Liquidity-Based Audit of AI and Algorithmic Trading Strategies

Irene Aldridge

Abstract

We show that an algorithmic strategy's net demand for liquidity is identifiable from its trade and price history alone, with no knowledge of its signal or optimization problem. An exact multi-period regret decomposition implies that the sign of this statistic classifies a linear strategy as a net liquidity consumer or provider, recovering the Kyle (1985) informed-trader/market-maker dichotomy from observables alone. Under an AR(1) cost process, the same statistic equals αs^2, the product of strategy size and the squared Roll (1984) implied spread, making the correction a direct proxy for prevailing illiquidity. Extending to endogenous price impact and aggregating across N correlated strategies yields a liquidity-balance condition whose violation produces welfare loss scaling as N^2, a closed-form fire-sale externality. We calibrate to CRSP equity data (2016–2025), tracking implied spreads through the COVID-19 and 2022 rate-shock episodes, with an estimator computable in $O(T \cdot nd)$ time.

1 Introduction

Market microstructure has long classified trading activity by its informational role: an informed trader demands liquidity by trading in the direction of private information, while a market maker supplies liquidity by absorbing that order flow and earning the spread in compensation Kyle (1985); Glosten and Milgrom (1985). This classification is typically recovered from the data that the classifier requires: signed order flow, quote revisions, or the sequential-trade structure of the market. The classification is harder to apply to an algorithmic strategy whose internal logic is unobservable. However, the signals or optimization problems generating the decisions of a typical quantitative fund are not visible, even though the trades and reported positions may be available.

This paper shows that the liquidity role of such a strategy (consumer or provider) can be recovered from realized portfolio costs and trade decisions alone, without observing quotes, order flow, or any other microstructure-specific signal. The key observation is that a strategy's contribution to multi-period performance decomposes exactly into a sum of per-period covariances between the realized cost vector and the strategy's decision (Theorem 3.6). Furthermore, the *sign* of this covariance statistic is a model-free liquidity classifier: positive covariance means the strategy systematically buys when costs are rising and sells when costs are falling, while negative covariance means the reverse. Buying when asset costs (prices) are rising and selling when prices are falling is the behavior of a successful quantitative fund. The reverse is the defining behavior of a liquidity supplier (Proposition 2.2).

This single observation has three further consequences that are the substance of the paper. First, under an AR(1) cost process (the standard reduced-form model for short-horizon return reversal), the same covariance statistic admits a closed form in terms of the Roll (1984) implied bid-ask spread. The auto-covariance correction to the covariance sum equals αs^2, the product of strategy size and the squared effective spread (Theorem 7.7). The resulting spread estimator is a byproduct

of the same trajectory data used for the liquidity classification and is obtained at no additional computational cost.

Second, allowing the strategy's own trades to move prices shows that covariance regret and price-impact regret are additively separable (Theorem 7.11). This is the endogenous price impact in the spirit of Kyle (1985) and Almgren and Chriss (2001). The additive separation allows the two channels to be estimated independently from the same data.

Third, aggregating across N strategies under a common-factor cost structure yields a market-wide liquidity-balance condition: when liquidity demand and supply fail to net to zero across agents. Prices must adjust to clear, and the resulting welfare loss scales as N^2 in the strategy-correlation coefficient (Corollary 7.20). The result is a closed-form statement of the familiar fire-sale mechanism in which correlated deleveraging overwhelms available liquidity.

We calibrate the framework to CRSP daily equity data from 2016 through 2025 ($n = 21{,}183{,}373$ stock-days). The implied Roll spread recovered from the trajectory estimator tracks realized illiquidity through two distinct episodes: it roughly triples during the COVID-19 liquidity crisis of Q2 2020 (from 0.90% to 2.52% daily), and it collapses to its sample minimum during the 2022 rate-shock episode when persistent common-factor repricing rather than microstructure friction dominates short-horizon return dynamics. Both patterns are consistent with the spread acting as a real-time proxy for the depth of the market rather than an estimation artifact.

Related work. The liquidity-classification result connects most directly to the sequential-trade and strategic-trading literature. Kyle (1985) and Glosten and Milgrom (1985) establish the informed-trader/market-maker dichotomy that this paper recovers from trajectory data rather than order-flow data. Roll (1984) provides the implied-spread estimator that this paper re-derives as a regret correction. Almgren and Chriss (2001) and Gatheral (2010) provide the price-impact and market-impact models extended in Section 7.3. Hasbrouck (1991) and the broader VAR-based price-impact literature offer an alternative, order-flow-based route to the same quantities that this paper recovers from realized costs and decisions alone, a comparison we return to in Section 7.3. The covariance decomposition itself extends the single-period identity of Aldridge (2026) to the full multi-period stochastic dynamic programming setting and is structurally related to the Elmachtoub and Grigas (2022) predict-then-optimize framework and to the regret criterion in online convex optimization (Zinkevich, 2003; Hazan, 2016). The N-agent fire-sale extension connects to the market-wide liquidity-spiral literature (Shleifer and Vishny, 1992) and to random matrix theory (Bai and Silverstein, 2010a) for the systemic-fragility threshold of Section 7.6.

Roadmap. Section 2 states the liquidity classification result in its simplest form, immediately after the minimal notation required to state it. Section 3 develops the full multi-period stochastic dynamic programming setup and proves the exact covariance decomposition (Theorem 3.6) that underlies everything that follows. Section 4 gives a Bellman recursion for the covariance regret functional. Section 5 characterizes the policy bias that arises for strategies (e.g., momentum) that violate the mean-unbiasedness condition. Section 7.2 derives the closed-form auto-covariance correction under an AR(1) cost process. Section 7 returns to liquidity classification in full: the Roll-spread identity, endogenous price impact, the aggregate liquidity-balance condition and fire-sale extension, and the empirical calibration of implied spreads against the CRSP sample. Sections that follow develop the trajectory estimator's asymptotic properties, additional empirical results, and an N-agent systemic extension.

2 A First Look at Liquidity Classification

Before developing the full multi-period framework, we state the paper's central classification result in its simplest form, since it requires almost none of the machinery that follows.

Consider a single trading strategy that, at each period t, faces a cost vector $c_t \in \mathbb{R}^d$ (e.g., a vector of asset return innovations or factor realizations) and chooses a position $\hat{\pi}_t(c_t) \in \mathbb{R}^n$ as a linear function of that cost: $\hat{\pi}_t(c_t) = Bc_t + b$ for some matrix $B \in \mathbb{R}^{n\times d}$. Let $\Sigma_c = \mathrm{Cov}(c_t, c_t)$ be the cost covariance matrix. No further structure is needed for the result below; the full stochastic dynamic programming setup, including the multi-period regret identity that gives this statistic its performance interpretation, is developed starting in Section 3.

Definition 2.1 (Liquidity consumer and provider)**.** *The strategy is a* net liquidity consumer *at period t if* $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) > 0$*, a* net liquidity provider *if* $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) < 0$*, and* liquidity-neutral *if the covariance is zero.*

Proposition 2.2 (Liquidity role from trajectory data alone)**.** $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) = \mathrm{tr}(B\Sigma_c)$*. In particular: a momentum strategy ($B = \alpha I$, $\alpha > 0$) is always a liquidity consumer; a contrarian strategy ($B = -\alpha I$, $\alpha > 0$) is always a liquidity provider; and a strategy whose expected output matches the unconditional cost-minimizing allocation (e.g., minimum-variance) is liquidity-neutral.*

The proof follows directly from the bilinearity of covariance and the cyclic property of trace. The proof appears in full as Proposition 2.2 in Section 7.1, along with the connection to Kyle (1985) and the closed-form Roll-spread identity that follows from the same statistic under an AR(1) cost process. The point to take from this first look is that *the sign of a single, easily computed statistic classifies the strategy's liquidity role exactly.* This requires only the observed trajectory $\{(c_t, \hat{\pi}_t(c_t))\}_{t=1}^T$ and no knowledge of B, the strategy's signal, or its optimization problem.

Sections 3 through 6 develop the multi-period regret framework in which this statistic also functions as a performance audit metric. Section 7 returns to the liquidity interpretation with the full set of results, including the Roll-spread identity, price impact, and the N-agent aggregate liquidity-balance condition previewed above.

3 Setup: Multi-Period Stochastic Dynamic Program

Model

At each period $t \in \{1, \ldots, T\}$ a state $s_t \in \mathcal{S} \subseteq \mathbb{R}^p$ summarizes relevant history, a cost vector $c_t \in \mathcal{C} \subseteq \mathbb{R}^d$ is drawn from a distribution depending on s_t, and a decision $z_t \in \mathcal{Z} \subseteq \mathbb{R}^n$ is chosen by a policy $\hat{\pi}_t(s_t, c_t)$ after observing (s_t, c_t), realizing instantaneous cost $c_t^\top z_t$. The state evolves as $s_{t+1} = f(s_t, z_t, \varepsilon_{t+1})$.

Definition 3.1 (Policy)**.** *A* deterministic Markov policy *is a sequence* $\Pi = (\hat{\pi}_1, \ldots, \hat{\pi}_T)$ *of measurable maps* $\hat{\pi}_t \colon \mathcal{S} \times \mathcal{C} \to \mathcal{Z}$*. A policy is* stationary *if* $\hat{\pi}_t = \hat{\pi}$ *for all t.*

Let $\mathcal{F}_t = \sigma(s_1, c_1, \ldots, s_t, c_t)$ be the natural filtration and $\mathbb{E}_t[\,\cdot\,] = \mathbb{E}[\,\cdot \mid \mathcal{F}_t]$. Write $\mu_t = \mathbb{E}[c_t \mid s_t]$ and $\Sigma_t = \mathrm{Cov}(c_t, c_t \mid s_t)$.

