Position Within Trend: The MAD-Markov Model for Calibrated Regime-Transition Forecasting

Lawson Arrington[1]

July 5, 2026

Abstract—Moving-average indicators are ubiquitous in practice but usually collapse price behavior into directional crossover events, discarding the dynamics of displacement from trend. This paper introduces Moving Average Distance (MAD), a novel displacement indicator. MAD measures the signed percentage distance between price and a moving-average reference; standardizing it by a trailing volatility estimate yields a series comparable across time. The standardized series is recast as a regime process by discretizing it into six ordered states. On this foundation, the MAD-Markov Model provides an interpretable semi-Markov framework for forecasting regime transitions. Analysis of the regime sequence reveals that dwell times are not memoryless: the hazard of leaving a regime declines with time already spent, a duration dependence that a first-order Markov chain cannot represent. The central result is that the continuous position of price within its current band—information the discrete chain discards—conditions the direction of the next transition in a strong, monotone, and well-calibrated manner. On 26 years of daily S&P 500 (SPY) data, this relationship replicates across three independent market eras and produces out-of-sample transition probabilities that lie on the calibration diagonal, improving Brier score and log-loss over a memoryless base-rate chain by approximately 10%. The contribution is an interpretable, calibrated model of how displacement regimes form, persist, and give way.

I. Introduction

Systematic trading strategies are shaped by a persistent tension between two opposing tendencies of asset prices: the tendency of trends to persist, exploited by momentum strategies, and the tendency of prices to revert toward a reference level. Momentum effects are among the most robust empirical regularities in asset pricing, documented across horizons, asset classes, and markets [1], [2], [3], while the reversion of extended prices is equally well established [4]. Trend persistence, however, is not uniform; it is regime dependent. Momentum strategies are most vulnerable near turning points, where a prevailing trend weakens, stalls, or reverses and persistence gives way to rapid mean reversion [5]. Recent work frames this explicitly as the problem of balancing slow-moving

---

[1] Disclaimer: MAD-Markov is an independent, self-directed research project developed exclusively on my personal time. It is not affiliated with, sponsored by, or conducted on behalf of the Department of the Army, the Department of War, or any U.S. Government entity.

trend persistence against fast, localized reversion [6]. The recurring question is therefore not merely whether price is trending, but which regime currently governs its behavior.

Moving-average systems are the most widely used tools for representing trend behavior in practice [7], [8]. In their standard form, however, they compress the price path into a sequence of directional crossover events: price lies above or below its average, or a fast average crosses a slow one. This representation discards the quantity that most directly measures the state of a trend—the displacement of price from that trend, and how the displacement evolves through time. Envelope-based methods such as Bollinger Bands [9] and Keltner Channels [10] partially address this by contextualizing deviation against a volatility scale, but they treat deviation as a static level or threshold event rather than as a dynamic state with its own transition behavior. The displacement process itself—how far price sits from trend, how long it remains there, and where it tends to move next—remains largely unmodeled as a state variable.

This paper takes displacement as the primary object of study and recasts it as an observable regime process. The starting point is Moving Average Distance (MAD), a displacement indicator measuring the signed percentage distance between price and a moving-average reference. Standardizing MAD by a trailing volatility estimate renders displacement comparable across time and across volatility environments, and the standardized series is discretized into six ordered regimes, ranging from extended displacement below trend to extended displacement above it. The resulting sequence of regimes admits a natural Markovian treatment, connecting to a long tradition of regime-switching models in which price dynamics are represented as movements among discrete states [11]. In contrast to that tradition, the regimes here are directly observable functions of price rather than latent states recovered by a filter—a deliberate choice that keeps every step of the model interpretable.

Modeling the regime sequence yields two findings. First, the process is not memoryless: the probability of leaving a regime declines with the time already spent in it, a duration dependence that a first-order Markov chain cannot capture and that motivates a semi-Markov treatment of regime dwell. Second, and central to this paper, the discrete regime label conceals information that strongly predicts the next transition. The continuous position of price within its current band—whether displacement sits near the boundary it is approaching or the one it has just crossed—conditions the direction of the next regime change in a strong, monotone, and well-calibrated fashion. Two observations occupying the same discrete regime can thus carry markedly different transition odds, a distinction the discrete chain discards by construction. This within-band position is the position within trend of the title.

