MAD-Manifold: Mapping and Testing the Phase-Space Flow of Moving-Average Distance

Lawson Arrington[1]

July 17, 2026

Abstract—Two companion studies established that a security's displacement from its moving average carries structure. Position within the displacement band forecasts the direction of the next regime transition; the composition of a displacement change—price moving versus the average catching up—converts into a tradeable entry gate. This paper asks the question both results imply: taken together, as position and velocity, what kind of object are displacement dynamics? Each security is embedded in the phase plane of standardized displacement and its five-day change, and every statistic is tested against autoregressive surrogates matched to the series' own persistence. The answer arrives in two halves. The orbit's geometry is unremarkable—its size and preferred radius are exactly what calibrated linear mean reversion produces. The traffic on the orbit is not. The circulation is clockwise, time-irreversible, and asymmetric; the rising and falling branches of the mean cycle fail to mirror one another ($p < 0.005$), irreversibility appears at two distinct timescales with opposite sign, and the asymmetry holds in both halves of three decades of SPY. Across 715 point-in-time S&P 500 constituents and 2.9 million daily observations, one flow field fits the whole universe—disjoint random halves of the cross-section yield fields with cosine similarity 0.99, the field is equally stable across decades, and 90% of individual names carry the same positive loop asymmetry where matched surrogates split evenly. Displacement dynamics evolve on a shared, stationary surface whose shape linear models explain and whose flow they cannot.

[1] *Disclaimer: MAD-Manifold 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

Two companion studies examined a security's displacement from its moving average, one coordinate at a time. The first established that position within the displacement band forecasts the direction of the next regime transition, monotonically and with calibrated probabilities [1]. The second established that velocity—specifically, the composition of a displacement change, price moving versus the average catching up—identifies which regime transitions are genuine, and converts that distinction into a trading gate that survives a paired significance test across the index's constituents [2]. Each result treats one projection of the same underlying object. This paper studies the object: displacement dynamics taken whole, as a trajectory through position and velocity jointly.

Plotting position against velocity is the oldest instrument of dynamics—the phase portrait, in which a system's behavior becomes geometry: equilibria appear as points, oscillations as closed orbits, and the rules of motion as a vector field [3]. A remarkable pair of results guarantees that the instrument applies far beyond mechanics: the state space of a dynamical system can be reconstructed, up to smooth deformation, from delayed observations of a single measured variable [4], [5]. Finance encountered these ideas in the late 1980s, when a literature searched stock returns for low-dimensional chaos and strange attractors [6], [7]; the search ended ambiguously, defeated by noise and nonstationarity, and phase-space methods largely left the field with it [8]. The question this paper asks is deliberately humbler than the one that failed. Not: are returns secretly deterministic? But: does a specific, already-validated indicator—standardized moving-average displacement—organize into a stable geometric structure when embedded with its own velocity? A humble question, but a falsifiable one.

The discipline comes from two tools the chaos literature left behind. First, surrogate data testing [9]: every statistic reported here is scored against ensembles of autoregressive surrogates fitted to each series' own persistence, so that "there is a loop" can never be mistaken for evidence—linear mean reversion also draws loops. What linear processes cannot do is break time symmetry: any Gaussian linear process is statistically identical run forward or backward, so genuine cycle asymmetry—fast declines and slow recoveries, corridors traversed in one direction—is a signature no autoregression can forge [10]. Asymmetry of exactly this kind is documented in macroeconomic aggregates [11]; whether it lives in the phase plane of a technical indicator, name by name across an equity universe, is the empirical question here.

The answer arrives in two halves, and the split is the finding. The shape of the orbit—its size, its preferred radius—is exactly what calibrated linear mean reversion produces; on geometry alone, the null survives. The flow on that orbit does not: the circulation is clockwise and time-irreversible, its rising and falling branches fail to mirror one another at $p < 0.005$, the irreversibility appears at two distinct timescales with opposite sign, and every one of these properties holds in both halves of three decades. Pooling 715 point-in-time S&P 500 constituents—2.9 million daily steps—one flow field fits the entire universe: disjoint random halves of the cross-section yield fields agreeing to cosine similarity 0.99, and ninety percent of individual names carry the same positive loop asymmetry where phi-matched surrogates split evenly. In the vocabulary of the machine-learning manifold hypothesis [12], the cross-section of displacement dynamics concentrates on a single low-dimensional structure—and this paper maps the traffic on it.

The remainder proceeds as follows. Section II reviews the phase-space toolkit and situates the contribution. Section III builds the constructions: the displacement stack inherited from the companions, the phase-plane embedding, and Moving Average Area, the signed area under displacement that the loop tests generalize. Section IV specifies the test battery and its null models. Section V describes the data. Section VI reports results for the index and the cross-section. Section VII discusses limitations, and Section VIII concludes.