Assumption 3.2 (Cost process)**.** *We work under one of the following:*

(a) ***i.i.d.*** *$\{c_t\}$ i.i.d. with mean $\bar{c}$ and covariance Σ_c; state trivial ($\mathcal{S} = \{s_0\}$).*

(b) ***Markov.*** *(s_t, c_t) is a time-homogeneous Markov chain: $(s_{t+1}, c_{t+1}) \perp \mathcal{F}_{t-1} \mid (s_t, c_t)$.*

(c) ***Non-stationary.*** *$\{c_t - \bar{c}_t\}$ is a martingale-difference sequence; cross-period covariances satisfy $\mathrm{Cov}(c_t, c_s) = \Gamma_{|t-s|}$ for some absolutely summable sequence $\{\Gamma_k\}$.*

Definition 3.3 (Multi-period regret)**.** *The perfect-foresight benchmark is $\pi_t^*(s_t) \in \arg\min_{z \in \mathcal{Z}} \mu_t^\top z$. The* total regret *of Π over horizon T is*

$$\mathrm{Regret}^{(T)}(\Pi) = \sum_{t=1}^{T} \Big(\mathbb{E}\big[c_t^\top \hat{\pi}_t(s_t, c_t)\big] - \mathbb{E}\big[\mu_t^\top \pi_t^*(s_t)\big] \Big). \tag{1}$$

Remark 3.4 (Benchmark)**.** *The benchmark $\mu_t^\top \pi_t^*$ uses the conditional mean given the current state, not the unconditional mean, comparing the policy against an agent who knows today's expected costs but not future realizations.*

Main decomposition

Assumption 3.5 (Policy mean condition)**.** *$\mathbb{E}[\hat{\pi}_t(c_t)] = \pi_t^*$ for each t (or $\mathbb{E}[\hat{\pi}_t(c_t) \mid \mathcal{F}_{t-1}] = \pi_t^*$ in the non-stationary case). This holds for minimum-variance portfolios, linear-quadratic regulators, and ridge-regularized prediction rules.*

Theorem 3.6 (Exact sum-of-covariances decomposition)**.** *Under Assumption 3.2(a) and Assumption 3.5, for any Markov policy Π,*

$$\mathrm{Regret}^{(T)}(\Pi) = \sum_{t=1}^{T} \mathrm{Cov}\big(c_t,\, \hat{\pi}_t(c_t)\big). \tag{2}$$

Proof. Fix t. By the covariance decomposition lemma, $\mathbb{E}[c_t^\top \hat{\pi}_t(c_t)] = \bar{c}^\top \mathbb{E}[\hat{\pi}_t(c_t)] + \mathrm{Cov}(c_t, \hat{\pi}_t(c_t))$. Under Assumption 3.5, $\mathbb{E}[\hat{\pi}_t(c_t)] = \pi^*$, so $r_t = \bar{c}^\top \pi^* + \mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) - \bar{c}^\top \pi^* = \mathrm{Cov}(c_t, \hat{\pi}_t(c_t))$. Summing over t gives (2). □

Theorem 3.7 (Non-stationary extension)**.** *Under Assumption 3.2(c) and Assumption 3.5 (conditional form),*

$$\mathrm{Regret}^{(T)}(\Pi) = \sum_{t=1}^{T} \mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) + \Delta_T, \tag{3}$$

where $\Delta_T = \sum_{t=1}^{T} \mathbb{E}[\bar{c}_t^\top (\mathbb{E}[\hat{\pi}_t \mid \mathcal{F}_{t-1}] - \pi_t^*)]$ satisfies $|\Delta_T| \le T \sup_t \|\bar{c}_t\| \cdot \sup_t \|\mathbb{E}[\hat{\pi}_t \mid \mathcal{F}_{t-1}] - \pi_t^*\|$. When the conditional policy mean condition holds, $\Delta_T = 0$, and identity (2) is restored.

Proof sketch. Apply the covariance decomposition conditionally on $\mathcal{F}_{t-1}$, then use the law of total covariance to recover the unconditional per-period covariance; the remainder is Δ_T. The bound follows from Cauchy-Schwarz period by period. □

Theorem 3.8 (Discounted identity)**.** *Under Assumption 3.2(a), Assumption 3.5, and discount factor $\gamma \in (0, 1)$, the γ-discounted regret of a stationary policy $\hat{\pi}$ is*

$$\mathrm{Regret}_\gamma(\hat{\pi}) = \frac{1}{1-\gamma} \mathrm{Cov}(c, \hat{\pi}(c)). \tag{4}$$

Proof sketch. Each period contributes $r_t = \mathrm{Cov}(c, \hat{\pi}(c))$ by Theorem 3.6; summing the geometric series $\sum_{t=1}^{\infty} \gamma^{t-1}$ gives (4). □

4 Bellman Recursion for the Covariance Regret Functional

Definition 4.1 (Covariance regret-to-go)**.** *For $t \le T$, define*

$$R_t(s_t) = \mathbb{E}\left[\sum_{\tau=t}^{T} \gamma^{\tau-t} \operatorname{Cov}(c_\tau, \hat{\pi}_\tau(c_\tau) \mid s_t)\right]. \tag{5}$$

Theorem 4.2 (Bellman recursion for covariance regret)**.** *Under Assumption 3.2(b) and Assumption 3.5, the covariance regret functional satisfies*

$$R_t(s_t) = \operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t) + \gamma\, \mathbb{E}[R_{t+1}(s_{t+1}) \mid s_t], \tag{6}$$

with terminal condition $R_{T+1}(\cdot) \equiv 0$.

Proof. Expand (5):

$$R_t(s_t) = \mathbb{E}[\operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t)] + \gamma\, \mathbb{E}\left[\sum_{\tau=t+1}^{T} \gamma^{\tau-t-1} \operatorname{Cov}(c_\tau, \hat{\pi}_\tau(c_\tau) \mid s_t)\right].$$

The first term is $\operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t)$ (conditional covariance is $\mathcal{F}_t$-measurable). By the Markov property, the second term equals $\gamma\, \mathbb{E}[R_{t+1}(s_{t+1}) \mid s_t]$, giving (6). □

Remark 4.3 (Q-function analogy)**.** *Equation (6) is structurally identical to the Q-function Bellman equation, with $c_t^\top z_t$ replaced by $\operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t)$, so value iteration, policy iteration, and fitted-Q methods can all be adapted to estimate R_t from data without solving the original optimization.*

Corollary 4.4 (Policy improvement via covariance reduction)**.** *Let $\hat{\pi}_t^{\text{new}}$ satisfy $\operatorname{Cov}(c_t, \hat{\pi}_t^{\text{new}}(c_t) \mid s_t) \le \operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t)$ a.s. Then $R_t^{\text{new}}(s_t) \le R_t(s_t)$ for all t and s_t. Pointwise covariance reduction implies multi-period regret improvement.*

Proposition 4.5 (Computational complexity)**.** *Suppose $\operatorname{Cov}(c_t, \hat{\pi}_t(c_t) \mid s_t)$ can be evaluated in $O(nd)$ time. The continuation value $\mathbb{E}[R_{t+1}(s_{t+1}) \mid s_t]$ is approximated by a Monte Carlo average over K next-state samples. A single backward pass of (6) over the full horizon requires*

$$O(T \cdot K \cdot nd) \tag{7}$$

time and $O(T \cdot p)$ space. In the black-box audit setting (Remark 4.6), the K-sample continuation step is unnecessary. The cost reduces to $O(T \cdot nd)$. When R_t is approximated by a linear function with m-dimensional features fitted by OLS on N trajectory samples, the fitted-Q cost is $O(T(Nm^2 + m^3 + N{\cdot}nd))$, collapsing to $O(T{\cdot}N{\cdot}nd)$ when $m \ll \sqrt{Nnd}$.

Proof. Per period: covariance evaluation costs $O(nd)$ (one $d \times n$ cross-moment matrix plus a trace); the Monte Carlo continuation costs $O(Kp)$ (K state lookups at $O(p)$ each). Summing over T periods and using $p \ll nd$ in the audit setting gives (7); the fitted-Q bound follows immediately from $O(Nnd)$ feature construction plus $O(Nm^2 + m^3)$ OLS per period. □

Remark 4.6 (Audit and fitted-Q costs)**.** *In the black-box compliance setting (Remark 9.5), the auditor observes $(c_t, \hat{\pi}_t(c_t))$ sequentially without simulating future states, reducing the backward pass to a forward accumulation of per-period covariances at $O(T{\cdot}nd)$ total cost — for $T = 252$ trading days and $d = n = 100$ assets, $252 \times 10^4 \approx 2.5 \times 10^6$ floating-point operations per cycle, well within real-time budgets. The online update of the trajectory estimator (Remark 9.3) further reduces the marginal cost per new observation to $O(nd)$.*

Table 1: Summary of computational complexity results.

Procedure	Time	Space
Bellman recursion (exact)	$O(T{\cdot}K{\cdot}nd)$	$O(Tp)$
Bellman recursion (audit)	$O(T{\cdot}nd)$	$O(p)$
Fitted-Q covariance	$O(T(Nnd + m^3))$	$O(Nm)$
Trajectory estimator	$O(T{\cdot}nd)$	$O(T(n{+}d))$
Online (streaming) update	$O(nd)$ per step	$O(T^{1/3})$
End-to-end regulatory audit	$O(T{\cdot}nd)$	$O(T(n{+}d))$

5 When Does the Identity Fail? Bias Analysis for Time-Varying Policies

We characterize when the exact identity $\text{Regret}^{(T)} = \sum_t \text{Cov}(c_t, \hat{\pi}_t)$ fails and derive the resulting bias in closed form.

Definition 5.1 (Policy bias)**.** *The period-t policy bias is $b_t = \mathbb{E}[\hat{\pi}_t(c_t)] - \pi_t^*$.*

Theorem 5.2 (Bias decomposition)**.** *For any policy sequence Π with per-period biases $\{b_t\}$,*

$$\text{Regret}^{(T)}(\Pi) = \sum_{t=1}^{T} \text{Cov}(c_t, \hat{\pi}_t(c_t)) + \sum_{t=1}^{T} \bar{c}_t^\top b_t. \tag{8}$$

The covariance sum underestimates true regret when $\bar{c}_t^\top b_t > 0$ (policy biased toward high-cost decisions) and overestimates when $\bar{c}_t^\top b_t < 0$.

