concentrated leaders was actively costly.

Third, our universe was assembled by filtering on contemporaneous market capitalisation and is therefore populated by securities that performed well over the test period by construction. In a universe of known survivors, a rule that reduces exposure to the highest-performing constituents operates at a structural disadvantage. This does not invalidate the comparison, since the bias applies uniformly across all variants, but it does make the experimental setting unfavourable to defensive conditioning specifically.

8.3 Remaining limitations

Two further limitations are worth noting, although we regard both as contributing factors rather than primary causes.

Even after the architecture change, Architecture Two turns over 965% per annum against 357% for the minimum-volatility sleeve held alone and 0% for the passive alternatives. At our assumed 15–25 basis point one-way costs this represents a drag of roughly 1.5 to 2.5 percentage points annually relative to passive holding, which is not the dominant effect but is certainly material given how small the differences in Table 4 are.

We also fixed $K = 2$ for the reasons given in Section 3, and a two-state specification may well be too coarse in the sense that the transitional states identified by three-state specifications could carry the informative content. That said, Table 2 indicates a well-behaved fitted model, so we do not think the specification is where the problem lies.

9 Directions for Further Research

9.1 Testing state dependence before building a portfolio

The most valuable modification we can suggest is also the least costly one. Before constructing any portfolio, the return *differential* between the two candidate sleeves should be regressed on the regime probability:

$$r_t^{\mathrm{mom}} - r_t^{\mathrm{def}} = \alpha + \beta\, \pi_{t-1} + \varepsilon_t. \tag{10}$$

If β is not reliably negative then no exploitable state dependence exists and no orchestration rule can create any. This diagnostic would have established in minutes what the full backtest established at considerably greater cost, and we would recommend it as a precondition for any regime-conditioned design.

9.2 Extensions to the design

Given Section 8.2, the highest-value change is to rotate between assets whose returns genuinely differ across states. Candidates include a long–short momentum factor in place of a long-only momentum book, or a multi-asset rotation spanning equities, government bonds and commodities, which is the setting in which Nystrup et al. (2015) and Guidolin and Timmermann (2007) report regime conditioning to be valuable. Rotation within long-only equity sleeves drawn from an overlapping universe may simply constitute too narrow a decision for a market-level signal to inform. A related alternative is to abandon rotation altogether and use the regime probability to scale aggregate portfolio exposure instead. This would employ the signal for the quantity it directly measures, namely expected volatility, rather than requiring it to inform a cross-sectional decision about which it carries no information, and it would circumvent the sleeve-similarity problem entirely.

We think that the universe itself should also be rebuilt from point-in-time index membership. A broad universe would eliminate the survivorship bias, give cross-sectional ranking adequate scope and reduce sleeve overlap, awhile being a precondition for generalising any conclusion beyond the 41 securities studied here.

Finally, two smaller refinements are worth pursuing in our opinion. Estimating over $K \in \{2, 3, 4\}$ with information criteria, together with relaxing the diagonal covariance restriction, would test the specification choices of Section 3, since three-state specifications separating calm, transitional and crisis regimes are common in the literature and may isolate the transitional periods in which conditioning is most informative. On the cost side, Novy-Marx and Velikov (2016) identify the buy/hold spread as the most effective single cost-mitigation technique, and our no-trade band is only a simplified implementation of it. A full specification with distinct entry and exit thresholds, combined with monthly rather than weekly rebalancing, should reduce the residual drag discussed above.

10 Conclusion

In this paper we designed and implemented a regime-conditioned equity rotation strategy in two architectures and assessed both of them against a pre-committed holdout with three families of control.

Of the three hypotheses stated in Section 4, two are supported. Volatility states are identifiable and persistent, as the fitted model separates cleanly, implies expected durations of 108 and 38 trading days, and operates out of sample without look-ahead. The conditioning is also implementable within transaction costs, since replacing a discrete liquidating switch with a continuous slider improved annualised return by 32.1 percentage points and reduced turnover by a factor of 3.5, with the netting mechanism of Equation (8) behaving exactly as derived.

The hypothesis that failed is unfortunately the one on which the strategy depended. The regime signal did not contribute measurably, as Architecture Two does not separate from a static blend of its own sleeves or from a placebo, and no variant outperformed a size-matched passive benchmark. We attribute this principally to a property of the sleeve pair rather than of the signal, since over the test period the momentum and minimum-volatility sleeves earned returns only 26 basis points apart and thus left very little for a rotation rule to exploit. Unfortunately, this was compounded by a universe too narrow and too overlapping to support cross-sectional selection, and by a sample period whose dominant return driver was size concentration rather than volatility state.

The framework and the infrastructure we built remain sound and reusable. Therefore, what requires revision is the decision the signal is asked to inform, and the sequence in which that decision is validated, because exploitable state dependence should be established before a portfolio is constructed around it rather than afterwards.

Reproducibility

All code is available at `https://github.com/AxelTurinPlessia/Regime_based_trading_strategy`. The regime model is estimated in `HMM.ipynb`, which writes the one-step-ahead forecasts consumed by every strategy. `all_equity_beta_slider.py` builds both architectures, runs the tuning and holdout evaluations, and regenerates every table and figure in Section 7 directly from the computed results rather than from transcribed numbers. The tuning and holdout windows are set in a single configuration block and the holdout is never used for parameter selection. Price panels are not redistributed but `dataset_maker.py` rebuilds them from the underlying source.

References

Ang, A., and Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies*, 15(4), 1137–1187.

Baker, M., Bradley, B., and Wurgler, J. (2011). Benchmarks as limits to arbitrage: understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1), 40–54.

Blitz, D., and van Vliet, P. (2007). The volatility effect: lower risk without lower return. *Journal of Portfolio Management*, 34(1), 102–113.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.

Daniel, K., and Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics*, 122(2), 221–247.

De Bondt, W. F. M., and Thaler, R. (1985). Does the stock market overreact? *Journal of Finance*, 40(3), 793–805.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.

Frazzini, A., and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1), 1–25.

Guidolin, M., and Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11), 3503–3544.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). . . . and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.

Hong, H., and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, 54(6), 2143–2184.

Jegadeesh, N., and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91.

Korajczyk, R. A., and Sadka, R. (2004). Are momentum profits robust to trading costs? *Journal of Finance*, 59(3), 1039–1082.

Lehmann, B. N. (1990). Fads, martingales, and market efficiency. *Quarterly Journal of Economics*, 105(1), 1–28.

Nagel, S. (2012). Evaporating liquidity. *Review of Financial Studies*, 25(7), 2005–2039.

Novy-Marx, R., and Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1), 104–147.

Nystrup, P., Hansen, B. W., Madsen, H., and Lindström, E. (2015). Regime-based versus static asset allocation: letting the data speak. *Journal of Portfolio Management*, 42(1), 103–109.

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257–286.

Sullivan, R., Timmermann, A., and White, H. (1999). Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, 54(5), 1647–1691.