

A Regime-Aware, Free-Boundary Framework for Statistical Arbitrage:

Kalman-Filtered Synthetic Baskets, EGARCH Volatility, Markov Regime Detection, and Optimal Stopping via Hamilton-Jacobi-Bellman

Berk Bektaş*

Independent Researcher

September 11, 2026

Abstract

We propose a statistical arbitrage framework for equity pairs and small synthetic baskets that replaces the conventional fixed-window z -score trading rule with a sequence of estimated stochastic models, each addressing a specific weakness of the standard approach. A synthetic long/short basket is constructed from Johansen cointegration loadings, restricted to the two largest-magnitude long and short weights per cluster on stability and significance grounds. The basket's hedge ratio is tracked with a Kalman filter rather than re-estimated by rolling OLS. The resulting spread's conditional volatility is modeled with an asymmetric EGARCH(1,1) process. A two-state Markov-switching model, informed by a small set of documented macro structural breaks (the 2020 COVID-19 liquidity shock and the 2024 AI-capex repricing / yen-carry unwind), gates trading activity to periods in which the spread is statistically behaving as mean-reverting. Entry and exit levels are not fixed multiples of a rolling standard deviation; they are the free boundaries of an Ornstein-Uhlenbeck optimal stopping problem, solved via its Hamilton-Jacobi-Bellman variational-inequality formulation with smooth-pasting conditions. A monthly seasonality overlay, shrunk toward a common prior via an empirical-Bayes (James-Stein) estimator and filtered through Benjamini-Hochberg false discovery rate control, acts as a veto rather than an independent signal. Capital is allocated across simultaneously active spreads via a mean-variance efficient frontier (GMV/MSR), scaled by an EGARCH-based dynamic Value-at-Risk. We present the full mathematical derivation, a from-scratch numerical implementation (Kalman recursion, EGARCH quasi-MLE, Hamilton-filter EM, Engle-Granger rolling cointegration, and a `scipy.integrate.solve_bvp`-based free-boundary solver), and a synthetic validation study confirming that the regime filter correctly de-weights the two targeted structural-break windows.

1 Introduction

Classical pairs trading rules—enter when a rolling z -score of the spread exceeds ± 2 , exit at 0—conflate three distinct sources of uncertainty into a single ad hoc threshold:

1. uncertainty about the current hedge ratio,

2. uncertainty about the spread's current volatility regime, and
3. uncertainty about whether the underlying cointegrating relationship still holds at all.

A fixed z -score band is optimal only under the counterfactual assumption that all three are constant. We decompose the problem into its constituent stochastic models and replace the z -score with the solution to an explicit optimal stopping problem, so that every parameter in the trading rule has an economic interpretation and a comparative-static derivative.

2 Universe and Synthetic Basket Construction

For a cluster of n economically related tickers with log-price vector $p_t \in \mathbb{R}^n$, we estimate a Johansen (1991) VECM:

$$\Delta p_t = \alpha \beta' p_{t-1} + \sum_{i=1}^k \Gamma_i \Delta p_{t-i} + \varepsilon_t, \quad (1)$$

and test $\text{rank}(\alpha \beta') = r \geq 1$ via the trace statistic. When a cointegrating vector β is found, each component β_i is the loading of ticker i in the stationary linear combination $\beta' p_t$. We select the synthetic long basket as the two tickers with the largest positive, temporally stable loadings, and the synthetic short basket as the two with the largest-magnitude negative, temporally stable loadings, where “stable” requires the sign and rank of β_i to persist across overlapping estimation windows. This operationalizes the notion of “the two names the model trusts most on each side” as a testable statistical statement rather than a discretionary choice, and generalizes directly to the single-pair case ($n = 2$) used throughout our implementation and validation.

3 Dynamic Hedge Ratio: Kalman Filter

We reject rolling-OLS re-estimation of the hedge ratio in favor of a local-level state-space model:

$$\beta_t = \beta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, Q), \quad (2)$$

$$y_t = x_t' \beta_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, R), \quad (3)$$

*E-mail: berkbektas36@gmail.com

where y_t is the long basket's log-price and x_t the short basket's. The Kalman recursion gives the minimum-MSE filtered $\hat{\beta}_{t|t}$ and, crucially, the one-step-ahead innovation $\varepsilon_{t|t-1} = y_t - x_t' \hat{\beta}_{t|t-1}$, which becomes the spread series feeding every downstream model. Q and R are estimated by maximizing the Gaussian prediction-error decomposition log-likelihood:

$$\log L(Q, R) = -\frac{1}{2} \sum_t \left[\log(2\pi S_t) + \frac{\varepsilon_t^2}{S_t} \right], \quad S_t = x_t' P_{t|t-1} x_t + R, \quad (4)$$

over a grid followed by a bounded local refinement.

