EPFL

A Regime-Aware Equity Timing Strategy Derived from the Marketron Model

Project Summary, TFA Quant Research Track, EPFL Spring 2026

Pablo Habib - David Attias

1 Project Motivation

This project studies one of the most recent and unconventional modeling directions in quantitative finance: the Marketron model of [Halperin and Itkin, 2025], a physics-inspired framework that replaces the standard single-regime view of equity markets with a nonlinear latent-state system featuring memory, metastability, and endogenous regime transitions. While geometric Brownian motion (GBM) remains the classical benchmark for equity modeling, it imposes Gaussian log-returns and therefore misses central empirical features of equity markets: negative skewness, excess kurtosis, asymmetric drawdowns, and persistent stress regimes.

More broadly, the model is connected to recent work on market inelasticity and nonlinear price impact, especially the ideas developed by [Gabaix and Koijen, 2021] and [Bouchaud, 2021], which suggest that investor flows can have persistent and structurally nonlinear effects on asset prices.

The core idea behind our project is that equity markets may be better described as evolving in a state-dependent energy landscape rather than around a single equilibrium. This opens an attractive research direction for quantitative finance: instead of forcing markets into a one-regime diffusion, one can model latent memory, nonlinear transitions, and regime shifts in a structurally interpretable way.

Our work is both recent and practical. It is recent because the underlying model was proposed only in 2025; it is practical because we do not stop at theory. We implement the full Marketron system, calibrate it on real S&P 500 data, infer its latent states via particle filtering, and extend it into a concrete equity timing strategy. This gives the project a complete pipeline from model to estimation to portfolio application.

2 Model Overview

The Marketron model describes the market as a particle moving in a deformable energy landscape. It combines three coupled state variables:

- $x_t = \log(S_t/S^\star)$, the observed log-price, referenced to $S^\star = 1000$;
- y_t, a latent memory variable capturing the accumulated effect of past investor flows;
- θ_t, a slowly varying latent signal with Ornstein–Uhlenbeck dynamics.

The system evolves according to the coupled stochastic differential equations

$$\begin{aligned}
dx_t &= \left[v\, f(\theta_t, t) + \eta - c\, y_t\, V_M'(x_t)\right] dt + \sigma\, dW_t, \\
dy_t &= \left[v\, h(\theta_t, t) + \mu(\bar{y} - y_t) - c\, V_M(x_t)\right] dt + \sigma_y\, d\widetilde{W}_t, \\
d\theta_t &= k(\hat{\theta} - \theta_t)\, dt + \sigma_z\, dZ_t.
\end{aligned}$$

This structure is what makes the model compelling. The drift of the price is not driven by a single constant trend, but by the interaction of a latent signal, a memory state, and a nonlinear potential. The resulting energy landscape can exhibit multiple extrema, interpreted as the *Good*, *Bad*, and *Ugly* market regimes: two metastable states separated by a tipping region. In contrast to exogenous regime-switching models, transitions here emerge endogenously from the dynamics.

For a quant audience, this is attractive for two reasons. First, it gives a structural explanation for non-Gaussian return features. Second, it naturally suggests regime-aware trading rules derived from latent-state inference rather than from ad hoc indicators.

3 Empirical Protocol and Calibration

We use 298 monthly S&P 500 observations from January 2000 to October 2024. The model is calibrated under a strict in-sample / out-of-sample protocol: parameters are estimated only on January 2000–December 2018 (228 months), while January 2019–October 2024 (70 months) is reserved as a genuine hold-out period.

Because only the price is observed, latent states are recovered through a bootstrap particle filter with 1500 particles, systematic resampling, and Seidel discretization. The 18 model parameters are estimated by minimizing an in-sample moment-matching objective across 5 horizons (2, 5, 10, 15, 18 years) and 4 moments (mean, volatility, skewness, kurtosis), augmented with two penalties:

- a shape penalty, enforcing preservation of the multi-extrema Good/Bad/Ugly structure;
- a bias penalty, discouraging systematic directional prediction bias.

We benchmarked several optimization pipelines. The retained combination, CMA-ES + L-BFGS-B, clearly dominated the alternatives both in objective value and computational efficiency:

Pipeline	Objective value	Runtime
DE + L-BFGS-B	7.336	∼55 min/run
CEopt + DE	2.878	∼109 min/run
CMA-ES + L-BFGS-B	2.035	∼2 min

This benchmark is an important part of the project: it shows that the contribution is not only conceptual, but also computational and methodological.

