MAD-Velocity: Constructing and Evaluating a Moving-Average-Distance Trading Signal

Lawson Arrington[1]
July 10, 2026

Abstract— A companion study, the MAD-Markov Model, established that a security's continuous position within its displacement band forecasts the direction of its next regime transition in a strong, monotone, and well-calibrated manner; it stopped short of asking whether that structure can be traded. This paper extends the work from forecasting to trading, and profitability hinges on a distinction the forecast never required: whether displacement is changing because price is moving, or because the average is catching up to it. A first-order decomposition separates each change in Moving Average Distance (MAD) into a price-driven component and a moving-average artifact. The composition inverts with position in the band. Deep-oversold "improvements" are roughly two-thirds artifact, and at the re-crossing of the trend the ratio flips—roughly two-thirds genuine price movement. That decomposition converts directly into a signal. Gating trend re-crossing entries on price dominance beats each name's own buy-and-hold, measured per day of capital deployed and gross of costs, on 59.8% of 672 point-in-time S&P 500 constituents (2016–2025, delisted members included) against 55.8% for the ungated rule (paired test, p = 0.021). On the index, an anticipatory variant that enters below trend, before the re-cross confirms, outperforms three decades of buy-and-hold on deployed capital (13.4% versus 10.9% annualized; in-trade Sharpe 0.74 versus 0.65) while invested 69% of days. The companion's calibrated transition forecast adds nothing beyond identifying the state. The composition and timing of the entry, not the forecast, carry the signal.

[1] *Disclaimer: MAD-Velocity 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.*

I. Introduction

Systematic equity strategies inherit a persistent tension between two well-documented tendencies of asset prices. Securities that have performed well continue, on average, to perform well over horizons of months [1], while over longer horizons prices display reversal, with extreme past losers overtaking past winners [2]. The momentum effect extends beyond the cross-section: an asset's own past return forecasts its future return across broad asset classes [3]. Recent work makes the tension explicit by engineering strategies that blend slow momentum with fast reversion [4] or learn the blend directly with attention-based architectures [5]. Any rule built on a moving average inherits this tension in miniature: the same indicator that confirms a trend is, at short horizons, a measure of the stretch that reversal corrects.

Moving averages are among the oldest objects in this literature. Simple moving-average rules were shown to carry genuine predictive content well before the modern factor era [6], and pattern-based technical analysis was subsequently given a rigorous statistical footing [7]. In the cross-section, moving-average signals are profitable in a manner not subsumed by momentum [8]; most directly relevant here, the distance between price and its moving average—the same primitive underlying this work—predicts equity returns with economically large alphas, an effect attributed to investor anchoring [9]. This literature also carries a standing caution: with enough rules tried, some succeed by chance [10], a caution the present evaluation is designed around.

A companion study, the MAD-Markov Model [11], contributes the representation this paper trades. Moving Average Distance (MAD)—the percentage displacement of price from its 20-day average, standardized by its own trailing volatility and discretized into six ordered regimes—yields a state whose continuous position forecasts the direction of the next regime transition in a strong, monotone, and well-calibrated manner under strictly walk-forward estimation. That study established forecastability and deliberately stopped short of evaluation: knowing which transition comes next is not the same as knowing that trading it profits.

This paper takes up that question and finds it pivots on a quantity the forecast never required: the composition of displacement change. MAD can rise for two very different reasons—price rising toward its average, or the average falling toward price as old highs leave the window. A first-order decomposition separates every change in MAD into a price-driven component and a moving-average artifact, and the mixture inverts with position in the band: apparent improvements deep below trend are roughly two-thirds artifact, while re-crossings of the trend from below invert the ratio—roughly two-thirds genuine price movement. The decomposition is this paper's central construction—the "velocity" of MAD-Velocity—and it functions not as a forecaster but as a qualifier, identifying which regime transitions are real.