II. Background: The Phase-Space Toolkit

A. Phase portraits and delay embedding

A phase portrait renders dynamics as geometry. For a mechanical system the axes are position and momentum; equilibria appear as fixed points, sustained oscillations as closed orbits, and the rule of motion as a vector field whose streamlines the trajectories follow [3]. The instrument would be useless for markets if it required knowing the true state variables—but it does not. The embedding theorems establish that delayed copies of a single smooth observable reconstruct the underlying state space up to smooth deformation [4], [5]. The pair used throughout this paper—standardized displacement and its own k-day change—is the minimal such embedding: (z, Δz) is a linear transformation of the delay pair (z_t, z_{t-k}), so whatever structure survives in this plane is structure of the dynamics, not an artifact of coordinates.

B. The chaos episode, and its lesson

Finance has visited phase space before. A late-1980s literature computed correlation dimensions and related statistics on stock returns, asking whether markets hide low-dimensional deterministic chaos [6]. The rigorous close of that literature found evidence of nonlinearity but not of determinism: dimension estimates are fragile under noise, and returns are too nonstationary to hold still for them [7], [8]. Two lessons survive the episode. Methodological: do not estimate what cannot be estimated at financial noise levels; test specific geometric properties against explicit null models instead. And practical: raw prices and returns are the wrong observable. Standardized displacement is not—MAD is stationary by construction, bounded in practice, and volatility-normalized, which is precisely the stability the 1980s program lacked in its raw material.

C. Surrogates and the arrow of time

The null-model discipline is the surrogate data method [9]: fit the null process to the observed series, simulate an ensemble, compute every statistic on real and surrogate data alike, and let the ensemble supply the distribution under the null. The null here is first-order autoregression matched to each series' own persistence—the mathematical form of "it is just mean reversion." One theorem gives the method its teeth: a stationary Gaussian linear process is time-reversible—statistically identical run forward or backward [13]. Time-irreversibility is therefore not merely evidence against a particular null but against the entire linear Gaussian class, and it has an economic pedigree: business-cycle aggregates fall faster than they rise [11], and irreversibility tests formalize the observation [10]. The loop asymmetry at the center of this paper is a phase-plane expression of the same arrow of time.

D. The manifold hypothesis, and this paper's position

Machine learning supplies the cross-sectional frame: the manifold hypothesis holds that high-dimensional data concentrate near low-dimensional structure, recoverable from the data themselves [12]. Its econometric translation here is concrete—715 securities, each tracing its own path in the same normalized plane, either scatter idiosyncratically or concentrate on one shared field. Prior work has studied single-series dynamics or applied dimension reduction to return panels; the contribution of this paper is the combination: a validated, stationary indicator as the observable, explicit autoregressive nulls for every claim, time-irreversibility as the discriminating physics, and cross-sectional replication—the same field from disjoint halves of the universe—as the manifold test. What hangs on the outcome deserves a sentence: if disjoint sets of companies trace one field, the dynamics belong to the market's machinery rather than to any particular stock—and every rule built on displacement, the companions' included, has been navigating that geometry without a map. The geometry is asked to prove itself the way the companions' trading claims were against a null, out of sample, and name by name.

III. The Constructions

A. The displacement stack

The observable is inherited unchanged from the companions [1], [2]. 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,

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

and displacement is standardized by its own trailing variability, the sample standard deviation over a window of W = 255 trading days,

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

giving the standardized displacement

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

Every quantity in (1)–(4) is strictly trailing, and z is stationary by construction and comparable across time and across assets—the properties Section II.B identified as prerequisites for phase-space work on financial data.

B. The phase-plane embedding

The velocity coordinate is the k-day change in standardized displacement, with k = 5 throughout, matching the decomposition horizon of the companion [2]:

$$v_t = z_t - z_{t-k} \tag{5}$$

Each security is embedded as the pair (z, v), with each coordinate divided by its own sample standard deviation so that every security occupies the same normalized plane and cross-sectional pooling is meaningful. By Section II.A this is a two-dimensional delay embedding: (z, v) is a linear transformation of (z_t, z_{t-k}), so the plane inherits the reconstruction guarantees of the embedding theorems. One mechanical property must be declared before any test is stated. When v is positive, z is by definition rising, so trajectories move rightward in the upper half-plane and leftward in the lower—every mean-reverting series circulates clockwise in this plane, real or simulated. Circulation is therefore never treated as evidence. The tests of Section IV measure properties of the circulation that a time-reversible process cannot possess.

One terminological boundary should be drawn explicitly. The companion [2] used "velocity" for the composition of a displacement change—the share of a MAD move attributable to price rather than to the moving average. The velocity coordinate here is the total change itself, undecomposed and standardized, because the embedding theorems of Section II.A require the full motion of the observable, whatever its cause. The two are complements rather than competitors: v measures how fast displacement moves; the companion's price share records who moved it. Coloring the present plane by that share is the natural next construction, and Section VII returns to it.

