Prices at Three Speeds

Aryan Ayyar*

First version: October 2025. This version: July 2026.

Abstract

Return predictability is usually measured at one horizon, even though prices can adjust at several speeds. I trace adjacent-return autocovariance from one day to three years in a 2002–2024 panel of Indian firms traded on both the NSE and BSE. The curve is small at short horizons, strongly positive over 3–8 months, and negative over 12–36 months. A dynamic permanent–transitory benchmark with one continuation rate and one reversal rate can cross zero at most once; the full curve therefore counts at least three economic clocks even when their exact rates are weakly identified. The model endogenizes the third clock as trend-chasing pressure that later unwinds. Same-firm venue gaps close within days, whereas the positive return band persists for months and is strongest where trading is costly and analyst coverage is low. Any account of price adjustment must therefore accommodate at least three distinct speeds.

Keywords: price impact, informed trading, momentum, reversal, attention, limits of arbitrage, price discovery.

JEL Classification: G12, G14, D82, D83.

*Fellow-in-Residence, Manipal Academy of Higher Education, Bengaluru (https://www.aryanayyar.com). The empirical analysis uses the daily equity data described in Section VI. All remaining errors are my own.

Disclosure Statement

The author has no conflicts of interest to disclose, no relevant or material financial interests that relate to the research described in this paper, and received no external funding for this research. The empirical analysis uses the daily equity data described in Section VI.

Return predictability changes character with the horizon. A security can reverse over one interval, continue over another, and reverse again later. Those signs are usually studied separately. Their ordered path reveals how many adjustment processes shape prices. I call each distinct adjustment or decay rate a *clock*. The horizon curve can count them even when their individual speeds are weakly identified.

Many Indian companies provide an unusual laboratory for that question. They end each trading day with one price on the National Stock Exchange (NSE) and another on the Bombay Stock Exchange (BSE). The cash-flow claim is the same; the order books are not. The BSE–NSE gap reveals how quickly a venue-specific discrepancy disappears. The return series, traced from one day to three years, reveals how predictability changes with the horizon. One security provides both a venue clock and a horizon curve.

The horizon curve is the paper's central object. In the 2002–2024 panel, which contains 5,690 regular sessions and a fixed universe of 1,613 listings across 822 companies, adjacent-return autocovariance is small at short horizons. It then rises sharply: the normalized 3–8 month band is roughly 10, and every component horizon is positive. At 12–36 months the curve is negative, with a band estimate of roughly -77. The same medium-positive, long-negative pattern appears in a 2017–2024 common-universe check and in a symmetric NSE–BSE diagnostic. Looking at one conventional horizon would miss the anatomy.

That anatomy restricts economic models. A permanent price component that adjusts gradually creates continuation. A transient liquidity component that decays creates reversal. Yet a model with one continuation rate and one reversal rate can cross zero at most once, for any parameter values and even when the underlying shocks are correlated (Proposition 4). The pattern places a lower bound on the number of adjustment rates needed to produce momentum and reversal.

The missing rate arises from the trading response to delayed interpretation. Public order flow can contain information before that information is fully reflected in executable prices. This delayed interpretation creates a medium-horizon trend. Traders enter to exploit it, but their own demand produces slow-moving price pressure. When their positions unwind, the same capital that chased continuation generates long-run reversal. Trend-capital entry and unwinding generate the third clock endogenously. Medium-horizon momentum and long-horizon reversal are two stages of one mechanism.

Costly interpretation and trading sustain continuation in equilibrium. Quote setters choose how often to refresh their interpretation of public order flow. Arbitrageurs enter when the permanent component recorded in transactions lags the public filter. Free entry pins the stationary variance of uninterpreted information to $4c_A\psi_A$, four times execution cost times attention cost (Theorem 1). The residual continuation is therefore set by the costs of arbitrage.

Cross-venue staleness would require the BSE–NSE gap to survive for months. In 2.57 million balanced same-firm observations, that gap closes with a pooled half-life of about 2.4 trading days, and the median company closes the gap faster still. A cross-venue price gap lasts days; the positive return band lasts months. This days-versus-months comparison bounds cross-venue staleness while allowing the common price to differ from fundamentals.

The cross-section then asks whether the medium clock is largest where the model says arbitrage should be most expensive. It is. The high-minus-low band is positive across quoted-spread, price-impact, and illiquidity groups; the price-impact contrast is about 16 normalized points. Analyst coverage moves in the opposite direction, with a high-minus-low difference of about -10. These are descriptive mechanism tests, because some proxy histories begin after 2017, but all four directions line up with costly interpretation.

The paper is closest to work that explains several return horizons within one mechanism. Hong and Stein (1999), Vayanos and Woolley (2013), and Luo et al. (2021) all connect continuation to a later correction through, respectively, gradual diffusion and trend trading, predictable fund flows, and skepticism about others' information. Jegadeesh et al. (2025) model the transition from short-term reversal to longer-term momentum using adjacent-return covariances, while Hendershott et al. (2022) show why multiple attention frequencies matter for price dynamics. The distinct question here is topological: what does the full horizon curve reveal about the minimum number and economic origin of price-adjustment rates? The same-security venue bound and the free-entry amplitude restriction give that question empirical content.

Sign changes identify the number of clocks more sharply than their individual speeds. Long-horizon inference is low-powered: a conservative eight-block calendar test finds no significant band at the five percent level. Structural rate estimates are also boundary-sensitive. The empirical sections emphasize the signed profile, predeclared bands, the venue speed bound, and comparative statics.

I. Prices, Information, and Liquidity

A. From the filter to the tape

Suppose public order flow has already revealed part of the asset's value while the transaction price still reflects an earlier interpretation of that flow. The price inferable from public history then differs from the price on the tape. The model separates those two prices from temporary liquidity pressure through four objects. The asset has terminal value V, and time is either discrete, $t = 0, 1, 2, \ldots$, or continuous, $t \in [0, T]$:

1. N_t: the instantaneous Bayesian efficient price, equal to the order-flow filter $\mathbb{E}[V \mid \mathcal{F}_t^Y]$;
2. M_t: the measured permanent component of the transaction price, common to market participants, which may adjust toward N_t with delay;
3. L_t: the transient liquidity component, possibly with several subcomponents;
4. $P_t = M_t + L_t$: the observed transaction price.

The gap $N_t - M_t$ is the stock of information that is inferable from public order flow but not yet incorporated in the measured permanent price. It is zero in the instantaneous rational-expectations benchmark. When interpretation takes time, M_t approaches N_t at a finite rate; Section II derives that rate from a costly refresh decision.

The decomposition maps directly into adjacent-return covariance. Gradual changes in M_t generate positive covariance between adjacent returns: information continues to enter the tape. Changes in L_t generate negative covariance: temporary price pressure reverses. At any horizon, the dominant force determines the sign of return predictability.

Permanent and transitory impact are jointly determined in the equilibrium foundations of Appendix A. In the one-period Kyle economy of Proposition 7, greater temporary price impact reduces informed trading, makes order flow less informative, and changes permanent impact. Proposition 8 gives the same joint determination in continuous time. Orthogonality provides a tractable benchmark for the covariance algebra.

B. The frictionless filtering benchmark

Let the terminal value satisfy $V \sim N(0, \Sigma_0)$. Noise order flow is Brownian: $dU_t = \sigma_u \, dB_t$. The informed trader chooses a trading rate q_t, so total order flow is

$$dY_t = q_t \, dt + \sigma_u \, dB_t. \tag{1}$$

Market makers observe $\mathcal{F}_t^Y = \sigma(Y_s : s \le t)$. The instantaneous Bayesian efficient price is $N_t = \mathbb{E}[V \mid \mathcal{F}_t^Y]$, with conditional variance $\Sigma_t = \mathrm{Var}(V \mid \mathcal{F}_t^Y)$. Conjecture a linear informed trading rule $q_t = \beta_t(V - N_t)$. Because $\mathbb{E}[V - N_t \mid \mathcal{F}_t^Y] = 0$, total order flow is already the public innovation process.

Lemma 1 (Continuous-time permanent impact)**.** *Given the linear trading rule, the permanent price-impact coefficient is the filtering coefficient $z_t = \beta_t \Sigma_t / \sigma_u^2$, so that $dN_t = z_t \, dY_t$, and the conditional variance evolves according to $\dot{\Sigma}_t = -z_t^2 \sigma_u^2$.*

In the Kyle–Back economy, market depth is constant, $z_t \equiv z = \sqrt{\Sigma_0}/(\sigma_u \sqrt{T})$, with $\Sigma_t = \Sigma_0(1 - t/T)$. All information is incorporated by T. But N_t is a martingale, so $\mathrm{Cov}(N_{t+h} - N_t, \, N_{t+2h} - N_{t+h}) = 0$. Instantaneous Bayesian price discovery cannot generate momentum by itself. The parameter $\sigma_N = z_t \sigma_u$ measures the local information-flow intensity.

Temporary impact $c_t q_t$ is the cost faced by an informed trader inside the trading problem. The transient state L_t is the temporary component observed in transaction prices, arising from liquidity provision, inventory compensation, and order-processing costs. The autocovariance analysis treats σ_N, ν, σ_L, and κ as local structural coefficients generated by the trading, attention, and liquidity environment. Appendix A gives equilibrium content to $\sigma_N = z_t \sigma_u$. Once N_t is a martingale, continuation requires M_t to lag behind it. The gap is public, however, and therefore invites a trade. What prevents that trade from erasing the delay?

II. The Price of Interpretation

If attention and execution were free, arbitrage would close the gap immediately. Quote setters instead choose how often to recompute the public filter, while arbitrageurs decide whether the remaining information gap justifies the attention and execution costs required to trade it. Over the horizons studied, the filtered efficient price has constant innovation intensity under Assumption 2, $dN_t = \sigma_N \, dW_t^N$.

A. Why quotes remain stale

A quote setter can observe order flow without continuously solving the inference problem hidden inside it. A continuum of setters $i \in [0,1]$ posts permanent quote components $m_{i,t}$. Setter i observes the public order-flow history but computes the filter N_t only at refresh times, which arrive at a chosen Poisson rate ν_i; computing at rate ν_i costs $\psi\nu_i$ per unit time, with $\psi > 0$. Upon refresh, the setter resets $m_{i,t} = N_t$; between refreshes the quote is unchanged. Stale quotes are picked off: traders who track N_t trade against setter i with intensity proportional to the mispricing, generating an expected loss flow $\varphi \, \mathbb{E}[(N_t - m_{i,t})^2]$ with picking-off intensity $\varphi > 0$. Section II.B endogenizes the pickers; here φ is a primitive.

Lemma 2 (Stationary staleness). *If setter i refreshes at Poisson rate ν_i, independent of W^N, the stationary expected squared staleness is $\mathbb{E}[(N_t - m_{i,t})^2] = \sigma_N^2/\nu_i$.*

Proposition 1 (Equilibrium attention). *Each setter minimizes $\varphi\sigma_N^2/\nu_i + \psi\nu_i$. The optimum is unique and symmetric:*

$$\nu^* = \sigma_N \sqrt{\varphi/\psi}, \tag{2}$$

increasing in information intensity and picking-off risk and decreasing in attention cost. The implied stationary variance of the uninterpreted stock is $\sigma_N^2/(2\nu^*) = \frac{1}{2}\sigma_N\sqrt{\psi/\varphi}$.

Lemma 3 (Aggregation). *Let* $M_t = \int_0^1 m_{i,t}\,di$. *If all setters refresh at rate* ν, *then* $dM_t = \nu(N_t - M_t)\,dt$.

The first implication is a hump. When interpretation is nearly free, $\psi \to 0$ and $\nu^* \to \infty$: quotes track the filter and continuation vanishes. When interpretation is prohibitively costly, $\nu^* \to 0$: almost nothing enters the permanent price within the horizon, so continuation again vanishes. Momentum is largest between those extremes.

The second implication is about matching a speed to a horizon. For a fixed h, the delayed-information covariance derived in Section III is maximized at $\nu h \approx 1.256$, the unique root of $2xe^{-x} = 1 - e^{-x}$. The same economy can therefore look efficient at one horizon and predictable at another. Momentum is strongest where the interpretation clock resonates with the return horizon.

B. Arbitrage and the information gap

Stale quotes are profitable to trade against because the information gap they create is public. Traders who compute it accelerate updating, so entry, attention, and execution costs jointly determine how much remains. The stock of unincorporated information,

$$H_t = N_t - M_t, \tag{3}$$

is a functional of the public order-flow history. Public, predictable arbitrage flow is filtered out of the innovation process, while free entry determines the variance of H_t.

Formally, let arbitrageurs submit an aggregate flow a_t adapted to $\mathcal{F}_t^Y$, so total order flow is $dY_t = q_t\,dt + a_t\,dt + \sigma_u\,dB_t$.

Lemma 4 (Filtering invariance under arbitrage flow). *For any publicly observed,* $\mathcal{F}^Y$*-predictable arbitrage flow* a_t *satisfying the usual integrability and pathwise-uniqueness conditions, the innovation process* $dI_t = dY_t - \mathbb{E}[q_t + a_t \mid \mathcal{F}_t^Y]\,dt = dY_t - a_t\,dt$ *is unchanged, and* *so are N_t, z_t, and Σ_t. Arbitrage flow moves measured prices; the efficient price and the speed of learning about V are unchanged.*

Arbitrage flow moves measured prices: executions against stale quotes force the executed setters to update, so arbitrage pressure passes through to the measured permanent component.

Assumption 1 (Pass-through of arbitrage flow)**.** *The measured permanent component obeys $dM_t = \nu_0(N_t - M_t)\,dt + \zeta a_t\,dt$, where ν_0 is the autonomous refresh speed of Proposition 1 and $\zeta \in (0, 1]$ is the fraction of arbitrage executions that trigger an update of the executed quote.*

A continuum of identical risk-neutral arbitrageurs may enter. An entrant pays attention cost ψ_A per unit time to track N_t, and trading at rate a against the measured price incurs quadratic execution costs, so the flow objective is $\max_a\, a(N_t - M_t) - c_A a^2 = \max_a\, aH_t - c_A a^2$, with $c_A > 0$. Individual optimality gives $a = H_t/(2c_A)$ and flow rents $H_t^2/(4c_A)$; with entrant mass A, aggregate flow is $a_t = \gamma H_t$ with $\gamma = A/(2c_A)$, so by Assumption 1 the effective interpretation speed is $\nu = \nu_0 + \zeta\gamma$.

Proposition 2 (Free entry into information-gap arbitrage)**.** *Suppose $\sigma_N^2 > 8c_A\psi_A\nu_0$. Then the free-entry equilibrium has positive arbitrage mass $A^* = \big(\sigma_N^2 - 8c_A\psi_A\nu_0\big)/(4\psi_A\zeta)$, effective interpretation speed*

$$\nu = \frac{\sigma_N^2}{8c_A\psi_A}, \tag{4}$$

and stationary variance of the information gap $\mathrm{Var}(H_t) = \sigma_N^2/(2\nu) = 4c_A\psi_A$. If $\sigma_N^2 \le 8c_A\psi_A\nu_0$, no arbitrageur enters and $\nu = \nu_0$, $\mathrm{Var}(H_t) = \sigma_N^2/(2\nu_0)$.