The evaluation is deliberately stringent. Signals are strictly causal and executed at the next open; performance is measured per day of capital deployed rather than against a fully invested calendar, so a low-exposure rule is neither flattered nor penalized by its time in cash; and the cross-section is built from point-in-time index membership, including delisted members, so the test carries the failures as well as the survivors. On 672 S&P 500 constituents over 2016–2025, gating trend re-crossing entries on price dominance raises the share of names beating their own buy-and-hold from 55.8% to 59.8% (McNemar $p = 0.021$)—the only refinement that survives a paired significance test. On the index itself, an anticipatory variant that enters below trend, before the re-cross confirms, outperforms buy-and-hold across three decades on deployed capital. The companion's calibrated forecast, by contrast, adds nothing beyond identifying the state—an instructive negative result in itself.

The remainder of the paper proceeds as follows. Section II situates the contribution in the momentum, reversal, and technical-analysis literatures. Section III reviews the MAD construction and regime representation. Section IV develops the velocity decomposition and the signals built on it. Section V describes the data, the point-in-time universe, and the measurement conventions. Section VI reports results on the index and the cross-section. Section VII discusses limitations, and Section VIII concludes.

II. Background and Related Work

A. Momentum and reversal

The empirical case for momentum is broad. Winners over the prior three to twelve months continue to outperform in the cross-section [1], and an asset's own past return forecasts its future return across markets and asset classes [3]. Reversal brackets momentum on both sides: at horizons of a week to a month, past returns predict opposite future returns [12], and at horizons of years, extreme losers overtake extreme winners [2]. The two families are less rivals than regimes of a single object, and recent work treats them as such—switching between slow momentum and fast reversion at statistically detected changepoints [4] or learning regime-adaptive blends with attention architectures [5]. The rule developed here operates at reversal's home horizon—positions are held for roughly a month—but demands a momentum-style confirmation before entering. It is, in effect, a blend of the two families expressed as a single transparent rule.

B. Moving-average signals

Within technical analysis, the moving average is the workhorse. Brock, Lakonishok, and LeBaron provided the first rigorous evidence that simple moving-average rules carry predictive content [6]; Lo, Mamaysky, and Wang gave chart-based methods a statistical foundation [7]. The case is not merely empirical: under uncertainty about the return process, allocating on a moving-average signal is theoretically justified and close to optimal for plausible parameterizations [13]. Technical indicators, moving-average distances among them, forecast the equity risk premium and capture information that macroeconomic predictors miss [14]. Moving-average timing likewise generates cross-sectional profits that momentum does not explain [8]. Closest to the present work, Avramov, Kaplanski, and Subrahmanyam show that the distance between short- and long-run moving averages predicts equity returns with economically large alphas, attributing the effect to anchoring investors tethered to the long-run average under-react to information that moves price away from it [9]. That literature studies the level of displacement. The present paper's object is the change in displacement, and specifically its composition—how much of a moving-average signal is price moving, and how much is the average catching up.

C. Machine learning and the case for transparent rules

A parallel literature applies machine learning to the same signals, from deep momentum networks [4], [5] to comprehensive treatments of machine learning in empirical asset pricing [15]. The gains are real but not unconditional: when the structure of a series is plain, shallower models can outperform deeper ones [16], and the more rules a researcher tries, the more will succeed by luck [10]. Transparency also carries value that predictive accuracy does not capture. A closed-form rule fails legibly—when its edge decays, the component that decayed can be identified and dated—whereas a learned model's edge can rarely be interrogated, a limitation the deep-learning literature itself acknowledges in its turn toward interpretable architectures [5]. A transparent rule can also be tied to a mechanism: every term in the decomposition of Section IV has a mechanical meaning, so performance can be attributed to an economic story rather than to a fitted surface. Those observations shape this paper's design. The signal is a closed-form decomposition with one threshold, the variant set is small and fully disclosed, and refinements are accepted only on paired significance tests against the base rule.

D. The companion study and the gap

The companion MAD-Markov study [11] established that position within the displacement band forecasts the direction of the next regime transition—monotonically, and with calibrated probabilities—under strictly walk-forward estimation. What it left open is whether that structure converts to trading value. This paper answers that question, and the answer carries a twist worth stating in advance: the tradeable content turns out to live not in the forecast but in the composition and timing of the transitions themselves. The calibrated probability, applied as an entry filter, either starves the rule of exposure or restates what the state already implies.