C. Moving Average Area

The loop statistics of Section IV integrate velocity around the orbit. Moving Average Area is the complementary integral of position—the signed area under the displacement curve over a rolling window of length L (with daily steps, a sum):

$$\mathrm{MAA}_t = \sum_{i=0}^{L-1} \mathrm{MAD}_{t-i} \tag{6}$$

Because MAD is signed, each window splits into a positive lobe A^+, the area contributed by days above trend, and a negative lobe A^-, the magnitude contributed by days below, with MAA = $A^+ - A^-$. The bounded ratio

$$\mathrm{bal} = \frac{A^+ - A^-}{A^+ + A^-} \tag{7}$$

runs from +1 for a window spent entirely above trend to −1 for one entirely below, independent of scale. Area connects the rolling view to the orbital one. Over one complete cycle of the displacement process—one upward zero-crossing of z to the next—the net standardized area is the discrete counterpart of a contour integral, and a time-reversible process balances its lobes in expectation. Per-cycle net area is thus the simplest expression of the asymmetry the phase-plane tests formalize, with one caution attached: positive drift in the underlying contributes positive area on its own, so Section VI reads the per-cycle statistic descriptively and lets the loop-asymmetry test, which is immune to location, carry the inference.

IV. The Test Battery and Its Nulls

A. The conditional mean loop

The orbit is summarized by its two branches. Binning the plane by position, the up branch u(x) is the mean velocity among rising observations at each position, and the down branch d(x) the mean among falling ones. Two integrals summarize the pair:

$$A_{\text{loop}} = \int (u(x) - d(x))\, dx \tag{8}$$

$$S_{\text{loop}} = \int (u(x) + d(x))\, dx \tag{9}$$

The first is the area of the mean orbit—its size. Every mean-reverting process, linear or not, produces a positive A, so size alone proves nothing. The second is the orbit's asymmetry, and it is the paper's central statistic: for any time-reversible process the down branch is the mirror image of the up branch, d(x) = −u(x), and S is identically zero. A nonzero S states that the system travels different roads in its two directions.

B. Time-irreversibility

The same property is tested without reference to the plane. For horizon k,

$$\mathrm{TR}_k = \frac{E\left[(z_t - z_{t-k})^3\right]}{\left(E\left[(z_t - z_{t-k})^2\right]\right)^{3/2}} \tag{10}$$

is the standardized skewness of k-day displacement changes—zero for every stationary Gaussian linear process [13], and the workhorse of irreversibility testing in economics [10]. It is computed at k=1, 5, 10, and 20, so that asymmetry is measured at horizons from a day to a month rather than asserted at one.

C. Radial drift and rotation

In polar coordinates, r measures distance from the plane's center and θ the angle. The radial drift profile $E[\Delta r \mid r]$ fingerprints an attracting orbit—outward drift near the center, inward beyond, a zero crossing at the preferred radius r^*—and the rotation statistics record the mean wrapped angular step and its consistency. Both are reported against the surrogate ensemble rather than as standalone tests, because a mean-reverting null also drifts toward a preferred radius and also circulates; the question these statistics answer is whether the observed geometry differs from the null's, not whether it exists.

D. The field and its similarity

The empirical vector field is the plane gridded into cells, each carrying the mean one-step motion of the trajectories passing through it, retained only where the cell is well populated. A cell's flow is called coherent when its mean vector exceeds twice its standard error; at pooled sample sizes in the millions this criterion saturates for real and surrogate data alike, so it serves as a floor, not a discriminator. The informative field statistics are similarities: two fields are compared by the cosine between their concatenated cell vectors, computed across halves of time (is the field stationary?) and across repeated random halves of the universe of names (is the field shared?). The second comparison is the operational form of the manifold hypothesis.

E. The surrogate ensemble, and what each test can conclude

Every null is first-order autoregression fitted to the series under test—the lag-one autocorrelation sets the persistence, the innovation variance is matched, and ensembles are simulated at equal length. For the index series the ensemble is 200 surrogates; for the cross-section, matched surrogates are generated per name, and every statistic is computed identically on real and simulated data. Reported p-values are two-sided ensemble frequencies. Table I states, before any result, what the null produces for each statistic and what a rejection therefore means—the reading protocol is fixed in advance. The choice of null is deliberate. First-order autoregression is the mathematical form of the competing explanation—that displacement is nothing but persistent mean reversion—fitted to each series rather than asserted, so a rejection dismisses exactly the story a skeptic would tell. The null omits heavy tails and volatility clustering by design: those enrichments make a series wilder without making it directional, and the battery's central statistics measure odd-moment, signed asymmetry that no sign-symmetric process supplies, however clustered its variance. What survives the ensemble survives for dynamical reasons, not distributional ones.

Table I. The test battery: null behavior and interpretation.

Statistic	**Under the AR(1) null**	**A pass means**
A (loop area, eq. 8)	positive — the null draws orbits	nothing alone; context only
S (loop asymmetry, eq. 9)	zero — branches mirror	genuine directional asymmetry
TR_k (eq. 10)	zero — reversible	irreversibility at horizon k
E[Δr | r] and r*	exists — the null has a radius	descriptive, read against the null profile
rotation consistency	positive — the null circulates	excess organization if above the ensemble
cell coherence	saturates at pooled sample sizes	floor only, not evidence
cosine, time halves	high if stationary	field stability
cosine, name halves	—	one field shared across the universe

V. Data

A. Index series

The index study uses SPY, the S&P 500 exchange-traded fund, over its full history: daily closes from January 1993 through July 2026, adjusted for splits and dividends. The indicator stack of Section III consumes roughly one trading year of warm-up (the 255-day σ window), leaving 8,149 valid observations of z from February 1994 onward. All SPY-level statistics, and their two hundred surrogate ensembles, are computed on this series.