These properties are established on twenty-six years of daily data for the S&P 500 (SPY). The position-conditioned transition relationship is monotone in every regime and replicates across three independent market eras spanning the dot-com decline, the post-crisis expansion, and the COVID-19 and inflation shocks. Out of sample, the resulting transition probabilities lie on the calibration diagonal and improve the Brier score of a memoryless base-rate chain by approximately ten percent, with a comparable gain in log-loss. The scope of these claims is deliberately narrow: the MAD-Markov Model forecasts the evolution of the regime state, not returns. Its purpose is to describe displacement dynamics faithfully and verifiably, and the empirical evaluation holds it to that standard rather than to a trading objective.

The contributions of this paper are fourfold. First, it introduces Moving Average Distance (MAD), a novel displacement indicator, together with a discretization that turns it into an ordered, observable regime state. Second, it characterizes the empirical dynamics of that state, showing that regime dwell times exhibit duration dependence and are better described as semi-Markov than Markov. Third, it establishes and validates the central result—that within-band position calibrates the direction of the next regime transition—out of sample, across market eras, and against a base-rate benchmark. Fourth, it assembles these components into the MAD-Markov Model, an interpretable and fully reproducible framework for forecasting regime transitions. The remainder of the paper is organized as follows. Section II situates MAD relative to prior work on moving-average indicators, regime-switching models, and momentum. Section III defines the indicator and its regime representation. Section IV develops the MAD-Markov Model, including its semi-Markov dwell structure and position-conditioned transition probabilities. Section V describes the data, Section VI reports the empirical results, and Section VII discusses scope and limitations before Section VIII concludes.

II. Background and Related Work

The empirical study of price continuation and reversal provides the backdrop for this work. Cross-sectional momentum, in which recent winners continue to outperform recent losers, and time-series momentum, in which an asset's own past return predicts its future return, are among the most durable regularities in asset pricing [1], [2], and momentum coexists with value as a pervasive source of return variation across markets [3]. Countervailing evidence documents reversal, whereby prices that have moved to extremes tend to revert [4]. These effects are not contradictory but horizon- and regime-dependent, and their interaction is most consequential near turning points, where momentum strategies are prone to sharp losses [5], [6]. This literature motivates treating the prevailing regime, rather than direction alone, as the object worth modeling.

Moving averages are the canonical tool for summarizing trend, and the statistical properties of moving-average trading rules have been studied extensively [7], [8]. Volatility-based envelopes, most prominently Bollinger Bands [9] and Keltner Channels [10], extend the idea by scaling deviation from a moving average against a measure of dispersion, making overextension comparable across assets and time. These constructions inform the indicator developed here. They share, however, a common limitation: deviation from trend is treated as a level to be thresholded—overbought, oversold, inside or outside a band—rather than as a persistent state whose evolution and transitions can themselves be modeled. Like these envelopes, MAD is standardized against a volatility scale; unlike them, the standardized value is treated as such a state rather than as a threshold to be crossed.

The term Moving Average Distance has a specific precedent that warrants careful distinction. Avramov, Kaplanski, and Subrahmanyam [13] define a moving-average distance as the gap between short- and long-horizon moving averages of price—specifically the 21- and 200-day averages—and show that it predicts equity returns in the cross-section, an effect they attribute to anchoring. The construction and purpose of MAD in this paper differ on every axis. Where their measure is the distance between two moving averages, MAD is the distance between price itself and a single moving-average reference; where theirs enters directly as a cross-sectional predictor of returns, MAD is volatility-standardized, discretized into an ordered set of regimes, and modeled as a transition process whose object of prediction is the regime state rather than return. The shared name reflects a shared intuition—that distance from a moving average is informative—but the indicators, their representation, and their use are distinct.

Representing asset dynamics as movement among discrete states has a long history, beginning with Markov regime-switching models in which returns are governed by an unobserved state that evolves as a Markov chain [11]. The present approach differs in two respects. First, the states are observable: each regime is a deterministic function of standardized MAD, so no latent variable need be filtered and every classification is transparent. Second, the transition structure is richer than a first-order Markov chain. The dwell time in a regime is found to depend on its current duration—the hazard of leaving declines the longer a regime has persisted—placing the process in the semi-Markov family, in which holding times may follow arbitrary distributions rather than the geometric law implied by a memoryless chain [15]. Section IV develops this structure explicitly.

A recent and influential line of work applies deep learning to time-series momentum and regime adaptation, learning trend estimation and position sizing directly from data and, in more recent variants, incorporating changepoint detection to respond to regime shifts [6]. The Momentum Transformer extends this with an attention mechanism that captures dependencies across multiple timescales while offering a degree of interpretability [12].