Theorem 1 (Equilibrium momentum bound)**.** *In the entry region of Proposition 2, the delayed-information continuation component of Section III satisfies, for every horizon $h > 0$,*

$$C_h^I = \mathrm{Var}(H_t)\left(1 - e^{-\nu h}\right)^2 \le 4c_A\psi_A, \tag{5}$$

with the bound approached as $h \to \infty$. Under the benchmark orthogonality of Section III, adjacent-return autocovariance therefore satisfies $\mathrm{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) \le 4c_A\psi_A$ at every *horizon. As $c_A \to 0$ or $\psi_A \to 0$, $\nu \to \infty$, the measured price converges to the filter, and continuation vanishes: the model recovers the martingale benchmark.*

Corollary 1 (Amplitude-horizon separation)**.** *In the entry region, the asymptotic amplitude of the continuation component, $\mathrm{Var}(H_t) = 4c_A\psi_A$, does not depend on the information-flow intensity σ_N; σ_N governs only the speed $\nu = \sigma_N^2/(8c_A\psi_A)$ at which the amplitude is approached. Two assets with equal trading costs but different information intensities display equally sized continuation at long horizons, reached at different horizons.*

The bound gives the continuation a price. Arbitrage proceeds until the marginal entrant's expected rents cover both attention and execution costs. The cost product $c_A\psi_A$ pins the amplitude; if either cost disappears, so does continuation. Corollary 1 then separates two empirical roles. Costs determine how much continuation can remain, while information flow determines where it appears on the horizon curve. The remaining question is how often that continuation can exchange dominance with temporary liquidity reversal as the horizon grows.

III. What Two Clocks Can Explain

The two-clock benchmark places one adjustment rate on each side of this contest: gradual permanent information against decaying liquidity pressure. Their relative magnitudes may reverse as the horizon changes. Can they reverse more than once?

A. Clock 1: delayed permanent information

Let M_t be the permanent component that enters observed returns. By Lemma 3 and Proposition 2, it adjusts toward the instantaneous Bayesian efficient price N_t at the equilibrium speed $\nu > 0$:

$$dM_t = \nu(N_t - M_t)\,dt. \tag{6}$$

With $H_t = N_t - M_t$ and locally constant filtering intensity $dN_t = \sigma_N\,dW_t^N$,

$$dH_t = \sigma_N\,dW_t^N - \nu H_t\,dt. \tag{7}$$

Assumption 2 (Stationary local reduction)**.** *The autocovariance analysis is local and excludes horizons near terminal revelation. Over the return horizons under study, σ_N, ν, σ_L, and κ are constant reduced-form equilibrium coefficients, and all transient states are initialized in their stationary distributions.*

The stationary variance is $\text{Var}(H_t) = \sigma_N^2/(2\nu)$, which Proposition 2 equates to $4c_A\psi_A$ in the entry region.

Lemma 5 (Delayed permanent continuation)**.** *For horizon $h > 0$,*

$$\text{Cov}(M_{t+h} - M_t,\, M_{t+2h} - M_{t+h}) = \frac{\sigma_N^2}{2\nu}(1 - e^{-\nu h})^2. \tag{8}$$

B. Clock 2: transient liquidity pressure

Let L_t be a transient liquidity state,

$$dL_t = -\kappa L_t\, dt + \sigma_L\, dW_t^L, \tag{9}$$

with $\kappa > 0$ and $\text{Var}(L_t) = \sigma_L^2/(2\kappa)$.

Lemma 6 (Transient liquidity reversal)**.** *For horizon $h > 0$,*

$$\text{Cov}(L_{t+h} - L_t,\, L_{t+2h} - L_{t+h}) = -\frac{\sigma_L^2}{2\kappa}(1 - e^{-\kappa h})^2. \tag{10}$$

C. The two-clock horizon curve

The observed transaction price is $P_t = M_t + L_t$, and horizon-h returns are $\Delta_h P_t = P_{t+h} - P_t$.

Assumption 3 (Benchmark orthogonality)**.** *The stationary permanent-backlog process H_t and the stationary transient liquidity process L_t are independent.*

Theorem 2 (Price-impact composition and return autocovariance)**.** *Under Assumptions 2 and 3, the adjacent horizon-h return autocovariance is*

$$\text{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) = \frac{\sigma_N^2}{2\nu}(1 - e^{-\nu h})^2 - \frac{\sigma_L^2}{2\kappa}(1 - e^{-\kappa h})^2. \tag{11}$$

Return continuation obtains if and only if the delayed permanent information component exceeds the transient liquidity component.

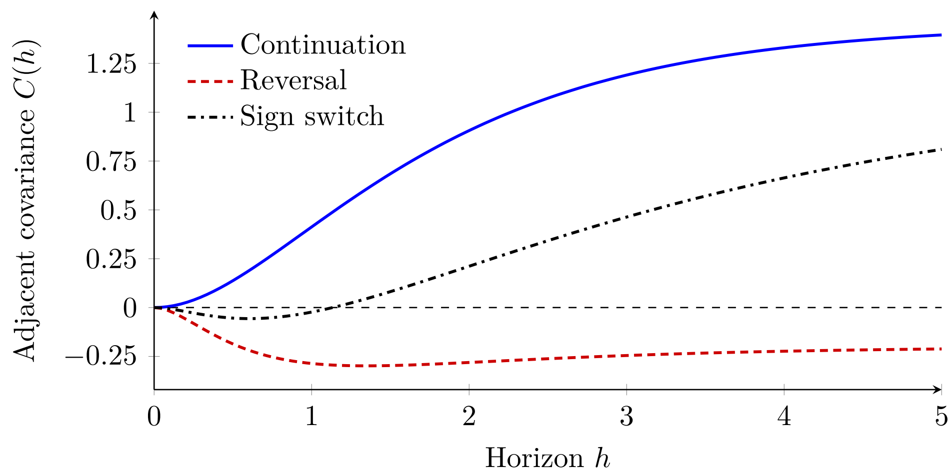

Figure 1. Model-implied horizon curves in the two-clock benchmark. The figure plots $C(h) = \frac{\sigma_N^2}{2\nu}(1-e^{-\nu h})^2 - \frac{\sigma_L^2}{2\kappa}(1-e^{-\kappa h})^2$ from Theorem 2 for three illustrative parameter sets. The figure is illustrative and uses no data.

Corollary 2 (Sign surface)**.** *If* $\sigma_L^2 > 0$, *then*

$$\mathrm{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) > 0 \quad \Longleftrightarrow \quad \frac{\sigma_N^2}{\sigma_L^2} > \frac{\nu}{\kappa}\left(\frac{1-e^{-\kappa h}}{1-e^{-\nu h}}\right)^2. \tag{12}$$

The sign condition compares two economic quantities: information power relative to liquidity power, and a threshold that changes with the two speeds and the horizon. Proposition 2 makes both sides equilibrium objects. The limiting cases give the intuition. If $\nu \to \infty$, information is incorporated immediately and the continuation term disappears. If $\nu \to 0$, almost no permanent information enters within the horizon and continuation again disappears. Predictability is largest between those extremes.

The horizon itself changes the contest. At short horizons the covariance is approximately $\frac{1}{2}(\sigma_N^2\nu - \sigma_L^2\kappa)h^2$; at long horizons it approaches $\sigma_N^2/(2\nu) - \sigma_L^2/(2\kappa)$. Thus the same asset can reverse at one horizon and continue at another. A single-horizon momentum statistic obscures the variation that distinguishes the mechanisms.

Figure 1 shows the possibilities. Depending on the relative power and speed of information and liquidity, the two-clock economy produces continuation, reversal, or one transition between them. Appendix Figure A1 plots the corresponding threshold $r^*(h)$.

D. The restriction under correlated shocks

Orthogonality provides the benchmark. Allowing correlated shocks tests whether the single-crossing result survives the joint determination of permanent and transitory impact.

Assumption 4 (Correlated shocks)**.** *The Brownian motions driving H_t and L_t have constant instantaneous correlation $\rho \in [-1, 1]$: $d\langle W^N, W^L\rangle_t = \rho\, dt$.*

Lemma 7 (Cross-covariances)**.** *Under Assumptions 2 and 4, $\Gamma_{HL}(0) \equiv \mathrm{Cov}(H_t, L_t) = \rho\sigma_N\sigma_L/(\nu + \kappa)$, and for $s \geq 0$, $\mathrm{Cov}(H_t, L_{t+s}) = \Gamma_{HL}(0)e^{-\kappa s}$, $\mathrm{Cov}(L_t, H_{t+s}) = \Gamma_{HL}(0)e^{-\nu s}$.*

Proposition 3 (General price-impact composition)**.** *Under Assumptions 2 and 4,*

$$\mathrm{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) = (1 - e^{-\nu h})^2 \left[\frac{\sigma_N^2}{2\nu} + \Gamma_{HL}(0)\right] - (1 - e^{-\kappa h})^2 \left[\frac{\sigma_L^2}{2\kappa} + \Gamma_{HL}(0)\frac{\nu}{\kappa}\right]. \quad (13)$$

Setting $\rho = 0$ recovers Theorem 2 exactly.

Remark 1 (Local effect of correlation)**.** *Differentiating the implied sign threshold with respect to ρ at $\rho = 0$ gives $2\kappa\sigma_N(\sigma_L^2 - \sigma_N^2)/\big((\nu + \kappa)\sigma_L^3\big)$. A small positive correlation pushes the boundary toward continuation when transient variance dominates and toward reversal when information variance dominates; correlated information and liquidity consumption has no fixed directional effect on the sign of predictability.*

E. The counting restriction

Both the orthogonal and correlated economies reduce to the same functional form. The horizon curves in Theorem 2 and Proposition 3 can be written as $f(h) = A(1 - e^{-\nu h})^2 - B(1 - e^{-\kappa h})^2$ with $A, B > 0$ constant in h.

Proposition 4 (Single crossing)**.** *Let $A, B > 0$ and $\nu, \kappa > 0$ with $\nu \neq \kappa$. Then $f(h) = A(1 - e^{-\nu h})^2 - B(1 - e^{-\kappa h})^2$ has at most one zero on $(0, \infty)$.*

Corollary 3 (The two-clock restriction)**.** *Under Theorem 2, and under Proposition 3 whenever the liquidity coefficient in (13) is positive, the return-autocovariance curve changes sign* *at most once. Thus a reversal–continuation–reversal curve requires another adjustment rate for all values of $\sigma_N, \nu, \sigma_L, \kappa$ and ρ.*

The single-crossing result has a sharp implication. A large literature documents reversal at very short horizons (Lehmann, 1990; Jegadeesh, 1990), momentum at intermediate horizons (Jegadeesh and Titman, 1993), and reversal over multi-year horizons (De Bondt and Thaler, 1985). A two-clock model can match pieces of that sequence. The full path leaves one structural timescale unaccounted for.

Remark 2 (Discrete-time special case)**.** *The familiar discrete-time boundary is the one-period analogue of Theorem 2. If a private signal shock s_t enters the measured permanent component over two periods, $\Delta M_t = K\mu s_t + \varepsilon_t$ and $\Delta M_{t+1} = K(1-\mu)s_t + \varepsilon_{t+1}$, and the transitory component is $\ell_t = Au_t$ with u_t i.i.d., then $\mathrm{Cov}(R_t, R_{t+1}) = \tau_s\mu(1-\mu) - A^2\sigma_u^2$ with $\tau_s = K^2\,\mathrm{Var}(s_t)$. The hump in informed participation μ is the discrete shadow of the hump in interpretation speed, and Proposition 1 supplies its equilibrium version: the hump in the attention cost ψ.*

IV. How Momentum Creates a Third Clock

A third exponential is easy to add. The harder question is what economic state could persist at that rate and why it should appear in equilibrium. Trend capital supplies both answers.

A. Why a third rate is necessary

Let the transient component be a sum of $K \geq 1$ mutually independent transient factors, $L_t = \sum_{k=1}^{K} L_{k,t}$, each Ornstein-Uhlenbeck with decay κ_k and intensity $\sigma_{L,k}$, independent of H_t. Then

$$\mathrm{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) = \frac{\sigma_N^2}{2\nu}(1-e^{-\nu h})^2 - \sum_{k=1}^{K} \frac{\sigma_{L,k}^2}{2\kappa_k}(1-e^{-\kappa_k h})^2, \tag{14}$$

With $K = 2$, the curve can change sign twice. For example, $\nu = 0.5$, $A = 1.4$, $\kappa_1 = 4$, $B_1 = 0.15$, $\kappa_2 = 0.1$, $B_2 = 2$ yields sign pattern negative-positive-negative with roots $h_1 \approx 0.79$ and $h_2 \approx 15.6$. A third timescale is therefore both necessary, by Proposition 4, and sufficient for the reversal–momentum–reversal sequence.

Finite-speed interpretation creates a trend worth chasing. Trend-capital entry turns the resulting trading pressure into the slow state.

B. Trend capital and slow unwinding

Trend-chasers use a coarse price signal. A mass $g \geq 0$ holds a per-unit-mass position equal to an exponentially weighted moving average s_t of past price changes,

$$ds_t = \omega_0 \left(dP_t - s_t \, dt\right), \tag{15}$$

with lookback rate $\omega_0 > 0$. Positions build after a price rise and decay at rate ω_0 when the trend stalls. These traders do not observe N_t, M_t, or H_t; the moving average s_t is their entire information set. That coarseness makes the strategy cheap to run, but it also makes its capital slow to exit. Their aggregate flow $g \, ds_t$ is absorbed by slow-moving liquidity providers who charge a price concession proportional to the inventory they carry, as in models of slow-moving capital: the resulting transient price component is

$$L_{2,t} = \eta \, g \, s_t \equiv p \, s_t, \qquad p \equiv \eta g \geq 0, \tag{16}$$

where $\eta > 0$ is the inventory-pressure coefficient. The transaction price is now $P_t = M_t + L_t + p \, s_t$. By Lemma 4, this publicly predictable flow leaves N_t, z_t, and Σ_t unchanged. The third clock is a measured-price state; the Bayesian value process remains unchanged.

The signal chases the price, and the price contains the signal's own pressure. That feedback sounds nonlinear economically but resolves linearly in the state system.

Lemma 8 (Feedback reduction)**.** *Suppose $\omega_0 p < 1$. Then (15) with $P_t = M_t + L_t + p \, s_t$ is equivalent to*

$$ds_t = \omega \left(dM_t + dL_t\right) - \omega \, s_t \, dt, \qquad \omega \equiv \frac{\omega_0}{1 - \omega_0 p}, \tag{17}$$

so the effective trend rate ω exceeds the lookback rate ω_0 and increases in the trend-chasing mass. If $\omega_0 p \geq 1$, no stationary linear equilibrium exists.