B. Cross-sectional universe

The cross-section is inherited from the companion study, which constructed it for the trading evaluation and documents its provenance in full [2]: every S&P 500 constituent between January 2016 and December 2025 as recorded in a point-in-time membership history [14], including members later delisted through failure or acquisition, with renamed and predecessor symbol lineages stitched so that each company's price history is continuous under its historical ticker. That construction was built to defeat survivorship bias in a trading test; it serves the same purpose here—the geometry is asked to hold for the failures as well as the survivors. For each of the 728 universe tickers, twenty years of daily bars are drawn from a commercial archive (polygon.io); names with insufficient history for the indicator warm-up are dropped, leaving 715 usable series and 2.89 million pooled daily steps in the phase plane.

One deliberate difference from the companion deserves note. Membership flags are not applied here: the manifold statistics use each name's full history, in and out of the index, because the object under test is the dynamics of a series, not the returns of a constituent-gated strategy—and restricting to membership windows would discard most of each name's trajectory without changing what the trajectory is. The universe definition is inherited purely so that the two papers describe the same companies.

Two data conventions matter less here than they did in the companion. Prices are split-adjusted without dividend reinvestment; MAD is a ratio of price to its own average, so any consistent price scaling cancels and the adjustment convention is immaterial to z. And a handful of merged files splice symbol-reuse eras of unrelated companies; each splice contributes one anomalous transition among several thousand per name, is diluted further by per-series standardization, and is disclosed rather than repaired.

C. Reproducibility

The pipeline is released as four scripts. The Moving Average Area and closed-loop tests—including every SPY result—reproduce from free data with no credentials; the universe pull and the cross-sectional field require a user-supplied polygon.io key. The appendix maps scripts to outputs.

VI. Results

A. Moving Average Area

Area behaves as the constructions anticipate. At short windows it is largely a smoothed level—the correlation between MAA and MAD is 0.89 at $L = 5$—but the redundancy decays to 0.24 by $L = 60$: long-window area carries content the level does not. The positive lobe exceeds the negative at every window, and the share of net-positive windows rises from 65% at $L = 5$ to 72% at $L = 60$—the drift of the underlying made visible as area. Per complete cycle (472 upward zero-crossings of z, median length 12 trading days), net standardized area averages +3.29 with a mean to standard error ratio of 3.7 and 55.5% of cycles positive—read descriptively, per the caution of Section III.C, with the location-immune loop statistic left to carry the inference. Two distributional details sharpen the picture. Short windows are one-sided—73% of five-day windows sit almost entirely in a single lobe, a share that falls to 11% by $L = 60$ as windows begin to straddle full cycles—and per-cycle area is heavy-tailed: the median cycle nets just +0.25 against the +3.29 mean, so most cycles roughly balance their lobes while a minority of long above-trend excursions carry nearly all of the surplus. A forward-return probe echoes the companion's index finding [2]: the deepest-negative balance quintile—windows spent hard below trend—precedes twenty-day returns near +1.8%, against roughly +0.5% for balanced windows.

B. The index loop

Table II reports the battery on SPY against 200 matched surrogates. The two statistics the protocol designated non-discriminating behave exactly as designated: the loop area matches the ensemble ($p = 0.92$), and the preferred radius is the null's almost to the third decimal (1.12 versus 1.13, $p = 0.63$). The orbit's shape is linear mean reversion's. The designated discriminators reject. The loop asymmetry is +0.467 against a surrogate mean of +0.004—no ensemble member approached it—and irreversibility appears at the five-day and twenty-day horizons with opposite signs: at the weekly scale displacement falls faster than it recovers ($TR_5 = -0.17$), at the monthly scale the skew inverts ($TR_{20} = +0.18$), and the ten-day horizon sits at the crossover between the two regimes. Rotation is clockwise, as mechanics requires, but more consistent than the ensemble produces ($p = 0.005$). The asymmetry holds in both halves of the sample, +0.33 and +0.58, and the radial drift profile traces the limit-cycle fingerprint—outward below $r \approx 1.1$, inward above—while the surrogates trace the same one: per the protocol, the annulus is not the finding. Figure 1 shows the portrait.

Table II. SPY closed-loop battery vs. 200 AR(1) surrogates (1994–2026).