These architectures are powerful and, in the latter case, deliberately interpretable, yet they remain data-intensive return-prediction systems whose internal state is learned rather than defined. The MAD-Markov Model occupies a different point on the trade-off: it forgoes the flexibility of a learned representation in favor of a fully specified, observable state and a transition model whose every probability can be read, audited, and checked for calibration. Evidence that parsimonious models can rival far more complex ones on structured time series [14] supports this emphasis on transparency over capacity.

III. Moving Average Distance and Regime Representation

This section defines Moving Average Distance and the regime representation built upon it. The construction proceeds in two steps—a signed measure of displacement from trend and a volatility standardization that renders displacement comparable across time—after which the standardized series is discretized into an ordered set of regimes.

A. Definition. Let P_t denote the closing price of a security on trading day t, and let SMA_t denote its w-day simple moving average,

$$\mathrm{SMA}_t = \frac{1}{w}\sum_{i=0}^{w-1} P_{t-i} \tag{1}$$

Moving Average Distance is the signed percentage displacement of price from this reference,

$$\mathrm{MAD}_t = 100 \cdot \frac{P_t - \mathrm{SMA}_t}{P_t} \tag{2}$$

so that $\mathrm{MAD}_t > 0$ when price trades above its moving average and $\mathrm{MAD}_t < 0$ when it trades below. Expressing displacement relative to price removes dependence on price level.

A fixed percentage displacement is not equally meaningful in all conditions: a two-percent gap from trend is routine when markets are volatile and extreme when they are calm. MAD is therefore standardized by a trailing estimate of its own volatility. Let

$$\sigma_t = \mathrm{std}(\mathrm{MAD}_{t-W+1}, \quad \dots, \quad \mathrm{MAD}_t) \tag{3}$$

be the sample standard deviation of MAD over the preceding W trading days. The standardized MAD is

$$z_t = \frac{\mathrm{MAD}_t}{\sigma_t} \tag{4}$$

expressing displacement in units of its trailing standard deviation. Every window in (1)–(4) is strictly trailing, so both MAD_t and z_t depend only on prices observed up to and including day t; the indicator is causal and free of look-ahead. Throughout, w = 20 trading days (approximately one trading month) and W = 255 trading days (approximately one trading year).

B. Regime representation. Standardized MAD is discretized into six ordered regimes that partition displacement into one-standard-deviation bands. The regime R_t on day t is assigned as follows:

Regime R_t	Condition	Interpretation
+3	$z_t > 2$	extended above trend
+2	$1 < z_t \leq 2$	moderately above
+1	$0 < z_t \leq 1$	mildly above
−1	$-1 < z_t \leq 0$	mildly below
−2	$-2 < z_t \leq -1$	moderately below
−3	$z_t \leq -2$	extended below

The regimes are ordinal: sign encodes the side of trend and magnitude the degree of displacement, so the six states form a natural ordering from deepest below-trend (−3) to most-extended above-trend (+3). No zero state is defined; the boundary $z_t = 0$ separates the mildly-below regime (−1) from the mildly-above regime (+1) and corresponds to price crossing its moving average. The discretization is deliberately coarse—six states distinguish mild from extreme displacement on either side of trend while leaving each state populated enough to estimate its transition behavior reliably, a point revisited in Section VI.

The regime sequence $\{R_t\}$ is the object modeled in the remainder of the paper. Because each regime is a deterministic function of standardized MAD, and MAD is itself a causal function of price, the state is fully observable: no latent variable is estimated, and the classification of any day is reproducible exactly from its price history.

IV. The MAD-Markov-Model

The regime sequence $\{R_t\}$ defined in Section III is modeled as a discrete-state stochastic process. Two questions organize the model: when the state changes, and in which direction. This section treats each in turn—beginning with the one-step transition structure, refining it to account for duration dependence in regime dwell, and finally introducing the position-conditioned transitions that constitute the model's central mechanism.