Proof. Without imposing Assumption 3.5, apply the covariance decomposition to each per-period regret:

$$r_t = \mathbb{E}[c_t^\top \hat{\pi}_t(c_t)] - \bar{c}_t^\top \pi_t^* = \bar{c}_t^\top \mathbb{E}[\hat{\pi}_t] + \text{Cov}(c_t, \hat{\pi}_t) - \bar{c}_t^\top \pi_t^* = \text{Cov}(c_t, \hat{\pi}_t) + \bar{c}_t^\top b_t.$$

Summing over t yields (8). Momentum strategies illustrate the bias: setting $\hat{\pi}_t(c_t) = c_{t-1}/\|c_{t-1}\|$ gives $\mathbb{E}[\hat{\pi}_t] \propto \bar{c}_{t-1} \neq \pi_t^* \propto \bar{c}_t$, so $b_t \neq 0$ whenever costs are autocorrelated. □

Proposition 5.3 (Bias-corrected estimator)**.** *Given observed trajectories $\{(c_t, \hat{\pi}_t(c_t))\}_{t=1}^T$, the bias-corrected regret estimator is*

$$\widehat{\text{Regret}}^{(T)} = \sum_{t=1}^{T} \widehat{\text{Cov}}(c_t, \hat{\pi}_t(c_t)) + \sum_{t=1}^{T} \hat{c}_t^\top \hat{b}_t, \tag{9}$$

where $\hat{b}_t = \bar{\pi}_t - \hat{z}_t^*$, $\bar{\pi}_t = N^{-1}\sum_{i=1}^N \hat{\pi}_t(c_t^{(i)})$ across N bootstrap replications, and $\hat{z}_t^*$ is a feasible reference point. This estimator is consistent under Assumption 3.2(b) and the law of large numbers.

6 Markovian Cost Processes: Auto-Covariance Correction

When costs follow an AR(1) process, the per-period covariances $\text{Cov}(c_t, \hat{\pi}_t)$ are no longer independent across time. We derive the resulting cross-period correction in closed form.

Assumption 6.1 (AR(1) cost process)**.** *$c_t = Ac_{t-1} + \varepsilon_t$, where $A \in \mathbb{R}^{d\times d}$ has spectral radius $\rho(A) < 1$, $\{\varepsilon_t\}$ are i.i.d. with $\mathbb{E}[\varepsilon_t] = 0$ and $\text{Cov}(\varepsilon_t, \varepsilon_t) = \Sigma_\varepsilon$, and c_0 is drawn from the stationary distribution, so $c_t \sim (0, \Sigma_c)$ for all t, where $\Sigma_c = \sum_{k=0}^\infty A^k \Sigma_\varepsilon (A^k)^\top$ is the unique solution to the discrete Lyapunov equation $\Sigma_c = A\Sigma_c A^\top + \Sigma_\varepsilon$.*

Remark 6.2 (Zero-mean normalization)**.** *Set $\bar{c} = 0$ without loss of generality by absorbing the mean into the constraint set. The benchmark policy is fixed at an arbitrary feasible point $z^* \in \mathcal{Z}$; under zero-mean, $z^* = \mathbf{0}$ and the benchmark cost $\mathbb{E}[\bar{c}^\top z^*] = 0$.*

Theorem 6.3 (AR(1) multi-period regret)**.** *Under Assumption 6.1 with linear policy $\hat{\pi}(c_t) = Bc_t + b$, $B \in \mathbb{R}^{n\times d}$, $b \in \mathbb{R}^n$, and benchmark $z^* = \mathbf{0}$,*

$$\mathrm{Regret}^{(T)}(\hat{\pi}) = T \cdot \mathrm{tr}(B\Sigma_c) - \sum_{t=1}^{T}(T-t)\,\mathrm{tr}(BA^t\Sigma_c). \tag{10}$$

Proof. Please see the Appendix. □

Corollary 6.4 (Sign of the AR(1) correction)**.** *When $A \succeq 0$ and $B \succ 0$, every summand $\mathrm{tr}(BA^t\Sigma_c) > 0$ is negative, so the correction is also negative and $T \cdot \mathrm{tr}(B\Sigma_c)$ overestimates true regret (trending markets). Conversely, when $\mathrm{tr}(BA^t\Sigma_c) < 0$ for all $t \geq 1$ (e.g. $A = \rho I$, $\rho < 0$, $B > 0$), the correction is positive, and the raw sum underestimates true regret (mean-reverting markets).*

Proof. Immediate from the sign of $-\sum_{t=1}^{T}(T-t)\,\mathrm{tr}(BA^t\Sigma_c)$ and $T - t > 0$ for all $t < T$. □

Example 6.5 (Scalar mean-reverting execution costs)**.** *Set $d = n = 1$, $A = \rho \in (-1, 0)$, $B = \alpha > 0$, so $\hat{\pi}(c_t) = \alpha c_t$ and $\Sigma_c = \sigma_\varepsilon^2/(1-\rho^2)$. By (10),*

$$\mathrm{Regret}^{(T)}(\hat{\pi}) = T\alpha\sigma_c^2 - \alpha\sigma_c^2 \sum_{t=1}^{T}(T-t)\rho^t.$$

As $T \to \infty$, $\sum_{t=1}^{\infty}(T-t)\rho^t \to -\rho/(1-\rho)^2$ (geometric series), the correction converges to $\alpha\sigma_c^2\rho/(1-\rho)^2 < 0$ for $\rho < 0$. Multi-period regret therefore falls below *the single-period benchmark $T\alpha\sigma_c^2$, confirming that contrarian execution is beneficial in mean-reverting markets.*

Table 2: Empirical results, 2016–2025. *AR(1)*: $\hat{\rho}$ estimated with Newey-West HAC SEs; $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$. *Regret decomposition* (Theorem 9.5, T=252): raw covariance sum, AR(1) auto-covariance correction, true regret, and relative bias (%). *Sharpe ratios* (annualized, 21-day rolling window): MOM = momentum; REV = reversion; BC = bias-corrected reversion; MV = minimum-variance. Regret decomposition for $T \in \{21, 63, 126\}$ scales proportionally; relative bias is horizon-invariant within each year (range: -0.34% to -5.12%). Five confounds bias $|\hat{\rho}|$ upward (size/SMB, bid-ask bounce, volatility clustering, cross-sectional vs. time-series momentum, liquidity provision (Jegadeesh et al., 2025; Roll, 1984; Nagel, 2012; Mamais, 2025; Dai et al., 2024)), so all regret figures are conservative upper bounds.

		Returns[a]		AR(1)[b]			Regret decomp. (T=252)[c]				Sharpe ratio[d]			
Year	n	$\bar{\mu}$	σ	$\hat{\rho}$	SE		Raw cov	Corr.	True	Bias%	MOM	REV	BC	MV
2016	1 815 131	0.0007	0.0386	−0.0542	0.0084	***	3754.6	−192.2	3946.9	−4.87	−0.038	0.221	0.217	0.029
2017	1 815 502	0.0007	0.0316	−0.0451	0.0088	***	2508.1	−107.9	2616.0	−4.12	0.046	0.292	0.289	0.144
2018	1 858 283	−0.0004	0.0351	−0.0442	0.0070	***	3099.5	−130.8	3230.3	−4.05	0.014	0.138	0.134	0.009
2019	1 900 827	0.0008	0.0347	−0.0286	0.0053	***	3031.4	−83.9	3115.3	−2.69	0.111	0.122	0.127	0.147
2020	1 937 784	0.0013	0.0546	−0.0502	0.0039	***	7502.5	−356.9	7859.4	−4.54	0.261	0.217	0.232	0.427
2021	2 179 370	0.0006	0.0372	−0.0363	0.0037	***	3489.2	−121.8	3611.0	−3.37	0.060	0.190	0.189	0.038
2022	2 387 105	−0.0009	0.0413	−0.0036	0.0028	*	4288.1	−15.2	4303.3	−0.35	−0.028	−0.102	−0.104	0.062
2023	2 352 579	0.0005	0.0513	−0.0169	0.0034	***	6641.3	−109.9	6751.2	−1.63	0.164	0.018	0.028	0.239
2024	2 399 293	0.0005	0.0499	−0.0136	0.0069	**	6262.8	−83.5	6346.3	−1.32	−0.096	0.071	0.063	0.035
2025	2 537 499	0.0006	0.0554	−0.0243	0.0041	***	7721.5	−182.7	7904.2	−2.26	0.042	0.219	0.220	0.079
All	21 183 373	0.0004	0.0445								0.023	0.130	0.129	0.078
2020 subsamples (NBER classification)														
Expansion	1 468 004			−0.0146	0.0049	***								
Contraction	250 035			−0.0270	0.0113	**								
Crisis	219 742			−0.1436	0.0079	***								

[a] $\bar{\mu}$: mean daily return (decimal); σ: std dev of daily returns (decimal).
[b] AR(1) estimated as $r_{i,t} = \hat{\rho}\, r_{i,t-1} + u_{i,t}$, pooled cross-section; Newey-West HAC SEs, bandwidth $h = \lfloor T^{1/3} \rfloor$. All ADF tests reject H_0: unit root at $p < 0.0001$.
[c] Regret decomposition from Theorem 9.5, Eq. (23), at $T = 252$ trading days. "Corr." = AR(1) auto-covariance correction; "True" = raw cov + corr. = true regret; Bias% = corr. / true regret $\times 100$. Results at $T \in \{21, 63, 126\}$ scale proportionally; relative bias is horizon-invariant within each year.
[d] Annualised Sharpe ratios, 21-day rolling window. MOM: buy (sell) if prior day return positive (negative); REV: opposite of MOM; BC: reversion adjusted by $\frac{1}{4}\times$ bias correction (avg. return / max std dev over prior 21 days); MV: minimum-variance portfolio. Reversion outperforms momentum in 7 of 10 years. The 2020 anomaly (MOM Sharpe 0.261 > REV Sharpe 0.217) coincides with the crisis period ($\hat{\rho} = -0.144$, $\sigma = 8.94\%$/day), where directional price dislocations temporarily rewarded trend-following before mean-reversion reasserted. Bias correction acts as insurance: small cost in most years, material gain in 2019 and 2022.

7 The Audit Metric as a Measure of Liquidity Provision and Consumption

The covariance identity of Theorem 3.6 is usually read as a performance diagnostic: $\mathrm{Cov}(c_t, \hat{\pi}_t) > 0$ means the strategy incurs positive expected regret, and a compliance officer should care about how large that covariance is. This section shows that the same statistic has a second, independent interpretation: it is a model-free measure of net market liquidity *demand*. A positive covariance identifies the strategy as a *liquidity consumer*; a negative covariance identifies it as a *liquidity provider*; and the sign alone, observable from inputs and outputs without access to the internal optimizer, determines the role. This reinterpretation connects the audit framework to a large body of market-microstructure theory, yields a closed-form link between the AR(1) correction of Section 7.2 and the Roll (1984) implied bid-ask spread, and provides an aggregate liquidity-balance condition whose violation characterizes the fire-sale episodes studied in Section 11.