III. Moving Average Distance and Regime Representation

A. Construction

Following the companion paper [11], let P denote the closing price. The trend reference is the w-day simple moving average, with w = 20 throughout:

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

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

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

positive when price trades above its average and negative when below. Because a fixed displacement is not equally meaningful across volatility regimes, MAD is standardized by its own trailing variability, the sample standard deviation over a window of W = 255 trading days,

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

and the standardized displacement

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

expresses displacement in units of its trailing standard deviation, comparable across time and across assets. The standardization allows a single rule to serve an entire cross-section. Every quantity in (1)–(4) is strictly trailing, so MAD and z depend only on information available at time t; a signal computed at one day's close is executable at the next day's open with no revision and no look-ahead.

B. Regime classification

The standardized series is discretized into six ordered regimes at thresholds ±1 and ±2, with no neutral state, so that every observation carries both a direction (above or below trend) and a magnitude (mild, moderate, or extreme):

Table I. Regime definitions.

Regime	Condition	Interpretation
+3	$z > 2$	extended above trend
+2	$1 < z \le 2$	above trend
+1	$0 < z \le 1$	just above trend (P > SMA)
−1	$-1 < z \le 0$	just below trend (P ≤ SMA)
−2	$-2 < z \le -1$	below trend
−3	$z \le -2$	deeply oversold

Discretization replaces a continuous signal with an interpretable state sequence while retaining, through z, the within-band position that the transition forecast exploits.

C. Walk-forward transition forecast

For each regime visit, the companion's procedure estimates the probability that the visit's next transition is upward, denoted P(up). Interior regimes fit a logistic function of the within-band position z, using only visits that resolved before the current day. P(up) enters the variant set of Section IV in two roles: as an entry filter on confirmed crossings, and as the trigger of an anticipatory entry that acts before the crossing confirms.

IV. The MAD-Velocity Signal

A. The velocity decomposition

A regime transition is, mechanically, a change in MAD—and that change is not a single thing. Writing (2) as MAD = 100·(1 − SMA/P), a first-order expansion of a change over k days separates it into a price term and a moving-average term:

$$\Delta \mathrm{MAD}_t \approx \left(\frac{100 \cdot \mathrm{SMA}_t}{{P_t}^2}\right) \cdot \Delta P_t - \left(\frac{100}{P_t}\right) \cdot \Delta \mathrm{SMA}_t \tag{5}$$

where $\Delta X_t = X_t - X_{t-k}$, with k = 5 trading days throughout. The first term is positive when price rises—displacement improving because the security is genuinely recovering. The second is positive when the moving average falls, typically because old highs are leaving the w-day window, so displacement can “improve” while price is flat or still declining. The signed share of the change attributable to price is then

$$s_t = \frac{c_t^P}{|c_t^P| + |c_t^{MA}|} \tag{6}$$

where the numerator is the price term of (5) and the denominator sums the magnitudes of both terms. The share lies in [−1, +1]: values near +1 mark a price-dominated move, values near −1 a moving-average artifact. This quantity—the velocity of MAD-Velocity—asks of every displacement change the question the level cannot answer: is price doing the moving, or is the average catching up?

B. The composition inverts with position in the band

Measuring s at each regime transition across three decades of SPY shows that the composition is not random; it inverts with position in the band:

Table II. Price vs. moving-average composition of regime transitions (SPY, 1993–2026).

Transition	Events	Price-driven	MA-driven	Mean share
−1 → +1 (re-cross above trend)	450	65%	35%	+0.25
−2 → −1 (deep "improvement")	251	35%	65%	−0.23

A deep-oversold improvement is, roughly two-thirds of the time, the average rolling over rather than the price recovering, while a re-crossing of the trend is roughly two-thirds genuine price movement. The warning for rule design is direct: a naive "buy the improving regime" rule deep in the band is, more often than not, buying an artifact. The opportunity is equally direct: even at the re-cross, one crossing in three is hollow, and (6) identifies which—strictly causally, from trailing quantities alone.