A modeling caveat we surface explicitly rather than hide: if no intercept is included and Q is left unconstrained, a state estimated against a strongly trending regressor can absorb (rather than merely track) the slow, mean-reverting component of the spread itself, leaving a near-white-noise innovation and destroying the very signal the OU stage needs. We therefore regularize Q 's admissible range to reflect the prior that hedge ratios are economically slow-moving (driven by index weights, capital structure, sector beta) relative to the spread's own reversion speed, and we omit an intercept from the state vector entirely, leaving long-run-mean estimation to the OU stage in Section 7 where it belongs statistically.

4 Conditional Volatility: EGARCH(1,1)

Let r_t denote the (demeaned) Kalman innovation. We model:

$$\log \sigma_t^2 = \omega + \beta_v \log \sigma_{t-1}^2 + \alpha_v (|z_{t-1}| - \mathbb{E}|z_{t-1}|) + \gamma_v z_{t-1}, \quad (5)$$

where $z_t = r_t / \sigma_t$, following Nelson (1991). The asymmetry term γ_v is economically meaningful for spreads: a sharp negative innovation (an abrupt re-convergence) is typically followed by a choppy, elevated-volatility consolidation phase distinct from the aftermath of an equally sized widening. σ_t estimated here is the sole volatility input to the OU diffusion coefficient in Section 7—it is *not* a rolling standard deviation, and unlike a rolling window it responds asymmetrically and with an estimated persistence β_v rather than a fixed lookback length.

5 Regime Detection: Markov-Switching with Structural-Break Priors

We fit a two-state Gaussian Markov-switching mean model on r_t via the Hamilton (1989) filter combined with the Baum-Welch EM algorithm:

$$r_t | S_t = k \sim \mathcal{N}(\mu_k, \sigma_k^2), \quad P(S_t = j | S_{t-1} = i) = p_{ij}, \quad (6)$$

with state 1 interpreted as “mean-reverting” (low $|\mu_1|$, moderate σ_1) and state 2 as “trending/broken.” A known limitation of regime-switching estimation on a few years of daily data is that genuinely rare, macro-driven dislocations are underweighted by a purely data-driven transition matrix—the base rate implied by the sample simply does not contain enough

such events to estimate their persistence reliably. We address this by treating two windows as **informative priors rather than re-discovered features**:

- **2020-02-20 to 2020-04-30 (acute phase 2020-03-09 to 2020-03-23):** The COVID-19 demand shock produced the fastest 30% S&P 500 drawdown on record and a VIX peak near 82.7 (2020-03-16), followed by the Federal Reserve's move to the zero lower bound (2020-03-03, 2020-03-15), unlimited quantitative easing and corporate credit facilities (2020-03-23), and the CARES Act (2020-03-27). Cross-sectional equity correlations converge toward one during the crash leg, and a liquidity-driven, valuation-agnostic trend regime persists into 2020–2021—both effects structurally break, rather than merely widen, historical cointegrating relationships.
- **2024-01-01 to 2024-03-31 and 2024-08-01 to 2024-09-30:** The first window reflects the AI-capex repricing that sharply widened dispersion among hyperscaler and semiconductor names following NVIDIA's 2024-02-22 earnings report, rendering pre-2023 cointegrating vectors stale for AI-exposed pairs. The second reflects the yen-carry unwind following the Bank of Japan's 2024-07-31 rate hike and a soft U.S. payrolls print, which produced the largest single-day VIX point spike on record (2024-08-05, intraday high near 65)—a cross-asset liquidity event unrelated to any individual pair's fundamentals—compounded by the first Federal Reserve rate cut in four years (2024-09-18), which resets the discount-rate regime entering the HJB problem's ρ parameter in Section 7.

During these windows the transition matrix's persistence toward the trending state is raised by a fixed Bayesian nudge before the E-step, rather than being left to the EM algorithm to infer from a handful of historical examples. The output used downstream is the smoothed probability $P(S_t = 1 | \text{data})$, applied as a continuous multiplicative gate rather than a hard binary switch, which materially reduces whipsaw relative to a thresholded regime call.