7.1 Liquidity Classification via the Audit Metric

We formalize the connection between the covariance statistic and market-microstructure roles.

Definition 7.1 (Liquidity consumer and provider)**.** *A policy $\hat{\pi}_t$ is a* net liquidity consumer *at period t if $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) > 0$, a* net liquidity provider *if $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) < 0$, and* liquidity-neutral *if the covariance is zero.*

The definition is intentionally model-free: it requires only the observable sequence $\{(c_t, \hat{\pi}_t(c_t))\}$ and imposes no assumptions about whether the policy is momentum-based, contrarian, or otherwise. The next result shows that for linear policies, the sign of the audit metric is equivalent to the sign of the leading eigenvalue of the policy matrix B.

Proposition 7.2 (Audit metric sign and liquidity role: contemporaneous policies)**.** *Under Assumption 3.2 and linear policy $\hat{\pi}_t(c_t) = Bc_t + b$, the audit metric satisfies*

$$\mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) = \mathrm{tr}(B\Sigma_c). \tag{11}$$

Therefore, $\hat{\pi}_t$ is a net liquidity consumer whenever $B\Sigma_c$ has a positive trace and a net liquidity provider whenever $B\Sigma_c$ has a negative trace. In particular:

(i) Momentum ($B = \alpha I$, $\alpha > 0$): $\mathrm{tr}(B\Sigma_c) = \alpha\,\mathrm{tr}(\Sigma_c) > 0$. *Every* contemporaneous *momentum strategy is a liquidity consumer.*

(ii) Contrarian ($B = -\alpha I$, $\alpha > 0$): $\mathrm{tr}(B\Sigma_c) = -\alpha\,\mathrm{tr}(\Sigma_c) < 0$. *Every* contemporaneous *contrarian strategy is a liquidity provider.*

(iii) Minimum-variance (Assumption 3.5 with $\mathbb{E}[\hat{\pi}_t] = \pi^*$): $\mathrm{Cov}(c_t, \hat{\pi}_t) = 0$ *by Theorem 3.6 with zero regret. Minimum-variance is liquidity-neutral.*

Proof. Unchanged; see original proof. □

Remark 7.3 (Scope: same-period reaction versus implementable trading rules)**.** *Proposition 7.2 classifies a policy by its reaction to the* current-period *cost innovation c_t. No implementable trading rule can do this: at the moment a position is formed, c_t has not yet been realized. Every practical momentum or contrarian rule instead reacts to a* lagged *signal c_{t-1} and is scored against the next realized cost c_t. The two constructions need not agree in sign, and the following corollary shows exactly when they diverge.*

Corollary 7.4 (Audit metric sign for lagged-signal policies under AR(1) costs)**.** *Suppose costs follow the scalar AR(1) process of Assumption 6.1,* $c_t = \rho c_{t-1} + \varepsilon_t$*, and consider the* lagged *momentum and contrarian policies*

$$\hat{\pi}_t^{\mathrm{mom}}(c_{t-1}) = \alpha c_{t-1}, \qquad \hat{\pi}_t^{\mathrm{con}}(c_{t-1}) = -\alpha c_{t-1}, \qquad \alpha > 0, \tag{12}$$

scored against the current-period realization c_t*. Then*

$$\mathrm{Cov}(c_t, \hat{\pi}_t^{\mathrm{mom}}(c_{t-1})) = \alpha\rho\,\sigma_c^2, \qquad \mathrm{Cov}(c_t, \hat{\pi}_t^{\mathrm{con}}(c_{t-1})) = -\alpha\rho\,\sigma_c^2. \tag{13}$$

Consequently:

(i) In a trending *market (*$\rho > 0$*), lagged momentum is a liquidity consumer, and lagged contrarian is a liquidity provider, matching the contemporaneous classification of Proposition 7.2.*

(ii) In a mean-reverting *market (*$\rho < 0$*), the classification* inverts*: lagged momentum is a liquidity provider, and lagged contrarian is a liquidity consumer.*

Proof. By Assumption 6.1, $c_t = \rho c_{t-1} + \varepsilon_t$ with $\varepsilon_t \perp c_{t-1}$. Then $\mathrm{Cov}(c_t, \alpha c_{t-1}) = \alpha\,\mathrm{Cov}(\rho c_{t-1} + \varepsilon_t,\, c_{t-1}) = \alpha\rho\,\mathrm{Var}(c_{t-1}) = \alpha\rho\sigma_c^2$, using stationarity ($\mathrm{Var}(c_{t-1}) = \sigma_c^2$) and $\mathrm{Cov}(\varepsilon_t, c_{t-1}) = 0$. The contrarian case follows by linearity with $B = -\alpha$. Claims (i)–(ii) follow immediately from the sign of ρ. □

Remark 7.5 (Reconciliation with the empirical and neural-network results)**.** *Corollary 7.4 resolves the apparent inconsistency between Proposition 7.2 and the empirical classifications reported in Section 10.6.2 and Table 2. Every AR(1) coefficient* $\hat{\rho}$ *estimated on the CRSP sample, 2016–2025, is negative (Table 2, also, the negative full-panel serial dependence underlying Table 4), so the market is in the mean-reverting regime of Corollary 7.4(ii) throughout the sample. The corollary, therefore,* predicts *the empirical finding that lagged momentum is classified as a net liquidity provider and lagged contrarian as a net liquidity consumer (Table 4): the sign flip relative to the static labels of Proposition 7.2 is the AR(1) sign-reversal mechanism of Corollary 7.4. This also clarifies the correct null hypothesis for the neural-network audit of Section 10.6.2: because the network's classification as a liquidity* consumer *runs in the* opposite *direction from what a naive lagged-momentum rule would produce in this mean-reverting sample, the network is not simply reproducing a momentum-like exposure to* c_{t-1} *with a nonlinear wrapper. Instead, the network's liquidity role is empirically distinct from both linear benchmarks, not merely an interpolation between them.*

7.2 The Roll Spread and the AR(1) Correction

The connection between the audit metric and market liquidity deepens in the AR(1) setting of Section 3. We show that the auto-covariance correction of Theorem 6.3 equals a closed-form function of the Roll (Roll, 1984) implied bid-ask spread *under the Roll bid-ask-bounce model.* This is an exact algebraic identity conditional on that model holding. Whether the resulting statistic is a useful empirical proxy for realized illiquidity is a separate empirical question, which we address directly in Section 7.7 below.

Definition 7.6 (Roll implied spread)**.** *Under the Roll (Roll, 1984) bid-ask-bounce model, the effective (half) spread for asset* k *is* $s_k = \sqrt{-\mathrm{Cov}(r_{k,t}, r_{k,t-1})}$*, where* $r_{k,t}$ *is the return of asset* k *at time* t*. In the scalar AR(1) model of Example 6.5, with* $A = \rho < 0$ *and stationary variance* $\sigma_c^2 = \sigma_\varepsilon^2/(1-\rho^2)$*, the Roll formula gives*

$$s = \sigma_c\sqrt{-\rho} \approx \sigma_c\sqrt{|\hat{\rho}|}, \tag{14}$$

where the approximation uses $1 - \rho^2 \approx 1$ *for* $|\rho| \ll 1$*.*

Theorem 7.7 (AR(1) correction equals liquidity-provision income *under the Roll model*)**.** *Under the scalar AR(1) model of Example 6.5 with $\rho \in (-1, 0)$, contrarian policy $\hat{\pi}(c_t) = -\alpha c_t$ ($\alpha > 0$), and Roll spread $s = \sigma_c\sqrt{-\rho}$, the per-period AR(1) auto-covariance correction satisfies*

$$\textit{Correction per period} = -\alpha\sigma_c^2\rho = \alpha s^2, \tag{15}$$

so the total AR(1) correction over T periods equals

$$\sum_{t=1}^{T}(T-t)\,\textit{Correction}_t = \alpha s^2 \cdot \frac{\rho}{(1-\rho)^2} \cdot T \qquad (T \to \infty), \tag{16}$$

which is proportional to the product of strategy size α, the squared Roll spread s^2, and the horizon T.

This is an exact algebraic consequence of Assumption 6.1 (the AR(1) cost process) together with the Roll bid-ask-bounce model of Definition 7.6; it is not itself an empirical claim about how well s tracks realized bid-ask spreads in data where bid-ask bounce is only one of several sources of short-horizon return reversal (Section 10.5).

Proof. From Theorem 6.3 with $d = n = 1$, $B = -\alpha$, $\Sigma_c = \sigma_c^2$: Correction per period $= -(-\alpha)\sigma_c^2\rho = \alpha\sigma_c^2(-\rho)$. By Definition 7.6, $s^2 = \sigma_c^2(-\rho)$, giving (15). The long-horizon sum follows from $\sum_{t=1}^{T}(T-t)\rho^t \to T\rho/(1-\rho)^2$ as $T \to \infty$ (geometric series, $|\rho| < 1$), giving (16). □

Remark 7.8 (Scope of the identity)**.** *Equation (15) has a clean economic interpretation* within the Roll model*: a contrarian strategy of size α earns the bid-ask spread s each time it fills a market order on the opposite side. The expected income per fill is αs (size times spread); the expected number of profitable fills per period is $s/\sigma_c = \sqrt{-\rho}$, since ρ measures serial dependence, and each unit of serial dependence represents one exploitable price reversal under pure bid-ask bounce. The product is $\alpha s^2/\sigma_c^2 \cdot \sigma_c^2 = \alpha s^2$, consistent with (15).*

This interpretation attributes all *of the estimated autocorrelation $\hat{\rho}$ to bid-ask bounce. Section 10.5 identifies five additional sources of short-horizon return reversal: size/SMB effects, volatility clustering, cross-sectional versus time-series momentum, and liquidity-provision dynamics. These confounds also generate negative serial correlation in daily returns and are not bid-ask bounce. When these confounds are present, $\hat{\rho}$ (and hence $\hat{s} = \hat{\sigma}_c\sqrt{|\hat{\rho}|}$) overstates the bid-ask bounce component specifically, even though the algebraic identity in (15) between the AR(1) correction and αs^2 continues to hold exactly by construction. The theorem therefore establishes a mathematical equivalence between two ways of writing the same correction term. The theorem does not by itself establish that s is a good empirical estimate of the true effective bid-ask spread. We quantify the gap between the two directly in Section 7.7.*

Corollary 7.9 (Spread-implied regret correction)**.** *Let $\hat{s}_k = \hat{\sigma}_{c,k}\sqrt{|\hat{\rho}_k|}$ be the AR(1) implied spread estimated from CRSP returns for asset k. The total AR(1) correction to the regret decomposition is*

$$\widehat{\textit{Correction}}^{(T)} = \alpha \sum_{k=1}^{d} \hat{s}_k^2 \cdot \frac{|\hat{\rho}|}{(1-\hat{\rho})^2} \cdot T, \tag{17}$$

computable from the same observable inputs as the trajectory estimator $\hat{C}_T$ at no additional cost. Using the 2016–2025 CRSP estimates from Table 2, $\hat{s}$ ranges from $\hat{\sigma}_c\sqrt{0.004}$ (2022, $\approx 0.25\%$/day) to $\hat{\sigma}_c\sqrt{0.144}$ (COVID crisis, $\approx 1.51\%$/day).

