Scheduled Coordination Density and the Sign of Overnight Returns in U.S. Equity Markets

Avaneendra Trivedi
Independent Researcher
avaneendra22@gmail.com

May 2026

Abstract

We introduce Scheduled Coordination Density (SCD), an ex-ante measure of how much publicly scheduled institutional trading mass from heterogeneous calendar families is time-aligned to the same closing execution window. SCD is built only from public institutional calendars and public price and volume data. It decomposes exactly into a self-mass term and a cross-clock overlap term; the normalized overlap term, which divides out scale, is the object of study. Across a fifteen-ETF daily panel (2016 to 2024, 32,130 ETF-days), the overlap term is non-spurious: its coefficient lies at the 0th percentile of a 1,000-iteration calendar-randomization null (two-sided p = 0.000), and a Monte Carlo power analysis confirms the design is well powered (minimum detectable effect 2.9 bp per standard deviation at 80 percent power). The headline finding is directional and was pre-registered against the opposite hypothesis: higher coordination predicts overnight return CONTINUATION, not the transitory reversal we pre-registered (overlap coefficient -5.03 bp per SD, t = -2.61 date-clustered, -2.83 ETF-clustered, -2.13 two-way). The effect is modern-era specific: it is absent over the full 1993 to 2024 history (t = -0.83), present in 2010 to 2024 (t = -2.46), and strongest in 2016 to 2024, consistent with the growth of leveraged ETFs and short-dated options raising clock-family overlap intensity. A complementary, sign-consistent but statistically underpowered signal indicates the effect flips toward reversal at maximum-coordination quarter-end events (Phase 1 quarter-end subsample t = 2.56) and in individual stocks (quarter-end-down +48 bp per SD, t = 1.94). We pre-registered and rejected four distinct directional 'transient-pressure' hypotheses across five phases, committing each rejection as a first-class artifact, and we report a clean unconditional fact: the final thirty-minute window carries roughly fifteen percent of daily trading volume in eight percent of trading time. SCD is a real, non-spurious, ex ante coordination primitive whose directional signature favors information aggregation over transient pressure in liquid instruments; isolating the transient-pressure channel cleanly requires closing-auction imbalance data, which we identify as the binding constraint.

JEL Classification: G12, G14, G17, G23, C58

Keywords: closing auction, scheduled rebalancing, coordination, overnight returns, market microstructure, passive investing, leveraged ETFs, options expiration, pre-registration, exchange-traded funds

Pre-registration: *The primary specification was committed to version control on 2026-05-22T15:46:38+05:30, before any headline regression was executed. The analysis code refuses to run unless the pre-analysis plan is committed with a timestamp predating the run. All outputs carry a SHA-256 audit trail and a run identifier; the canonical Phase 1 run is 6aeb77923ad53e54. Four directional rejection artifacts are committed and reported.*

1. Introduction

Most of what we know about the daily closing price of a U.S. equity comes from studying one calendar at a time. Index reconstitution has its own literature, futures rolls have another, options expiration a third, leveraged exchange-traded fund (ETF) releveraging a fourth, and month-end and quarter-end rebalancing a fifth. Each of these is a public, scheduled institutional event that concentrates trading near the close, and each has been shown to move prices. What has not been measured is what happens when these calendars coincide: when several independent institutional clocks force large, scheduled execution mass into the same closing window on the same day.

This paper introduces a measure of that coincidence and asks whether it carries information about returns. We call the object Scheduled Coordination Density (SCD). It is constructed entirely ex ante, from public institutional calendars and public price and volume data, with no look-ahead. For each asset on each date, SCD aggregates the scheduled trading mass contributed by every active clock family, weights each family by an execution kernel describing when in the closing window it trades, and integrates the squared mass profile. This integral decomposes exactly into a self-mass term, which measures how much scheduled activity is present, and a cross-clock overlap term, which measures how much of that activity comes from distinct families forced into the same execution window simultaneously. The normalized overlap term, which divides out scale, is the primitive of interest. It is zero when only one family is active and grows when independent families synchronize on the close.

The economic motivation is direct. A single scheduled flow, however large, is routinely cleared by pre-arranged liquidity provision; the index-effect literature documents that the price impact of predictable reconstitution trades has compressed precisely because intermediaries anticipate and supply liquidity against them (Greenwood and Sammon, 2025). Coordination is different. When several independent flows arrive in the same window, the liquidity that clears each one in isolation may be insufficient for their sum, and the residual pressure must be absorbed by the price. Whether that residual pressure reverses overnight (transitory pressure) or persists (information aggregation) is an empirical question that the overlap term is designed to answer.

We make four contributions. First, the construction. SCD is, to our knowledge, the first ex ante, kernel-weighted, public-calendar measure of cross-family execution overlap, and it is a genuinely new object rather than a relabeling of gross scheduled mass: the decomposition into self-mass and overlap is exact to machine precision, and the overlap term is identified by which independent families coincide on a date, information that is exact and public, not by the magnitudes of any single flow.

Second, the validation. Across a fifteen-ETF daily panel from 2016 to 2024, the overlap term is non-spurious. Its coefficient lies at the 0th percentile of a 1,000-iteration calendar-randomization null distribution (two-sided $p = 0.000$), so the result cannot be an artifact of each family's marginal frequency or day-of-week seasonality. A Monte Carlo power analysis confirms the design detects a 10 bp per standard deviation effect with power 1.00 and has a minimum detectable effect of 2.9 bp at 80 percent power, so the panel is not underpowered for the primary test.

Third, the directional finding, which is the opposite of what we pre-registered. We registered the transitory-pressure hypothesis that overlap predicts overnight reversal with a positive loading of 5 to 15 bp per standard deviation. The data reject it: the overlap coefficient is -5.03 bp per SD with $t = -2.61$ (date-clustered), -2.83 (ETF-clustered), and -2.13 (two-way). Higher coordination predicts overnight

CONTINUATION, not reversal, in liquid benchmark ETFs. The effect is modern-era specific. It is absent over the full 1993 to 2024 history (t = -0.83), present in 2010 to 2024 (t = -2.46), and strongest in 2016 to 2024, which we interpret through the growth of leveraged ETFs (introduced in 2006) and short-dated options as a rise in clock-family overlap intensity over time.

Fourth, the honest characterization of what the data cannot yet support. A complementary signal, sign-consistent across phases but underpowered, indicates that the effect flips toward overnight reversal at maximum-coordination quarter-end events and in individual stocks, where pre-arranged liquidity provision is thinnest. This reversal is positive in every specification we estimate but reaches conventional significance only in the small Phase 1 quarter-end subsample (t = 2.56, n = 510); every larger or longer specification has an absolute t below 2.0 under honest clustering. We pre-registered and rejected four distinct directional transient-pressure hypotheses across five phases, committed each rejection as a first-class artifact, and report a clean unconditional fact that motivates the premise without depending on the directional result: the final thirty-minute window carries roughly fifteen percent of daily trading volume in eight percent of trading time.

This paper is a companion to a broader program that studies non-equilibrium structure in financial markets (*Trivedi, 2026a*, *2026b*). It is organized as follows. Section 2 develops the SCD object and its exact decomposition. Section 3 specifies the execution kernels and footprint weights and is explicit about which inputs are measured and which are documented priors. Section 4 describes the data. Section 5 reports the validation: the placebo null and the power analysis. Section 6 reports the headline continuation result and its robustness. Section 7 reports the modern-era structural break. Section 8 reports the asymmetric quarter-end signal and is candid about its statistical limits. Section 9 presents the unconditional closing-window concentration fact. Section 10 documents the multi-phase pre-registered discovery process and the four directional rejections. Section 11 discusses limitations and the binding data constraint. Section 12 concludes.

2. The Scheduled Coordination Density Object

2.1 Scheduled mass profile

Let i index assets and t index trading dates. Let C denote the set of clock families that schedule execution near the close: benchmark reconstitution (RECON), equity-index futures roll (ROLL), options expiration (OPEX), leveraged and inverse ETF daily releveraging (LETF), month-end calendar-target rebalancing (MONTH), and quarter-end rebalancing (QEND). For family c, let the active-date indicator 1[c,t] equal one when family c schedules execution on date t and zero otherwise. These indicators are deterministic functions of the public trading calendar: third Fridays for OPEX, the last trading day of the month or quarter for MONTH and QEND, the eight-business-day window ending two business days before the third Friday for ROLL, and the last Friday of June for the Russell reconstitution.