Write $s_t = s_t^M + s_t^L$, separating the part of the trend signal loaded on permanent returns from the part loaded on temporary pressure. The distinction will determine whether trend-following earns or loses. All second moments are exponential mixtures; the following lemma collects the terms needed for the horizon curve, with $d \geq 0$ throughout.

Lemma 9 (Cascade covariance functions)**.** *Under Assumption 2, Assumption 3, and (17), with $\omega \notin \{\nu, \kappa\}$:*

$$\mathrm{Cov}(s_t^M, H_{t+d}) = \frac{\omega\sigma_N^2}{2(\nu+\omega)}e^{-\nu d}, \qquad \mathrm{Cov}(s_t^L, L_{t+d}) = \frac{\omega\sigma_L^2}{2(\kappa+\omega)}e^{-\kappa d}, \tag{18}$$

$$\mathrm{Cov}(H_t, s_{t+d}^M) = \frac{\omega\sigma_N^2\left[2\nu e^{-\omega d} - (\nu+\omega)e^{-\nu d}\right]}{2(\nu-\omega)(\nu+\omega)}, \qquad \mathrm{Cov}(s_t^M, s_{t+d}^M) = \frac{\nu\omega\sigma_N^2\left[\nu e^{-\omega d} - \omega e^{-\nu d}\right]}{2(\nu^2-\omega^2)}, \tag{19}$$

$$\mathrm{Cov}(L_t, s_{t+d}^L) = \frac{\omega\sigma_L^2\left[(\kappa+\omega)e^{-\kappa d} - 2\omega e^{-\omega d}\right]}{2(\kappa-\omega)(\kappa+\omega)}, \qquad \mathrm{Cov}(s_t^L, s_{t+d}^L) = \frac{\omega^2\sigma_L^2\left[\kappa e^{-\kappa d} - \omega e^{-\omega d}\right]}{2(\kappa^2-\omega^2)}. \tag{20}$$

In particular $\mathbb{E}[s_t H_t] = \omega\sigma_N^2/(2(\nu+\omega)) > 0$ and $\mathbb{E}[s_t L_t] = \omega\sigma_L^2/(2(\kappa+\omega)) > 0$: the trend position loads positively on uninterpreted information and positively on transient liquidity pressure. The first loading is the source of trend-following profits; the second is the source of trend-following losses.

The feedback economy has a three-term horizon representation, with one squared-exponential term for each economic clock.

Theorem 3 (Three-clock representation)**.** *Under the assumptions of Lemma 9, the adjacent horizon-h return autocovariance of $P_t = M_t + L_t + p\, s_t$ is exactly*

$$C(h) = W_\nu\left(1 - e^{-\nu h}\right)^2 + W_\omega\left(1 - e^{-\omega h}\right)^2 + W_\kappa\left(1 - e^{-\kappa h}\right)^2, \tag{21}$$

with weights

$$W_\nu = \sigma_N^2 \, \frac{\nu^2(1+\omega p)^2 - \omega^2}{2\nu\,(\nu^2-\omega^2)}, \tag{22}$$

$$W_\omega = -\,\frac{p\,(2+\omega p)}{2}\left[\frac{\nu^2\sigma_N^2}{\nu^2-\omega^2} - \frac{\omega^2\sigma_L^2}{\kappa^2-\omega^2}\right], \tag{23}$$

$$W_\kappa = -\,\sigma_L^2 \, \frac{\kappa^2(1+\omega p)^2 - \omega^2}{2\kappa\,(\kappa^2-\omega^2)}. \tag{24}$$

At $p = 0$ the weights reduce to $W_\nu = \sigma_N^2/(2\nu)$, $W_\omega = 0$, $W_\kappa = -\sigma_L^2/(2\kappa)$, recovering Theorem 2 exactly.

The effective trend speed ω is the third rate and measures how quickly trend capital unwinds. The weight on that clock, W_ω, is negative if and only if $\nu^2\sigma_N^2/(\nu^2-\omega^2) > \omega^2\sigma_L^2/(\kappa^2-\omega^2)$, the condition under which the signal's information component dominates its bounce-noise exposure. In that region, trend-chasing converts continuation into slow reversal. All three weights depend on the same six primitives $(\sigma_N, \sigma_L, \nu, \kappa, \omega, p)$ and obey the cross-equation restrictions developed in Section V.

Corollary 4 (Long-run plateau)**.** $\lim_{h\to\infty} C(h) = W_\nu + W_\omega + W_\kappa = \frac{\sigma_N^2}{2\nu} - \frac{\sigma_L^2}{2\kappa} - p\left[\frac{\nu\sigma_N^2}{\nu+\omega} + \frac{\omega\sigma_L^2}{\kappa+\omega}\right] - p^2\left[\frac{\nu\omega\sigma_N^2}{2(\nu+\omega)} + \frac{\omega^2\sigma_L^2}{2(\kappa+\omega)}\right]$. *For fixed rates, the plateau is strictly decreasing in the trend-chasing pressure p: any positive mass of trend-chasers strictly deepens long-run reversal, at first order in mass.*

Remark 3 (Ultra-short horizons)**.** *As $h \to 0$, $\partial C(h)/\partial p\big|_{p=0} = h^2\omega\left[\frac{\nu^2\sigma_N^2}{\nu+\omega} - \frac{\sigma_L^2(\kappa^2+\kappa\omega+\omega^2)}{\kappa+\omega}\right] + O(h^3)$. When bounce noise is fast and sizable, the bracket is negative: trend-chasers, who cannot distinguish bounce from information, chase the bounce and deepen very-short-horizon reversal as well. The same friction that lets them survive (coarse information) makes them amplify both ends of the curve.*

Proposition 5 (Endogenous reversal-momentum-reversal)**.** *There exist parameters with $p > 0$ for which $C(h)$ in (21) changes sign exactly twice on $(0, \infty)$, with pattern negative-positive-negative, while the same parameters with $p = 0$ produce a single crossing. One*

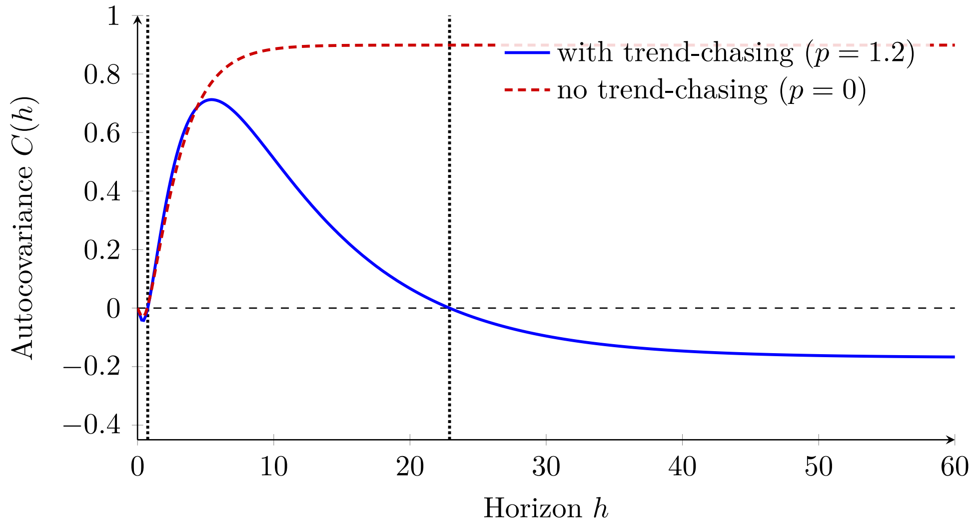

Figure 2. The endogenous third clock. The figure plots the three-clock representation (21) for $\nu = 0.5$, $\kappa = 4$, $\omega = 0.12$, $\sigma_N = 1$, $\sigma_L = 0.9$ (Proposition 5). With trend-chasing pressure $p = 1.2$ (solid), the curve is negative below $h_1 \approx 0.75$, positive between h_1 and $h_2 \approx 22.9$, and negative beyond: short-run reversal, intermediate momentum, long-run reversal. With $p = 0$ (dashed), the same economy single-crosses and continuation persists at all long horizons. Vertical dotted lines mark h_1 and h_2. The figure is illustrative and uses no data.

such parameterization is $\nu = 0.5$, $\kappa = 4$, $\omega = 0.12$, $\sigma_N = 1$, $\sigma_L = 0.9$, $p = 1.2$, *for which* $(W_\nu, W_\omega, W_\kappa) \approx (1.3276, -1.3641, -0.1325)$, *the crossings are* $h_1 \approx 0.75$ *and* $h_2 \approx 22.9$, *and the long-run plateau is* ≈ -0.169; *at* $p = 0$ *the curve crosses once and the plateau is* $+0.899$.

Figure 2 shows the three rates at work. The fast clock κ captures order-processing and inventory pressure in the sense of Roll (1984); it produces reversal below h_1. The interpretation clock ν dominates between h_1 and h_2, producing continuation. The slow trend clock ω then takes over. Because the trend signal loads mainly on information at these parameters, its weight is negative: positions accumulated during the momentum region unwind and pull the measured price back. Long-run reversal is the equilibrium shadow cast by intermediate momentum.

C. Who chooses the trend-chasing mass?

For the slow state to appear in equilibrium, trend followers must find entry profitable. A trend-chaser's expected flow profit is $\pi(\omega) = \omega\big[\nu\sigma_N^2/2(\omega+\nu) - \kappa\sigma_L^2/2(\omega+\kappa)\big]$ (Proposition 9). The position earns when it follows uninterpreted information and loses when it follows bounce. Profit rises while faster chasing loads more heavily on delayed information, peaks, and falls once bounce dominates. With a participation cost, free entry selects a stable interior mass in the regime plotted in Figure 2. Trend-chasers earn positive net profits there even as their aggregate pressure produces the later reversal.

When the information component is strong enough, entry becomes a strategic complementarity and moves the economy toward the fragility bound in Lemma 8, providing the model's connection to crash-prone momentum (Daniel and Moskowitz, 2016). The entry margin lies outside the empirical tests; Appendix A contains the formal proposition, profit figure, and crash interpretation.

D. What ties the three clocks together

The three exponentials share the same primitives. The information and liquidity processes that create the trend signal also determine the profitability of chasing it. The interpretation rate $\nu = \sigma_N^2/(8c_A\psi_A)$ and the amplitude bound $4c_A\psi_A$ share the same execution and attention costs. The trend rate ω is the unwinding speed of trend capital; its weight W_ω depends on $p(2 + \omega p)$ and on the information and liquidity clocks it chases. The bounce parameters (κ, σ_L) determine both W_κ and, through trend-following losses, the equilibrium mass p.

Changing one clock generally changes the incentives and weights of the others. These shared primitives impose cross-equation restrictions on the three-exponential representation.

V. What the Curve Identifies: Counting before Timing

Two crossings already reveal that two rates are insufficient. Recovering the three rates themselves is harder: nearby exponentials can produce almost the same finite-sample curve. This difference separates counting from timing. Under regularity conditions, a smooth curve identifies three rates, three weights, and the underlying primitives; its sign changes require less structure. Both inferences use sample autocovariances at a grid of horizons.

A. What the horizon curve identifies

Proposition 6 (What return autocovariances identify)**.** *Suppose the adjacent autocovariance curve satisfies* (21) *on an open interval of horizons, with distinct positive rates, no rate equal to twice another, and nonzero weights.*

1. *The three rates and three weights are uniquely determined by the curve. In the reversal-momentum-reversal regime, the labels are determined by the signs: the positive-weight rate is ν, and the smaller and larger negative-weight rates are ω and κ.*

2. *Suppose in addition that $\omega < \nu < \kappa$. Define, for $x \geq 1$,*

$$\Phi(x) = W_\omega + \frac{x-1}{2\omega}\left[\frac{2\nu^3 W_\nu}{\nu^2 x - \omega^2} + \frac{2\kappa\omega^2 W_\kappa}{\kappa^2 x - \omega^2}\right]. \tag{25}$$

   If $W_\nu/|W_\kappa| \geq \nu\omega^2(\kappa^2-\omega^2)/\big(\kappa^3(\nu^2-\omega^2)\big)$, then Φ is strictly increasing on $[1,\infty)$ and has the unique root $x^* = (1+\omega p)^2$, and the primitives are recovered in closed form:

$$p = \frac{\sqrt{x^*}-1}{\omega}, \qquad \sigma_N^2 = \frac{2\nu(\nu^2-\omega^2)\,W_\nu}{\nu^2 x^* - \omega^2}, \qquad \sigma_L^2 = -\frac{2\kappa(\kappa^2-\omega^2)\,W_\kappa}{\kappa^2 x^* - \omega^2}. \tag{26}$$

3. *In the entry region of Proposition 2, the product of execution and attention costs is identified from returns alone: $c_A\psi_A = \sigma_N^2/(8\nu)$.*

Proposition 6 establishes population identification. In this panel, however, the fitted rates lie near the parameter boundary. The empirical analysis relies on the predeclared band signs, the same-firm venue speed bound, and the friction and attention comparative statics. Appendix D documents the boundary behavior.

B. Monte Carlo lessons

The Monte Carlo illustrates the gap between clock counting and clock timing. Appendix C simulates the parameterization of Proposition 5 for a single asset and for pooled panels of 100 and 500 assets, with 204 replications per design and forty years of approximately daily prices (Table AI, Figure A3). A single series recovers only the fast rate and the first crossing. Pooling recovers the interpretation and liquidity rates, their variance intensities, and the cost product $c_A \psi_A$ much more tightly. The slow clock remains difficult: even with 500 assets, ω, W_ω, and the second crossing have wide intervals, and only 68% of replications detect both crossings. The simulation predicts where the empirical claims should be modest: the second crossing is the hardest feature to locate.

C. Predictions taken to the data

The theory reaches the data through four predictions.

1. *Counting.* By Proposition 4 and Corollary 3, a model with one continuation speed and one reversal speed produces at most one sign change in the horizon curve, whatever its parameters and whatever the shock correlation. More than one crossing counts at least three components.

2. *Amplitude.* Theorem 1 bounds continuation by $4c_A\psi_A$. In the cross-section, the momentum region should strengthen with execution-cost proxies, such as spreads, price impact, and illiquidity, and with attention-cost proxies such as inverse analyst coverage. Amplitudes far above any plausible cost product would reject the interpretation channel outright.

3. *Separation.* By Corollary 1, information intensity governs the horizon at which continuation peaks, while trading and attention costs govern its asymptotic size.

4. *Long reversal.* The endogenous trend-chasing clock turns medium-horizon continuation into a later reversal. The empirical implication is a positive 3–8 month band followed by a negative 12–36 month band.

VI. Two Prices for One Firm

Before attaching the medium clock to delayed interpretation, the same-firm design confronts a simpler possibility: the horizon pattern may inherit stale prices from two venues. That account requires the NSE–BSE gap itself to survive on a monthly horizon. Many Indian

Table I. Main empirical sample. The table reports the panel and estimand used in the main empirical tables. A listing is a company–venue security. Symmetric cross-venue rows retain companies with one eligible NSE listing and one eligible BSE listing.