C. Trading rules

All variants share one architecture: long or flat, a single position, no leverage, no short leg. Signals are evaluated at the close of day t−1 and executed at the open of day t; positions exit at the next open after the regime rolls back from +2 to +1, the displacement extension that marks the move as spent. No stop-loss is used, and results are gross of transaction costs (Section V). The variants differ only in the entry:

Table III. Variant set.

Variant	Entry condition	Role
Buy-and-hold	hold every day	benchmark
Confirmed re-cross	regime crossed −1 → +1	base rule
Velocity-gated re-cross	confirmed cross and s ≥ 0.50	the decomposition as an entry gate: price-dominated crossings only
Anticipatory (θ = 0.50 / 0.60 / 0.70)	state −1 and P(up) ≥ θ	enter before the cross, on the forecast
Anticipatory (unfiltered)	state −1, forecast ignored	control: isolates the forecast's contribution

The set is small and fully disclosed [10]. The confirmed re-cross is the base rule; the velocity-gated variant applies the decomposition of (5)–(6) as an entry condition; the anticipatory ladder asks whether the companion's forecast can buy entry into the bounce before confirmation; and the unfiltered control settles whether any anticipatory value comes from the forecast or from the state itself. Each refinement is accepted or rejected on a paired comparison against the base rule across the same names (Section VI).

V. Data

A. Index series

The index study uses SPY, the S&P 500 exchange-traded fund, over its full history: daily opening and closing prices from January 1993 through July 2026, adjusted for splits and dividends. The indicator stack warms up over the first year (the 255-day σ window of (3)); the backtest begins at the first valid regime day in February 1994 and spans 8,144 sessions.

B. Cross-sectional universe and point-in-time membership

The cross-section comprises every S&P 500 constituent between January 2016 and December 2025, as recorded in a point-in-time membership history [17]. Membership is reconstructed day by day from index snapshots; 728 distinct tickers touch the index during the window—roughly 525 in any calendar year, with 19 to 29 additions and removals annually. For each name, twenty years of daily bars are drawn from a commercial archive (polygon.io), providing indicator warm-up well before the evaluation window. A name qualifies once it has at least 252 member-days with a formed regime; 672 names qualify. Trading is membership-gated: a position may be entered or held only while the name is in the index, and each name's benchmark is its own buy-and-hold over those same member-days.

C. Survivorship and symbol continuity

A survivorship-free cross-section must carry the failures. The universe includes members later delisted through failure or acquisition—Silicon Valley Bank and First Republic among them—precisely the names a survivor-only universe silently drops. Two data pathologies required explicit handling. First, the membership history records renamed companies under their current symbols, while the price archive stores each era under the symbol that actually traded; fourteen lineages (Priceline→Booking and Dr Pepper Snapple→Keurig Dr Pepper among them) were stitched by mapping predecessor symbols, with the earlier symbol taking priority on overlapping dates. Second, ticker reuse—one symbol, two unrelated companies—is displaced by the same priority rule. After stitching, 99.2% of member-days carry a formed regime. The single residual is Arconic, at 77.5% coverage, whose corporate split leaves no clean predecessor symbol; it is retained with partial coverage rather than dropped.

D. Measurement conventions

Two conventions could measure a rule that is sometimes in cash. Full-calendar annualization charges every cash day against the strategy, which anchors the evaluation to the benchmark's always-invested calendar and conflates signal quality with capital utilization. This paper measures performance per day of capital deployed:

$$\text{deployed CAGR} = W^{\frac{252}{D}} - 1 \tag{7}$$

where W is terminal wealth and D the number of days a position is actually held, entry through exit inclusive. The in-trade Sharpe ratio is computed on in-position days only, at a zero risk-free rate; return per deployed day is reported as 10,000·ln(W)/D basis points. Trade consistency is reported alongside—completed round trips, win rate, days per trade, and return per trade—so that a large annualized figure earned over a handful of days cannot masquerade as an edge. Cross-sectional aggregates report both mean and median across names, and every claim of improvement between variants is tested with a paired McNemar comparison [18] on the per-name beat-own-buy-and-hold outcome. The test fits the claim being made. Each name serves as its own control—the same stock, the same window, the same member-days under both variants—so market-wide fortune cancels; the binary outcome gives each name one vote, so no single outsized winner can carry the verdict; and only names whose outcome flips between variants are informative, since under the null hypothesis that a refinement is worthless, flips are equally likely in either direction. All results are gross of transaction costs; entries and exits occur at next-day opens.