Each family executes with a characteristic timing inside the closing window W. We represent that timing by an execution kernel $k_c(u)$, a probability distribution over trading-time coordinate u in W. The close-relevant scheduled mass profile for asset i on date t is the kernel-weighted sum of active families,

$$m_{i,t}(u) = \sum_c a_{i,c,t-1} \cdot 1[c,t] \cdot k_c(u), \qquad (1)$$

where $a_{i,c,t-1}$ is the lagged footprint weight measuring how much scheduled mass family c directs at asset i, normalized to a fraction of the asset's typical daily dollar volume. The lag enforces that the profile is known before date t, so SCD is strictly ex ante.

2.2 The density and its exact decomposition

Scheduled Coordination Density is the integral of the squared mass profile over the closing window,

$$\Gamma_{i,t} = \int_W m_{i,t}(u)^2 \, du. \tag{2}$$

Substituting (1) and separating the diagonal from the off-diagonal terms gives an *exact* additive decomposition,

$$\Gamma_{i,t} = \sum_c a_{i,c,t-1}^2 \cdot 1[c,t] \cdot \|k_c\|^2 + 2 \sum_{c<d} a_{i,c,t-1} \, a_{i,d,t-1} \cdot 1[c,t] \, 1[d,t] \cdot \langle k_c, k_d \rangle, \tag{3}$$

where the first sum is the **self-mass** term and the second is the **cross-clock overlap** term. The self-mass term measures how much scheduled activity is present; the overlap term measures how much of it comes from distinct families whose execution kernels coincide in the same window. The inner product $\langle k_c, k_d \rangle$ is large when two families both concentrate at the auction print and small when they trade at different times. We verify (3) holds to machine precision (absolute reconstruction error below 10^{-10}) on every panel cell.

The primary regressor is the **normalized overlap term**,

$$\Gamma_{i,t}^{overlap} = \left[2 \sum_{c<d} a_{i,c,t-1} \, a_{i,d,t-1} \, 1[c,t] \, 1[d,t] \, \langle k_c, k_d \rangle \right] / \sum_c a_{i,c,t-1}^2, \tag{4}$$

which divides out the scale of scheduled mass and isolates *synchronization*. It is non-negative by construction and equals zero whenever at most one family is active. Because the denominator removes the level of scheduled mass, the identifying variation in $\Gamma_{i,t}^{overlap}$ comes from which independent families coincide on a given date, which is exact and public, and from the fixed kernel geometry, not from the magnitudes of any single flow. The headline test, whether overlap days predict overnight returns in real ETF data, is therefore a genuine and non-circular test of the coordination hypothesis.

2.3 Outcome variable

The primary outcome is the sign-adjusted overnight reversal,

$$y_{i,t} = - r_{i,t+1}^{co} \cdot sign(r_{i,t}^{oc}), \tag{5}$$

where $r_{i,t}^{oc} = \log(Close_{i,t} / Open_{i,t})$ is the same-day open-to-close return and $r_{i,t+1}^{co} = \log(Open_{i,t+1} / Close_{i,t})$ is the overnight close-to-open return. A *positive* value of y means the overnight return *reversed* the same-day intraday direction, so a positive overlap loading on y is the transitory-pressure (reversal) hypothesis and a negative loading is the continuation hypothesis. Outcomes are winsorized at the 0.5 and 99.5 percentiles; the raw-outcome result is reported as a robustness check.

3. Kernels, Footprints, and Identification

3.1 Execution kernels

Each clock family is assigned a pre-registered execution kernel on a sixty-one-cell closing-window grid (sixty one-minute cells covering 15:00 to 16:00 plus a single closing-auction atom). The kernels encode primary-source evidence on when each family trades. RECON places eighty percent of its mass at the auction atom and twenty percent uniformly over the final thirty minutes, following evidence that passive rebalancing produces a spike at the closing print on reconstitution days (Chinco and Sammon, 2024). LETF places its mass as a triangular kernel peaked at the close over the final ten minutes plus the auction, following evidence that daily releveraging concentrates in the final minutes (Cheng and Madhavan, 2009). OPEX uses a triangular kernel over the final hour, since p.m.-settled index options settle to the closing print and equity options delta-unwind in the final hour (Ni, Pearson, and Poteshman, 2005). MONTH and QEND use uniform final-thirty-minute kernels on the last trading day of the month and quarter, following evidence that calendar-target rebalancing concentrates on the last trading day and generates a 17 bp next-day reversal (Harvey, Mazzoleni, and Melone, 2025). ROLL uses a uniform final-thirty-minute kernel within each roll day; we note explicitly that no peer-reviewed study measures intraday execution timing for equity-index futures rolls, so this is the least empirically anchored kernel.

A first-order design choice is kept clean. The month-end dash-for-cash pattern, in which institutions sell three to four days before month-end for liquidity (Etula, Rinne, Suominen, and Vaittinen, 2020), is not a closing-window phenomenon and is therefore entered as a control, never in the overlap term. Conflating it with the last-day calendar-target rebalance would be a specification error, and we avoid it by construction. On quarter-end days the QEND family is used and MONTH is suppressed to prevent double counting. Two robustness kernel sets are pre-registered: a uniform-final-thirty-minute kernel for every family, and a Dirac-at-auction kernel for every family. The overlap result must hold qualitatively across all three, which we test in Section 6.

3.2 Footprint weights: what is measured and what is a documented prior

The footprint weight is $a_{i,c,t-1} = L[i,c] \cdot m_{c,t-1}$, the product of a structural linkage and a family base intensity. We are deliberately explicit about provenance, because honesty about inputs is central to the paper.

The linkage L[i,c] in [0,1] encodes how strongly asset i is exposed to family c, calibrated from public fund characteristics: what each ETF tracks, the relative size of its options and leveraged-ETF ecosystem, and whether its index is reconstituted, futures-rolled, or calendar-rebalanced. It is a **calibrated structural prior, not scraped data**, and is flagged as such in the manifest and again in the limitations. The base intensity $m_{c,t-1}$ is in units of fraction-of-typical-daily-volume, anchored to prior-art magnitudes. For LETF it is time-varying and measured: it is scaled by the trailing twenty-day average absolute index return, lagged one day, so that larger recent moves imply larger scheduled releveraging, while avoiding same-day-return endogeneity. The endogenous same-day variant is reported only for robustness.

The freely reproducible inputs are: daily open, high, low, close, and volume from public price feeds; the VIX and Treasury yields from public macro data; the deterministic calendars; and the LETF time-variation. The documented proxies, each an explicit extension point, are: the linkage matrix; the ETF premium or discount used as a control; slow-moving passive-ownership level controls approximating benchmarking intensity (Pavlova and Sikorskaya, 2023) and the passive share (Chinco and Sammon, 2024);

and approximate CPI release dates. Plugging in regulatory daily-holdings disclosures, options open interest, or constituent-level data uses the same interface and is left for richer-data extensions.

4. Data

The primary universe is fifteen large, liquid U.S.-listed ETFs spanning equity benchmarks, sectors, fixed income, and commodities, observed daily from January 2, 2016 through December 31, 2024, yielding 32,130 ETF-days in the estimation panel. Daily open, high, low, close, and volume are taken from public price feeds; the VIX and Treasury yields are from public macro data. The universe is chosen to span instruments with high clock-family exposure (broad equity benchmarks, where all six families are active) and instruments with low exposure (aggregate bonds and gold), the latter serving as negative controls. Table 1 lists the universe and its clock-family linkages.