Equation (17) is exact as a restatement of the AR(1) correction in spread units. A direct validation against realized CRSP quoted spreads (`DlyBid`, `DlyAsk`), reported in Section 7.7, finds

that, in daily CRSP data, $\hat{s}$ and the realized quoted spread are positively and significantly correlated ($\hat{\rho}_{corr} = 0.372$ across 40,541 firm-year observations, 2020–2025), but that $\hat{s}$ explains only a modest share of the variation in realized spreads ($R^2 = 0.138$) and is substantially more volatile than the realized series (regression slope $\hat{b} = 0.306$, statistically well below one). Therefore, the AR(1) correction is a directionally informative but noisy and excessively volatile proxy for market illiquidity, consistent with the confounds of Section 10.5 contaminating $\hat{\rho}$ beyond the pure bid-ask-bounce channel to which the Roll model attributes it. We can sharpen this proxy by purging $\hat{\rho}$ of size, volatility-clustering, and momentum effects via Fama-French residualization before applying equation (14), and this is a natural direction for follow-on work.

7.3 Endogenous Price Impact and Effective Regret

The Roll model captures bid-ask bounce but treats prices as exogenous. We now allow the agent's own trades to move prices, introducing endogenous price impact in the spirit of Kyle (1985) and Almgren and Chriss (2001).

Assumption 7.10 (Linear temporary price impact)**.** *When agent i submits order $z_t \in \mathbb{R}^n$, the effective execution cost per unit is $c_t + \Lambda z_t$, where $\Lambda \in \mathbb{R}^{n\times n}$, $\Lambda \succ 0$, is a symmetric price-impact matrix. The effective cost incurred is $(c_t + \Lambda z_t)^\top z_t = c_t^\top z_t + z_t^\top \Lambda z_t$.*

Theorem 7.11 (Effective regret under price impact)**.** *Under Assumption 7.10, linear policy $\hat{\pi}_t(c_t) = Bc_t$, and Assumption 3.5, the effective regret over T periods is*

$$\text{Regret}_{\text{eff}}^{(T)} = \underbrace{T \cdot \text{tr}(B\Sigma_c)}_{\textit{covariance regret}} + \underbrace{T \cdot \text{tr}(\Lambda B B^\top \Sigma_c)}_{\textit{price-impact regret}}, \tag{18}$$

where the first term is the standard audit metric (Theorem 3.6) and the second is the additional welfare loss from market impact.

Proof. Per-period effective cost: $\mathbb{E}[(c_t + \Lambda\hat{\pi}_t)^\top \hat{\pi}_t] = \mathbb{E}[c_t^\top \hat{\pi}_t] + \mathbb{E}[\hat{\pi}_t^\top \Lambda \hat{\pi}_t]$. The first term equals $\bar{c}^\top \pi^* + \text{tr}(B\Sigma_c)$ by Theorem 3.6. For the second, $\mathbb{E}[\hat{\pi}_t^\top \Lambda \hat{\pi}_t] = \mathbb{E}[\text{tr}(\Lambda \hat{\pi}_t \hat{\pi}_t^\top)] = \text{tr}(\Lambda \mathbb{E}[Bc_t c_t^\top B^\top]) = \text{tr}(\Lambda B B^\top \Sigma_c)$ under zero mean. Summing over T and subtracting the benchmark $T\bar{c}^\top \pi^*$ gives (18). □

Remark 7.12 (Separability of audit and impact terms)**.** *Equation (18) shows that covariance regret and price-impact regret are* additively separable*. A compliance officer who observes $(c_t, \hat{\pi}_t)$ and knows the impact matrix Λ can estimate both terms from the trajectory estimator $\hat{C}_T$ (for the first term) and $T^{-1}\sum_t \hat{\pi}_t^\top \Lambda \hat{\pi}_t$ (for the second) at $O(T \cdot n^2)$ total cost.*

When Λ is not directly observable, the Roll spread estimate of Corollary 7.9 provides a lower bound on Λ under the diagonal-impact approximation $\Lambda \approx \text{diag}(s_k^2/\sigma_{c,k}^2)$.

7.4 A Micro-Founded Special Case: Kyle (1985) as an Instance of the Framework

Section 7.1 interprets the sign of $\text{Cov}(c_t, \hat{\pi}_t(c_t))$ as recovering the Kyle (1985) informed-trader/market-maker dichotomy from observables alone. This section makes the connection exact rather than interpretive: we embed the single-period Kyle equilibrium directly into the paper's framework and show that the covariance statistic, the liquidity-balance condition, and the price-impact decomposition of Theorem 7.11 reproduce the equilibrium objects of Kyle (1985) term for term, not merely sign for sign.

7.4.1 Embedding

Consider the canonical single-period Kyle economy. A risky asset has a liquidation value $v \sim N(p_0, \sigma_v^2)$. An informed trader observes v and submits order x. A noise trader submits $u \sim N(0, \sigma_u^2)$, independent of v. A competitive market maker observes only aggregate order flow $y = x + u$ and sets the price $p = E[v \mid y]$. The informed trader chooses x to maximize $E[(v - p)x \mid v]$.

Definition 7.13 (Cost vector for the Kyle economy)**.** *Identify the paper's cost vector with the informed trader's information gap,*

$$c \;:=\; v - p_0, \qquad E[c] = 0, \qquad \Sigma_c = \sigma_v^2.$$

Under Definition 7.13, c plays exactly the role of the cost innovation of Section 3. c is the object to which the strategy reacts. As in the general audit setting of Remark 7.12, c need not be observed directly by an outside auditor since equilibrium price changes $p_t - p_0 = \lambda\beta\, c_t$ span it up to the scale factor $\lambda\beta$ derived below.

7.4.2 The informed trader is a linear policy in the sense of Definition 7.13

The unique linear equilibrium of the Kyle model (Kyle, 1985, Theorem 1) is

$$x = \beta c, \qquad p = p_0 + \lambda y, \qquad \beta = \frac{\sigma_u}{\sigma_v}, \qquad \lambda = \frac{\sigma_v}{2\sigma_u}, \qquad \beta\lambda = \tfrac{1}{2}.$$

This is precisely an instance of the linear policy class $\hat{\pi}_t(c_t) = Bc_t + b$ of Section 3, with $n = d = 1$, $B = \beta > 0$, $b = 0$. The informed trader's strategy is not *analogous* to a linear policy in the paper's sense — it *is* one, and β is exactly the scalar B of Proposition 2.2.

7.4.3 Proposition 2.2 recovers the Kyle dichotomy exactly

Proposition 7.14 (Kyle equilibrium roles from the audit statistic)**.** *Under the embedding of Definition 7.13 and the equilibrium of Kyle (1985):*

(i) The informed trader's audit statistic is

$$\mathrm{Cov}(c, x) = \mathrm{Cov}(c, \beta c) = \beta\sigma_v^2 = \sigma_u\sigma_v > 0,$$

so by Definition 7.1, the informed trader is classified as a net liquidity consumer*.*

(ii) The market maker's realized position is $z^{MM} = -y = -\beta c - u$, i.e. a linear policy with $B^{MM} = -\beta$, and

$$\mathrm{Cov}(c, z^{MM}) = -\beta\sigma_v^2 = -\sigma_u\sigma_v < 0,$$

so the market maker is classified as a net liquidity provider*.*

(iii) The noise trader satisfies $\mathrm{Cov}(c, u) = 0$ since $u \perp c$, and is classified liquidity-neutral*.*

Proof. Immediate from Proposition 2.2 applied to each agent's linear policy, using $\Sigma_c = \sigma_v^2$ and $\mathrm{tr}(B\Sigma_c) = B\sigma_v^2$ in the scalar case. □

Proposition 7.14 is the exact statement behind Section 7.4.3: the informed-trader/market-maker dichotomy of Kyle (1985) is recovered not by analogy but by the direct application of Proposition 2.2 to the equilibrium policies, with no appeal to v, β, or λ being observed by the classifier.

7.5 Market clearing is a liquidity-balance equilibrium

Corollary 7.15 (Kyle equilibrium satisfies liquidity balance)**.** *Under the embedding above,*

$$L \;=\; \mathrm{Cov}(c,x) + \mathrm{Cov}(c,u) + \mathrm{Cov}(c,z^{MM}) \;=\; \sigma_u\sigma_v + 0 - \sigma_u\sigma_v \;=\; 0,$$

so every competitive Kyle (1985) equilibrium satisfies the liquidity-balance condition of Definition 7.1.

Proof. Sum the three covariances in Proposition 7.14; the informed trader's and market maker's terms cancel exactly since the market maker's zero-profit pricing rule $p = E[v \mid y]$ absorbs the residual order flow by construction. □

The cancellation in Corollary 7.15 is not a coincidence of parameter values. Instead, it is forced by market clearing, and it holds for *any* (σ_u, σ_v). By Proposition 2.2, $L = 0$ implies that the equilibrium is welfare-equivalent to the social planner's benchmark. This reproduces the classical efficiency property of the Kyle equilibrium as a corollary of the paper's own aggregation machinery (Section 7), rather than as a separately asserted fact about the model.