One convention requires disclosure. The cross-sectional archive is split-adjusted without dividend reinvestment. Benchmarks and strategies are computed from the same series, so every comparison is internally consistent; but buy-and-hold, always invested, forgoes more dividend accrual than a partial-exposure rule, so cross-sectional beat-rates carry a modest upward bias—on the order of half a percentage point of annualized return at typical exposure differences—and should be read accordingly. The index series is total-return adjusted and does not share this bias.

VI. Results

A. The index: entry timing

Table IV reports all variants on SPY over 8,144 sessions. The headline is the anticipatory entry: at θ = 0.50 it outperforms buy-and-hold on every deployed measure—13.4% versus 10.9% annualized on deployed capital, in-trade Sharpe 0.74 versus 0.65, 4.98 versus 4.12 basis points per day—while invested only 69% of the time and with a shallower maximum drawdown (−49.7% versus −55.2%). It is the only variant family to do so. The one column buy-and-hold wins, total return, is the exposure difference itself: the benchmark compounds across all 8,144 sessions while the anticipatory entry compounds across 69% of them; on the days capital is actually at risk—the measure of (7)—the ordering reverses. The confirmed re-cross, by contrast, trails the benchmark per deployed day (3.72 basis points): by the time displacement has re-crossed the trend, the bounce it confirms is largely spent. Waiting for confirmation costs more than it protects.

Table IV. SPY, all variants, February 1994 – July 2026 (8,144 sessions, gross).

Variant	Expo %	Trades	Total ret.	Depl. CAGR	Sharpe	bps/day	Win %	Max DD
Buy-and-hold	100.0	—	2,760%	10.93%	0.646	4.12	—	−55.2%
Confirmed re-cross	60.5	159	524%	9.82%	0.602	3.72	74.2	−46.7%
Velocity-gated	44.1	115	152%	6.72%	0.441	2.58	70.4	−61.3%
Anticipatory (0.50)	**69.1**	**153**	**1,548%**	**13.38%**	**0.741**	**4.98**	**75.8**	**−49.7%**
Anticipatory (0.60)	68.2	152	1,356%	12.92%	0.719	4.82	75.7	−49.7%
Anticipatory (0.70)	63.8	141	967%	12.17%	0.678	4.56	75.9	−57.9%
Anticipatory (unfiltered)	70.0	154	1,473%	12.96%	0.717	4.84	75.3	−48.1%

The forecast ladder is nearly flat: raising θ from 0.50 to 0.70 lowers deployed return monotonically, and the unfiltered control—entering on every −1 day with no forecast at all—performs within a fraction of the filtered variant. The θ = 0.50 filter skips the handful of −1 days on which the model predicts further deterioration, worth a small trim and nothing more; the entry's value comes from the state, not from the probability attached to it. The velocity gate, finally, is the worst performer on the index—a result flagged here and taken up in Section VI.D.

Trade-level consistency spans eras (Table V): the anticipatory entry wins at least 71% of its trades in each of four sub-periods covering 1994–2026, on 30 to 43 trades per era, with the average trade never negative—and it out-earns the confirmed entry per trade in every era. Median holding periods are roughly a month (29 sessions anticipatory, 21 confirmed).

Table V. Era consistency on SPY (completed trades, by entry-date era).