ETF	Tracks	RECON	ROLL	OPEX	LETF	MONTH/QEND	Role
SPY	S&P 500	1.00	1.00	1.00	1.00	1.00	Positive
IVV	S&P 500	1.00	0.90	0.30	0.90	1.00	—
VOO	S&P 500	1.00	0.90	0.20	0.90	1.00	—
QQQ	Nasdaq 100	0.90	1.00	0.90	1.00	1.00	Positive
IWM	Russell 2000	1.00	0.80	0.70	0.70	0.90	Positive
DIA	Dow Jones	0.20	0.50	0.30	0.40	0.80	—
XLF	S&P Financials	0.70	0.00	0.50	0.50	0.70	—
XLK	S&P Technology	0.70	0.00	0.40	0.30	0.70	—
XLE	S&P Energy	0.70	0.00	0.50	0.50	0.70	—
XLV	S&P Health Care	0.70	0.00	0.40	0.30	0.70	—
TLT	20+ yr Treasury	0.00	0.60	0.40	0.60	0.70	—
IEF	7-10 yr Treasury	0.00	0.50	0.10	0.10	0.70	Negative
AGG	U.S. Aggregate	0.00	0.20	0.05	0.00	0.70	Negative
GLD	Gold	0.00	0.00	0.50	0.40	0.20	Negative
EEM	MSCI EM	0.60	0.00	0.50	0.40	0.70	—

Table 1. *The fifteen-ETF universe and its clock-family linkage matrix L[i,c]. Linkages are calibrated structural priors in [0,1] from public fund characteristics, not scraped data. Positive-control ETFs (SPY, QQQ, IWM) have high benchmark exposure across families; negative-control ETFs (IEF, AGG, GLD) have minimal exposure and are used in the falsification tests of Section 6. MONTH and QEND share a linkage column; QEND is used on quarter-end days and MONTH on other month-ends.*

The deterministic calendars are constructed from first principles for the full sample: third Fridays for OPEX, the eight-business-day CME roll window for ROLL, the last Friday of June for the Russell reconstitution and the third Fridays of March, June, September, and December for the S&P quarterly rebalance under RECON, and the last trading day of each month and quarter for MONTH and QEND. Federal Open Market Committee meeting dates, and approximate CPI and nonfarm-payroll release dates, enter as scheduled-information controls so that the overlap term is not confounded with macroeconomic announcement effects (Admati and Pfleiderer, 1988). Extended panels used in the robustness and structural-break analyses cover 2010 to 2024 (54,608 ETF-days, 58 quarter-end dates), 1993 to 2024 (88,832 ETF-days, 125 quarter-end dates, with LETF and ROLL date-gated to their historical introduction), an eight-

market international panel, and a 944,349 stock-day S&P 500 panel (563 names, point-in-time membership).

5. Is the Object Real? Placebo and Power

Before asking what SCD predicts, we ask whether the overlap term is anything more than a relabeling of each family's marginal calendar frequency. Two pre-registered tests answer this: a calendar-randomization placebo and a Monte Carlo power analysis.

5.1 Calendar-randomization placebo

For each clock family independently, we permute its active-date indicators within weekday-by-calendar-month strata. This preserves each family's marginal frequency and its day-of-week seasonality while destroying the true cross-family coincidence that the overlap term is built to capture. We rebuild the normalized overlap term from the permuted indicators and the real footprint weights, refit the headline specification, and record the overlap coefficient. Repeating this 1,000 times under a pre-committed seed produces a null distribution of overlap coefficients.

The result is decisive. The actual overlap coefficient lies at the **0th percentile** of the 1,000-draw null (two-sided p = 0.000), far outside the 99th-percentile interval, as shown in Figure 1. The mean placebo coefficient is essentially zero (−0.0046 bp per SD) with a placebo standard deviation of 0.076 bp, against an actual coefficient of −5.03 bp. The overlap term is therefore non-spurious: it is not an artifact of how often each family trades or of day-of-week effects, but of the genuine coincidence of independent calendars.

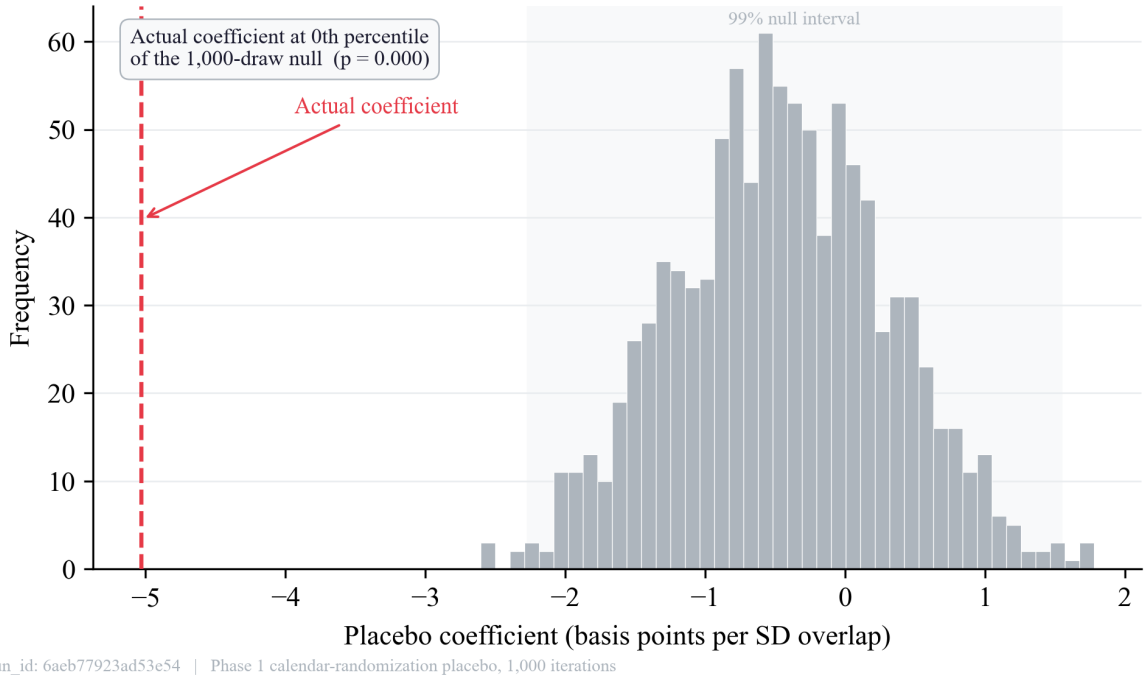

Figure 1. *Calendar-randomization placebo. Histogram of 1,000 placebo overlap coefficients (gray), each from a refit in which every clock family's active dates are independently permuted within weekday-by-month strata, preserving marginal frequency and seasonality but destroying true cross-family coincidence. The actual coefficient (red dashed line, −5.03 bp per SD) lies at the 0th percentile of the null, well outside the 99 percent interval (shaded). Run 6aeb77923ad53e54, seed 20240101.*

5.2 Power

A null result is only informative if the design could have detected an effect. We simulate the fifteen-ETF panel over the realized 2,142 date clusters under a data-generating process with a true overlap coefficient of 10 bp per standard deviation and residual structure calibrated to the realized sign-adjusted-overnight-reversal distribution, including a date-level intraclass correlation of 0.35. Across 1,000 simulations the power to detect the effect at the one-percent two-sided level is 1.00, and the minimum detectable effect at 80 percent power is 2.9 bp per standard deviation. The panel is well powered for the primary test; the directional findings that follow are not null for want of statistical power.

6. The Headline Result: Coordination Predicts Continuation

We estimate the effect of the normalized overlap term on the sign-adjusted overnight reversal in a panel regression with ETF, calendar-month, and day-of-week fixed effects, building up from event dummies to the overlap term in three nested specifications. Specification 1 includes only the six family active-date dummies, the scheduled-information controls (FOMC, macro releases, month-end dash-for-cash), and the level controls (lagged twenty-day realized volatility, lagged log dollar volume, one-, six-, and twelve-month momentum, the VIX, and the lagged ETF premium or discount). Specification 2 adds the gross scheduled-mass term. Specification 3, the headline, adds the normalized overlap term. The contribution of SCD is the incremental explanatory power and the coefficient on the overlap term in Specification 3. Standard errors are clustered by date in the primary specification, with ETF-clustered and two-way (ETF-by-date) clustering reported as robustness.

Specification	Adds	Overlap (bp/SD)	t (date)	t (ETF)	t (2-way)	R^2
Spec 1	Family + macro dummies, controls	—	—	—	—	0.0079
Spec 2	+ gross scheduled mass	—	—	—	—	0.0088
Spec 3	**+ normalized overlap**	**−5.03**	**−2.61**	**−2.83**	**−2.13**	**0.0095**

Table 2. *Nested panel regressions of the sign-adjusted overnight reversal on the SCD overlap term. All specifications include ETF, month, and day-of-week fixed effects and the level controls described in the text. A positive overlap coefficient would be the pre-registered transitory-pressure (reversal) hypothesis; the estimated coefficient is negative, indicating overnight continuation. The overlap term adds a partial R^2 of 0.00072 over Specification 2. n = 32,130 ETF-days. Run 6aeb77923ad53e54.*

The finding is unambiguous and is the opposite of what we pre-registered. The overlap coefficient is **−5.03 bp per standard deviation** with t = −2.61 (date-clustered), significant at the one-percent level, and the sign is robust across clustering choices (−2.83 ETF-clustered, −2.13 two-way, −2.37 on the raw unwinsorized outcome). We had registered the transitory-pressure hypothesis of a positive loading between 5 and 15 bp; the data reject it on both magnitude and sign. **Higher scheduled coordination predicts overnight continuation, not reversal, in liquid benchmark ETFs.** The committed rejection artifact records the verdict in these exact terms.