Statistic	Observed	Surrogate mean	p (two-sided)	Protocol verdict
A (loop area)	+5.08	+5.28	0.92	null-like, as designated
S (loop asymmetry)	+0.467	+0.004	< 0.005	reject: asymmetric
TR_1	+0.013	−0.002	0.66	null-like
TR_5	−0.171	0.000	< 0.005	reject: irreversible (weekly)
TR_{10}	+0.017	−0.001	0.70	crossover
TR_{20}	+0.182	0.000	< 0.005	reject: irreversible (monthly)
rotation consistency	0.172	0.140	0.005	excess organization
r*	1.119	1.130	0.63	null-like, as designated

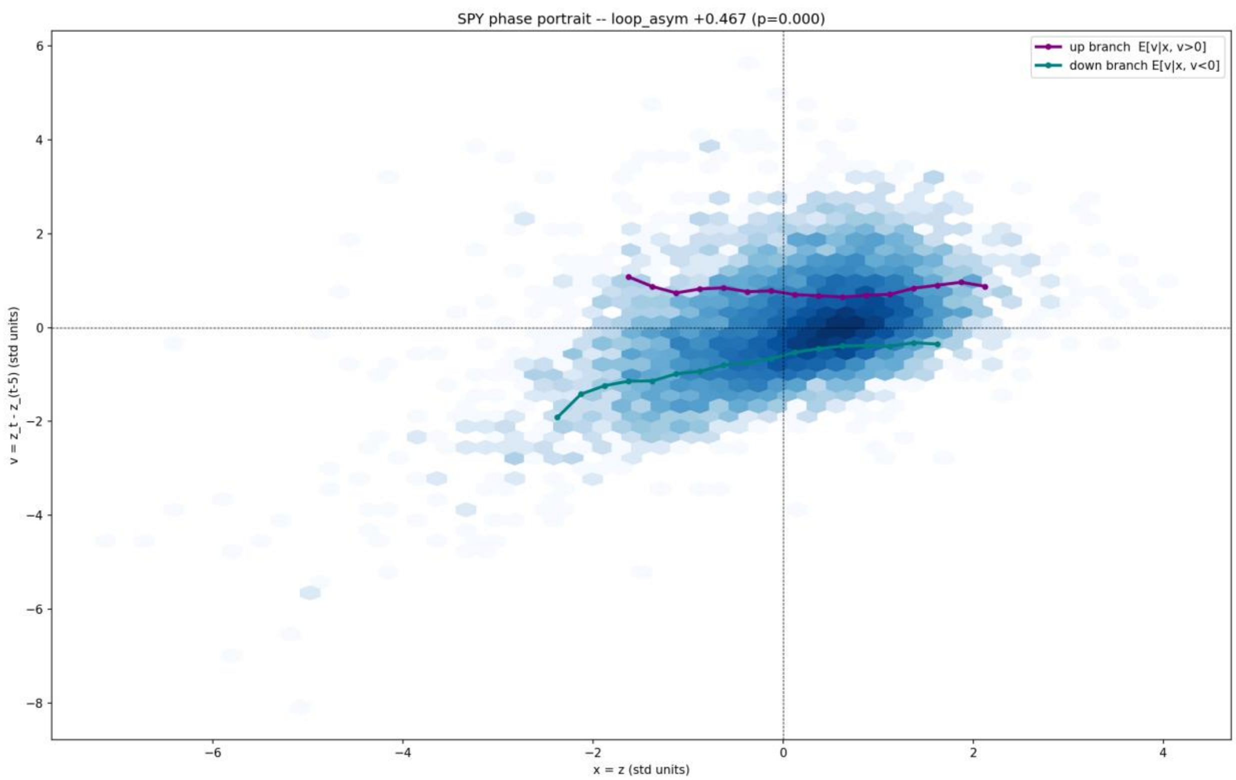

Figure 1. SPY phase portrait: trajectory density with the conditional mean loop. The up and down branches fail to mirror—the loop asymmetry of Table II.

C. The cross-section

Table III reports the flow-manifold verdicts on 715 names and 2.89 million pooled steps. The coherence floor saturates for the real field and the surrogate universe alike, exactly as Section IV.D anticipated at these sample sizes; it is reported and set aside. The similarities decide. The field built from the first half of each name's history and the field built from the second agree to cosine 0.988; fields built from disjoint random halves of the universe agree to 0.993 on average and never below 0.990 in twenty splits—roughly 357 companies per side, drawing the same vector field each time. Name by name, 89.5% of the 708 series with sufficient data carry positive loop asymmetry (mean +0.232, t = +35), where 3,551 surrogates matched to each name's own persistence split 50.2% to 49.8%. The two-timescale signature replicates across the universe: the five-day statistic is negative in 70% of names, the twenty-day positive in 79%. One caution attaches: names share a market factor, so the t-statistic overstates the independent evidence; the ninety-versus-fifty contrast carries the conclusion without it. The index's asymmetry, +0.47, is roughly twice the per-name mean—aggregation cancels idiosyncratic noise and concentrates the common cycle, making the index the cleanest window on the shared loop. Figure 2 shows the pooled field; Figure 3, the per-name distribution against its null.

Table III. Flow-manifold verdicts: 715 names, 2.89M steps, 2006–2026.