A. Transition structure. At the coarsest level, the evolution of $\{R_t\}$ is summarized by its one-step transition matrix,

$$P_{jk} = \Pr(R_{t+1} = k \mid R_t = j), \;\; j, k \in S \tag{5}$$

where S = {−3, −2, −1, +1, +2, +3} is the ordered regime set. Two structural features follow from the construction of the regimes. First, the matrix is strongly diagonal-dominant: because standardized MAD is persistent, a regime is far more likely to repeat than to change on any given day. Second, off-diagonal mass concentrates on adjacent regimes; since the regimes are ordered bands of a continuous quantity, price rarely skips a band between consecutive days, and the transition matrix is therefore approximately tridiagonal. Because same-day persistence dominates (5), it is largely uninformative about the question of interest—where the state goes when it does change. That question is captured by the embedded, or jump, chain, obtained by conditioning on a change occurring:

$$J_{jk} = \Pr(R_{t+1} = k \mid R_t = j, R_{t+1} \neq j), \;\; J_{jj} = 0 \tag{6}$$

which records only genuine transitions. For the interior regimes, the jump chain concentrates on the two adjacent states, so a transition from a regime is, in effect, a choice between moving to the regime above and the regime below. Predicting that choice is the object of Section IV-C.

B. Dwell times and duration dependence. The timing of transitions is governed by the distribution of dwell times. Let D denote the number of consecutive trading days in a single visit to a regime. A first-order Markov chain implies that D is geometric: the probability of leaving on any day is constant, independent of how long the regime has already persisted. This memoryless property is testable through the discrete hazard,

$$h_j(d) = \Pr\big(D_j = d \mid D_j \geq d\big) \tag{7}$$

the probability of leaving on day d given survival to day d; under a Markov chain the hazard is flat in d. Empirically (Section VI), the hazard declines with d: the longer a regime has persisted, the less likely it is to end. Regime dwell is therefore over-dispersed relative to the geometric law, and the expected remaining time in a regime increases with time already spent. Duration dependence of this kind cannot be represented by a first-order Markov chain and places the process in the semi-Markov family, in which a visit is described by the jump chain together with a dwell-time distribution that may take any form [15]. The model retains the empirical hazard directly rather than imposing a parametric dwell law, so that the estimated time-to-change reflects the observed duration dependence.

C. Position-conditioned transitions. The transition matrix (5) and jump chain (6) treat the regime label as the complete state. Yet the label is a coarsening of the continuous standardized MAD: within the band of an interior regime, z_t ranges over a full standard-deviation interval, and its position within that interval is discarded by the discrete representation. The central hypothesis of the MAD-Markov Model is that this discarded position is informative about the direction of the next transition. Formally, for a day in an interior regime, let $Y \in \{\text{up}, \text{down}\}$ denote the direction of the next transition—up if the next distinct regime lies above the current one in the ordering, down if below. The position-conditioned model specifies

$$\Pr(Y = \text{up} \mid R_t = j, z_t) = \frac{1}{1 + \exp(-\alpha_j - \beta_j z_t)} \tag{8}$$

a logistic function of standardized MAD with a regime-specific intercept α and slope β. A positive slope encodes the hypothesis directly: the higher z_t sits within the band—the closer displacement is to the upper boundary—the more likely the next transition is upward. When β = 0, (8) collapses to the unconditional jump probability and the model reduces to the memoryless chain; the empirical content of the model is the extent to which β departs from zero. The intercept and slope are estimated per regime from historical visits. Because the level of the transition bias drifts with the broader market while the position gradient is stable across eras (Section VI), the intercept is fit on a trailing window so that the baseline adapts, while the slope captures the structural relationship. All estimation uses only data preceding the day being predicted, so the resulting probabilities are causal and admit honest out-of-sample evaluation.

D. The forecast. Together, (7) and (8) constitute the forecast. Given the current state—the regime, its standardized MAD z_t, and the elapsed duration d—the model returns a calibrated probability that the next transition is upward, from (8), and an expected time until that transition, from the duration-conditioned dwell implied by (7). The direction forecast refines the jump chain with within-band position; the timing forecast refines the geometric dwell with duration dependence. Section VI evaluates both against memoryless baselines using proper scoring rules.

V. Data

The empirical analysis uses daily closing prices for the SPDR S&P 500 ETF (SPY), a highly liquid proxy for the U.S. large-cap equity market, retrieved from Yahoo Finance [16]. The sample runs from July 2000 through July 2026, comprising 6,537 trading days. SPY is chosen for its long, continuous history and its exposure to a broad index rather than to firm-specific events; evaluating a regime model over a single episode would be uninformative, whereas this sample spans the deflation of the dot-com bubble, the 2008 financial crisis, the post-crisis expansion, the COVID-19 shock of 2020, and the 2022 inflation and rate-hiking drawdown.