Regular sessions	5,690
Eligible listings	1,613
Eligible companies	822
Symmetric NSE–BSE company pairs	791
Minimum price screen	INR 20
Drift buffer around target	21 sessions
Predeclared horizons	1, 2, 3, 5, 10, 16, 21, 32, 42, 63, 84, 126, 168, 210, 252, 378, 504, 630, 756

companies reveal that persistence directly. Their cash-flow claim is the same, while their order books and closing prices are venue-specific.

A. Data and estimand

The core sample is a quality-deduplicated same-firm panel from 2002 through 2024; the replication archive records the input-file checksums. The common regular-session calendar contains 5,690 sessions, and the fixed eligibility universe contains 1,613 listings across 822 companies. Returns are raw daily log returns. A return enters only when the immediately preceding declared-session close is at least INR 20. Missing returns are never imputed; sparse excluded-session returns are folded into the next included return when available.

The estimator preserves calendar time. For each horizon h, it forms two adjacent, complete h-session blocks. If a return is missing or fails a screen, the target is invalid; the procedure never joins observations across the gap and calls them consecutive. Drift is estimated from two calendar blocks outside the target, separated by a 21-session buffer, and the left and right assignments are symmetrized. The listing-level covariance is normalized by daily return variance. Predeclared horizons are first averaged within a listing and only then across listings or companies. Table I summarizes the design.

B. A same-firm speed bound

The required persistence is absent. Across 2.57 million balanced same-firm observations, the pooled daily AR(1) coefficient of the BSE–NSE log-price gap is 0.749, implying a half-life of 2.40 trading days. The median company coefficient is only 0.078. Mechanical disagreement between the two order books disappears far too quickly to account for a positive return band that persists for 3–8 months.

VII. The Horizon Curve in the Data

Venue disagreement disappears in days. Figure 3 follows the return curve on its own scale, from one day to three years. No structural fit enters the plot.

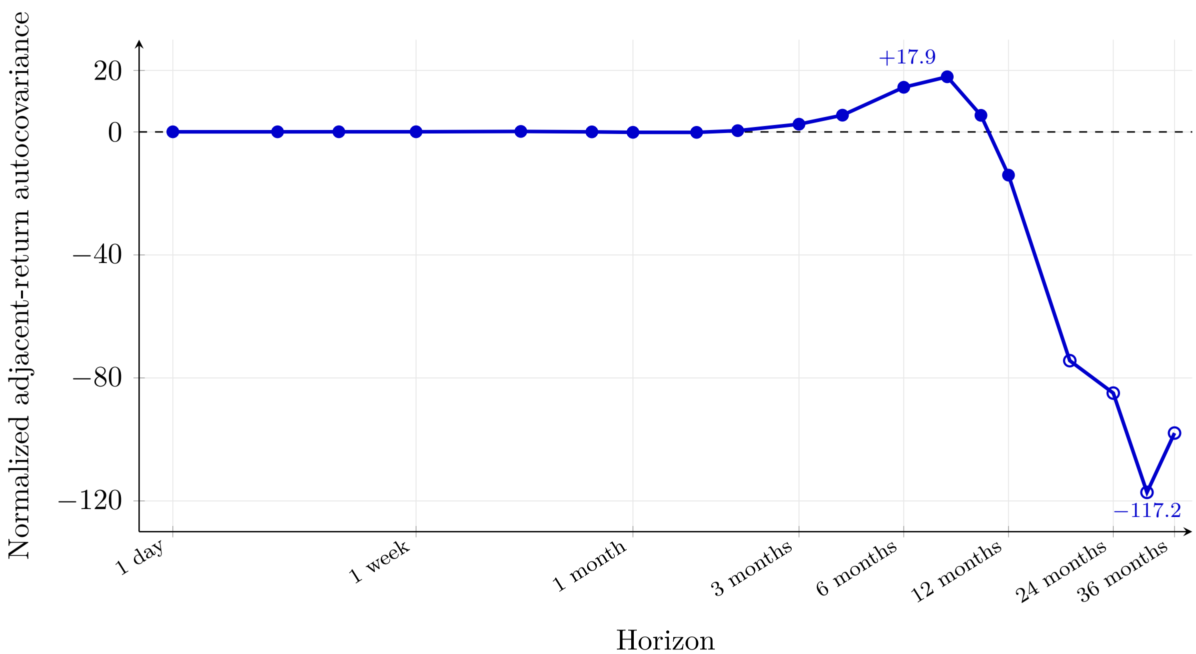

Figure 3. The empirical horizon curve. The figure plots the equal-listing mean of adjacent-return autocovariance at each of 19 predeclared horizons in the 2002–2024 panel, normalized by each listing's daily return variance. The horizontal axis is logarithmic so that days, months, and years remain visible on one scale. Filled markers denote horizons with at least eight effective calendar blocks; open markers denote longer horizons for which calendar inference is descriptive. The line connects model-free point estimates and is not a structural fit.

Table II. Horizon bands. Bands are arithmetic means of the predeclared horizons within listing. The normalized scale divides each listing's covariance by its daily return variance. Calendar HAC and company cluster columns are one-way diagnostics for the same point estimate. The long band has too few independent calendar blocks for calendar-time inference.

Band	*Horizons*	*Mean raw*	*Norm.*	*Listings*	*Effective blocks*
Micro	1,2,3,5,10	0.000065	0.068	1,613	90.30
Transition	16,21,32,42	0.000067	0.039	1,613	44.56
Medium clock	63,84,126,168	0.008909	10.115	1,612	14.68
Long reversal	252,378,504,630,756	-0.050516	-76.519	1,131	3.34

Table III. Common-universe and cross-venue robustness. The table reports normalized band estimates. The fixed early top-500 row uses the same company identifiers selected for the pre-2017 backfill. The expanded late row uses every company with a late-period primary moment. The cross-venue row symmetrizes NSE-left/BSE-right and BSE-left/NSE-right.

Design	*Micro*	*3–8m*	*12–36m*
2017–2024 fixed early top-500	0.070	13.401	-157.537
2017–2024 expanded late universe	0.057	13.005	-168.651
Symmetric NSE–BSE diagnostic	0.077	7.979	-137.344

A. From days to years

Table II traces the ordered profile. The short bands are close to zero, the 3–8 month band is economically large and positive, and the 12–36 month band is negative. Inference becomes less precise as the number of independent long-horizon blocks shrinks.

The normalized estimate rises from 2.513 at 63 sessions to 17.936 at 168 sessions, with positive values at 84 and 126 sessions in between. Every reported horizon in the long region is negative, from -14.049 at 252 sessions to values between roughly -74 and -117 thereafter. The model-free moments alone trace the curve.

B. Does the pattern depend on the universe or venue?

Two sample features could produce the shape spuriously: the universe expands after 2017, and one exchange could dominate the common price. Neither does. In the 2017–2024 period, the fixed early universe and the expanded late universe produce the same medium-positive, long-negative ordering. A symmetric cross-venue diagnostic, which pairs one eligible NSE listing with one eligible BSE listing for each company, does so as well.

Table IV. Friction, attention, and the 3–8 month band. Entries are high-minus-low differences in the normalized 3–8 month band. Equal-company rows first collapse the outcome and proxy to company level. Positive spread, impact, and illiquidity effects mean stronger medium-clock continuation in higher-friction names. The negative analyst-count effect means weaker continuation where attention is higher.

Proxy	*Equal listing*	*95% CI*	*Equal company*	*95% CI*
Quoted spread cost	4.787	[0.268, 9.306]	8.240	[1.783, 14.697]
Price-impact cost	16.021	[9.825, 22.217]	17.743	[11.112, 24.374]
Illiquidity	5.132	[0.452, 9.812]	9.958	[3.266, 16.650]
Analyst count	-9.917	[-15.568, -4.267]	-9.896	[-15.574, -4.218]

The remaining concern is calendar-time dependence. The exact-block audit partitions the common calendar into eight long blocks, calibrates size under simulated panels with the actual support masks, and selects the Ibragimov–Mueller group-t rule using simulated null size. With only eight blocks, none of the bands is significant at the five percent level. The low power at long horizons defines the claim boundary, while the point estimates agree with the main estimates.

VIII. Friction, Attention, and the Medium Clock

The curve has counted a medium clock. The mechanism now makes a harder cross-sectional prediction. Medium-horizon continuation should be larger when trading against the price-information gap is expensive and smaller when public information is easier to process. Quoted spreads, price impact, illiquidity, and analyst coverage provide four comparisons.

Each proxy is collapsed to a listing median and sorted without splitting tied values. Standard errors cluster listings by company and condition on the estimated return moments. The estimates are descriptive; mapping them into structural costs would require additional identification.

All four comparisons point in the predicted direction. The medium clock is larger among stocks with wider quoted spreads, greater price impact, and higher illiquidity. Analyst count moves the other way, as an attention proxy should. The largest contrast is price impact: high-impact listings have a 3–8 month band about 16 normalized points above low-impact listings, and the equal-company difference is 17.7. Together, the trading-cost and attention comparisons support the same mechanism.

Some proxy histories begin in 2017, so the estimates remain descriptive despite the stable signs in a common-window sensitivity. The weak monthly-momentum bridge and exploratory event study carry no identifying weight.

IX. Relation to the Literature

Four literatures meet here: horizon-dependent predictability, filtering of order flow, temporary price pressure, and feedback trading. Their different adjustment speeds jointly restrict the shape of the horizon curve.

A. Momentum, reversal, and the horizon dimension

The empirical ingredients are well established but usually measured separately. Jegadeesh (1990) and Lehmann (1990) document short-horizon reversal, Jegadeesh and Titman (1993) establish intermediate-horizon momentum, and De Bondt and Thaler (1985) document long-horizon reversal. Lee and Swaminathan (2000) show that trading volume predicts both the magnitude and persistence of momentum and the timing of its later reversal. More recent work makes horizons explicit: Medhat and Schmeling (2022) show that the one-month sign varies with turnover, Dai et al. (2024) separate the magnitude and persistence of liquidity-provision reversals, and Jegadeesh et al. (2025) model and test the transition from short reversal to longer momentum using adjacent-return covariances.

Several theories already connect continuation and correction. Daniel, Hirshleifer, and Subrahmanyam (1998) obtain continuing overreaction and eventual reversal from overconfidence and self-attribution. Hong and Stein (1999) show how gradual information diffusion creates an opportunity that simple momentum traders amplify into overshooting. Vayanos and Woolley (2013) generate momentum and reversal through inertial fund flows, and Luo et al. (2021) generate both from investors' skepticism about information produced by others. Kelly, Moskowitz, and Pruitt (2021) show that time-varying conditional risk exposures can explain a sizable share of both momentum and long reversal. The minimality theorem adds a lower bound on the number of adjustment modes consistent with the whole curve.

B. Filtering, interpretation, and attention

The model begins from the microstructure view that prices are filters. In Kyle (1985), informed trading and market-maker inference jointly determine price impact; Back (1992) gives the continuous-time version, and Back and Baruch (2004) connect the Kyle and Glosten–Milgrom settings. Glosten and Milgrom (1985) show how adverse selection creates a spread. The present model retains that Bayesian benchmark but distinguishes the information inferable from public order flow from the permanent price currently recorded in transactions.

That distinction connects to sticky information (Mankiw and Reis, 2002), infrequent rebalancing (Bogousslavsky, 2016), and limits of arbitrage (Shleifer and Vishny, 1997). Holden and Subrahmanyam (2002) study costly private-information acquisition and serial correlation, while Cespa and Vives (2012) study dynamic trading and public information. Hendershott et al. (2022) show that several Poisson attention frequencies are needed to fit price, inventory, and retail-flow dynamics. Here the costly object is narrower: interpreting information already present in public order flow. Quote setters choose the refresh rate, and free entry prices the continuation that remains.

C. Liquidity, price pressure, and state dependence

The permanent–transitory decomposition follows the price-impact literature. Glosten and Harris (1988) separate adverse-selection impact from inventory and order-processing effects; Hasbrouck (1991) measures the information content of trades; Madhavan, Richardson, and Roomans (1997) emphasize unexpected order flow; and Huang and Stoll (1997) decompose the spread. Roll (1984) provides the short-reversal benchmark, and the decaying transient state is close to Obizhaeva and Wang (2013). Sadka (2006) directly connects a Glosten–Harris liquidity decomposition to momentum and post-earnings drift. The present analysis links permanent and transitory rates to changes in the sign of predictability across horizons.

The empirical interpretation must also distinguish stock-level trading costs from aggregate liquidity states. Campbell, Grossman, and Wang (1993) and Llorente et al. (2002) use volume to separate information from risk-sharing trades, while Chordia and Subrahmanyam (2004) study order imbalance and short-horizon predictability. Nagel (2012) interprets short-run reversal returns as compensation for liquidity provision. Hameed, Kang, and Viswanathan (2010) show that market declines and tight funding reduce liquidity; Avramov, Cheng, and Hameed (2016) find momentum profits are larger after liquid aggregate market states; and Daniel and Moskowitz (2016) show that momentum crashes following market declines and during rebounds. These aggregate results confine the cross-sectional interpretation to stock-level cost and attention mechanisms.

D. Feedback, flows, and same-security designs

The endogenous third clock belongs to the literature on positive-feedback trading and slow-moving capital. De Long et al. (1990) study destabilizing feedback trading; Hong and Stein (1999) show how momentum traders arbitrage underreaction into overreaction; and Vayanos and Woolley (2013) derive continuation and later reversal from predictable fund flows. The slow absorption of trend capital is also in the spirit of Grossman and Miller (1988) and Duffie (2010). Here, feedback trading creates a third squared-exponential term whose weight is tied to the information and liquidity clocks it chases.

Finally, Chui, Subrahmanyam, and Titman (2022) use different share classes of the same firms to hold cash flows fixed while studying investor clientele. The NSE–BSE setting applies the same-security comparison to the persistence of venue discrepancies. Investor allocation across venues remains endogenous, so the design identifies the lifespan of a venue-specific price gap. This controlled comparison produces the days-versus-months bound.

X. Conclusion

Return predictability is a curve, not a single number. In the Indian same-firm, two-venue panel, that curve can be followed from one trading day to three years. Its short bands are small, its 3–8 month region is strongly positive, and its 12–36 month region is strongly negative. A conventional momentum horizon would reveal only the middle of that path.

The counting theorem turns the path into an economic restriction. One continuation rate and one reversal rate can change dominance at most once, for any parameters and even with correlated shocks. The two crossings in the Indian curve identify at least three timescales. Boundary-sensitive estimates leave their individual speeds unresolved.

The crossing count alone leaves the economic states unnamed. Same-firm BSE–NSE gaps narrow that ambiguity: their pooled half-life is 2.4 trading days, and the median-company half-life is below one day. The venue price gap closes in days; the medium band lasts months. In the cross-section, the 3–8 month band is larger when spreads, price impact, and illiquidity make arbitrage costly, and smaller when analyst coverage makes public information easier to interpret. The two comparisons point toward costly interpretation while leaving the speed of fundamental learning unresolved.