Verdict	Value	Reference
cell coherence	100% of cells	surrogate universe: 100% (floor; set aside)
field stability (time halves)	cosine +0.988	—
shared field (name halves)	cosine +0.993 (min +0.990)	20 random splits
loop asymmetry per name	89.5% positive; mean +0.232 (t = +35)	surrogates: 50.2% positive
TR_5 per name	negative in 70.3%	null: 50%
TR_{20} per name	positive in 79.0%	null: 50%

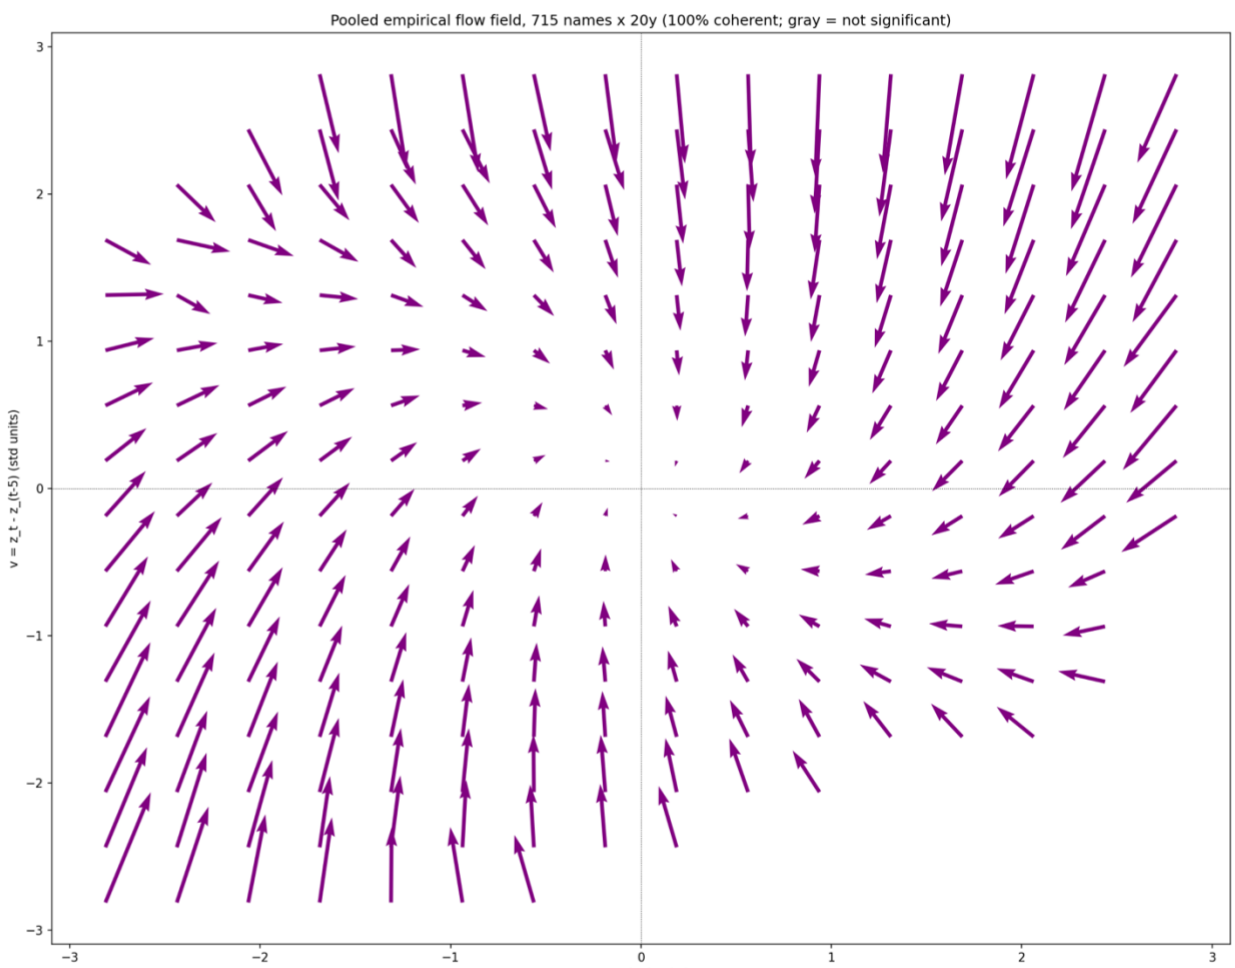

Figure 2. The pooled empirical flow field, 715 names over two decades. The circulation is clockwise; folding the field through the origin does not map it to itself—the visible form of the loop a symmetry.