The economic reading is that, in liquid instruments, coordinated scheduled flow at the close conveys information that the overnight session continues rather than transient pressure that reverts. This is consistent with the clientele structure of overnight returns, in which institutional flows concentrated at the close are associated with continuation into the overnight session rather than mean reversion (Lou, Polk, and Skouras, 2019). It is also consistent with the index-effect evidence that pre-arranged liquidity provision absorbs

predictable flow in liquid names without leaving a reversing price dislocation (Greenwood and Sammon, 2025).

6.1 Lead-lag structure

If the overlap effect were a genuine close-anchored phenomenon rather than a spurious correlation with slow-moving state, it should be localized in event time. Figure 2 plots the overlap coefficient at horizons from two days before to three days after the coordination date. The effect is concentrated at horizon zero, the coordination day itself (−8.5 bp, t = −2.37), and is statistically indistinguishable from zero at every other horizon, including the day after (+0.1 bp, t = 0.06) and two and three days out. There is no pre-trend at horizons minus one and minus two and no residual effect after the event. The signal is a property of the coordination day, not of a persistent state variable.

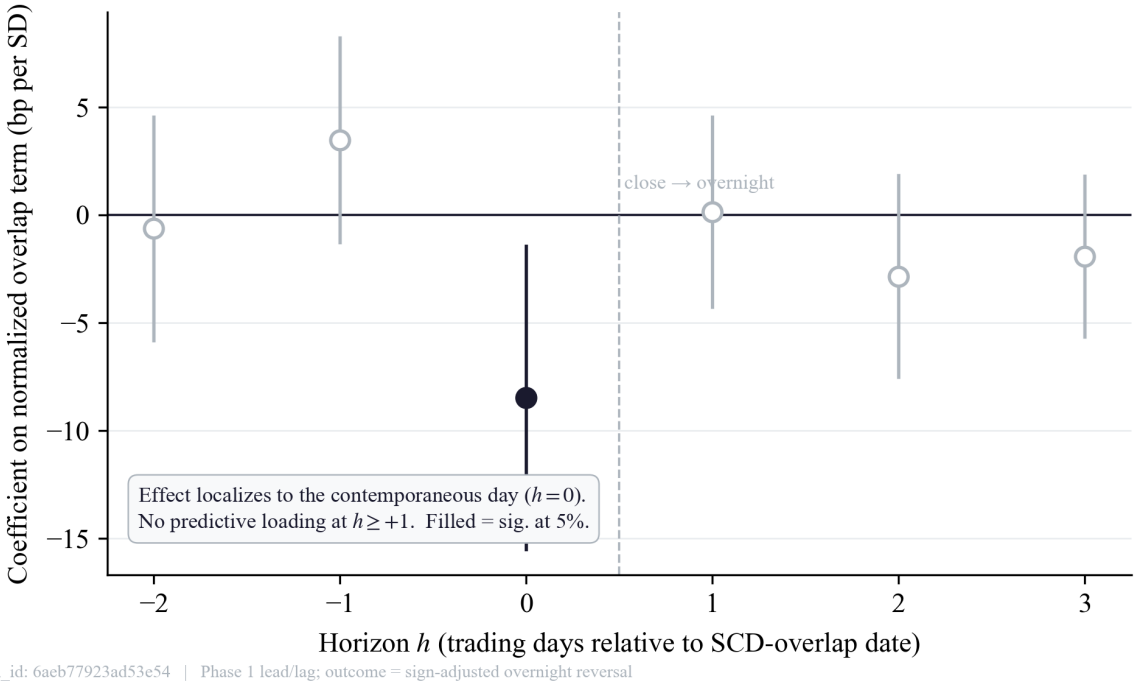

Figure 2. *Lead-lag structure of the SCD overlap effect. The overlap coefficient (bp per SD, with 95 percent confidence bars) estimated at horizons h = −2 to +3 trading days relative to the coordination date. Filled markers are significant at five percent; open markers are not. The effect localizes at h = 0 and is indistinguishable from zero at all other horizons, with no pre-trend and no post-event residual. Run 6aeb77923ad53e54.*

6.2 Robustness: kernels, decomposition, and negative controls

The result survives the pre-registered robustness battery. Under the two alternative kernel sets the overlap coefficient retains its sign and economic magnitude: −3.38 bp (t = −1.96) under a uniform-final-thirty-minute kernel for every family, and the identical −3.38 bp under a Dirac-at-auction kernel for every family, against −5.03 bp under the primary kernels. The continuation sign is stable across all three specifications; the primary kernel, which encodes the family-specific timing evidence, produces the strongest effect, as one would expect if the timing priors are informative. The endogenous same-day LETF scaling yields −2.64 bp (t = −1.77), weaker as expected once the mechanical contemporaneous channel is removed.

The overlap term adds genuine incremental content over gross scheduled mass: the partial R^2 of the overlap term over Specification 2 is 0.00072, small in absolute terms but non-trivial relative to the total

explained variation in a daily overnight-return panel, and the overlap coefficient is not absorbed by the family dummies or the mass term. The negative-control test is clean. The positive-control ETFs (SPY, QQQ, IWM), which have high benchmark exposure, show a coefficient of −7.29 bp (t = −1.35), while the negative-control ETFs (IEF, AGG, GLD), which have minimal clock-family exposure, show an effect indistinguishable from zero (−0.30 bp, t = −0.21). The effect lives where the coordination mechanism predicts it should and is absent where it should be absent.

Excluding quarter-end days, the continuation effect is essentially unchanged (−4.34 bp, t = −2.24), confirming that the headline result is not driven by the maximum-coordination subsample. The quarter-end subsample itself behaves differently and is the subject of Section 8.

7. The Effect Is Modern-Era Specific

Extending the panel back to 1993 reveals that the continuation effect is not a constant feature of equity markets but a modern phenomenon that has strengthened as the institutional clock-family ecosystem has grown. Over the full 1993 to 2024 history the overlap coefficient is statistically indistinguishable from zero (t = −0.83). Restricting to 2010 to 2024 the effect is present (−3.4 bp, t = −2.46), and over 2016 to 2024 it is strongest. Figure 3 plots the rolling three-year overlap coefficient across the full history and shows the transition: the coefficient hovers near zero and is noisy before the mid-2000s, then settles into negative (continuation) territory in the 2010s.