6 Rolling Cointegration as an Independent Early-Warning Layer

Independently of the regime filter, we re-run an Engle-Granger two-step test (OLS followed by an augmented Dickey-Fuller test on the residual) over a trailing six-month window at every date, producing a cointegration confidence score $c_t = 1 - \hat{p}_t$. This is deliberately redundant with the Markov filter: the latter detects volatility/mean-shift anomalies in the *spread's time series*, while rolling cointegration directly re-tests the *cross-sectional relationship* between the legs. A relationship can fail one test without failing the other (e.g., a slow fundamental divergence that does not manifest as an acute return-distribution shift), so both gates must pass for a trade to be permitted.

7 Replacing the Z-Score: Optimal Stopping via Hamilton-Jacobi-Bellman

The filtered, EGARCH-conditioned spread is modeled as an Ornstein-Uhlenbeck (OU) process:

$$dS_t = \kappa(\theta - S_t)dt + \sigma_t dW_t, \quad (7)$$

with (κ, θ) estimated by exact discrete-time conditional MLE—equivalently, OLS on the AR(1) representation:

$$S_{t+1} = \theta(1 - e^{-\kappa\Delta t}) + e^{-\kappa\Delta t}S_t + \varepsilon_{t+1}, \quad (8)$$

and σ_t supplied period-by-period by the EGARCH filter of Section 4.

7.1 Exit Problem

Conditional on holding a long-spread position, the holder chooses a stopping time to maximize expected discounted proceeds net of a round-trip cost c :

$$V(s) = \sup_{\tau} \mathbb{E}_s[e^{-\rho\tau}(S_{\tau} - c)]. \quad (9)$$

In the continuation region $s < b^*$, the Hamilton-Jacobi-Bellman equation for a time-homogeneous diffusion reduces to the linear ODE:

$$\kappa(\theta - s)V'(s) + \frac{1}{2}\sigma^2 V''(s) - \rho V(s) = 0, \quad (10)$$

with **value matching** $V(b^*) = b^* - c$ and **smooth pasting** $V'(b^*) = 1$ at the free boundary b^* , and a regularity condition ruling out the explosive solution branch as $s \rightarrow -\infty$. We impose the latter numerically by truncating the state space at ± 6 stationary standard deviations ($\sigma/\sqrt{2\kappa}$) from θ and applying a Neumann condition there, solving the resulting two-point boundary value problem with `scipy.integrate.solve_bvp`, and locating b^* by a one-dimensional root search on the smooth-pasting residual $V'(b; b) - 1$.

7.2 Entry Problem

Symmetrically, conditional on being flat, the holder chooses when to buy the spread:

$$U(s) = \sup_{\tau} \mathbb{E}_s[e^{-\rho\tau}(V(S_{\tau}) - S_{\tau} - c)], \quad (11)$$

solving the same homogeneous ODE on $s > d^*$ with matching $U(d^*) = V(d^*) - d^* - c$ and smooth pasting $U'(d^*) = V'(d^*) - 1$.

This construction yields $d^* < \theta < b^*$ with the following comparative statics, each verified numerically in Section 11:

- $\partial b^*/\partial \sigma > 0$, $\partial d^*/\partial \sigma < 0$: higher EGARCH volatility widens the band—a rising- σ_t day should not trigger the same trade a low- σ_t day would;
- $\partial(b^* - d^*)/\partial c > 0$: transaction costs must be recovered before the boundary is reached;

- The band narrows as κ (mean-reversion speed) rises or ρ (opportunity cost of capital, reset by events such as the September 2024 Fed cut) rises, since faster reversion or a higher cost of waiting both argue for realizing gains sooner.

In production, boundaries are recomputed on a coarse grid of plausible σ_t values and looked up rather than re-solved at every tick.

8 Seasonality as a Veto, Not a Signal

For each ticker we compute calendar-month compounded returns across all available years and their standard errors. Raw month-by-month means are unreliable for two compounding reasons: a 20-year seasonal claim is a 20-observation estimate at best, and several names in the target universe have far shorter histories, making the raw estimate almost pure noise. We apply an empirical-Bayes shrinkage:

$$\hat{m}_{i,\text{shrunk}} = w_i \hat{m}_i + (1 - w_i) \bar{m}, \quad w_i = \frac{\tau^2}{\tau^2 + \text{se}_i^2}, \quad (12)$$

which automatically pulls short-history, high-standard-error estimates toward the cross-sectional prior. Because a full ticker-by-month grid constitutes several hundred simultaneous hypothesis tests, we further apply Benjamini-Hochberg false discovery rate control across the entire grid (not ticker-by-ticker, which would understate the multiplicity problem) and zero out any monthly effect that does not survive. The surviving, shrunk seasonal tilt is used only to veto an entry that would contradict it; it never independently triggers a trade.