Moving Average Distance, its standardization, and the regime assignment are computed as defined in Section III. Because standardization applies a 255-day trailing window on top of the 20-day moving average, the earliest observations have no defined regime and are excluded; 6,264 classified trading days remain, beginning in August 2001. Across this sample the six regimes are populated unevenly—the mildly-above-trend regime is by far the most common and the two extreme regimes the least—a distribution reported in Section VI.

Two evaluation protocols are used. For out-of-sample scoring, the classified sample is split chronologically, with the earliest 60% used to estimate model parameters and the most recent 40% reserved for evaluation, so that no future information enters any estimate. For the replication analysis, the sample is instead partitioned into three contiguous, equal-length eras—spanning approximately 2001–2009, 2009–2018, and 2018–2026—so that the position-conditioned relationship can be estimated independently within each. Because these eras encompass markedly different macroeconomic environments, agreement across them provides a stringent test of whether the relationship is structural rather than an artifact of any single period.

VI. Results

This section reports the empirical behavior of the regime process and evaluates the position-conditioned model, proceeding from the coarsest description to the finest: regime frequencies and one-step transitions, the dwell-time structure, the position-conditioned transition probabilities that constitute the central finding, and an out-of-sample evaluation of the resulting forecasts.

A. Regime frequencies and transition structure. Table I reports the unconditional frequency of each regime over the 6,264 classified days. The distribution is skewed toward above-trend states, which together account for roughly two-thirds of the sample, consistent with the upward drift of the index over the period; the mildly-above-trend regime alone occupies 46% of days, while the two extreme regimes are rare. Table II reports the one-step transition matrix (5). The two features anticipated in Section IV-A are evident: self-transition probabilities range from 0.50 to 0.79, so a regime is far more likely to repeat than to change on any given day, and off-diagonal mass is almost entirely confined to adjacent regimes, leaving the matrix effectively tridiagonal. Because same-day persistence dominates, the jump chain (6), reported in Table III, is more informative about the direction of change. Its structure is distinctly centripetal: from every below-trend regime the probable move is upward (0.70 from moderately-below, 0.62 from mildly-below, and with certainty from extended-below), and from every above-trend regime the probable move is downward (0.90 from moderately-above, and unity from extended-above). The mildly-above regime is the pivot, splitting 0.54 down and 0.43 up. This is the mean reversion of standardized MAD expressed at the level of regime transitions.

Table I. Regime frequencies

Regime	Interpretation	Days	Share
−3	extended below	260	4.2%
−2	moderately below	578	9.2%
−1	mildly below	1,439	23.0%
+1	mildly above	2,894	46.2%
+2	moderately above	1,001	16.0%
+3	extended above	92	1.5%

Table II. One-step transition matrix, P(next | current)

	−3	**−2**	**−1**	**+1**	**+2**	**+3**
−3	0.65	0.30	0.04	—	—	—
−2	0.13	0.50	0.35	0.02	—	—
−1	0.01	0.13	0.62	0.24	0.00	—
+1	0.00	0.01	0.12	0.79	0.09	0.00
+2	—	—	0.00	0.26	0.71	0.03
+3	—	—	—	—	0.29	0.71

Table III. Jump chain, P(next distinct regime | current)

	−3	**−2**	**−1**	**+1**	**+2**	**+3**
−3	—	0.88	0.12	—	—	—
−2	0.27	—	0.70	0.03	—	—
−1	0.02	0.35	—	0.62	0.01	—
+1	0.00	0.03	0.54	—	0.43	0.00
+2	—	—	0.01	0.90	—	0.09
+3	—	—	—	—	1.00	—

B. Dwell times and duration dependence. Table IV summarizes the dwell-time distribution for each regime. Visits are short on average—median dwell is one to three days—but the distributions are right-skewed, the mildly-above regime reaching a maximum visit of 46 trading days. The critical question is whether these dwell times are memoryless. Under a first-order Markov chain the discrete hazard (7) would be constant in duration; it is not. Table V reports the empirical hazard of leaving a regime as a function of days already spent, for the three most populous regimes. In each case the hazard declines with duration—for the mildly-above regime it falls from 0.26 in the first three days to 0.15 beyond fifteen.