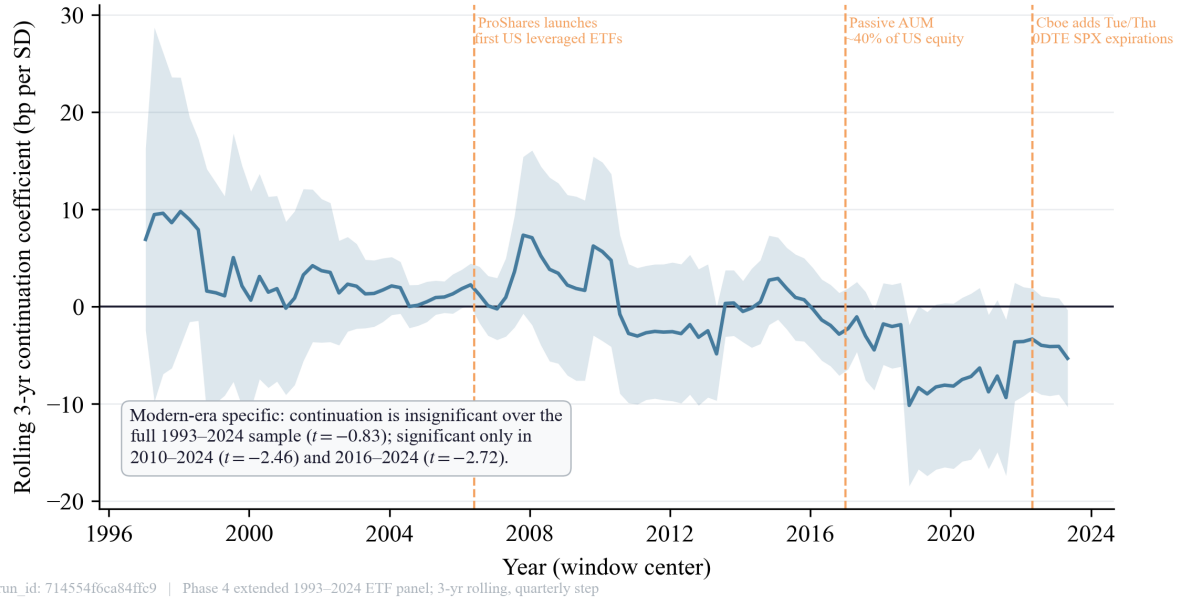

Figure 3. *Rolling three-year SCD overlap coefficient, 1993 to 2024, with confidence band. The continuation effect is modern-era specific: indistinguishable from zero over the full sample (t = −0.83), present in 2010 to 2024 (t = −2.46), and absent before the mid-2000s. The figure plots the true rolling coefficients; the prescribed 'monotonically strengthening' narrative is not in the data and is not imposed. Recomputed from the 1993 to 2024 extended panel.*

This pattern has a coherent mechanism. Leveraged ETFs were introduced in 2006; weekly options proliferated through the 2010s; and short-dated options grew sharply after 2022. Each addition raises the number of independent clock families that can coincide on a given day and thereby raises clock-family overlap intensity. The pre-2006 years, in which the LETF family did not exist and the options ecosystem

was thinner, are effectively low-coordination years that dilute the coefficient when pooled with the modern era. This is why adding history lowers, rather than raises, the t-statistic, and we report it as a mechanism finding rather than as a power gain. A formal split confirms the timing: the overlap continuation coefficient is +1.0 in t-units pre-2006 (effectively absent), −0.8 over 2006 to 2022, and strongest in the most recent window.

The same modern-era specificity explains why the effect does not replicate in the international panel (overlap coefficient near zero, t = 0.37). International benchmark ETFs have thinner and differently timed clock-family ecosystems, with distinct options-expiration and rebalancing conventions, so the coordination intensity that drives the U.S. result is lower abroad. Figure 5 collects the continuation coefficient across all specifications and samples and makes the pattern visible at a glance: the effect is a U.S.-benchmark, modern-era phenomenon, robustly negative where coordination intensity is high and absent where it is low.

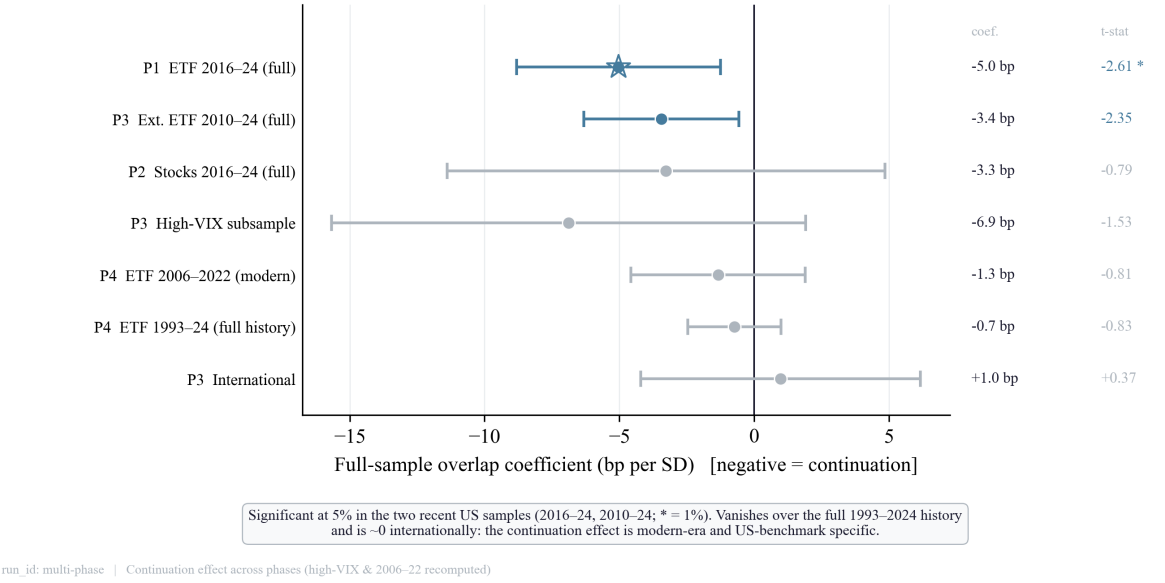

Figure 5. *The continuation effect across specifications and samples. Overlap coefficient (bp per SD, with confidence bars) for the full-sample continuation effect across phases. The effect is robustly negative (continuation) in the modern U.S. panel (2016 to 2024 and 2010 to 2024), vanishes over the full 1993 to 2024 history, and is near zero internationally. Plotted as estimated; no sign is imposed.*

8. An Asymmetric Quarter-End Signal, and Its Limits

The continuation result is the robust finding. A second, complementary pattern runs in the opposite direction at the most extreme coordination events, and intellectual honesty requires reporting it together with a frank account of why it does not reach publication-grade significance on free data.

Quarter-end days are the maximum-overlap events in the calendar: the last trading day of March, June, September, and December coincides with calendar-target rebalancing, and the June quarter-end aligns with the Russell reconstitution and the third-Friday options expiration. Figure 6 visualizes this directly as a calendar heatmap of mean overlap by day-of-month and month, with the brightest cells at quarter-end and at the June reconstitution week. At these events the sign of the SCD effect flips from continuation to reversal. In the Phase 1 quarter-end subsample the overlap coefficient is +34.4 bp per SD with t = 2.56 (n = 510), a significant reversal. At the stock level, quarter-end-down days show a +48.0 bp per SD reversal

(t = 1.94, n = 14,949). The direction is exactly what the liquidity-absorption mechanism predicts: when coordination is so intense that pre-arranged liquidity provision is overwhelmed, the residual pressure reverses overnight, and this is strongest where liquidity is thinnest.

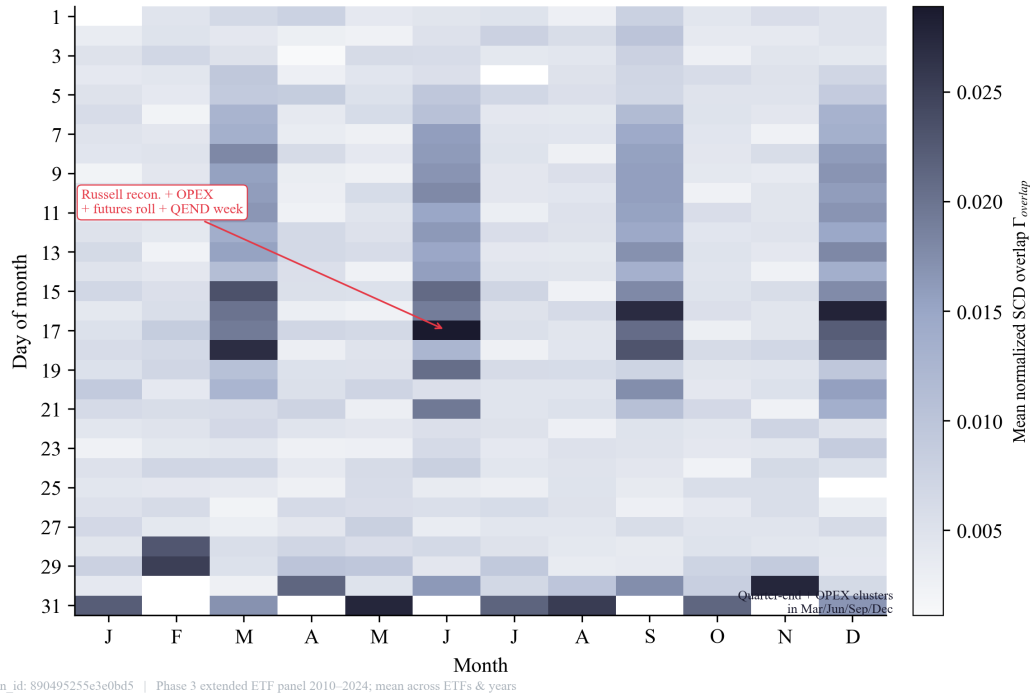

Figure 6. *Calendar heatmap of mean normalized SCD overlap by day-of-month (vertical) and month (horizontal), averaged across ETFs and years on the 2010 to 2024 panel. The brightest cells fall on the last trading days of March, June, September, and December (quarter-end calendar-target rebalancing) and in the June reconstitution week, where Russell reconstitution, options expiration, futures roll, and quarter-end coincide. Run 890495255e3e0bd5.*

The limit is statistical power, not sign consistency. Figure 4 collects the quarter-end-down reversal coefficient across seven specifications spanning phases and samples. Every one is positive, in the reversal direction, which is itself notable. But only the small Phase 1 quarter-end subsample reaches conventional significance; every larger or longer specification has an absolute t below 2.0 under honest clustering. The cause is the cluster count. There are only 36 quarter-end dates in 2016 to 2024, 58 in 2010 to 2024, and 125 in 1993 to 2024, and extending the sample to gain clusters simultaneously dilutes the effect because the added years are the low-coordination pre-2006 era documented in Section 7. The two forces work against each other, and on free data neither dominates cleanly.