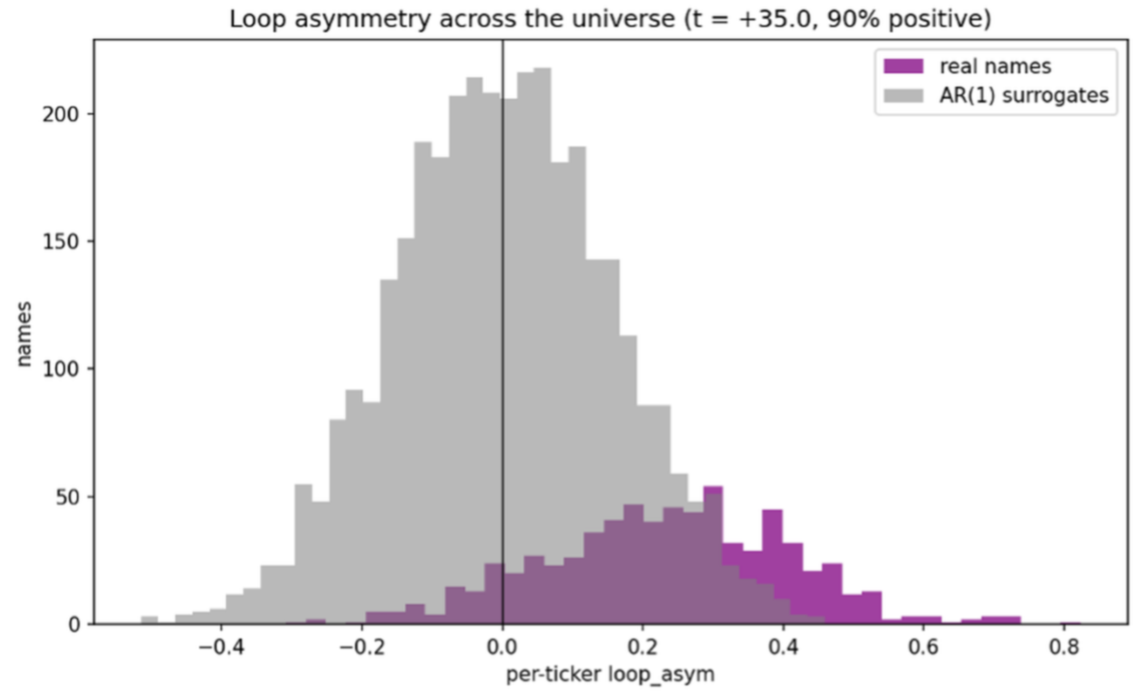

Figure 3. Per-name loop asymmetry against the surrogate ensemble. The distribution sits bodily to the right of the null.

D. What passed, and what did not

Read against the protocol of Table I, the results divide cleanly. Everything the null can produce, it produced: the orbit's size, its preferred radius, its saturated coherence. Everything the null cannot produce, the data produced anyway: a loop whose branches do not mirror, irreversibility at two horizons with a sign inversion between them, rotation more organized than mean reversion supplies—each stable across decades, and each replicated in roughly nine of every ten names against surrogates that split evenly. The shape of the orbit belongs to linear mean reversion. The traffic on it does not, and the traffic is the same in every corner of the universe tested.

VII. Discussion and Limitations

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

The results establish three things. Displacement dynamics carry a circulation that is asymmetric, time-irreversible, and organized beyond what matched linear mean reversion produces—properties that are stable across decades and replicated in roughly nine of ten names. The asymmetry has structure of its own: two horizons of opposite sign, a weekly scale that falls faster than it recovers and a monthly scale that inverts. And the geometry that hosts this flow is ordinary—the annulus, its radius, its size are all the null's—so the finding is located precisely in the flow and nowhere else.

Three things the results do not establish deserve equal billing. They do not establish trading utility: no test in this paper shows that conditioning on the phase-plane state improves a forecast or a portfolio over conditioning on position alone, and the companions' results [1], [2] set exactly the bar—paired tests, name by name—that such a claim would have to clear. They do not establish a manifold in the topological sense; the term is used operationally throughout, for a shared low-dimensional structure that disjoint halves of the universe reproduce, and nothing stronger is claimed. And they do not identify a mechanism. Anchoring, leverage, flow-driven feedback, and volatility dynamics are all candidates for why displacement travels different roads in its two directions; the tests here can reject reversibility, not adjudicate its cause.

B. Alternative explanations

Two deserve explicit statement. First, the cross-sectional replication is not 708 independent experiments—names share a market factor, and the t-statistic of 35 should be read as arithmetic, not as evidence strength; the comparison that carries the conclusion is nine-of-ten names positive against surrogates that split evenly, a contrast no plausible correction erases. Second, the surrogate class excludes asymmetric volatility models. A GARCH process with a leverage effect is not sign-symmetric and could in principle contribute odd-moment structure of the kind measured here; the present battery rejects the linear reversible class, and promoting the null to asymmetric-volatility alternatives is the natural sharpening—a harder gauntlet this paper's framework can run unchanged. Parameter dependence is disclosed rather than explored: $k = 5$, $w = 20$, and $W = 255$ are inherited from the companions, not tuned, and every statistic faces surrogates with identical windows and identical overlap, so the p-values compare like with like.