4 Main Results

1. Strong statistical fit beyond GBM

The calibrated Marketron reproduces the in-sample S&P 500 return structure closely across all horizons, including:

- the sign change in mean returns,
- the increasingly negative skewness,
- the rise in excess kurtosis at medium and long horizons.

On the aggregate in-sample moment criterion, the result is striking:

	Marketron	GBM
In-sample moment MSE	0.032	8.644
Relative improvement	269× better	

This is a key quantitative result. The gain is not driven by mean and volatility alone, but by the model's ability to capture higher-order moments that GBM cannot reproduce by construction.

2. Stable out-of-sample latent-state inference

A common weakness of rich state-space models is that they may overfit in-sample and fail when rolled forward. Here, the particle filter remains remarkably stable on unseen data:

$$\text{Train RMSE} = 0.0202, \qquad \text{Test RMSE} = 0.0204, \qquad \text{ratio} = 1.010.$$

This near-equality between train and test error indicates that latent-state inference generalizes cleanly beyond the calibration window. In other words, the model does not merely fit the past; it continues to track the observed market consistently out of sample.

3. A usable regime-aware trading overlay

We then extend the model by transforming its filtered latent state into a trading signal through the model-implied drift

$$\mu_x(t) = v\, f(\hat{\theta}_t, t) + \eta - c\, \hat{y}_t\, V_M'(\hat{x}_t).$$

This signal is translated into two allocation rules: a binary *long/flat* strategy, which stays invested in the S&P 500 when $\mu_x(t) > 0$ and moves to cash otherwise, and a *sized* strategy, which scales exposure continuously with the strength of the signal (both net of 5 bps transaction costs). Empirically, the long/flat rule turns out to be the more convincing specification, with the clearest downside-protection benefit out of sample. The main results are:

	Long/Flat	Buy & Hold
In-sample Sharpe	0.084	0.212
Out-of-sample Sharpe	0.864	0.831
Out-of-sample Max Drawdown	−20.1%	−24.8%

The interpretation is deliberately cautious. The strategy is *not* a robust alpha machine: it underperforms buy-and-hold in-sample. However, it does produce a slightly higher Sharpe ratio out of sample and, more importantly, a materially smaller maximum drawdown. Robustness checks across seven practical variants (threshold shifts, implementation lag, quarterly rebalancing, added signal noise) confirm that this defensive profile is persistent rather than accidental.

5 Why This Project Stands Out

We believe the project is appealing for quantitative research selection for four reasons:

- Topicality: It builds directly on a 2025 model, rather than revisiting a mature and over-explored framework.
- Cross-disciplinary originality: The project imports ideas from physics (energy landscapes, metastability, barrier crossing) into a concrete equity modeling problem.
- End-to-end implementation: We do not just discuss the theory: we implement the SDE system, filtering layer, calibration engine, benchmark comparison, and trading application.
- Honest empirical assessment: The project includes a real out-of-sample protocol, robustness checks, and a balanced interpretation of strengths and limitations.

6 Limitations and Future Work

The project also opens several natural directions for extension. First, the out-of-sample window is still short (70 monthly observations), so stronger statistical claims would require either a longer horizon or repeated walk-forward experiments. Second, the calibrated specification behaves more like a conditional filtering model than a realistic autonomous simulator, which suggests separating estimation-oriented and simulation-oriented parameter choices in future work.

More broadly, several research extensions are especially promising:

- Walk-forward recalibration: sequentially updating parameters over time rather than fixing them once.
- Richer benchmark set: comparing against stochastic-volatility and regime-switching models, not only GBM.
- Uncertainty-aware trading: integrating particle dispersion and latent-state confidence directly into position sizing.
- Higher-frequency extensions: testing whether the same latent-state architecture retains value at weekly or daily frequencies.
- Multi-asset generalization: extending the framework beyond single-index timing toward cross-asset regime allocation.

7 Conclusion

This project shows that a recent physics-inspired model can be turned into a serious quantitative research pipeline. The Marketron is not only conceptually original; it also delivers strong empirical results: a much better fit to higher-order return moments than GBM, stable latent-state inference out of sample, and a regime-sensitive trading signal with meaningful downside protection. In our view, the most interesting contribution is precisely this combination of novel modeling, rigorous empirical validation, and portfolio-level applicability.

References

[Bouchaud, 2021] Bouchaud, J.-P. (2021). The inelastic market hypothesis: A microstructural interpretation.

[Gabaix and Koijen, 2021] Gabaix, X. and Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. Research Paper 20-91, Swiss Finance Institute.

[Halperin and Itkin, 2025] Halperin, I. and Itkin, A. (2025). Marketron games: Self-propelling stocks vs dumb money and metastable dynamics of the Good, Bad and Ugly markets.