Equivalently, dwell dispersion exceeds the geometric benchmark implied by the self-transition probability: for the mildly-above regime the observed standard deviation is 5.4 days against a geometric prediction of 4.2, and the longest visit, 46 days, far exceeds what a memoryless process would produce. Persistence begets further persistence; the process is semi-Markov, as anticipated in Section IV-B, and the expected time to the next transition increases with time already spent.

Table IV. Dwell times (completed visits)

Regime	Visits	Median (days)	Mean	Max
−3	90	2	2.9	14
−2	287	1	2.0	15
−1	549	2	2.6	19
+1	614	3	4.7	46
+2	294	2	3.4	17
+3	27	2	3.4	17

Table V. Hazard of leaving, by days already in regime

Regime	1–3 d	4–7 d	8–15 d	>15 d
+1 mildly above	0.26	0.19	0.15	0.15
−1 mildly below	0.40	0.36	0.24	—
+2 moderately above	0.32	0.25	0.25	—

C. Position-conditioned transitions. The preceding results treat the regime label as the state. The central result is that the continuous position within the band—discarded by that label—strongly conditions the direction of the next transition. For each interior regime, the upward-transition probability is estimated as a function of standardized MAD within the band, following (8). The relationship is strong, monotone, and consistent. Within the mildly-above regime, the probability that the next transition is upward rises from 0.18 when standardized MAD sits in the lowest quartile of the band to 0.56 in the highest—a swing of nearly forty points driven entirely by within-band position. The same pattern holds in every interior regime: 0.46 to 0.81 in the mildly-below regime, 0.02 to 0.31 in the moderately-above, and 0.55 to 0.98 in the moderately-below. Two days assigned to the same discrete regime can thus carry markedly different transition odds, to which the discrete chain is blind. Two properties make the relationship trustworthy rather than incidental. First, it replicates: the upper panels of Figure 1 plot the upward-transition probability against within-band position separately for the three independent eras of Section V, and the curves rise monotonically in every regime and every era—the level shifts modestly with the market's prevailing bias, but the gradient is invariant. Second, it is calibrated out of sample: fitting on the earlier portion and applying to the held-out remainder, the lower panel of Figure 1 shows predicted probabilities lying on the diagonal against realized frequencies across the full range. The gradient is therefore both structural and numerically honest out of sample.

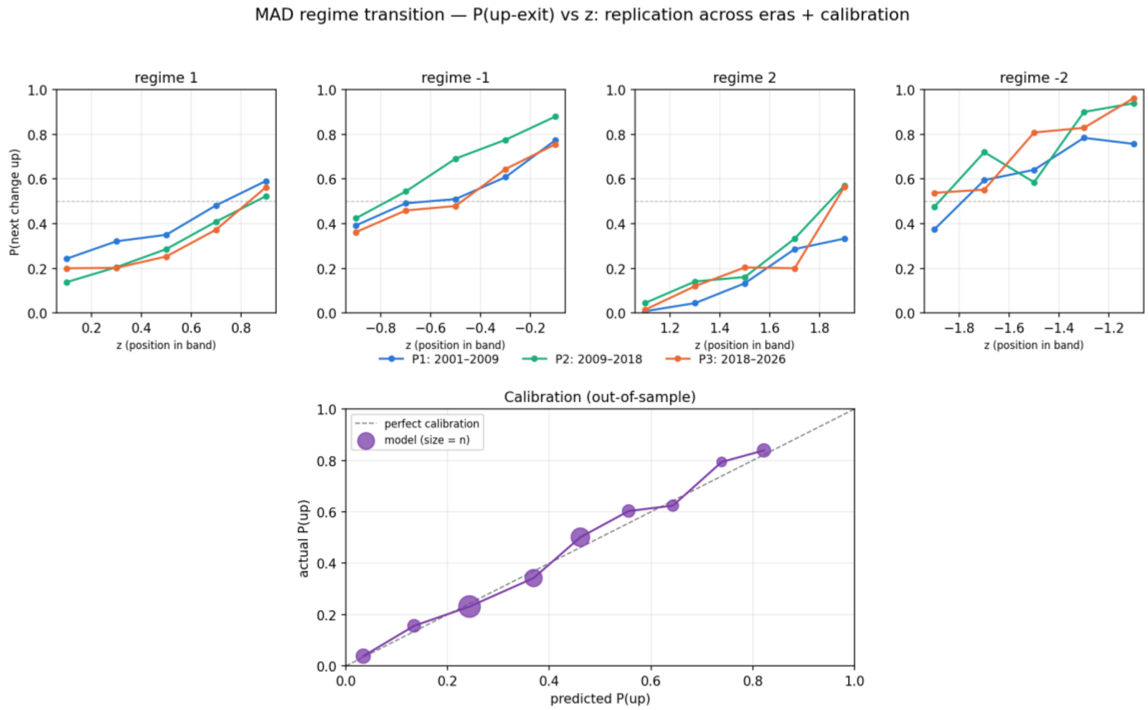

Figure 1. (Upper) Upward-transition probability versus within-band standardized MAD, by regime, in three independent eras; the gradient is monotone in every panel. (Lower) Out-of-sample calibration of predicted versus realized upward-transition probability.