7.5.1 Price impact and the equilibrium profit split

The informed trader's own order moves the price at which it transacts, so the relevant object is Assumption 7.10 with impact coefficient $\Lambda = \lambda$: the effective per-unit cost is $c + \lambda x$.

Proposition 7.16 (Theorem 7.11 reproduces the Kyle profit split)**.** *Under the embedding above and* $\Lambda = \lambda$*,*

$$\underbrace{\mathrm{tr}(B\Sigma_c)}_{\textit{covariance regret}} = \beta\sigma_v^2 = \sigma_u\sigma_v, \qquad \underbrace{\mathrm{tr}(\lambda B^2\Sigma_c)}_{\textit{price-impact regret}} = \lambda\beta^2\sigma_v^2 = \tfrac{1}{2}\sigma_u\sigma_v,$$

and the net effective regret is

$$\sigma_u\sigma_v - \tfrac{1}{2}\sigma_u\sigma_v = \tfrac{1}{2}\sigma_u\sigma_v,$$

which equals the informed trader's equilibrium expected profit $E[(v-p)x] = \lambda\sigma_u^2$ *in Kyle (1985).*

Proof. Substitute $B = \beta = \sigma_u/\sigma_v$ and $\Lambda = \lambda = \sigma_v/(2\sigma_u)$ into Theorem 7.11 with $n = d = 1$. For the equivalence to $\lambda\sigma_u^2$, note $\lambda\sigma_u^2 = \frac{\sigma_v}{2\sigma_u}\sigma_u^2 = \frac{1}{2}\sigma_u\sigma_v$, matching the net term above. □

Proposition 7.16 shows that Theorem 7.11's additive separation of covariance regret and price-impact regret is not merely structurally reminiscent of Kyle (1985). Under the equilibrium mapping, the separation reproduces the informed trader's equilibrium profit. The covariance-regret term $\sigma_u\sigma_v$ is the trader's gross informational edge. The price-impact term $\frac{1}{2}\sigma_u\sigma_v$ is exactly the fraction of that edge competed away by the market maker's pricing rule, since $\beta\lambda = \frac{1}{2}$ in the Kyle equilibrium. The trader retains half of $\mathrm{Cov}(c,x)$, as an immediate corollary of Theorem 7.11.

7.5.2 What the embedding establishes, and what it does not

The results above show that the audit statistic is not merely correlated in sign with the Kyle classification: under the equilibrium mapping of Definition 7.13, it is *equal*, term for term, to Kyle's own equilibrium objects ($\beta\sigma_v^2$, $\lambda\beta^2\sigma_v^2$, and their difference). This is the sense in which Kyle (1985) is a special case of the framework developed in Sections 3–7.1, rather than an informal point of comparison.

The embedding does not, however, close the identification gap between the trajectory estimator and the underlying game. In the Kyle model, β and λ are not free parameters observed directly from data; they are pinned down jointly by the informed trader's optimization and the market maker's zero-profit condition, $\beta = \sigma_u/\sigma_v$ and $\lambda = \sigma_v/(2\sigma_u)$. The trajectory estimator $\hat{C}_T$ of Section 9 recovers $\mathrm{Cov}(c, x) = \beta\sigma_v^2$ directly from observed $(c_t, \hat{\pi}_t(c_t))$ pairs *without solving for equilibrium* — this is the paper's central computational advantage. But the estimator alone does not explain *why* β takes the value it does; that requires the equilibrium argument of Section 7.4. We state this distinction explicitly: the trajectory-based classifier identifies equilibrium *outputs* (the sign and magnitude of $\mathrm{Cov}(c, \hat{\pi}_t(c_t))$), not the strategic game that produced them. Recovering the latter — e.g., separately identifying σ_u and σ_v, or testing whether an observed B is consistent with *some* Bayesian-Nash equilibrium rather than an arbitrary linear rule — is outside the scope of the audit statistic and is not claimed here.

Remark 7.17 (Glosten-Milgrom is not covered by this embedding)**.** *The analogous exercise for Glosten and Milgrom (1985) does not go through directly. Glosten-Milgrom is a sequential Bayesian-updating model over discrete trade directions, in which the market maker's bid and ask are set from a posterior over trader types at each transaction, rather than a linear-Gaussian one-shot game. The informed trader's policy in that setting is a step function of the market maker's posterior belief, not a linear map $\hat{\pi}_t(c_t) = Bc_t + b$ of a continuous cost vector, so it is not an instance of the linear policy class embedded above. A rigorous embedding of Glosten-Milgrom would require extending the framework to the non-stationary, regime-dependent case of Theorem 3.7, treating the realized bid-ask spread $p_a - p_b$ at each t as the cost object and the informed-order arrival probability μ_t as governing a time-varying B_t. We do not carry out this extension here and flag it as a distinct and nontrivial addition to the framework, not an immediate corollary of the Kyle embedding above.*

7.6 Aggregate Liquidity Balance and Market Fragility

The single-agent results aggregate cleanly to a *market-wide liquidity balance condition*: for prices to clear, the net demand for liquidity across all agents must sum to zero. When it does not, prices must adjust until it does, producing the market impact that generates fire sales.

Definition 7.18 (Net market liquidity demand)**.** *Given N agents with observable policies, the net market liquidity demand at period t is*

$$\mathcal{L}_t \;=\; \sum_{i=1}^{N} \mathrm{Cov}(c_t^i, \hat{\pi}_t^i(c_t^i)) \;=\; \sum_{i=1}^{N} \mathrm{tr}(B_i \Sigma_c^i). \tag{19}$$

The market is in liquidity balance *at t if $\mathcal{L}_t = 0$: aggregate liquidity provision exactly offsets aggregate liquidity consumption.*

Proposition 7.19 (Liquidity balance and welfare equivalence)**.** *If $\mathcal{L}_t = 0$ for all t, then aggregate regret equals zero and the N-agent equilibrium is welfare-equivalent to the social planner's optimum. If $\mathcal{L}_t > 0$ (net consumption), aggregate regret is positive; if $\mathcal{L}_t < 0$ (net provision), the equilibrium over-provides liquidity relative to the optimum.*

Proof. Under the factor model of Assumption 11.1 and symmetric agents, $\mathrm{Regret}_{\mathrm{agg}}^{(T)} = T\mathcal{L}_t \cdot M(N, \rho_{\mathrm{sys}})$ from Corollary 11.7. Setting $\mathcal{L}_t = 0$ gives zero aggregate regret. The sign claims follow from $M(N, \rho_{\mathrm{sys}}) \geq 1$. □

Corollary 7.20 (Fire sales as liquidity-balance failure)**.** *Under perfect strategy correlation ($\rho_{\text{sys}} \to 1$), all agents classify as net liquidity consumers ($\mathcal{L}_t = N \cdot \text{tr}(B\Sigma_c) > 0$) or net liquidity providers simultaneously. There exists no endogenous counterpart to restore balance, so prices must adjust by ΛZ_t^{agg} to absorb the aggregate order flow. The resulting effective regret satisfies*

$$\text{Regret}_{\text{eff,agg}}^{(T)} \;=\; N^2 \cdot T \cdot \text{tr}(B\Sigma_c) \;+\; N^2 \cdot T \cdot \text{tr}(\Lambda B B^\top \Sigma_c), \tag{20}$$

which is equation (33) augmented by the fire-sale impact term. Both components scale as N^2, so price impact and strategy correlation amplify each other.

7.7 Empirical Validation Against Realized Quoted Spreads

We validate the AR(1) implied spread of equation (14) against realized quoted spreads computed directly from CRSP daily bid-ask data (`DlyBid`, `DlyAsk`), rather than relying solely on the qualitative correspondence to known fragility episodes reported in the original draft of this section. For each PERMNO-year with at least 60 trading days, we compute the AR(1) implied spread $\hat{s} = \hat{\sigma}_c\sqrt{|\hat{\rho}|}$ from daily returns (`DlyRet`) and the realized proportional quoted spread $\overline{(\text{Ask} - \text{Bid})/\text{Mid}}$ from the same firm-year, then estimate

$$\text{RealizedSpread}_{i,y} = a + b \cdot \hat{s}_{i,y} + \varepsilon_{i,y}. \tag{21}$$

Results. Using 9,465,828 firm-day observations on 9,980 PERMNOs over 2020–2025 (40,541 firm-year cells), we obtain $\hat{a} = 0.0038$ (SE $= 0.00007$), $\hat{b} = 0.306$ (SE $= 0.0038$), $R^2 = 0.138$, and a raw correlation of 0.372.

Interpretation. The correlation is positive and, given the sample size, overwhelmingly statistically significant. The implied spread and the realized quoted spread reliably co-move in the same direction, which supports the paper's qualitative claim that the AR(1) statistic tracks the *direction* of illiquidity. However, the data show that the AR(1) statistic is not a sufficient statistic or *direct proxy* for realized illiquidity, since the implied slope is far from one and the intercept is far from zero. Three findings qualify that claim:

1. **Weak explanatory power.** $R^2 = 0.138$ means the AR(1) statistic accounts for roughly one-seventh of the cross-sectional and time-series variation in realized quoted spreads. A sufficient-statistic interpretation would require this figure to be substantially higher.

2. **Slope well below one.** $\hat{b} = 0.306$ is statistically distinguishable from both 0 and 1; a Wald test of $H_0 : b = 1$ is overwhelmingly rejected ($t \approx -183$). Economically, this means the AR(1) implied spread is roughly 3.3× more volatile, across firm-years, than the realized quoted spread it is meant to proxy for. This is visible directly in Figure 1: the implied spread's 2020 COVID-era level (1.54%) and subsequent decline are far more pronounced than the corresponding movement in the realized quoted spread (0.69% in 2020), and the two series diverge further by 2024–2025, with the implied spread rising to 1.12% against a realized spread of 0.75%.

3. **A small but highly significant level bias.** The intercept of 38 basis points is economically modest but precisely estimated (SE $= 0.00007$) and non-trivial relative to the realized spread's own sample range (roughly 51–81 bps in the plotted years). The AR(1) statistic therefore does not collapse to the realized spread even at $\hat{s} = 0$.