The third clock links the two ends of the curve. Delayed interpretation makes trend-following profitable; trend capital then builds a slow price-pressure state whose eventual unwinding produces long reversal—momentum's equilibrium afterlife.

Timing the clocks directly will require longer panels or data on the states the model leaves latent: order flow and inference for the interpretation clock, and trend-follower positions for the slow unwinding clock. Such data could price the execution-cost and attention-cost product that bounds continuation and the trend-capital mass that converts it into reversal. The counting is done; the timing is the agenda.

Appendix A. Supporting Equilibrium Results

The main text uses a small set of local coefficients so that the horizon logic remains visible. This appendix supplies the equilibrium machinery behind them. It derives static and continuous-time trading foundations for the permanent–transitory decomposition, plots the information–liquidity boundary from Corollary 2, and prices the trend-chasing mass.

A. The static fixed point

The one-period economy makes the joint determination of permanent and transitory impact easiest to see. Start with a Kyle market and add an execution cost. Let m be the public prior mean of V, $\delta = V - m \sim N(0, \sigma_\delta^2)$, let an informed trader submit $x = \beta\delta$, and let noise traders submit $u \sim N(0, \sigma_u^2)$, so signed order flow is $y = x + u$. The market maker sets the permanent component by Bayesian updating, $M = m + zy$, and the transaction price includes temporary price impact, $P = m + (z + c)y$.

Proposition 7 (Static Kyle–Glosten–Harris fixed point)**.** *In the one-period linear-normal economy, a linear equilibrium with permanent impact z, temporary impact $c \geq 0$, and informed intensity β satisfies*

$$z = \frac{\beta\sigma_\delta^2}{\beta^2\sigma_\delta^2 + \sigma_u^2}, \qquad \beta = \frac{1}{2(z + c)}, \tag{27}$$

and total execution impact $\Lambda = z + c$ solves $4\sigma_u^2\Lambda^3 - 4c\sigma_u^2\Lambda^2 - \sigma_\delta^2\Lambda - c\sigma_\delta^2 = 0$ on the branch $\Lambda > c$, $z > 0$, $\beta > 0$. When $c = 0$, the equilibrium is Kyle's benchmark, $\beta = \sigma_u/\sigma_\delta$, $z = \sigma_\delta/(2\sigma_u)$.

The execution cost changes informed trading, which changes the information in order flow and hence permanent impact. Permanent and transitory components are jointly determined, while shock orthogonality provides the covariance benchmark used below.

B. Temporary execution cost in continuous time

Let the execution price faced by the informed trader be $P_t^{exec} = N_t + c_t q_t$, where $c_t > 0$ is deterministic. Define $G_t = V - N_t$, so $dG_t = -z_t q_t\, dt - z_t \sigma_u\, dB_t$. The informed trader

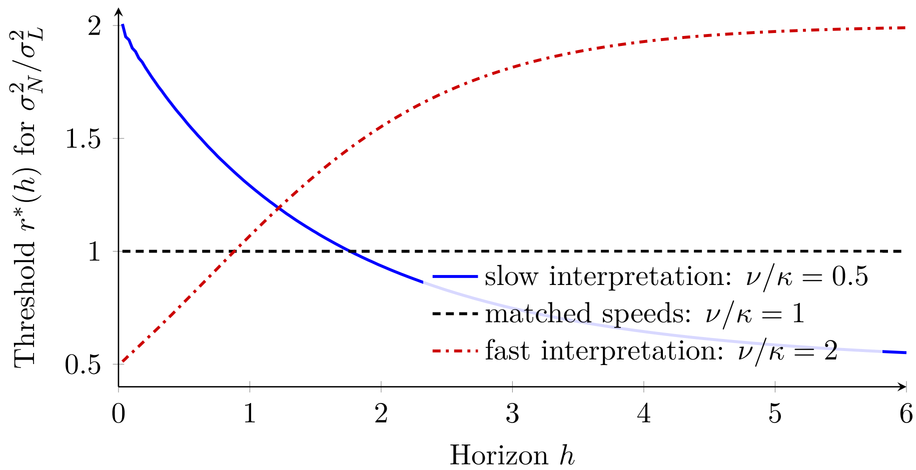

Figure A1. The information–liquidity sign boundary. The figure plots the threshold $r^*(h) = \frac{\nu}{\kappa}\left(\frac{1-e^{-\kappa h}}{1-e^{-\nu h}}\right)^2$ from Corollary 2, holding κ fixed and varying the interpretation speed ν. Continuation requires $\sigma_N^2/\sigma_L^2 > r^*(h)$. The figure is illustrative and uses no data.

maximizes expected cumulative trading profits; with quadratic value function $J(t, G) = A_t G^2 + B_t$, the Hamilton-Jacobi-Bellman equation yields the following system.

Proposition 8 (Continuous-time Kyle–Glosten–Harris system)**.** *With local temporary execution cost $c_t > 0$, a linear equilibrium is characterized by*

$$\beta_t = \frac{1 - 2A_t z_t}{2c_t}, \qquad z_t = \frac{\beta_t \Sigma_t}{\sigma_u^2}, \qquad \dot{\Sigma}_t = -z_t^2 \sigma_u^2, \qquad \dot{A}_t = -\frac{(1 - 2A_t z_t)^2}{4c_t}, \tag{28}$$

on the branch $1 - 2A_t z_t > 0$, with terminal condition $A_T = 0$.

C. The information–liquidity sign boundary

Figure A1 plots the threshold $r^*(h)$ of Corollary 2: continuation requires the information-to-liquidity power ratio to clear a horizon-adjusted speed threshold, which shifts with the interpretation speed ν.

D. Entry, complementarity, and fragility

The third clock is endogenous only if trend capital enters voluntarily. A trend-chaser's impact-adjusted flow profit per unit mass values external price changes at the previsible position and changes in the cohort's own linear pressure at the midpoint of its position adjustment:

$$d\Pi_t = s_t\, d(M_t + L_t) + \frac{p}{2}\, d(s_t^2) = s_t\, dP_t + \frac{p}{2}\, d[s]_t. \tag{29}$$

The second equality follows from Itô's formula. The midpoint term removes gains or losses produced solely by marking the cohort's own linear price pressure at one endpoint of a trade. In stationarity, the exact differential has zero expectation. Lemma 9 then leaves two economic terms: profits from following uninterpreted information and losses from following bounce.

Proposition 9 (Trend-following profits and entry)**.** *Under the assumptions of Lemma 9,*

$$\pi(\omega) \equiv \frac{\mathbb{E}[d\Pi_t]}{dt} = \nu\, \mathbb{E}[s_t H_t] - \kappa\, \mathbb{E}[s_t L_t] = \omega \left[\frac{\nu\sigma_N^2}{2(\omega+\nu)} - \frac{\kappa\sigma_L^2}{2(\omega+\kappa)} \right], \tag{30}$$

where $\omega = \omega_0/(1-\omega_0 p)$ increases in the mass g. Hence:

1. *If $\sigma_N^2 \le \sigma_L^2$ and $\nu\sigma_N^2 \le \kappa\sigma_L^2$, then $\pi \le 0$ for all ω: no trend-chasing survives.*

2. *If $\sigma_L^2 < \sigma_N^2 < (\kappa/\nu)\sigma_L^2$, then π is positive on $(0, \bar{\omega})$ and negative beyond, with $\bar{\omega} = \nu\kappa(\sigma_N^2 - \sigma_L^2)/(\kappa\sigma_L^2 - \nu\sigma_N^2)$: chasing faster captures relatively more bounce than information. With a participation cost $F \in (0, \max_\omega \pi)$ per unit mass, free entry has a stable interior solution at the larger root of $\pi(\omega(g)) = F$, where π is decreasing in mass.*

3. *If $\nu\sigma_N^2 > \kappa\sigma_L^2$ and $\sigma_N^2 > \sigma_L^2$, then $\pi > 0$ for every chasing speed, with $\pi(\infty) = (\nu\sigma_N^2 - \kappa\sigma_L^2)/2 > 0$: for small enough participation costs no finite mass extinguishes profits, and entry is limited only by capital or by the fragility bound of Lemma 8. If moreover $\nu\sigma_N \ge \kappa\sigma_L$, profits strictly increase in aggregate mass through the feedback multiplier: entry is a strategic complementarity.*

Remark 4 (Fragility and momentum crashes)**.** *By Lemma 8, the feedback multiplier $1/(1-\omega_0 p)$ diverges as $p \uparrow 1/\omega_0$: there is a finite trend-chasing mass beyond which no stationary linear equilibrium exists. In the complementarity regime of Proposition 9, entry pushes the market toward that bound, the trend clock accelerates ($\omega \to \infty$), and $|W_\omega|$ grows without* *limit. Daniel and Moskowitz (2016) emphasize crash-prone momentum. The model locates both ignition and fragility in a measurable object: the distance $1 - \omega_0 p$ between trend-chasing pressure and its fragility bound.*

Figure A2 shows the interior-entry regime. Profit initially rises with chasing speed because the signal captures more delayed information. It peaks and then falls as faster chasing loads increasingly on bounce. The participation-cost line selects the stable mass. Trend-chasers can therefore earn positive net profits and still create their own eventual reversal: individual profitability and aggregate destabilization coexist.

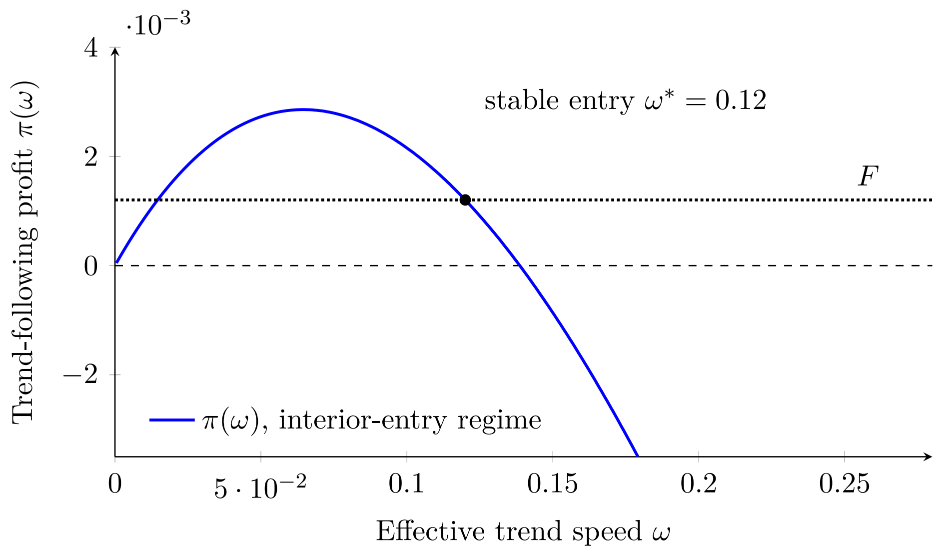

Figure A2. The trend-chasing entry margin. The figure plots expected trend-following profit per unit mass, equation (30), for $\nu = 0.5$, $\kappa = 4$, $\sigma_N = 1$, $\sigma_L = 0.9$ (the interior-entry regime of Proposition 9: profits are positive below $\bar{\omega} \approx 0.139$ and negative beyond). The dotted line is a participation cost F; the marked point is the stable free-entry equilibrium at $\omega^* = 0.12$, the value used in Proposition 5 and Figure 2. The figure is illustrative and uses no data.

Appendix B. Proofs

Proof of Proposition 7. The proof closes the two equilibrium conditions in (27). Because the market maker observes $y = \beta\delta + u$, joint normality gives

$$z = \frac{\mathrm{Cov}(\delta, y)}{\mathrm{Var}(y)} = \frac{\beta\sigma_\delta^2}{\beta^2\sigma_\delta^2 + \sigma_u^2}.$$

The informed trader takes total execution impact $\Lambda = z + c$ as given and maximizes $x(\delta - \Lambda x)$. The first-order condition gives $x = \delta/(2\Lambda)$, and hence $\beta = 1/(2\Lambda)$. Substitution into the market maker's coefficient gives

$$\Lambda - c = \frac{2\Lambda\sigma_\delta^2}{\sigma_\delta^2 + 4\Lambda^2\sigma_u^2} \quad \iff \quad 4\sigma_u^2\Lambda^3 - 4c\sigma_u^2\Lambda^2 - \sigma_\delta^2\Lambda - c\sigma_\delta^2 = 0.$$

The admissible branch has $\Lambda > c$, so $z = \Lambda - c > 0$ and $\beta > 0$. When $c = 0$, the positive root is $\Lambda = \sigma_\delta/(2\sigma_u)$, which gives $\beta = \sigma_u/\sigma_\delta$. $\square$

Proof of Lemma 1. Under $q_t = \beta_t(V - N_t)$, order flow has innovation representation $dY_t = \beta_t(V - N_t)\,dt + \sigma_u\,dB_t$ with $\mathbb{E}[q_t \mid \mathcal{F}_t^Y] = 0$. The Kalman-Bucy filter gives $dN_t = (\beta_t\Sigma_t/\sigma_u^2)\,dY_t$, so $z_t = \beta_t\Sigma_t/\sigma_u^2$, and the Riccati equation $\dot{\Sigma}_t = -\beta_t^2\Sigma_t^2/\sigma_u^2 = -z_t^2\sigma_u^2$. $\square$

Proof of Proposition 8. Let $G_t = V - N_t$, so $dG_t = -z_t q_t\,dt - z_t\sigma_u\,dB_t$, and conjecture $J(t, G) = A_t G^2 + B_t$. With flow payoff $qG - c_t q^2$, the HJB is

$$0 = \dot{A}_t G^2 + \dot{B}_t + \max_q \left\{(1 - 2A_t z_t)Gq - c_t q^2 + A_t z_t^2\sigma_u^2\right\}.$$

The maximand is strictly concave in q. Its first-order condition and the linear strategy $q_t = \beta_t G_t$ therefore give

$$q_t = \frac{1 - 2A_t z_t}{2c_t} G_t, \qquad \beta_t = \frac{1 - 2A_t z_t}{2c_t}.$$

Substituting the optimizer back into the HJB gives

$$0 = \left[\dot{A}_t + \frac{(1 - 2A_t z_t)^2}{4c_t}\right] G^2 + \dot{B}_t + A_t z_t^2\sigma_u^2.$$

The coefficient on G^2 yields the equation for $\dot{A}_t$; the constant coefficient determines B_t, which does not enter the trading rule. Lemma 1 supplies $z_t = \beta_t\Sigma_t/\sigma_u^2$ and $\dot{\Sigma}_t = -z_t^2\sigma_u^2$. The terminal condition $J(T, G) = 0$ implies $A_T = 0$, completing (28). $\square$

Proof of Lemma 2. For a stationary Poisson refresh process with rate ν_i, the age $A_{i,t}$ of the current quote is exponentially distributed with mean $1/\nu_i$, independent of W^N. Conditional on age a, the quote is $m_{i,t} = N_{t-a}$, and since N has independent increments with $\mathrm{Var}(N_t -$