D. Forecast evaluation. Table VI evaluates the model on the held-out sample against a memoryless base-rate benchmark—the jump chain of Table III, which assigns every day in a regime the same transition probability. Because the model's contribution is a sharper, calibrated probability rather than a different hard call, the natural metrics are proper scoring rules. The position-conditioned model reduces the Brier score from 0.204 to 0.184 and the log-loss from 0.589 to 0.541, improvements of roughly ten and eight percent. Against the uninformative benchmark of a constant one-half forecast, whose Brier score is 0.25, the base rate captures 18% of available skill and the position-conditioned model 26%, so within-band position supplies roughly a further third of attainable skill beyond the discrete state. As anticipated, the two models are nearly indistinguishable on hard direction accuracy, 71.8% versus 71.0%, because the sample's down-bias keeps the calibrated probability from crossing the decision boundary even as it moves substantially; the gain is in the quality of the probabilities, not the classification. Timing is the weakest component: the duration-conditioned expected time-to-change improves only marginally on a benchmark that ignores duration, a mean absolute error of 3.37 versus 3.43 trading days, reflecting the difficulty of the long right tail of dwell times. The model's value is thus concentrated in calibrated transition-direction probabilities—precisely the quantity the discrete chain cannot provide.

Table VI. Out-of-sample forecast evaluation (vs. memoryless base rate)

Metric	MAD-Markov	Base rate
Brier score	0.184	0.204
Log-loss	0.541	0.589
Direction accuracy	71.8%	71.0%
Timing MAE (days)	3.37	3.43

VII. Discussion and Limitations

The results establish two properties of the displacement-regime process and one modeling consequence. First, the process is semi-Markov rather than Markov: regime dwell times exhibit duration dependence, so a first-order chain misspecifies the timing of transitions. Second, the discrete regime label discards information that materially predicts the next transition; the continuous position of standardized MAD within its band conditions the direction of change in a monotone, well-calibrated, and era-stable manner. The consequence is that a transparent, fully observable model—the MAD-Markov Model—produces calibrated transition-direction probabilities that improve on a memoryless base rate out of sample. The value of the approach lies in this combination of interpretability and calibration: every probability is a legible function of price, and each is numerically honest against realized frequencies.

The scope of these claims is deliberately limited to the regime state. The MAD-Markov Model forecasts the direction and timing of transitions between displacement regimes; it is not a forecast of returns, and its evaluation is conducted entirely against regime-transition benchmarks rather than a market or trading objective. This distinction is intentional. Regime dynamics and return dynamics are related but not equivalent, and a calibrated forecast of the former does not imply an exploitable forecast of the latter; conflating the two is a recurrent source of overstatement in the technical-analysis literature. The contribution here is a faithful, verifiable description of how displacement regimes evolve, and the empirical section holds the model to precisely that standard.

Several limitations qualify the results. The analysis uses a single instrument, the S&P 500 ETF; while the position-conditioned relationship replicates across three independent eras of that series, its generality across other assets, sectors, and market-capitalization segments is not established here. Timing is the weakest component of the forecast: the duration-conditioned dwell model improves only marginally on an unconditional benchmark, reflecting the difficulty of predicting the long right tail of regime durations. The two extreme regimes are visited infrequently, so their transition estimates carry wider sampling uncertainty than those of the interior regimes and should be read as qualitative. Finally, the position-conditioned model adopts a simple logistic form; richer specifications of the within-band relationship, or of the dwell distribution, may sharpen the forecast.