Era	Trades (Ant.)	Win % (Ant.)	Avg trade (Ant.)	Win % (Conf.)	Avg trade (Conf.)
1994–2002	38	71%	+1.11%	66%	+0.41%
2003–2010	43	72%	+1.12%	68%	+0.57%
2011–2018	42	83%	+2.16%	82%	+1.62%
2019–2026	30	77%	+4.27%	86%	+3.42%

B. The cross-section: the velocity gate

Table VI reports medians across the 672 qualifying constituents, each measured against its own buy-and-hold over its own member-days. Every variant beats a majority of names on deployed return—on a universe that includes the delistings. The velocity-gated re-cross is the standout on quality: the highest beat-rate (59.8%), median deployed return (8.6%), in-trade Sharpe (0.449), and return per deployed day (3.29 basis points), achieved at the lowest exposure (50.6% of member-days). The anticipatory family captures the most wealth per name (median total return near 48%, versus 37% for the confirmed base) by holding nearly 80% of the time at slightly lower per-day quality. Among names with at least fifteen completed trades, the top of the edge distribution includes First Republic Bank—a +58 percentage-point deployed-return edge over its own buy-and-hold across 18 trades while a constituent—evidence that the point-in-time machinery handles failures as designed, trading them while they were members and releasing them when they left.

Table VI. Cross-section: 672 point-in-time S&P 500 constituents, 2016–2025 (medians across names, gross).

Variant	Expo %	Depl. CAGR	Sharpe	bps/day	Total ret.	Win %	Trades	% beat B&H
Buy-and-hold	100.0	6.5%	0.368	2.51	56%	—	—	—
Confirmed re-cross	64.6	7.1%	0.408	2.74	37%	70.5	37	55.8
Velocity-gated	50.6	8.6%	0.449	3.29	35%	71.4	31	59.8
Anticipatory (0.50)	79.0	6.8%	0.384	2.60	48%	68.0	37	58.5
Anticipatory (0.60)	77.8	7.0%	0.393	2.69	48%	68.0	37	59.2
Anticipatory (0.70)	76.2	7.4%	0.395	2.83	48%	68.1	36	59.8
Anticipatory (unfiltered)	79.4	6.8%	0.380	2.59	48%	67.6	37	58.5

C. Paired comparisons

Table VII reports McNemar tests [18] on the per-name beat-own-buy-and-hold outcome. One comparison survives: the velocity gate over the confirmed base, 77 names to 50 ($p = 0.021$). The anticipatory entry's advantage over the base is directional but not significant across names (99 to 81, $p = 0.21$). The unfiltered control against the filtered anticipatory variant is a dead tie—14 names each way—the cleanest statement in the paper that the forecast contributes nothing beyond identifying the state. The gate's improvement is also not a tail effect: it shifts the entire distribution of per-name edges rightward, roughly doubling the names with edges above +30 points per year (27 versus 13) while thinning the no-edge center.

Table VII. Paired McNemar comparisons (beat own buy-and-hold, deployed CAGR).

Comparison	A only	B only	χ^2	p	Verdict
Velocity-gated vs. confirmed	77	50	5.32	0.021	significant
Anticipatory (0.50) vs. confirmed	99	81	1.61	0.205	directional
Anticipatory (0.50) vs. velocity-gated	**102**	**111**	**0.30**	**0.584**	**indistinguishable**
Unfiltered vs. anticipatory (0.50)	14	14	0.04	0.850	tie
Anticipatory (0.50) vs. (0.70)	37	46	0.77	0.380	indistinguishable

D. The split, and what it means

The index and the cross-section disagree about which refinement matters, and the disagreement is itself informative. On the index—a diversified aggregate whose idiosyncratic moves cancel—short-horizon dynamics are dominated by reversion, so the profitable act is to buy weakness early; a gate that demands visible price strength filters toward the moves that are already spent, and the velocity gate is worst precisely where reversion rules. Individual names carry idiosyncratic momentum that the aggregate averages away, so the profitable act there is to demand that a re-cross be price-dominated before trusting it; the gate is the only significant refinement precisely where momentum lives. One indicator, two regimes of use: buy weakness on the aggregate; demand decisive strength on the names. The decomposition of Section IV is what makes the second half of that sentence executable.