$N_{t-a}) = \sigma_N^2 a$, $\mathbb{E}[(N_t - m_{i,t})^2] = \sigma_N^2\,\mathbb{E}[A_{i,t}] = \sigma_N^2/\nu_i$. $\square$

Proof of Proposition 1. The objective $\varphi\sigma_N^2/\nu + \psi\nu$ is strictly convex on $(0,\infty)$ with unique stationary point $\nu^* = \sigma_N\sqrt{\varphi/\psi}$ (second derivative $2\varphi\sigma_N^2/\nu^3 > 0$). Substituting into $\sigma_N^2/(2\nu^*)$ gives $\frac{1}{2}\sigma_N\sqrt{\psi/\varphi}$. $\square$

Proof of Lemma 3. Over $[t, t+dt]$, a fraction $\nu\,dt$ of setters refreshes to N_t while the rest keep their quotes, so $M_{t+dt} - M_t = \nu\,dt \int_0^1 (N_t - m_{i,t})\,di + o(dt) = \nu(N_t - M_t)\,dt + o(dt)$. $\square$

Proof of Lemma 4. Because a_t is publicly observed and predictable,

$$\mathbb{E}[dY_t \mid \mathcal{F}_t^Y] = \big(a_t + \beta_t \mathbb{E}[V - N_t \mid \mathcal{F}_t^Y]\big)\,dt = a_t\,dt.$$

Define the cumulative innovation by

$$I_t := Y_t - \int_0^t a_s\,ds, \qquad dI_t = q_t\,dt + \sigma_u\,dB_t.$$

The first definition makes I measurable with respect to Y. Conversely, pathwise uniqueness for the predictable feedback rule reconstructs $Y_t = I_t + \int_0^t a_s\,ds$ from the history of I. Thus $\mathcal{F}_t^I = \mathcal{F}_t^Y$. After subtracting the observable drift, the market maker faces exactly the benchmark observation equation, and the Kalman–Bucy filter depends on order flow only through I. Its conditional mean, impact coefficient, and conditional variance are therefore the same N_t, z_t, and Σ_t as without arbitrage flow. $\square$

Proof of Proposition 2. An individual arbitrageur solves $\max_a aH_t - c_A a^2$. Thus

$$a^*(H_t) = \frac{H_t}{2c_A}, \qquad \pi_A(H_t) = \frac{H_t^2}{4c_A}, \qquad \mathbb{E}[\pi_A(H_t)] = \frac{\mathrm{Var}(H_t)}{4c_A}.$$

With mass A, aggregate flow is γH_t, where $\gamma = A/(2c_A)$. Assumption 1 and Lemma 4 give

$$dH_t = \sigma_N\,dW_t^N - \nu H_t\,dt, \qquad \nu = \nu_0 + \zeta\gamma, \qquad \mathrm{Var}(H_t) = \frac{\sigma_N^2}{2\nu}.$$

When the implied mass is positive, free entry requires

$$\frac{\mathrm{Var}(H_t)}{4c_A} = \psi_A \quad\iff\quad \nu = \frac{\sigma_N^2}{8c_A\psi_A}.$$

Combining this equality with $\gamma = A/(2c_A)$ gives (4) and the stated A^*, which is positive exactly when $\sigma_N^2 > 8c_A\psi_A\nu_0$. Otherwise the rent at $A = 0$ is below the entry cost, so $A^* = 0$ and $\nu = \nu_0$. $\square$

Proof of Theorem 1 and Corollary 1. By Lemma 5, $C_h^I = \mathrm{Var}(H_t)(1 - e^{-\nu h})^2$, which is increasing in h with supremum $\mathrm{Var}(H_t)$, equal to $4c_A\psi_A$ in the entry region by Proposition 2. Under Assumption 3 the transient term in (11) is nonpositive, so the total is at most $C_h^I \leq 4c_A\psi_A$. As $c_A\psi_A \to 0$ with σ_N fixed, $\nu = \sigma_N^2/(8c_A\psi_A) \to \infty$ and $C_h^I \leq \mathrm{Var}(H_t) = 4c_A\psi_A \to 0$. Finally, $\mathrm{Var}(H_t) = 4c_A\psi_A$ and $\nu = \sigma_N^2/(8c_A\psi_A)$ establish the separation result. $\square$

Proof of Lemma 5. H_t is Ornstein-Uhlenbeck with $\Gamma_H(s) = \sigma_N^2 e^{-\nu|s|}/(2\nu)$. Since $M_{t+h} - M_t = \int_t^{t+h} \nu H_s\, ds$, $\mathrm{Cov}(\Delta_h M_t, \Delta_h M_{t+h}) = \nu^2 \int_0^h \int_h^{2h} \Gamma_H(b-a)\, db\, da = \frac{\sigma_N^2}{2\nu}(1 - e^{-\nu h})^2$. $\square$

Proof of Lemma 6. With $\Gamma_L(s) = \sigma_L^2 e^{-\kappa|s|}/(2\kappa)$, $\mathrm{Cov}(L_{t+h} - L_t, L_{t+2h} - L_{t+h}) = 2\Gamma_L(h) - \Gamma_L(0) - \Gamma_L(2h) = -\frac{\sigma_L^2}{2\kappa}(1 - e^{-\kappa h})^2$, using $2e^{-\kappa h} - 1 - e^{-2\kappa h} = -(1 - e^{-\kappa h})^2$. $\square$

Proof of Theorem 2 and Corollary 2. Write $\Delta_h X_t = X_{t+h} - X_t$. Since $P = M + L$, bilinearity gives the complete decomposition

$$\begin{aligned}\mathrm{Cov}(\Delta_h P_t, \Delta_h P_{t+h}) = {} & \mathrm{Cov}(\Delta_h M_t, \Delta_h M_{t+h}) + \mathrm{Cov}(\Delta_h M_t, \Delta_h L_{t+h}) \\ & + \mathrm{Cov}(\Delta_h L_t, \Delta_h M_{t+h}) + \mathrm{Cov}(\Delta_h L_t, \Delta_h L_{t+h}).\end{aligned}$$

Assumption 3 sets the two cross terms to zero. Lemmas 5 and 6 supply the first and fourth terms, respectively, yielding (11). For $\sigma_L^2 > 0$, dividing the continuation inequality by $\sigma_L^2(1 - e^{-\nu h})^2/(2\kappa) > 0$ gives (12) without reversing its sign. $\square$

Proof of Lemma 7. Write the stationary solutions as $H_t = \sigma_N \int_{-\infty}^t e^{-\nu(t-u)} dW_u^N$ and $L_t = \sigma_L \int_{-\infty}^t e^{-\kappa(t-u)} dW_u^L$. Then $\mathrm{Cov}(H_t, L_t) = \rho\sigma_N\sigma_L \int_0^\infty e^{-(\nu+\kappa)v} dv = \rho\sigma_N\sigma_L/(\nu+\kappa)$. For $s \geq 0$, $L_{t+s} = e^{-\kappa s}L_t + \int_t^{t+s}(\cdot)\, dW^L$ with the integral uncorrelated with H_t, so $\mathrm{Cov}(H_t, L_{t+s}) = e^{-\kappa s}\Gamma_{HL}(0)$; symmetrically for $\mathrm{Cov}(L_t, H_{t+s})$. $\square$

Proof of Proposition 3. Since $\Delta_h M_t = \nu \int_t^{t+h} H_u \, du$ and Lemma 7 gives the two directional cross-covariances decay at different rates,

$$\begin{aligned}
\operatorname{Cov}(\Delta_h M_t, \Delta_h L_{t+h}) &= \nu \Gamma_{HL}(0) \int_0^h \left[e^{-\kappa(h+r)} - e^{-\kappa r}\right] dr \\
&= -\frac{\nu}{\kappa} \Gamma_{HL}(0)(1 - e^{-\kappa h})^2, \\
\operatorname{Cov}(\Delta_h L_t, \Delta_h M_{t+h}) &= \nu \Gamma_{HL}(0) \int_0^h \left[e^{-\nu r} - e^{-\nu(h+r)}\right] dr \\
&= \Gamma_{HL}(0)(1 - e^{-\nu h})^2.
\end{aligned}$$

Adding these terms to (11) gives (13). The local effect in Remark 1 follows by substituting $\Gamma_{HL}(0) = \rho \sigma_N \sigma_L / (\nu + \kappa)$ into the two bracketed coefficients and differentiating their ratio at $\rho = 0$. □

Proof of Proposition 4. Let $x = e^{-\nu h} \in (0, 1)$, $r = \kappa/\nu \neq 1$, so $e^{-\kappa h} = x^r$ and $f(h) = 0$ iff

$$\phi(x) := \sqrt{A}(1 - x) - \sqrt{B}(1 - x^r) = 0,$$

where taking square roots is valid because both bracketed terms are positive for $x \in (0, 1)$. Extend ϕ continuously to $x = 1$, where $\phi(1) = 0$. Its derivative is

$$\phi'(x) = -\sqrt{A} + \sqrt{B}\, r\, x^{r-1}.$$

Because $r \neq 1$, the function x^{r-1} is strictly monotone on $(0, 1)$; hence ϕ' has at most one zero. If ϕ had two interior zeros $x_1 < x_2 < 1$, Rolle's theorem applied on $[x_1, x_2]$ and $[x_2, 1]$ would give two distinct zeros of ϕ', a contradiction. Thus ϕ, and therefore f, has at most one zero on the stated domain. □

Derivation of equation (14). Lemma 6 applies to each independent factor $L_{k,t}$; bilinearity and independence across k and from H_t eliminate all cross terms and give the sum in (14).

It remains to verify the two-crossing example. Denote its horizon curve by g and set $y = e^{-h/10} \in (0, 1)$. Then

$$g(h) = q(y) := \frac{7}{5}(1 - y^5)^2 - \frac{3}{20}(1 - y^{40})^2 - 2(1 - y)^2,$$

and expansion in descending powers gives

$$q(y) = -\frac{3}{20}y^{80} + \frac{3}{10}y^{40} + \frac{7}{5}y^{10} - \frac{14}{5}y^{5} - 2y^2 + 4y - \frac{3}{4}.$$

The nonzero coefficients have sign sequence $(-, +, +, -, -, +, -)$, which has four sign changes. Descartes' rule of signs therefore bounds the number of positive roots of q, counted with multiplicity, by four. Moreover, $q(1) = q'(1) = 0$, so the excluded point $y = 1$ already contributes a root of multiplicity at least two. Exact evaluation at three rational points gives

$$q(99/100) < 0, \qquad q(3/4) > 0, \qquad q(1/10) < 0,$$

with approximate values -0.0133, 0.5394, and -0.3700. Continuity therefore supplies two distinct roots in $(0, 1)$. Together with the double root at $y = 1$, these exhaust the four positive roots allowed by Descartes' rule. Hence the example has exactly two roots on $(0, \infty)$, numerically $h_1 \approx 0.7900$ and $h_2 \approx 15.6165$. $\square$

Proof of Lemma 8. Substituting $dP_t = dM_t + dL_t + p\, ds_t$ into (15) gives

$$(1 - \omega_0 p)\, ds_t = \omega_0(dM_t + dL_t) - \omega_0 s_t\, dt.$$

When $\omega_0 p < 1$, division by the positive coefficient yields (17) with $\omega = \omega_0/(1 - \omega_0 p) > 0$. At $\omega_0 p = 1$, the left side vanishes while the right side retains the Brownian innovations in $dM_t + dL_t$, so the identity cannot define a linear signal process. When $\omega_0 p > 1$, the implied ω is negative and the drift $-\omega s_t$ is explosive. Thus a stationary linear solution exists only on the branch $\omega_0 p < 1$. $\square$

Proof of Lemma 9. Write $s_t^M = \omega\nu \int_{-\infty}^{t} e^{-\omega(t-r)} H_r\, dr$ and $s_t^L = \omega \int_{-\infty}^{t} e^{-\omega(t-r)} dL_r$. Substituting the moving-average representations of H and L and applying Fubini gives kernel representations $s_t^M = \int_{-\infty}^{t} k_{s^M}(t-u)\, dW_u^N$ and $s_t^L = \int_{-\infty}^{t} k_{s^L}(t-u)\, dW_u^L$ with

$$k_{s^M}(d) = \frac{\sigma_N \omega \nu}{\omega - \nu}\left(e^{-\nu d} - e^{-\omega d}\right), \qquad k_{s^L}(d) = \frac{\sigma_L \omega}{\omega - \kappa}\left(\omega e^{-\omega d} - \kappa e^{-\kappa d}\right),$$

and $k_H(d) = \sigma_N e^{-\nu d}$, $k_L(d) = \sigma_L e^{-\kappa d}$. For any two kernels, Itô isometry gives $\mathrm{Cov}(X_t, Y_{t+d}) =$

$\int_0^\infty k_X(u)\, k_Y(u+d)\, du$ for $d \geq 0$. Evaluating these elementary exponential integrals for the six pairs yields (18)–(20); for instance $\mathrm{Cov}(s_t^M, H_{t+d}) = e^{-\nu d}\, \mathbb{E}[s_t^M H_t]$ follows from the Markov property of H, with $\mathbb{E}[s_t^M H_t] = \int_0^\infty k_{s^M}(u) k_H(u)\, du = \omega\sigma_N^2/(2(\nu+\omega))$. □

Proof of Theorem 3. The proof has three steps. First, decompose adjacent return covariance into the two independent Brownian blocks. Second, reduce every block to the squared-exponential basis $(1-e^{-rh})^2$. Third, collect the coefficients rate by rate.