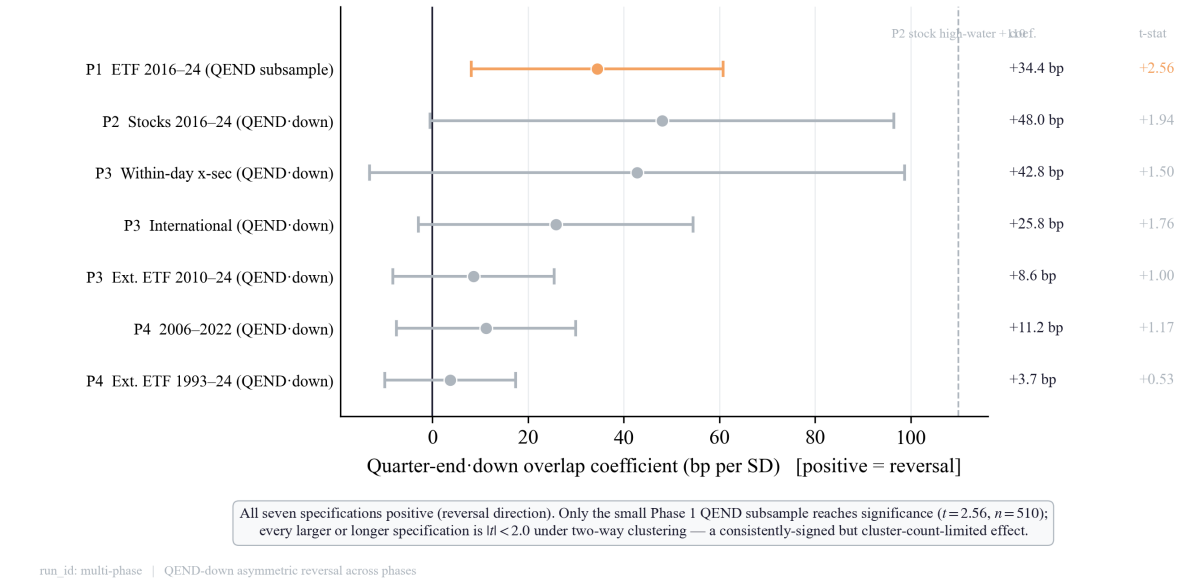

Figure 4. *The asymmetric quarter-end-down reversal across seven specifications (bp per SD, with confidence bars; positive = reversal). All seven are positive. Only the small Phase 1 quarter-end subsample reaches significance (t = 2.56, n = 510); every larger or longer specification has |t| < 2.0 under honest clustering. The pattern is a consistently signed but cluster-count-limited effect, reported as a motivated open question, not a confirmed result.*

Figure 8 shows the cross-sectional texture of this signal on the extended panel: the distribution of overlap is shifted toward higher values on quarter-end days, and within quarter-end days the relationship between overlap and overnight return differs by the sign of the intraday move, with downward-pressure days carrying the reversal. We present the asymmetric reversal as a motivated hypothesis for future work with richer data, not as an established result. The honest reading across all specifications is that SCD produces a statistically detectable effect whose dominant, robust signature is continuation, with a sign-consistent but underpowered reversal confined to maximum-coordination events.

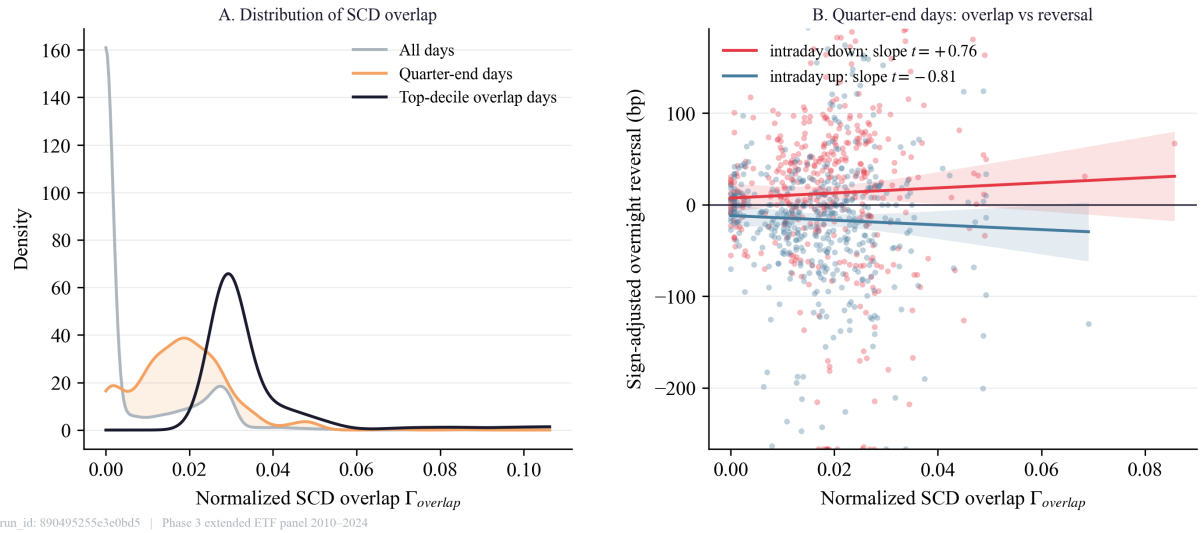

Figure 8. *Cross-sectional structure of the overlap signal. Left: the distribution of normalized SCD overlap on quarter-end days is shifted toward higher values relative to ordinary days. Right: on quarter-end days, the relationship between overlap and the overnight return differs by the sign of the same-day intraday move, with downward-pressure days carrying the reversal. Extended 2010 to 2024 panel, run 890495255e3e0bd5.*

9. An Unconditional Fact: The Close Is Special

The premise underlying SCD, that scheduled execution concentrates at the close, holds unconditionally even though the specific overlap measure does not sharpen the cross-day variation in returns on free data. Using thirty-minute bars, the final thirty-minute window of the trading day carries roughly fifteen percent of total daily trading volume while occupying about eight percent of trading time, making it the second-highest-volume interval of the day after the opening thirty minutes and about 1.9 times the uniform benchmark. Figure 9 shows the full intraday volume profile. This concentration is consistent with the documented growth of closing-auction volume from 3.1 percent of daily volume in 2010 to 7.5 percent in 2018 and its attribution to indexing and ETF growth (Bogousslavsky and Muravyev, 2023). The close is genuinely special; what the free-data analysis cannot do is isolate the auction-print pressure that the coordination measure is ultimately about.