These limitations point to natural extensions. The most direct is cross-sectional: applying the same construction to a broad universe of equities would test whether the position-conditioned relationship is a general property of displacement regimes rather than a feature of the aggregate index and would furnish a far larger sample of transitions in the sparsely-visited extreme states. A second direction is to condition the transition probability on information orthogonal to the price path—volume, breadth, or volatility state—which may carry predictive content the displacement series alone cannot. A third is to embed the calibrated regime forecast within an explicit risk or exposure framework; whether the regime signal confers an advantage in that setting is a separate empirical question, and one that—consistent with the scope maintained throughout—would require its own careful, out-of-sample evaluation.

VIII. Conclusion

This paper introduced Moving Average Distance, a displacement indicator measuring the signed percentage distance of price from a moving-average reference and developed it into an observable regime process: standardized by a trailing volatility scale and discretized into six ordered states. Modeling the resulting sequence, the MAD-Markov Model treats regime evolution as a semi-Markov process and conditions each transition on the continuous position of displacement within its band.

The empirical analysis, on twenty-six years of daily S&P 500 data, yielded a single central result: where price sits within its current regime—information the discrete state discards—predicts the direction of the next regime transition in a strong, monotone, and well-calibrated way. The relationship replicates across three independent market eras and produces out-of-sample transition probabilities that lie on the calibration diagonal, improving a memoryless base rate in Brier score and log-loss by roughly ten and eight percent. The dwell-time structure is likewise informative: regime persistence is duration-dependent, placing the process outside the first-order Markov family.

The contribution is an interpretable, calibrated account of how displacement regimes form, persist, and give way—a faithful model of market state rather than a forecast of returns. Whether the same position-conditioned structure holds across a broad cross-section of assets, and whether it can inform regime-aware risk management, are the questions this work leaves open.

References

[1] N. Jegadeesh and S. Titman, "Returns to buying winners and selling losers: Implications for stock market efficiency," The Journal of Finance, vol. 48, no. 1, pp. 65–91, 1993.

[2] T. J. Moskowitz, Y. H. Ooi, and L. H. Pedersen, "Time series momentum," Journal of Financial Economics, vol. 104, no. 2, pp. 228–250, 2012.

[3] C. S. Asness, T. J. Moskowitz, and L. H. Pedersen, "Value and momentum everywhere," The Journal of Finance, vol. 68, no. 3, pp. 929–985, 2013.

[4] W. F. M. De Bondt and R. Thaler, "Does the stock market overreact?," The Journal of Finance, vol. 40, no. 3, pp. 793–805, 1985.

[5] K. Daniel and T. J. Moskowitz, "Momentum crashes," Journal of Financial Economics, vol. 122, no. 2, pp. 221–247, 2016.

[6] K. Wood, S. Roberts, and S. Zohren, "Slow momentum with fast reversion: A trading strategy using deep learning and changepoint detection," arXiv:2105.13727, 2021. (Also in The Journal of Financial Data Science, 2022.)

[7] W. Brock, J. Lakonishok, and B. LeBaron, "Simple technical trading rules and the stochastic properties of stock returns," The Journal of Finance, vol. 47, no. 5, pp. 1731–1764, 1992.

[8] A. W. Lo, H. Mamaysky, and J. Wang, "Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation," The Journal of Finance, vol. 55, no. 4, pp. 1705–1765, 2000.

[9] J. Bollinger, Bollinger on Bollinger Bands. New York, NY, USA: McGraw-Hill, 2001.

[10] C. L. Keltner, How to Make Money in Commodities. Kansas City, MO, USA: The Keltner Statistical Service, 1960.

[11] J. D. Hamilton, "A new approach to the economic analysis of nonstationary time series and the business cycle," Econometrica, vol. 57, no. 2, pp. 357–384, 1989.

[12] K. Wood, S. Giegerich, S. Roberts, and S. Zohren, "Trading with the momentum transformer: An intelligent and interpretable architecture," arXiv:2112.08534, 2022.

[13] D. Avramov, G. Kaplanski, and A. Subrahmanyam, "Moving average distance as a predictor of equity returns," Review of Financial Economics, vol. 39, no. 2, pp. 127–145, 2021, doi: 10.1002/rfe.1118.

[14] D. Santoro, T. Ciano, and M. Ferrara, "A comparison between machine and deep learning models on high stationarity data," Scientific Reports, vol. 14, art. 19409, 2024, doi: 10.1038/s41598-024-70341-6.

[15] R. A. Howard, Dynamic Probabilistic Systems, Vol. II: Semi-Markov and Decision Processes. New York, NY, USA: Wiley, 1971.

[16] Yahoo Finance, "Historical market data," accessed via the yfinance Python package, 2026.