Expand $P_t = M_t + p\, s_t^M + L_t + p\, s_t^L$. By independence of W^N and W^L, the adjacent covariance splits into a W^N-block and a W^L-block:

$$C(h) = C_{MM} + p\,(C_{Ms^M} + C_{s^M M}) + p^2 C_{s^M s^M} + C_{LL} + p\,(C_{Ls^L} + C_{s^L L}) + p^2 C_{s^L s^L},$$

where each term is $\mathrm{Cov}(\Delta_h X_t, \Delta_h Y_{t+h})$ for the indicated pair. Three identities reduce every term to the squared-exponential basis. (i) For jointly stationary X, Y with $\Gamma_{XY}(d) \equiv \mathrm{Cov}(X_t, Y_{t+d}) = \sum_j c_j e^{-r_j d}$ for $d \geq 0$,

$$\mathrm{Cov}(\Delta_h X_t, \Delta_h Y_{t+h}) = 2\Gamma_{XY}(h) - \Gamma_{XY}(0) - \Gamma_{XY}(2h) = -\sum_j c_j \left(1 - e^{-r_j h}\right)^2,$$

using lags $\{0, h, 2h\} \geq 0$ and $2e^{-rh} - 1 - e^{-2rh} = -(1-e^{-rh})^2$. (ii) For stationary Z driven by W^N with $\Gamma_{HZ}(d) = \sum_j a_j e^{-r_j d}$,

$$\mathrm{Cov}(\Delta_h M_t, \Delta_h Z_{t+h}) = \nu \int_0^h \left[\Gamma_{HZ}(2h - v) - \Gamma_{HZ}(h - v)\right] dv = -\nu \sum_j \frac{a_j}{r_j} \left(1 - e^{-r_j h}\right)^2.$$

(iii) With $\Gamma_{ZH}(d) = \sum_j b_j e^{-r_j d}$,

$$\mathrm{Cov}(\Delta_h Z_t, \Delta_h M_{t+h}) = \nu \int_0^h \left[\Gamma_{ZH}(v) - \Gamma_{ZH}(h + v)\right] dv = +\nu \sum_j \frac{b_j}{r_j} \left(1 - e^{-r_j h}\right)^2.$$

Identity (i) applies to every pair of stationary components, including all terms in the L-block. Identities (ii) and (iii) apply only to cross terms involving the integrated component M. The pure M-term is $C_{MM} = \frac{\sigma_N^2}{2\nu}(1-e^{-\nu h})^2$ by Lemma 5. To make the table below checkable, consider the L-block cross term. From Lemma 9,

$$\Gamma_{Ls^L}(d) = \frac{\omega\sigma_L^2\left[(\kappa+\omega)e^{-\kappa d} - 2\omega e^{-\omega d}\right]}{2(\kappa^2-\omega^2)}, \qquad \Gamma_{s^L L}(d) = \frac{\omega\sigma_L^2}{2(\kappa+\omega)}e^{-\kappa d}.$$

Identity (i) therefore gives a rate-ω coefficient $\omega^2\sigma_L^2/(\kappa^2-\omega^2)$ and a rate-κ coefficient

$$-\frac{\omega\sigma_L^2(\kappa+\omega)}{2(\kappa^2-\omega^2)} - \frac{\omega\sigma_L^2}{2(\kappa+\omega)} = -\frac{\kappa\omega\sigma_L^2}{\kappa^2-\omega^2}.$$

Applying the same identities to the remaining terms gives the following coefficients on $(1-e^{-rh})^2$:

source	rate ν	rate ω	rate κ
C_{MM}	$\sigma_N^2/(2\nu)$	–	–
$p(C_{Ms^M} + C_{s^M M})$	$p\,\dfrac{\nu\omega\sigma_N^2}{\nu^2-\omega^2}$	$-p\,\dfrac{\nu^2\sigma_N^2}{\nu^2-\omega^2}$	–
$p^2 C_{s^M s^M}$	$p^2\,\dfrac{\nu\omega^2\sigma_N^2}{2(\nu^2-\omega^2)}$	$-p^2\,\dfrac{\nu^2\omega\sigma_N^2}{2(\nu^2-\omega^2)}$	–
C_{LL}	–	–	$-\sigma_L^2/(2\kappa)$
$p(C_{Ls^L} + C_{s^L L})$	–	$+p\,\dfrac{\omega^2\sigma_L^2}{\kappa^2-\omega^2}$	$-p\,\dfrac{\kappa\omega\sigma_L^2}{\kappa^2-\omega^2}$
$p^2 C_{s^L s^L}$	–	$+p^2\,\dfrac{\omega^3\sigma_L^2}{2(\kappa^2-\omega^2)}$	$-p^2\,\dfrac{\kappa\omega^2\sigma_L^2}{2(\kappa^2-\omega^2)}$

Summing each column and completing the square in $(1+\omega p)$ gives (22)–(24); for example the rate-ν column is $\sigma_N^2\left[(\nu^2-\omega^2) + 2p\nu^2\omega + p^2\nu^2\omega^2\right]/\left(2\nu(\nu^2-\omega^2)\right) = \sigma_N^2\left[\nu^2(1+\omega p)^2 - \omega^2\right]/\left(2\nu(\nu^2-\omega^2)\right)$. □

Proof of Corollary 4. Each basis function tends to one, so the plateau is $W_\nu + W_\omega + W_\kappa$. Collecting powers of p: the p^0 terms give $\sigma_N^2/(2\nu) - \sigma_L^2/(2\kappa)$; the p^1 terms give $\frac{\nu\omega\sigma_N^2 - \nu^2\sigma_N^2}{\nu^2-\omega^2} + \frac{\omega^2\sigma_L^2 - \kappa\omega\sigma_L^2}{\kappa^2-\omega^2} = -\frac{\nu\sigma_N^2}{\nu+\omega} - \frac{\omega\sigma_L^2}{\kappa+\omega}$; the p^2 terms give $-\frac{\nu\omega\sigma_N^2}{2(\nu+\omega)} - \frac{\omega^2\sigma_L^2}{2(\kappa+\omega)}$. Both p-coefficients are strictly negative. □

Proof of Remark 3. Since $(1-e^{-rh})^2 = r^2h^2 + O(h^3)$, the p-linear part of $C(h)$ has h^2-coefficient $\frac{\nu\omega\sigma_N^2}{\nu^2-\omega^2}\nu^2 - \frac{\nu^2\sigma_N^2}{\nu^2-\omega^2}\omega^2 + \frac{\omega^2\sigma_L^2}{\kappa^2-\omega^2}\omega^2 - \frac{\kappa\omega\sigma_L^2}{\kappa^2-\omega^2}\kappa^2 = \frac{\nu^2\omega\sigma_N^2}{\nu+\omega} - \frac{\omega\sigma_L^2(\kappa^2+\kappa\omega+\omega^2)}{\kappa+\omega}$, using $\kappa^3-\omega^3 = (\kappa-\omega)(\kappa^2+\kappa\omega+\omega^2)$. □

Proof of Proposition 5. The sign changes can be counted exactly. For the stated parameters, set $y = e^{-h/50} \in (0, 1)$. Because $(\nu, \omega, \kappa) = (25, 6, 200)/50$, equation (21) becomes

$$C(h) = q(y) := W_\nu(1 - y^{25})^2 + W_\omega(1 - y^6)^2 + W_\kappa(1 - y^{200})^2.$$

The parameterization gives $W_\nu > 0$, $W_\omega < 0$, $W_\kappa < 0$, and $W_\nu + W_\omega + W_\kappa < 0$. Consequently, the coefficients of q at powers $(400, 200, 50, 25, 12, 6, 0)$ have signs

$$(-, +, +, -, -, +, -),$$

respectively. There are four sign changes, so Descartes' rule bounds the number of positive roots by four, counted with multiplicity. Since every term in q is a squared difference, $q(1) = q'(1) = 0$; the excluded point $y = 1$ is a root of multiplicity at least two. Exact evaluation at rational points gives

$$q(199/200) < 0, \qquad q(19/20) > 0, \qquad q(1/2) < 0,$$

with approximate values -0.0359, 0.4650, and -0.1267. Continuity supplies two distinct roots in $(0, 1)$, and these roots together with the double root at $y = 1$ exhaust the Descartes bound. Thus $C(h)$ has exactly two roots on $(0, \infty)$, numerically $h_1 \approx 0.7498$ and $h_2 \approx 22.9057$, with the stated negative–positive–negative sign pattern.

At $p = 0$ the curve is $\frac{\sigma_N^2}{2\nu}(1 - e^{-\nu h})^2 - \frac{\sigma_L^2}{2\kappa}(1 - e^{-\kappa h})^2$ with both coefficients positive, so Proposition 4 applies: at most one crossing, and since the short-horizon coefficient $\frac{1}{2}(\sigma_N^2 \nu - \sigma_L^2 \kappa) = -1.37$ is negative while the plateau 0.899 is positive, continuity gives exactly one. $\square$

Proof of Proposition 9. Profit decomposition. Taking expectations in (29) and using stationarity gives

$$\pi(\omega) = \frac{\mathbb{E}[s_t\, dM_t]}{dt} + \frac{\mathbb{E}[s_t\, dL_t]}{dt} + \frac{p}{2}\frac{d\mathbb{E}[s_t^2]}{dt} = \nu\mathbb{E}[s_t H_t] - \kappa\mathbb{E}[s_t L_t].$$

The stochastic integral in $s_t\, dL_t$ has zero expectation because s_t is previsible; orthogonality gives $\mathbb{E}[s_t H_t] = \mathbb{E}[s_t^M H_t]$ and $\mathbb{E}[s_t L_t] = \mathbb{E}[s_t^L L_t]$. Substituting the two date-zero covariances from Lemma 9 yields (30).

Signs of profit and its derivative. For $\omega > 0$, multiply (30) by a positive denominator and define

$$Q(\omega) := \frac{2(\omega+\nu)(\omega+\kappa)}{\omega}\,\pi(\omega)$$
$$= \omega(\nu\sigma_N^2 - \kappa\sigma_L^2) + \nu\kappa(\sigma_N^2 - \sigma_L^2).$$

Thus $\operatorname{sgn}\pi(\omega) = \operatorname{sgn} Q(\omega)$. For the slope,

$$2\pi'(\omega) = \left(\frac{\nu\sigma_N}{\omega+\nu}\right)^2 - \left(\frac{\kappa\sigma_L}{\omega+\kappa}\right)^2.$$

Both ratios are positive, so the sign of π' is the sign of

$$D(\omega) := \nu\sigma_N(\omega+\kappa) - \kappa\sigma_L(\omega+\nu) = (\nu\sigma_N - \kappa\sigma_L)\omega + \nu\kappa(\sigma_N - \sigma_L).$$

Parameter regimes. (1) If $\sigma_N^2 \le \sigma_L^2$ and $\nu\sigma_N^2 \le \kappa\sigma_L^2$, both the intercept and slope of Q are nonpositive. Hence $\pi(\omega) \le 0$ for every $\omega > 0$.

(2) Suppose $\sigma_L^2 < \sigma_N^2 < (\kappa/\nu)\sigma_L^2$. The intercept of Q is positive and its slope is negative, so Q has the unique positive root

$$\bar{\omega} = \frac{\nu\kappa(\sigma_N^2 - \sigma_L^2)}{\kappa\sigma_L^2 - \nu\sigma_N^2}.$$

It follows that $\pi > 0$ on $(0, \bar{\omega})$ and $\pi < 0$ beyond it. These inequalities also imply $\kappa > \nu$ and $1 < \sigma_N/\sigma_L < \sqrt{\kappa/\nu} < \kappa/\nu$. Therefore $D(0) > 0$ and the slope of D is negative. The derivative π' changes sign exactly once: profit rises to a unique maximum and then falls. For $F \in (0, \max \pi)$, the larger solution of $\pi(\omega(g)) = F$ lies on this decreasing branch. Lemma 8 makes $\omega(g)$ strictly increasing, so that solution is stable under entry and exit.

(3) If $\nu\sigma_N^2 > \kappa\sigma_L^2$ and $\sigma_N^2 > \sigma_L^2$, both coefficients of Q are positive. Thus $\pi > 0$ for every $\omega > 0$, and

$$\lim_{\omega\to\infty} \pi(\omega) = \frac{\nu\sigma_N^2 - \kappa\sigma_L^2}{2} > 0.$$

If F lies below both this limit and $\inf_{\omega \ge \omega_0} \pi(\omega)$, no finite mass eliminates the entry rent. Finally, if $\nu\sigma_N \ge \kappa\sigma_L$, the slope of D is nonnegative while its intercept is strictly positive. Hence $\pi'(\omega) > 0$ for every $\omega > 0$; since $\omega(g)$ rises with mass, entry is a strategic complementarity. $\square$

Proof of Proposition 6. Step 1: recover the reduced-form curve. For a rate r,

$$(1-e^{-rh})^2 = 1 - 2e^{-rh} + e^{-2rh}.$$

Thus $C(h)$ is a finite linear combination of $\{1\} \cup \{e^{-\lambda h}\}$, with $\lambda \in \{\nu, 2\nu, \omega, 2\omega, \kappa, 2\kappa\}$. The assumptions that rates are distinct and no rate is twice another make these exponents distinct. Distinct exponentials are linearly independent on any open interval, so the exponent set and the associated coefficients are unique. For each recovered rate r_j, the coefficient on $e^{-r_j h}$ is $-2W_j$, while the coefficient on $e^{-2r_j h}$ is W_j. The paired coefficients distinguish the base rate from its double and recover the three weights. In the reversal-momentum-reversal regime, the positive weight labels ν, and the two negative weights are ordered as $\omega < \kappa$.

Step 2: reduce primitive recovery to one scalar equation. Let

$$x = (1+\omega p)^2, \qquad p(2+\omega p) = \frac{x-1}{\omega}.$$

Equation (22) gives

$$W_\nu = \frac{\sigma_N^2}{2\nu(\nu^2-\omega^2)}\big(\nu^2 x - \omega^2\big),$$

so

$$\frac{\sigma_N^2}{\nu^2-\omega^2} = \frac{2\nu W_\nu}{\nu^2 x - \omega^2}.$$

Similarly, (24) gives

$$W_\kappa = -\frac{\sigma_L^2}{2\kappa(\kappa^2-\omega^2)}\big(\kappa^2 x - \omega^2\big),$$

and therefore

$$\frac{\sigma_L^2}{\kappa^2-\omega^2} = -\frac{2\kappa W_\kappa}{\kappa^2 x - \omega^2}.$$

Substitute these two displays into (23):

$$\begin{aligned} W_\omega &= -\frac{1}{2}\,p(2+\omega p)\left[\nu^2\frac{\sigma_N^2}{\nu^2-\omega^2} - \omega^2\frac{\sigma_L^2}{\kappa^2-\omega^2}\right] \\ &= -\frac{x-1}{2\omega}\left[\frac{2\nu^3 W_\nu}{\nu^2 x-\omega^2} + \frac{2\kappa\omega^2 W_\kappa}{\kappa^2 x - \omega^2}\right]. \end{aligned}$$

Moving the right side to the left gives exactly $\Phi(x) = 0$ as defined in (25). Once its root x^* is known, (26) follows by the same two displayed equations and $p = (\sqrt{x^*} - 1)/\omega$.