VII. Discussion and Limitations

A. What the results establish—and what they do not

The results establish three things. On the index, the confirmed re-cross forfeits most of what the anticipatory entry captures: confirmation is expensive, and the expense is measurable per deployed day. On individual names, the velocity gate is a genuine refinement—the only variant change that survives a paired test—and its improvement comes from shifting the whole distribution of per-name edges, not from a handful of outliers. And throughout, the companion's calibrated forecast is redundant for trading: its value is exhausted by naming the state, a negative result this paper considers as informative as the positive ones.

Three things the results do not establish deserve equal prominence. Deployed return is not realized wealth: a rule invested half the time earns its per-day edge only while deployed and converting that quality into portfolio wealth requires breadth—capital redeployed across names whose signals fire asynchronously. That portfolio-construction exercise is deliberately out of scope. The entry-timing result is not settled cross-sectionally; it is directional (99 names to 81) and unproven at conventional significance. And the interpretation of the index-versus-names split in Section VI.D, while consistent with the reversion and momentum literatures, is offered post hoc—a hypothesis that predicts, testably, that the gate should refine other momentum-carrying cross-sections and the anticipatory entry should time other diversified aggregates.

B. Costs, turnover, and design choices

All results are gross of transaction costs. Turnover is moderate—roughly 31 to 37 completed round trips per name over the decade, with month-long holds—so at five basis points per side, cost drag is on the order of 0.3 to 0.6 percentage points of deployed annual return depending on exposure: enough to narrow the reported gaps, not to erase the central ones. Two design choices are disclosed as tested-and-rejected rather than silently absent. A fixed 8% stop-loss degraded every variant—for a reversion-family entry, a stop systematically sells at maximum stretch—and is omitted. A forecast filter applied at the +1 state (requiring continuation to +2, rather than the −1-state bounce forecast) collapsed exposure to a handful of days per name and manufactured large, annualized figures on trivial deployment; it is exactly the artifact the trade-consistency reporting of Section V.D exists to expose. The architecture is long-only, one position, unsized: deliberately austere, so that what is measured is the signal, not the wrapper around it.

C. Robustness and future work

The indicator parameters ($w = 20$, $W = 255$) are inherited from the companion unchanged, and the decomposition parameters ($k = 5$, $s \geq 0.50$) were fixed before the cross-sectional test; none were optimized against the results reported here. Their sensitivity is nonetheless unexplored, and the single significant comparison warrants family-level caution: at five paired tests, a Bonferroni adjustment puts the gate's $p = 0.021$ at the margin of the conventional threshold (0.105 family-wise). The gate was the mechanism-first hypothesis of Table II rather than a searched rule, which mitigates but does not eliminate the concern [10]; out-of-sample replication is the appropriate arbiter. Natural extensions follow directly: portfolio construction over the gated signal, replication of the split prediction on international indices and other cross-sections, a dividend-adjusted rerun to close the disclosed benchmark bias, and position sizing by state depth for the anticipatory entry.

VIII. Conclusion

The companion study established that position within a displacement band forecasts the direction of the next regime transition; this paper asked whether that structure trades. It does—but not through the forecast. The tradeable content of Moving Average Distance lives in the composition and timing of its changes. A first-order decomposition separates every change in MAD into a price-driven component and a moving-average artifact, and the distinction is not decoration: the mixture inverts across the band, and demanding a price-dominated re-cross is the only refinement of the base rule that survives a paired significance test across 672 point-in-time S&P 500 constituents—raising the share of names beating their own buy-and-hold from 55.8% to 59.8% while improving every per-day quality measure at lower exposure. On the index, the lesson is timing rather than composition: entering at the −1 state, before the re-cross confirms, outperforms three decades of buy-and-hold on every deployed measure at 69% exposure—while the calibrated transition probability, the companion's central object, adds nothing a state label does not already carry.

The two domains disagree, and the disagreement is the finding. Buy weakness on the aggregate. Demand decisive strength on the names. Both rules are executable because the decomposition turns one question into a measurable, strictly causal quantity: is price doing the moving, or is the average catching up? That question is not specific to MAD. Every moving-average-distance indicator conflates the two motions, and every rule built on one inherits the confusion until it is separated. This paper separates them and measures the result against the universe as it actually was, on the capital actually deployed.