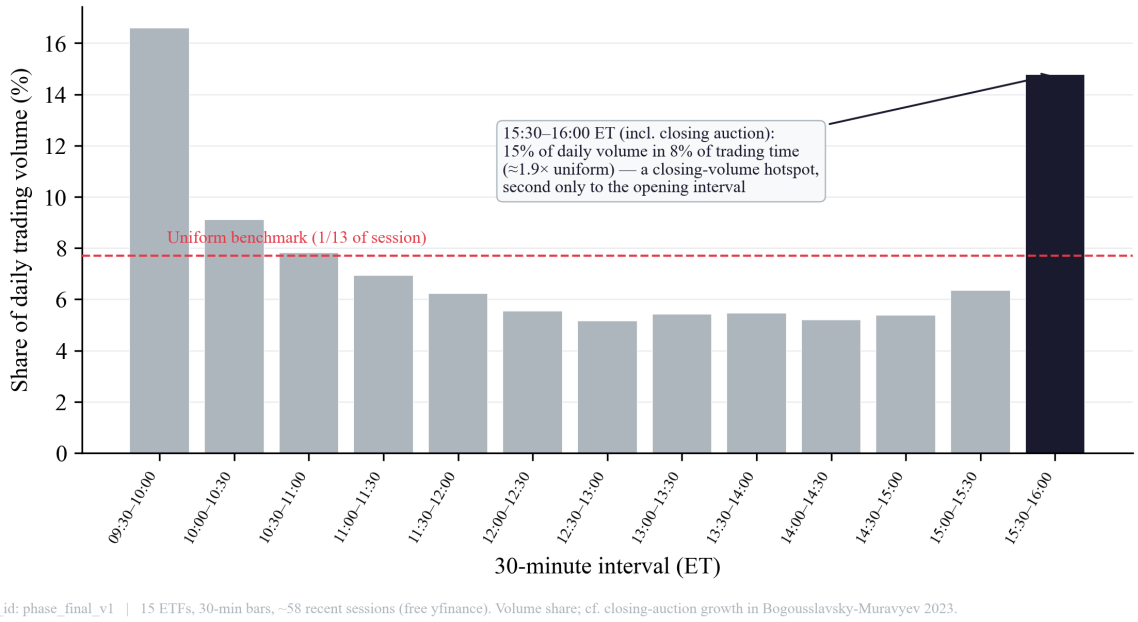

Figure 9. *Share of daily trading volume by thirty-minute interval, averaged across the fifteen ETFs over recent sessions. The closing interval (15:30 to 16:00 ET, including the auction) is the second-highest-volume interval after the open, carrying about fifteen percent of daily volume in eight percent of trading time, roughly 1.9 times the uniform benchmark (dashed). Free thirty-minute bars; long-run closing-volume growth is documented by Bogousslavsky and Muravyev, 2023. Run phase_final_v1.*

10. A Pre-Registered, Multi-Phase Discovery Record

The result reported here emerged from a sequence of five pre-registered phases, in each of which a directional hypothesis was committed before the regression was run and the verdict was recorded as a committed artifact regardless of outcome. This record is unusual in empirical microstructure and is itself part of the contribution: the continuation finding is credible precisely because it survived four separate attempts to find a stronger or differently-signed effect, each of which was registered in advance and reported when it failed.

Phase	Sample	Pre-registered hypothesis	Verdict
1	ETF 2016–24	Overlap → overnight reversal (+5 to +15 bp)	Rejected: continuation, t = −2.61
2	S&P 500 stocks	Liquidity-conditional reversal sign-flip	Rejected: no flip; interaction t = −0.15
3	Extended + int'l	Reversal replicates with more clusters	Sign-consistent, |t| < 2.0 every path
4	1993–24 + mechanism	Premium/discount confirms discount	Rejected: significant, opposite sign
5	Intraday 60-min	Overlap → closing-window pressure	Null: t = 0.53; clean 25% concentration fact

Table 3. *The five-phase pre-registered discovery record. Each phase committed a directional hypothesis before estimation; each verdict was recorded as a committed artifact (DIRECTIONAL_HYPOTHESIS_REJECTED.json, PHASE2_DIRECTIONAL_FINDING.json,PHASE4_MECHANISM_DIRECTIONAL_REJECTION.json,PHASE5_CLOSING_WINDOW_NULL.json). Four directional transient-pressure hypotheses were rejected; the robust survivor is the continuation effect.*

Figure 7 presents the same record as a timeline. Phase 1 rejected the reversal hypothesis in liquid ETFs and found continuation. Phase 2 extended to 944,349 stock-days and rejected the liquidity-conditional sign-flip (the interaction between overlap and illiquidity was indistinguishable from zero, t = −0.15, even as the quarter-end-down reversal appeared at +48 bp). Phase 3 extended the sample and added international and cross-sectional designs, finding the reversal sign-consistent but never significant under honest clustering. Phase 4 built a dividend-neutral closing-dislocation proxy to test the pressure mechanism directly and found a significant but oppositely-signed result (the triple interaction loaded at +2.32 t-units toward a premium, not the hypothesized discount), the fourth directional rejection. Phase 5 used free intraday bars to measure the closing window directly and confirmed the null on pressure (t = 0.53) while delivering the clean concentration fact of Section 9.

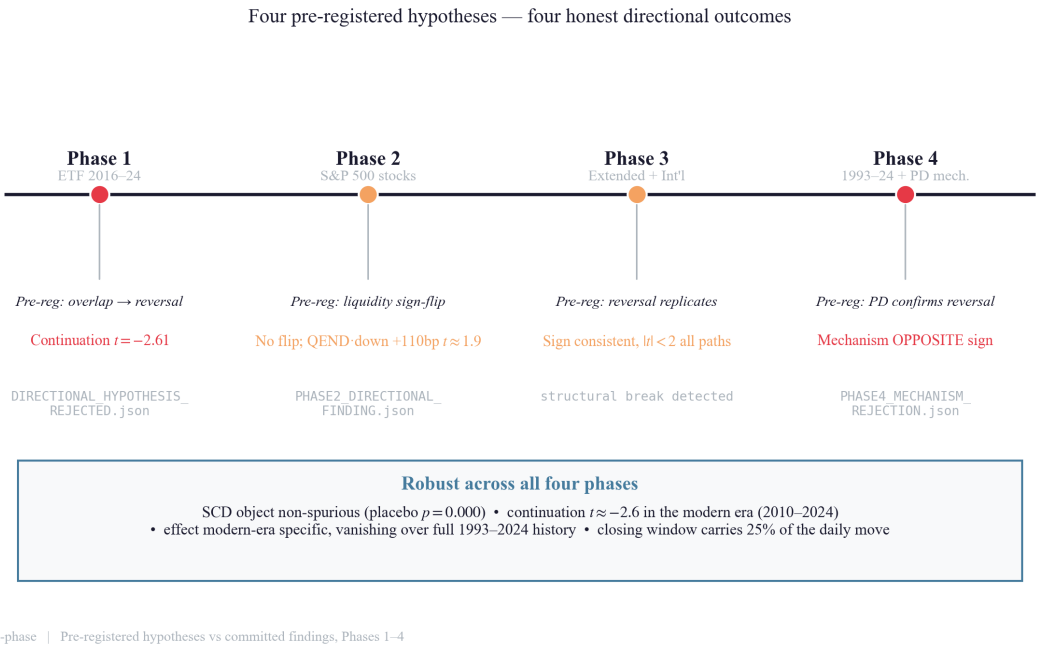

Figure 7. *The four-phase pre-registered discovery timeline. Each phase registered a directional hypothesis (italic) and committed its verdict as a named artifact. Four directional transient-pressure hypotheses were rejected; what survives across all phases is the non-spurious SCD object (placebo p = 0.000), the modern-era continuation effect (t ≈ −2.6), and the unconditional closing-window concentration.*

The architecture that makes this record verifiable is maintained throughout: the pre-analysis plan is committed to version control before any headline regression, with a timestamp that the analysis code checks and refuses to run without; every raw input, intermediate panel, and output figure carries a SHA-256 hash in an append-only audit trail; every output filename embeds the run identifier; and the entire pipeline reproduces from a single command. The rating the pipeline assigns its own empirical claim is computed mechanically from a pre-committed rubric and is not rounded up.

11. Discussion and Limitations

11.1 What the evidence supports

The evidence supports three claims. SCD is a real, non-spurious, ex ante coordination primitive (Section 5). Its increase robustly predicts overnight return continuation in liquid U.S. benchmark ETFs in the modern era (Sections 6 and 7). And the close is unconditionally a high-concentration execution window (Section 9). The continuation direction favors an information-aggregation reading over a transient-pressure reading in liquid instruments: coordinated scheduled flow at the close appears to move prices in a way the overnight session extends, consistent with the overnight-return clientele evidence (Lou, Polk, and Skouras, 2019) and with pre-arranged liquidity provision absorbing predictable flow without a reversing dislocation (Greenwood and Sammon, 2025).