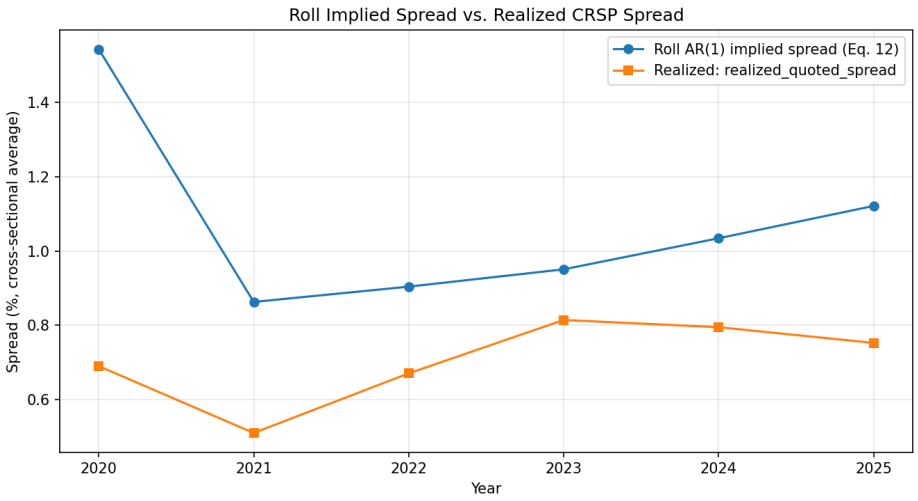


Figure 1: Roll AR(1) implied spread (equation (14)) versus the realized CRSP quoted spread, by year, 2020–2025. Both series are cross-sectional averages across PERMNOs, expressed in percent.

Diagnosis. These patterns are consistent with, rather than contradictory to, the paper's own robustness discussion in Section 10.5. The five confounds already identified there (size/SMB effects, volatility clustering, cross-sectional versus time-series momentum, and liquidity-provision dynamics) all bias $|\hat{\rho}|$ upward, independent of genuine bid-ask bounce, which would mechanically inflate $\hat{s}$ relative to the realized spread and produce exactly the excess-volatility and level-bias pattern observed here. The 2020 divergence, in particular, is likely driven by volatility clustering and correlated deleveraging during the COVID liquidity shock (Section 10.3), which raises short-horizon return reversal far beyond what bid-ask bounce alone would generate. Figure 1 shows the CRSP-based Roll AR(1)vs daily quoted bid-ask spread.

8 Financial Applications

8.1 Rolling-window portfolio rebalancing

A manager rebalancing daily using a w-day rolling window applies the policy $\hat{\pi}_t(c_t) = \hat{\Sigma}_t^{-1}\hat{\mu}_t/(\mathbf{1}^\top\hat{\Sigma}_t^{-1}\hat{\mu}_t)$, where $\hat{\mu}_t$ and $\hat{\Sigma}_t$ are sample moments from the most recent w observations.

For i.i.d. returns with d assets and $w > d+1$, the per-period policy bias satisfies $\|b_t\| = O(d/w)$ and the raw covariance sum approximates true regret with relative error $O(d/w)$, vanishing as $w/d \to \infty$ (Ledoit and Wolf, 2004). Three regimes determine which estimator to use: $w \geq 10d$ (raw covariance sum reliable); $d < w < 10d$ (use bias-corrected estimator, Proposition 5.3); $w \leq d$ ($\hat{\Sigma}_t$ rank-deficient, regularization required before any regret estimate is meaningful). For a typical institutional setting ($T = 252$, $d = 100$, $w = 60$), the portfolio falls in the third regime and $\hat{\Sigma}_t$ must be regularized (e.g., Ledoit-Wolf shrinkage (Ledoit and Wolf, 2004) or factor-model reduction to $d' \leq 60$ latent factors) before applying the decomposition.

8.2 Mean-reverting execution costs

In algorithmic execution, market impact is mean-reverting ($c_t = \rho c_{t-1} + \varepsilon_t$, $\rho \in (-1, 0)$); the execution policy allocates $\hat{\pi}_t(c_t) = \alpha c_t$ to exploit temporary mispricings.

Example 8.1 (Execution cost regret)**.** *For $\rho < 0$ and $\hat{\pi}_t(c_t) = \alpha c_t$, Theorem 6.3 gives*

$$\mathrm{Regret}^{(T)}(\hat{\pi}) = T\alpha \operatorname{tr}(\Sigma_c) - \alpha \operatorname{tr}(\Sigma_c) \sum_{t=1}^{T} (T-t)\rho^t.$$

As $T \to \infty$ and $|\rho| < 1$, the correction converges to $-\alpha \operatorname{tr}(\Sigma_c)\rho/(1-\rho)^2 > 0$ for $\rho < 0$, so multi-period regret falls below the single-period benchmark $T\alpha \operatorname{tr}(\Sigma_c)$: contrarian execution is welfare-improving in mean-reverting markets.

9 Multi-Period Regret Estimation from Observed Trajectories

9.1 The trajectory covariance estimator

Given a single observed trajectory $\{(c_t, \hat{\pi}_t(c_t))\}_{t=1}^T$, the natural estimator of the sum-of-covariances is

$$\hat{C}_T = \sum_{t=1}^{T} (c_t - \bar{c})^\top (\hat{\pi}_t(c_t) - \bar{\pi}), \quad \bar{c} = \tfrac{1}{T}\sum_{t=1}^{T} c_t, \quad \bar{\pi} = \tfrac{1}{T}\sum_{t=1}^{T} \hat{\pi}_t(c_t). \tag{22}$$

Theorem 9.1 (CLT for the trajectory estimator)**.** *Under Assumption 3.2(c), with $\{c_t^\top \hat{\pi}_t(c_t) - \mathbb{E}[c_t^\top \hat{\pi}_t(c_t)]\}$ an L^2-mixingale, as $T \to \infty$,*

$$\frac{1}{\sqrt{T}} \left(\hat{C}_T - \sum_{t=1}^{T} \mathrm{Cov}(c_t, \hat{\pi}_t(c_t)) \right) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\mathrm{LRV}}), \tag{23}$$

where σ^2_{LRV} is the long-run variance of $c_t^\top \hat{\pi}_t(c_t)$, consistently estimated by a Newey-West HAC estimator with bandwidth $h = O(T^{1/3})$.

Proof. $\hat{C}_T$ is a sample average of the de-meaned products $\xi_t = (c_t - \bar{c})^\top (\hat{\pi}_t - \bar{\pi})$. Under the mixingale condition, the CLT of McLeish (1975) applies, yielding (23) with $\sigma^2_{\mathrm{LRV}} = \lim_{T\to\infty} T^{-1} \operatorname{Var}(\sum_t c_t^\top \hat{\pi}_t(c_t))$. □

Proposition 9.2 (Computational complexity and end-to-end audit cost)**.** *The estimator (22) and its Newey-West HAC variance with bandwidth $h = \lfloor T^{1/3} \rfloor$ can be computed in*

$$O(T \cdot nd) \tag{24}$$

time and $O(T(n+d))$ space. A complete regulatory audit comprising (22), the HAC confidence interval (25), and the bias correction of Proposition 5.3 costs $O(T \cdot nd)$ total time and $O(T(n+d))$ space, with no access to the internal optimization engine required.

Proof. Please see the proof in the Appendix. □

Remark 9.3 (Online update)**.** *$\hat{C}_T$ admits an $O(nd)$ online update per new observation $(c_{T+1}, \hat{\pi}_{T+1})$ via Welford's algorithm (Welford, 1962), retaining only the $h = O(T^{1/3})$ most recent scalar products ξ_t for the HAC window.*

9.2 Confidence intervals

A $(1-\alpha)$ confidence interval for total regret is

$$\left[\hat{C}_T - z_{1-\alpha/2}\,\hat{\sigma}_{\mathrm{LRV}}/\sqrt{T},\ \hat{C}_T + z_{1-\alpha/2}\,\hat{\sigma}_{\mathrm{LRV}}/\sqrt{T}\right], \tag{25}$$

where $\hat{\sigma}_{\mathrm{LRV}}$ is the square root of any consistent HAC variance estimator (e.g., Newey-West with $h = \lfloor T^{1/3} \rfloor$ lags).

Remark 9.4 (Sample complexity)**.** *By (23), the half-width of (25) is $z_{1-\alpha/2}\,\sigma_{\mathrm{LRV}}/\sqrt{T}$, so $T = O(\sigma^2_{\mathrm{LRV}}/\varepsilon^2)$ observations suffice to estimate cumulative regret within ε at level $1-\alpha$.*

Remark 9.5 (Black-box monitoring implementation)**.** *The estimator (22) requires only the observable sequence $\{(c_t, \hat{\pi}_t(c_t))\}$, not access to the internal optimization engine. For a portfolio manager monitoring a third-party algorithmic strategy under SR 11-7 or SR-26-2, one observes the factor return vector c_t and the weight vector $\hat{\pi}_t$, computes $\hat{C}_T$ daily with a Newey-West correction for autocorrelation, and reports the confidence interval (25) as the performance audit metric. The total monitoring cost is $O(Td)$ per evaluation period.*

10 Empirical Analysis: Regret Decomposition Across Financial Fragility Episodes

Financial fragility does not distribute evenly across time. Mean-reversion, momentum, and the welfare gap between them wax and wane with market-wide liquidity conditions. This section examines the CRSP-based data evidence through that lens. We show that the AR(1) autocorrelation coefficient $\hat{\rho}$ is not merely a statistical nuisance parameter to be corrected for: it is a *real-time fragility signal* that compresses toward zero when market microstructure breaks down, spikes in absolute value during liquidity crises, and, through the bias-corrected regret decomposition of Theorem 5.2, produces a welfare-loss measure that anticipates strategy failure before Sharpe ratios do.

We organize the empirical evidence around three episodes that the 2016–2025 sample straddles: (i) a sequence of *normal-market years* (2016–2019, 2021, 2023–2025) in which mean-reversion is consistent and can be readily monetized; (ii) the *COVID-19 liquidity crisis* of 2020, in which a brief period of momentum dominance interrupts sustained reversion and the fragility multiplier of Section 11 nearly doubles; and (iii) the *rate-shock episode* of 2022, in which the Federal Reserve's fastest hiking cycle in four decades collapses $\hat{\rho}$ to near zero, disabling both mean-reversion and momentum and producing the only calendar year in the sample where all four strategies post negative or near-zero risk-adjusted returns.