Appendix A. Reproducibility

The full analysis pipeline is released as five Python scripts, together with every result file cited in this paper. Price data is not redistributed: the index series (Section V.A) rebuilds from a free public source, and the cross-sectional archive rebuilds from polygon.io under a user-supplied API key, with delisted names and predecessor symbols recovered as described in Section V.C.

Table A1. Analysis pipeline

Script	Output	Role
01_spy.py	results/01_spy_*.csv	SPY history, indicator stack, walk-forward P(up), decomposition, index backtest (Tables IV–V)
02_SP500_constituents.py	results/02_sp500_constituents.csv	Point-in-time membership 2016–2025 [17], freshness-guarded
03_data.py	data/raw/<TICKER>.csv	20-year daily bars per constituent, delisted and predecessor symbols included
04_data_prep.py	data/<TICKER>.csv, _prep_summary.csv	Per-ticker indicator stack, membership flags, coverage audit
05_backtest.py	results/05_*.csv	Cross-sectional backtest, paired McNemar tests (Tables VI–VII)

References

[1] N. Jegadeesh and S. Titman, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," Journal of Finance, vol. 48, no. 1, pp. 65–91, 1993.

[2] W. F. M. De Bondt and R. Thaler, "Does the Stock Market Overreact?," Journal of Finance, vol. 40, no. 3, pp. 793–805, 1985.

[3] 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.

[4] K. Wood, S. Roberts, and S. Zohren, “Slow Momentum with Fast Reversion: A Trading Strategy Using Deep Learning and Changepoint Detection,” Journal of Financial Data Science, vol. 4, no. 1, pp. 111–129, 2022.

[5] K. Wood, S. Giegerich, S. Roberts, and S. Zohren, “Trading with the Momentum Transformer: An Intelligent and Interpretable Architecture,” arXiv:2112.08534, 2022.

[6] W. Brock, J. Lakonishok, and B. LeBaron, “Simple Technical Trading Rules and the Stochastic Properties of Stock Returns,” Journal of Finance, vol. 47, no. 5, pp. 1731–1764, 1992.

[7] A. W. Lo, H. Mamaysky, and J. Wang, “Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation,” Journal of Finance, vol. 55, no. 4, pp. 1705–1765, 2000.

[8] Y. Han, K. Yang, and G. Zhou, “A New Anomaly: The Cross-Sectional Profitability of Technical Analysis,” Journal of Financial and Quantitative Analysis, vol. 48, no. 5, pp. 1433–1461, 2013.

[9] 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.

[10] R. Sullivan, A. Timmermann, and H. White, “Data-Snooping, Technical Trading Rule Performance, and the Bootstrap,” Journal of Finance, vol. 54, no. 5, pp. 1647–1691, 1999.

[11] L. Arrington, “Position Within Trend: The MAD-Markov Model for Calibrated Regime-Transition Forecasting,” independent research, July 2026.

[12] N. Jegadeesh, “Evidence of Predictable Behavior of Security Returns,” Journal of Finance, vol. 45, no. 3, pp. 881–898, 1990.

[13] Y. Zhu and G. Zhou, “Technical Analysis: An Asset Allocation Perspective on the Use of Moving Averages,” Journal of Financial Economics, vol. 92, no. 3, pp. 519–544, 2009.

[14] C. J. Neely, D. E. Rapach, J. Tu, and G. Zhou, “Forecasting the Equity Risk Premium: The Role of Technical Indicators,” Management Science, vol. 60, no. 7, pp. 1772–1791, 2014.

[15] S. Gu, B. Kelly, and D. Xiu, “Empirical Asset Pricing via Machine Learning,” Review of Financial Studies, vol. 33, no. 5, pp. 2223–2273, 2020.

[16] 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.

[17] fja05680, “S&P 500 Historical Components & Changes,” GitHub repository, accessed July 2026.

[18] Q. McNemar, “Note on the Sampling Error of the Difference Between Correlated Proportions or Percentages,” Psychometrika, vol. 12, no. 2, pp. 153–157, 1947.