C. Future work

The extensions rank themselves. First, the utility bar: condition the companions' entry signals on phase-plane state and test, with their own paired-comparison discipline, whether two coordinates beat one—the question that connects this geometry back to capital. Second, color the plane by the companion's price share, separating price-driven from average-driven traffic within the field; Section III.B promised this construction and the decomposition is already built [2]. Third, promote the null to asymmetric-volatility classes, and replicate the battery on other indices, other asset classes, and other markets—a shared field that survives international replication would considerably strengthen the manifold reading. Last, the exit side: the flow near the upper band is where positions come home, and the field's geometry there is unexplored territory with a standing customer.

VIII. Conclusion

When this line of research began, the Moving Average Distance was an indicator. The first companion showed that its position forecasts regime transitions; the second showed that the composition of its motion selects the tradeable ones; this paper asked what the whole object is—and the answer divides exactly in half. The shape of displacement dynamics is unremarkable: an attracting annulus of the same size and radius that calibrated linear mean reversion builds. The motion on that shape is not: a clockwise circulation whose branches refuse to mirror, irreversible at two horizons with opposite signs, more organized than any of its matched surrogates, unchanged across two decades, and drawn identically by disjoint halves of a 715-name universe. Nine of ten individual companies carry the asymmetry that every linear clone lacks.

The companions' results now read as consequences of the geometry. A confirmed re-cross enters late because the circulation is asymmetric—by the time displacement re-crosses its trend, the fast half of the cycle is already spent. The velocity gate works because genuine motion and frame drift are different flows on the same field. What remains unproven is stated plainly: no test here converts geometry into capital, and whether two coordinates beat one—in forecasts, in portfolios, name by name against the companions' own paired-test bar—is the open question. The market's displacement dynamics run on a shared, stationary surface. This paper measured the traffic. The next one puts the map to work.

Appendix A. Reproducibility

The full pipeline is released as four Python scripts, together with every result file and figure cited in this paper. Price data is not redistributed. The index-level results—Moving Average Area and the complete closed-loop battery, including Table II and Figure 1—reproduce from free data with no credentials. The cross-sectional results rebuild from polygon.io under a user-supplied API key, with the universe constructed as in Section V. Surrogate ensembles use a fixed random seed, so every p-value in the paper is reproducible exactly.

Table A1. Analysis pipeline (Python; pandas/numpy; yfinance for the index series, polygon.io for the cross-section).

Script	Output	Role
01_data.py	data/raw/<TICKER>.csv	Universe (point-in-time membership [14], rename and predecessor stitching) + 20-year daily bars (user key)
02_maa_spy.py	results/02_maa_spy.csv	SPY displacement stack + Moving Average Area: lobes, balance, per-cycle areas (free data)
03_loop_spy.py	results/03_loop_*.csv, fig_03_*	SPY phase portrait, closed-loop battery, 200 AR(1) surrogates (Table II, Figure 1; free data)
04_flow_manifold.py	results/04_*.csv, fig_04_*	Pooled cross-sectional field, similarity verdicts, per-name replication (Table III, Figures 2–3)

References

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

[2] L. Arrington, "MAD-Velocity: Constructing and Evaluating a Moving-Average-Distance Trading Signal," independent research, July 2026.

[3] S. H. Strogatz, Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering, 2nd ed. Boulder, CO: Westview Press, 2015.

[4] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, "Geometry from a Time Series," Physical Review Letters, vol. 45, no. 9, pp. 712–716, 1980.

[5] F. Takens, "Detecting Strange Attractors in Turbulence," in Dynamical Systems and Turbulence, Lecture Notes in Mathematics, vol. 898, Springer, pp. 366–381, 1981.

[6] J. A. Scheinkman and B. LeBaron, "Nonlinear Dynamics and Stock Returns," Journal of Business, vol. 62, no. 3, pp. 311–337, 1989.

[7] D. A. Hsieh, “Chaos and Nonlinear Dynamics: Application to Financial Markets,” Journal of Finance, vol. 46, no. 5, pp. 1839–1877, 1991.

[8] R. N. Mantegna and H. E. Stanley, An Introduction to Econophysics: Correlations and Complexity in Finance. Cambridge, UK: Cambridge University Press, 2000.

[9] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer, “Testing for Nonlinearity in Time Series: The Method of Surrogate Data,” Physica D, vol. 58, no. 1–4, pp. 77–94, 1992.

[10] J. B. Ramsey and P. Rothman, “Time Irreversibility and Business Cycle Asymmetry,” Journal of Money, Credit and Banking, vol. 28, no. 1, pp. 1–21, 1996.

[11] S. N. Neftci, “Are Economic Time Series Asymmetric over the Business Cycle?,” Journal of Political Economy, vol. 92, no. 2, pp. 307–328, 1984.

[12] J. B. Tenenbaum, V. de Silva, and J. C. Langford, “A Global Geometric Framework for Nonlinear Dimensionality Reduction,” Science, vol. 290, pp. 2319–2323, 2000.

[13] G. Weiss, “Time-Reversibility of Linear Stochastic Processes,” Journal of Applied Probability, vol. 12, no. 4, pp. 831–836, 1975.

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