9 Portfolio Construction and Risk Scaling

Across simultaneously active spreads, capital is allocated via the mean-variance efficient frontier—Global Minimum Variance (GMV) and Maximum Sharpe Ratio (MSR) endpoints, and the connecting frontier obtained by repeated variance minimization at fixed target returns (numerically equivalent to Markowitz’s Critical Line Algorithm’s turning-point curve). Given the well-documented fragility of mean-return inputs, we weight toward the GMV end in practice. Total exposure is further scaled by an EGARCH-implied parametric Value-at-Risk:

$$\text{VaR}_t(\alpha) = -(\mu_t + \sigma_t \Phi^{-1}(\alpha)), \quad (13)$$

and a volatility-targeting multiplier that caps leverage and shrinks exposure automatically as σ_t rises.

10 The Composite Decision Rule

The comprehensive execution signals and dynamic allocations satisfy:

$$\begin{aligned} \text{Signal}_{k,t} &= \mathbf{1}\{P(\text{mean-reverting})_t > \pi^*\} \\ &\quad \times \mathbf{1}\{c_t > \gamma^*\} \\ &\quad \times \mathbf{1}\{S_t \text{ crosses } b_t^* \text{ or } d_t^*\} \\ &\quad \times \mathbf{1}\{\text{seasonal tilt does not contradict}\}, \quad (14) \\ \text{Size}_{k,t} &= w_k^{\text{GMV/MSR}} \times \text{VolTargetScale}(\sigma_t) \\ &\quad \times P(\text{mean-reverting})_t. \quad (15) \end{aligned}$$

11 Numerical Validation (Synthetic Data)

Because live market data could not be retrieved in the execution environment used to build this reference implementation, we validate the pipeline on a synthetic long/short pair whose spread follows the OU process of Section 7 everywhere **except** inside the two structural-break windows of Section 5, where it is deliberately switched to a trending, variance-inflated process—i.e., we construct exactly the two dislocations the framework is designed to detect and confirm the filter responds as intended, rather than backtesting for a performance claim.

Across three break sub-windows, mean smoothed $P(\text{mean-reverting})$ collapsed from a baseline of approximately 0.85 outside any break window to below 0.02 inside every one, and the HJB band width increased roughly five-fold between the 20th and 80th percentile of realized EGARCH σ_t . Rolling cointegration confidence declined modestly inside the break windows relative to baseline, consistent with its role as a slower-moving, complementary gate rather than the primary break detector. These results should be read only as a check that each module behaves as specified given a known ground truth; they are not, and should not be interpreted as, an estimate of live trading performance.

12 Limitations

The Kalman/EGARCH/OU chain, run against a synthetic single-factor trending regressor, is demonstrably capable of materially distorting the estimated mean-reversion half-life if the state-noise variance is left unregularized—we document this failure mode explicitly in Section 3 because it is a general risk in any Kalman-based hedge-ratio pipeline, not an artifact specific to this implementation.

The Engle-Granger p-values used for the rolling cointegration gate are computed from an interpolated approximation to MacKinnon critical values rather than an exact response-surface calculation, adequate for a relative, rolling confidence signal but not a substitute for exact inference where the precise p-value matters.

Regime-switching estimation on a handful of years of data is intrinsically data-poor for rare events, which motivated the informative-prior treatment of Section 5; extending or relocating those windows requires the same macro judgment applied here, not a purely statistical re-estimation.

Finally, every module in this framework carries estimation risk that compounds across stages; walk-forward, out-of-sample validation against real market data—not presented here—is a prerequisite before any capital allocation decision.

References

- [1] Bertram, W. (2010). Analytic solutions for optimal statistical arbitrage trading. *Physica A*, 389(11).
- [2] Dixit, A. K. (1993). *The Art of Smooth Pasting*. Harwood Academic Publishers.
- [3] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4).
- [4] Engle, R. F. & Granger, C. W. J. (1987). Co-integration and error correction: representation, estimation, and testing. *Econometrica*, 55(2).
- [5] Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2).
- [6] Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- [7] Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 59(6).
- [8] Leung, T. & Li, X. (2015). Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance*, 18(3).
- [9] Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica*, 59(2).
- [10] MacKinnon, J. G. (2010). Critical values for cointegration tests. *Queen's Economics Department Working Paper No. 1227*.