Step 3: show uniqueness of the scalar root. Throughout this step we use the regime ordering $\omega < \nu < \kappa$, $W_\nu > 0$, $W_\kappa < 0$, and $p \geq 0$. The last inequality implies $x = (1+\omega p)^2 \geq 1$. Differentiating Φ gives

$$\Phi'(x) = \frac{1}{\omega}\left[\frac{\nu^3 W_\nu(\nu^2 - \omega^2)}{(\nu^2 x - \omega^2)^2} + \frac{\kappa\omega^2 W_\kappa(\kappa^2 - \omega^2)}{(\kappa^2 x - \omega^2)^2}\right].$$

Because $W_\kappa < 0$, the second term is the only possible source of non-monotonicity. For $x \geq 1$,

$$\frac{\kappa^2 x - \omega^2}{\nu^2 x - \omega^2} > \frac{\kappa^2}{\nu^2},$$

since cross-multiplication leaves $\nu^2(\kappa^2 x - \omega^2) - \kappa^2(\nu^2 x - \omega^2) = \omega^2(\kappa^2 - \nu^2) > 0$. Hence

$$\frac{1}{(\kappa^2 x - \omega^2)^2} < \frac{\nu^4}{\kappa^4}\,\frac{1}{(\nu^2 x - \omega^2)^2}.$$

Using this bound on the negative term in $\Phi'(x)$,

$$\Phi'(x) > \frac{\nu^3(\nu^2 - \omega^2)}{\omega(\nu^2 x - \omega^2)^2}\left[W_\nu - |W_\kappa|\,\frac{\nu\omega^2(\kappa^2 - \omega^2)}{\kappa^3(\nu^2 - \omega^2)}\right],$$

The bracket on the right is nonnegative when

$$\frac{W_\nu}{|W_\kappa|} \geq \frac{\nu\omega^2(\kappa^2 - \omega^2)}{\kappa^3(\nu^2 - \omega^2)}.$$

The preceding strict inequality therefore gives $\Phi'(x) > 0$, including when the restriction holds with equality. Thus Φ has at most one root on $[1, \infty)$. The true primitive vector generated the observed curve, so its $x = (1 + \omega p)^2$ is a root in that interval. It is therefore the unique root.

Step 4: identify the cost product. In the entry region, Proposition 2 gives

$$\nu = \frac{\sigma_N^2}{8 c_A \psi_A}.$$

The objects ν and σ_N^2 have already been recovered, so

$$c_A \psi_A = \frac{\sigma_N^2}{8\nu}$$

Table AI. Monte Carlo recovery of the three-clock model. Prices are simulated from the exact discretization of the model of Section V at intervals of 0.05 months (approximately daily) over 480 months, with heterogeneous unobserved drifts; the estimator uses only pooled sample autocovariances of drift-adjusted returns at 18 horizons between 0.25 and 48 months. Entries are medians over 204 replications with 5th and 95th percentiles in brackets. The final row is the percentage of replications in which the fitted curve changes sign exactly twice.

	True	$N = 1$	$N = 100$	$N = 500$
ν (interpretation)	0.50	0.29 [0.05, 0.95]	0.49 [0.37, 0.56]	0.50 [0.45, 0.54]
κ (liquidity)	4.00	4.56 [0.65, 25.0]	4.01 [3.85, 4.18]	4.00 [3.93, 4.08]
ω (trend)	0.12	0.12 [0.01, 0.66]	0.11 [0.01, 0.36]	0.11 [0.06, 0.21]
W_ν	1.33	9.64 [0.89, 443]	1.37 [1.07, 7.47]	1.33 [1.16, 1.87]
W_κ	-0.13	-0.10 [-47.0, -0.00]	-0.13 [-0.14, -0.12]	-0.13 [-0.14, -0.13]
W_ω	-1.36	-28.0 [-4667, -0.00]	-2.73 [-85.2, -1.00]	-1.53 [-3.62, -1.00]
h_1 (first crossing)	0.75	0.77 [0.54, 1.06]	0.75 [0.74, 0.77]	0.75 [0.74, 0.76]
h_2 (second crossing)	22.9	7.0 [2.9, 14.0]	14.7 [8.6, 29.9]	18.7 [14.2, 26.9]
σ_N^2 (information)	1.00	0.96 [0.10, 62.2]	0.98 [0.71, 1.10]	0.99 [0.88, 1.05]
σ_L^2 (bounce)	0.81	0.35 [0.03, 41.5]	0.72 [0.49, 0.89]	0.79 [0.67, 0.86]
p (trend pressure)	1.20	3.99 [0.18, 50.4]	1.50 [0.73, 7.25]	1.31 [0.83, 2.93]
$c_A\psi_A$ (cost product)	0.25	0.39 [0.05, 27.2]	0.25 [0.17, 0.34]	0.25 [0.21, 0.28]
two crossings detected (%)	–	60	63	68

is identified from the return-autocovariance curve. □

Appendix C. A Monte Carlo Assessment of the Estimator

The simulation asks which part of the three-clock structure a realistic price history can recover. It uses the parameterization of Proposition 5, expressed in monthly units: $\nu = 0.5$, $\kappa = 4$, and $\omega = 0.12$, with crossings near two weeks and two years and a momentum peak near six months. Prices are drawn from the model's exact discretization at 0.05-month intervals, approximately daily, for 480 months, with heterogeneous unobserved drifts.

Table AI compares a single asset with pooled panels of 100 and 500 independent assets. Each design has 204 replications. Every replication uses the full Section V procedure: drift adjustment, a pooled curve with cross-sectional standard errors, grid and simplex searches over rates, sign-constrained weights, the h^2 bias regressor, and the structural back-out in Proposition 6.

The recovery follows the clocks' speed. A single asset's forty years of daily prices identify

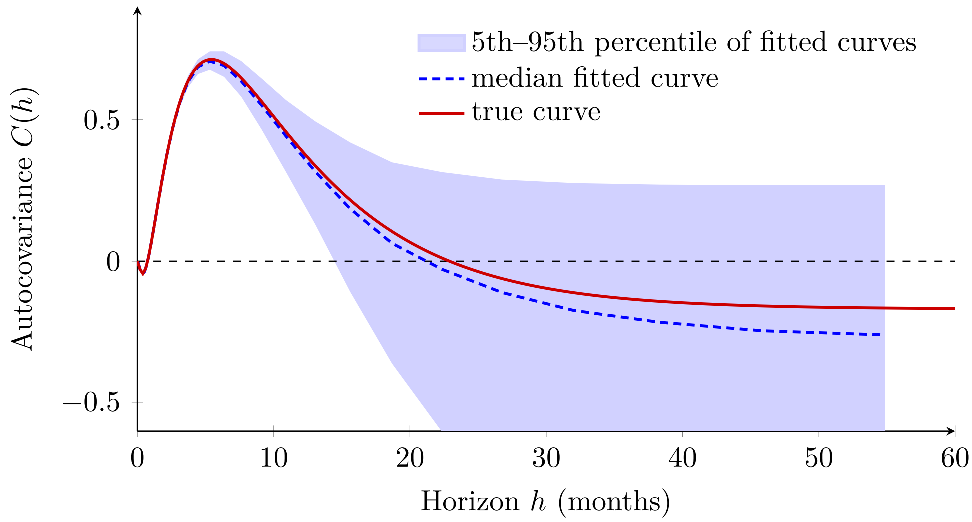

Figure A3. Monte Carlo distribution of fitted horizon curves. The figure shows the point-wise 5th–95th percentile band and median of the fitted three-clock curves across 204 replications of the $N = 500$ panel design of Table AI, against the true curve of Proposition 5. The fast reversal, the momentum hump, and the first crossing are recovered with little error; the slow clock's descent into long-run reversal is where sampling uncertainty concentrates.

the fast clock and first crossing, but little else. Pooling changes the picture: with 100 assets, the interpretation and liquidity rates, their weights, both variance intensities, and the cost product $c_A\psi_A$ are recovered tightly; with 500 assets, their medians match the truth to two digits. The cost product is among the best measured objects, with a 90% range of $[0.21, 0.28]$ around a truth of 0.25.

The slow clock remains the exception. Even in the 500-asset design, ω, W_ω, and the second crossing h_2 have wide bands, and only 68% of replications detect both crossings. Figure A3 shows where that uncertainty accumulates. The simulation supports the paper's empirical hierarchy: fast and medium components can be estimated, whereas the second crossing remains the hardest feature to estimate.

Independence makes the design favorable because common factors reduce the effective number of assets in actual panels. The estimator uses a simple equally spaced grid, without efficient cross-asset weighting or overidentification tests. The reported bands quantify information per effective independent asset; corresponding intervals in the Indian panel will generally be wider.

Appendix D. Boundary Diagnostics

In this panel, the fitted rates cannot be separated reliably because the estimates lie on a constraint boundary. The diagnostics explain why the main text omits individual structural rates, weights, and the curve-implied cost product.

The three-term fit uses 25,800 normalized instrument–horizon moments for 1,613 listings. Numerical stability requires a minimum ratio between adjacent fitted rates. In the full sample, at least one rate ratio lands exactly on that boundary; the same is true in every one of the 60 converged company-cluster bootstrap replications.

At a constraint boundary, the estimate responds to the constraint rather than to interior curvature of the objective. The fit can still reveal the number of components and the approximate locations of sign changes, but it is locally flat in the direction that would separate neighboring rates. A bootstrap pinned to the same boundary then understates uncertainty. The rates and weights are consequently boundary-contaminated, as is any object constructed from them, including the implied cost product $c_A\psi_A = \sigma_N^2/(8\nu)$ in Proposition 6(iii). Their full distributions are reported only in the Internet Appendix, with this qualification attached.

A short-grid diagnostic reaches the same conclusion from the opposite direction. Restricting the fit to the densely measured fast bands still does not separate the two faster rates; the objective again stops at the rate-ratio boundary. This is the sample analogue of Appendix C's Monte Carlo lesson: one well-measured region cannot time two neighboring clocks.

The model-free band signs remain intact. The representation diagnoses sign changes, and Table IV uses cross-sectional band moments insulated from boundary-sensitive rate estimates. The evidence identifies the component count more reliably than the individual speeds.

The Internet Appendix reports full bootstrap distributions for boundary rates and weights, Medhat–Schmeling double sorts, the short-grid separation failure, further subpe-

Table AII. Subperiod robustness and venue-gap speed. The table reports fitted sign changes and second-crossing locations by subperiod, together with company-cluster bootstrap evidence and same-firm BSE–NSE venue-gap persistence. The 2002–2008 window is sparse and has no usable bootstrap fit. The crossing locations are fitted-representation statistics; the venue-gap half-life is computed from the pooled daily AR(1) coefficient of the same-firm log-price gap.

Period	*Sign changes*	*Second crossing (months)*	*Bootstrap two-sign share*	*Gap half-life (days)*	*Median gap ρ*
2002–2008	0	–	–	0.41	0.070
2009–2016	2	5.35	90.0%	2.81	0.096
2017–2024	2	14.27	93.3%	1.36	0.069
2002–2024	2	13.13	93.3%	2.40	0.078

riod detail, venue-gap checks, and construction code.

References

Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Avramov, Doron, Si Cheng, and Allaudeen Hameed, 2016, Time-varying liquidity and momentum profits, *Journal of Financial and Quantitative Analysis* 51, 1897–1923.

Back, Kerry, 1992, Insider trading in continuous time, *Review of Financial Studies* 5, 387–409.

Back, Kerry, and Shmuel Baruch, 2004, Information in securities markets: Kyle meets Glosten and Milgrom, *Econometrica* 72, 433–465.

Bogousslavsky, Vincent, 2016, Infrequent rebalancing, return autocorrelation, and seasonality, *Journal of Finance* 71, 2967–3006.

Campbell, John Y., Sanford J. Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *Quarterly Journal of Economics* 108, 905–939.

Cespa, Giovanni, and Xavier Vives, 2012, Dynamic trading and asset prices: Keynes vs. Hayek, *Review of Economic Studies* 79, 539–580.

Chordia, Tarun, and Avanidhar Subrahmanyam, 2004, Order imbalance and individual stock returns: Theory and evidence, *Journal of Financial Economics* 72, 485–518.

Chui, Andy C. W., Avanidhar Subrahmanyam, and Sheridan Titman, 2022, Momentum, reversals, and investor clientele, *Review of Finance* 26, 217–255.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1885.

Daniel, Kent, and Tobias J. Moskowitz, 2016, Momentum crashes, *Journal of Financial Economics* 122, 221–247.

Dai, Wei, Mamdouh Medhat, Robert Novy-Marx, and Savina Rizova, 2024, Reversals and the returns to liquidity provision, *Financial Analysts Journal* 80, 122–151.

De Bondt, Werner F. M., and Richard Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379–395.

Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *Journal of Finance* 65, 1237–1267.

Glosten, Lawrence R., and Lawrence E. Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 21, 123–142.

Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.

Hameed, Allaudeen, Wenjin Kang, and S. Viswanathan, 2010, Stock market declines and liquidity, *Journal of Finance* 65, 257–293.

Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.

Hellwig, Martin F., 1980, On the aggregation of information in competitive markets, *Journal of Economic Theory* 22, 477–498.

Hendershott, Terrence, Albert J. Menkveld, Rémy Praz, and Mark Seasholes, 2022, Asset price dynamics with limited attention, *Review of Financial Studies* 35, 962–1008.

Holden, Craig W., and Avanidhar Subrahmanyam, 2002, News events, information acquisition, and serial correlation, *Journal of Business* 75, 1–32.

Hong, Harrison, Terence Lim, and Jeremy C. Stein, 2000, Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies, *Journal of Finance* 55, 265–295.

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.

Huang, Roger D., and Hans R. Stoll, 1997, The components of the bid-ask spread: A general approach, *Review of Financial Studies* 10, 995–1034.

Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *Journal of Finance* 45, 881–898.

Jegadeesh, Narasimhan, Jiang Luo, Avanidhar Subrahmanyam, and Sheridan Titman, 2025, Short-term reversals and longer-term momentum around the world: Theory and evidence, *Review of Financial Studies* 38, 3673–3728.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Kelly, Bryan T., Tobias J. Moskowitz, and Seth Pruitt, 2021, Understanding momentum and reversal, *Journal of Financial Economics* 140, 726–743.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Lehmann, Bruce N., 1990, Fads, martingales, and market efficiency, *Quarterly Journal of Economics* 105, 1–28.

Lee, Charles M. C., and Bhaskaran Swaminathan, 2000, Price momentum and trading volume, *Journal of Finance* 55, 2017–2069.

Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang, 2002, Dynamic volume-return relation of individual stocks, *Review of Financial Studies* 15, 1005–1047.

Lo, Andrew W., and A. Craig MacKinlay, 1990, When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175–205.

Luo, Jiang, Avanidhar Subrahmanyam, and Sheridan Titman, 2021, Momentum and reversals when overconfident investors underestimate their competition, *Review of Financial Studies* 34, 351–393.

Madhavan, Ananth, Matthew Richardson, and Mark Roomans, 1997, Why do security prices change? A transaction-level analysis of NYSE stocks, *Review of Financial Studies* 10, 1035–1064.

Mankiw, N. Gregory, and Ricardo Reis, 2002, Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics* 117, 1295–1328.

Medhat, Mamdouh, and Maik Schmeling, 2022, Short-term momentum, *Review of Financial Studies* 35, 1480–1526.

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

Obizhaeva, Anna A., and Jiang Wang, 2013, Optimal trading strategy and supply/demand dynamics, *Journal of Financial Markets* 16, 1–32.

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1127–1139.

Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309–349.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Verrecchia, Robert E., 1982, Information acquisition in a noisy rational expectations economy, *Econometrica* 50, 1415–1430.

Vayanos, Dimitri, and Paul Woolley, 2013, An institutional theory of momentum and reversal, *Review of Financial Studies* 26, 1087–1145.

Vives, Xavier, 2008, *Information and Learning in Markets: The Impact of Market Microstructure* (Princeton University Press, Princeton, NJ).