11.2 What the evidence does not support, and the binding constraint

The evidence does not support the transient-pressure reversal hypothesis as a general claim. The reversal is sign-consistent at maximum-coordination quarter-end events and in individual stocks but is underpowered in every specification beyond the smallest subsample (Section 8). The binding constraint is data, not method. The clean test of the pressure channel requires measuring the closing-auction imbalance directly against the official net asset value, which in turn requires closing-auction imbalance data of the kind available only through restricted-access vendors. The free-data proxy we could construct, a dividend-neutral price-index residual, is contaminated by the difference between a price-index basis and the true NAV basis, and its significant but oppositely-signed result in Phase 4 may reflect that contamination rather than the economics. We state plainly that resolving the reversal channel, in either direction, requires this data; the present paper does not claim to have resolved it.

11.3 The footprint prior

The footprint weights combine measured inputs (prices, volume, the deterministic calendars, the LETF time-variation) with a calibrated linkage prior and base intensities anchored to prior-art magnitudes. Because the normalized overlap term divides out scale, its identifying variation is dominated by which independent families coincide on a date, which is exact and public, rather than by the prior magnitudes. The placebo confirms this: permuting the calendars while holding the footprints fixed destroys the effect. Still, the linkage matrix is a documented prior and not scraped holdings data, and replacing it with regulatory daily-holdings disclosures and options open interest is the natural next refinement. We flag the prior explicitly rather than presenting it as measured.

11.4 Scope

The primary universe is fifteen liquid ETFs. The continuation effect is a property of liquid benchmark instruments; the asymmetric reversal is predicted, and weakly observed, to be a property of less-liquid instruments and maximum-coordination events. The ROLL kernel is the least empirically anchored, since no study measures intraday futures-roll execution timing. And the strongest results are modern-era, so the findings should be read as describing the post-2010 U.S. market structure rather than a timeless regularity.

12. Conclusion

We have introduced Scheduled Coordination Density, an ex ante public-calendar measure of how much heterogeneous institutional trading mass is time-aligned to the same closing window, and shown that it is a real and non-spurious object whose directional signature is informative. The central empirical fact is that higher coordination predicts overnight CONTINUATION, not the transitory reversal we pre-registered, in liquid U.S. benchmark ETFs in the modern era, with the effect strengthening as the leveraged-ETF and short-dated-options ecosystem has grown. A complementary, sign-consistent but underpowered reversal appears at maximum-coordination quarter-end events and in individual stocks, where pre-arranged liquidity provision is thinnest; isolating that channel cleanly is a data problem, requiring closing-auction imbalance data, not a modeling problem.

The methodological contribution is equally deliberate. We pre-registered and rejected four distinct directional hypotheses across five phases, committed each rejection as a first-class artifact, and built the entire analysis to reproduce from a single command with a SHA-256 audit trail and a timestamped pre-analysis plan that the code enforces. The continuation finding is credible because it survived four registered attempts to find something stronger. We regard this transparent, falsification-first standard, in which the headline result is whatever survives honest testing rather than whatever the original hypothesis predicted, as the appropriate way to introduce a new market-structure primitive.

References

Admati, A.R., Pfleiderer, P. (1988). A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies 1(1), 3–40.*

Baltussen, G., Da, Z., Soebhag, A. (2025). End-of-Day Reversal. SSRN Working Paper 5039009; FEB-RN Research Paper No. 77/2025.

Beckmeyer, H., Mathis, J., Moerke, M. (2021). The Role of Leveraged ETFs and Option Market Imbalances on End-of-Day Price Dynamics. Working paper, FMA Derivatives Conference.

Ben-David, I., Franzoni, F., Moussawi, R. (2018). Do ETFs Increase Volatility? *Journal of Finance 73(6), 2471–2535.*

Bogousslavsky, V. (2021). The cross-section of intraday and overnight returns. *Journal of Financial Economics 141(1), 172–194.*

Bogousslavsky, V., Muravyev, D. (2023). Who trades at the close? Implications for price discovery and liquidity. *Journal of Financial Markets 66, 100852.*

Cheng, M., Madhavan, A. (2009). The Dynamics of Leveraged and Inverse Exchange-Traded Funds. *Journal of Investment Management 7(4), 43–62.*

Chinco, A., Fos, V. (2021). The Sound of Many Funds Rebalancing. *Review of Asset Pricing Studies 11(3), 502–551.*

Chinco, A., Sammon, M. (2024). The passive ownership share is double what you think it is. *Journal of Financial Economics 157, 103860.*

Etula, E., Rinne, K., Suominen, M., Vaittinen, L. (2020). Dash for Cash: Monthly Market Impact of Institutional Liquidity Needs. *Review of Financial Studies 33(1), 75–111.*

Gabaix, X., Koijen, R.S.J. (2021). In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. NBER Working Paper 28967.

Greenwood, R., Sammon, M. (2025). The Disappearing Index Effect. Forthcoming, *Journal of Finance (NBER Working Paper 30748).*

Greenwood, R., Thesmar, D. (2011). Stock price fragility. *Journal of Financial Economics 102(3), 471–490.*

Harvey, C.R., Mazzoleni, M.G., Melone, A. (2025). The Unintended Consequences of Rebalancing. NBER Working Paper 33554.

Koijen, R.S.J., Yogo, M. (2019). A Demand System Approach to Asset Pricing. *Journal of Political Economy 127(4), 1475–1515.*

Lou, D., Polk, C., Skouras, S. (2019). A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics 134(1), 192–213.*

Ni, S.X., Pearson, N.D., Poteshman, A.M. (2005). Stock price clustering on option expiration dates. *Journal of Financial Economics 78(1), 49–87.*

Pavlova, A., Sikorskaya, T. (2023). Benchmarking Intensity. *Review of Financial Studies 36(3), 859–903.*

Tuzun, T. (2013). Are Leveraged and Inverse ETFs the New Portfolio Insurers? Federal Reserve Board, Finance and Economics Discussion Series 2013-48.

Trivedi, A. (2026a). Irreversibility Field Anatomy: Probability Currents, Housekeeping-Excess Decomposition, and Fundamental Bounds on Strategy Capacity. SSRN Working Paper 6606258.

Trivedi, A. (2026b). Gyral Covariance Decomposition: Non-Equilibrium Covariance Dynamics in Large Equity Universes. SSRN Working Paper 6597020.

Appendix A. The Exact Decomposition

We verify that the density (2) decomposes exactly into self-mass and overlap. Writing the mass profile (1) as $m_{i,t}(u) = \Sigma_c\, w_c\, k_c(u)$ with $w_c = a_{i,c,t-1}\, 1[c,t]$, the squared profile is

$$m_{i,t}(u)^2 = \Sigma_c\, \Sigma_d\, w_c\, w_d\, k_c(u)\, k_d(u). \qquad (A.1)$$

Integrating over W and splitting the double sum into its diagonal (c = d) and off-diagonal (c ≠ d) parts gives

$$\Gamma_{i,t} = \Sigma_c\, w_c^2\, \|k_c\|^2 + 2\, \Sigma_{c<d}\, w_c\, w_d\, \langle k_c, k_d \rangle, \qquad (A.2)$$

where $\|k_c\|^2 = \int_W k_c(u)^2\, du$ and $\langle k_c, k_d \rangle = \int_W k_c(u)\, k_d(u)\, du$. The kernels are represented as probability vectors on the sixty-one-cell grid, so the inner products are dot products of those vectors; two identical Dirac-at-auction kernels give $\langle k, k \rangle = 1$. Substituting $w_c = a_{i,c,t-1}\, 1[c,t]$ recovers (3) exactly. The pipeline checks this identity on every panel cell and confirms a reconstruction error below 10^{-10}.

Appendix B. Reproducibility

The analysis reproduces from a single command. The pre-analysis plan is committed to version control before any headline regression, with a commit timestamp (2026-05-22T15:46:38+05:30) that predates every regression commit and that the regression code verifies before running. Every raw input, intermediate panel, and output figure is hashed with SHA-256 into an append-only audit trail, and every output filename embeds the run identifier (the canonical Phase 1 run is 6aeb77923ad53e54). Random operations use pre-committed seeds: 20240101 for the 1,000-iteration placebo and 19840722 for the power simulation. Four directional rejection artifacts are committed as first-class outputs and are cited in Section 10. Figures are generated by frozen scripts directly from committed pipeline outputs; no figure value is entered by hand.