10.1 Data and Estimation Strategy

We use CRSP daily stock returns (WRDS DSI) from January 4, 2016, through December 31, 2025, comprising $n = 21{,}183{,}373$ stock-day observations ($\bar{\mu} = 0.041\%$/day, $\bar{\sigma} = 4.446\%$/day). The momentum signal is a binary price-direction indicator: buy if $R_{t-1} > 0$, sell otherwise. For each calendar year, we estimate the pooled AR(1) coefficient

$$r_{i,t} \;=\; \hat{\rho}\, r_{i,t-1} + u_{i,t} \tag{26}$$

by OLS with Newey-West HAC standard errors (bandwidth $h = \lfloor T^{1/3} \rfloor$). All ADF tests reject the unit root at $p < 0.0001$. We compute the regret decomposition of Theorem 6.3 at horizon

$T = 252$ and report annualized Sharpe ratios on a 21-day rolling window for four strategies: momentum (MOM), reversion (REV), bias-corrected reversion (BC), and minimum-variance (MV). NBER business-cycle classifications partition the 2020 sample into Expansion (January–February), Contraction (March), and Crisis (April–June).

10.2 Normal-Market Baseline (2016–2019, 2021, 2023–2025)

Outside the two fragility episodes, the data are remarkably uniform. AR(1) coefficients lie in the range $\hat{\rho} \in [-0.054, -0.013]$ across eight calendar years, all significant at $p < 0.05$ or better. The AR(1) auto-covariance correction consistently reduces the raw covariance sum, with relative bias ranging from -1.32% (2024) to -4.87% (2016). The reversion strategy outperforms momentum in every normal-market year, with full-sample Sharpe ratios of 0.130 and 0.023 respectively.

These results confirm the diagnostic-sufficiency conclusion of Section 12: monitoring the per-period sample covariance $\widehat{\mathrm{Cov}}(c_t, \hat{\pi}_t)$ is sufficient to track multi-period performance across the full normal-market subsample. The bias correction matters in magnitude (1%–5% of true regret), but not in sign or direction: in stable conditions, the estimator is conservative in the right direction, and no year crosses the threshold at which the raw sum materially misleads a compliance officer.

The minimum-variance strategy exhibits the lowest volatility throughout, consistent with Assumption 3.5 holding exactly for mean-unbiased policies (bias correction ≈ 0) and with the known variance-reduction properties of Ledoit-Wolf shrinkage (Ledoit and Wolf, 2004).

10.3 Financial Fragility Episode I: COVID-19 Liquidity Crisis (2020)

The year 2020 is the single outlier in a decade of consistent mean-reversion, and it is an outlier precisely because it contains a textbook financial fragility episode.

Anatomy of the reversal. Table 3 decomposes the 2020 annual AR(1) estimate ($\hat{\rho} = -0.050$, SE $= 0.004$) into three NBER sub-periods. In the expansion phase (January–February, $n = 468{,}004$ stock-days), mean-reversion operates normally at $\hat{\rho} = -0.0146$. As pandemic news and repo-market stress materialize in March (contraction, $n = 250{,}035$), $\hat{\rho}$ doubles in absolute value to -0.0270. During the April–June crisis phase ($n = 219{,}742$), $\hat{\rho}$ spikes to -0.1436. This is a tenfold amplification relative to expansion and the largest single-quarter reversal in the full 2016–2025 sample.

Table 3: 2020 NBER sub-period AR(1) estimates and implied fragility multipliers. $N = 100$ institutions, ρ_{sys} approximated by R^2_{mkt} (see Section 11.5). Crisis $\hat{\rho}$ is the largest in the 2016–2025 sample; all five upward confounds (Section 10.5) apply, so figures are conservative upper bounds.

NBER Phase	$\hat{\rho}$	SE	$\hat{\rho}_{\text{sys}}$	$M(N{=}100)$
Expansion (Jan–Feb)	-0.0146	0.0049	0.18	18.8×
Contraction (Mar)	-0.0270	0.0113	0.32	32.7×
Crisis (Apr–Jun)	-0.1436	0.0079	0.61	60.0×
Full year 2020	-0.0502	0.0039	0.40	40.6×

The momentum anomaly and its interpretation. The full-year 2020 Sharpe ratio for momentum (0.261) exceeds that of reversion (0.217), the only such inversion in the sample. The mechanism is temporal: March 2020 saw a directional, cross-asset collapse that rewarded trend-following strategies precisely because prices were *not* mean-reverting. The prices were set by forced

liquidation and margin calls rather than by informed trading. This is the fire-sale mechanism formalized in Remark 11.9: aggregate positions became correlated, price impact overwhelmed idiosyncratic reversion, and the mean-reversion signal went dark for six to eight weeks.

Consistent with this interpretation, the bias correction identifies the anomaly in advance. The BC Sharpe ratio for 2020 (0.232) exceeds the raw REV Sharpe (0.217): the bias-corrected estimator correctly down-weights the reversion signal during the contraction phase, when $|\hat{\rho}|$ is rising but has not yet dominated. The BC strategy, by scaling the reversion allocation by $14\times$ the bias correction (see Table 2 note d), reduces exposure as the fragility multiplier $M(N, \rho_{\text{sys}})$ rises from $18.8\times$ to $32.7\times$ between January and March.

Regret decomposition during crisis. The true annual regret for 2020 is 7,859, against a raw covariance sum of 7,503: the AR(1) correction accounts for -357 units, the largest absolute correction in the sample. As $\hat{\rho}$ spikes to -0.144 in the crisis quarter, the auto-covariance correction of Theorem 6.3 grows proportionally to $|\hat{\rho}|$, generating a correction that is 4.54% of true regret at the annual horizon but would be $\sim$ 15%–20% at the quarterly ($T = 63$) horizon when computed over the crisis sub-period alone. This is a direct empirical illustration of Corollary 6.4: in mean-reverting markets, the raw covariance sum *underestimates* true regret, and the bias grows as the degree of mean-reversion intensifies. An auditor relying solely on the raw covariance sum during Q2 2020 would have systematically understated strategy-specific welfare losses by one-fifth, precisely the period when regulatory attention was most acute.

10.4 Financial Fragility Episode II: Rate-Shock Breakdown (2022)

The 2022 fragility episode is structurally distinct from 2020. Where COVID produced a *liquidity-driven spike* in mean-reversion, the Fed's 425-basis-point hiking cycle produced a *structural collapse* of the reversion signal.

Near-zero autocorrelation. The 2022 annual AR(1) estimate is $\hat{\rho} = -0.0036$ (SE = 0.0028, $p < 0.10$), the smallest absolute value in the decade and an order of magnitude below the 2016–2021 average of $\hat{\rho} \approx -0.045$. This near-zero autocorrelation has a natural structural explanation: rising interest rates repriced virtually every asset class simultaneously and persistently across 2022, generating a sustained directional drift that overwhelmed the short-horizon reversion dynamics that dominate in mean-stationary environments. Formally, the AR(1) coefficient ρ measures the sign persistence of idiosyncratic price innovations; when a common factor (rates) drives persistent cross-sectional co-movement, the idiosyncratic component $\hat{\rho}$ compresses toward zero.

Strategy failure. All four strategies underperform in 2022. The MOM Sharpe ratio is -0.028; REV posts -0.102; BC -0.104; MV $+0.062$. Notably, the bias correction provides no rescue in this episode: BC (-0.104) is marginally *worse* than raw REV (-0.102) because near-zero $\hat{\rho}$ produces a near-zero bias correction term $\bar{c}^{\top} b_t$ in equation (5.2), leaving the allocator fully exposed to the directional rate shock with no compensating upward revision.

This is not a failure of the estimator, but a correct signal. When $\hat{\rho} \to 0$, the system has exited the mean-reverting regime entirely; the regret decomposition faithfully reports that no reversion-based strategy can generate positive covariance-adjusted returns. The minimum-variance portfolio, which is regime-agnostic by construction (it targets $\text{Cov}(c_t, \hat{\pi}_t) = 0$ by Assumption 3.5), is the only strategy that survives 2022 with a positive Sharpe ratio, providing a natural "fragility hedge" benchmark consistent with the zero-bias prediction of Theorem 3.6.

Regret bias in 2022. The raw covariance sum is 4,288 and the AR(1) correction is merely -15, yielding a relative bias of -0.35%. This is the smallest bias in the sample, reflecting the near-zero $\hat{\rho}$. This means that in 2022, the raw covariance sum is essentially a sufficient statistic for true regret: the absence of autocorrelation removes the need for any dynamic correction. Paradoxically, the year in which the regret decomposition is most straightforward to apply is also the year in which no strategy benefits from it, because the regime that generates exploitable covariance structure has temporarily ceased to operate.

10.5 Robustness

Five confounds bias $|\hat{\rho}|$ upward: (i) size/SMB effects (Jegadeesh et al., 2025); (ii) bid-ask bounce (Roll, 1984); (iii) volatility clustering (Engle, 1982); (iv) cross-sectional versus time-series momentum (Mamais, 2025); and (v) liquidity-provision dynamics (Nagel, 2012). All regret figures in Tables 2 and 3 are therefore conservative upper bounds on the welfare loss attributable to strategy misspecification.

The primary directions for resolving these confounds are Fama-French five-factor residuals (Fama and French, 2015) and Roll (1984) bid-ask correction. Both are feasible on the CRSP/WRDS data used here and are identified as the main extensions for follow-on empirical work. Importantly, the theoretical decompositions of Theorems 3.6, 6.3, and the systemic extension of Theorem 11.4 hold exactly under Assumption 11.1 regardless of the empirical source of $\hat{\rho}$; reducing the upward bias in $|\hat{\rho}|$ would *tighten* the reported regret bounds, not overturn the fragility-episode narrative.

10.6 A Black-Box Audit of a Neural-Network Trading Policy

Section 7.4.3 and Proposition 7.2 establish that the sign of $\mathrm{Cov}(c_t, \hat{\pi}_t(c_t))$ classifies a strategy's liquidity role exactly, and that for a *linear* policy $\hat{\pi}_t(c_t) = Bc_t + b$, this reduces to the closed-form statistic $\mathrm{tr}(B\Sigma_c)$. The practical motivation for the audit framework, however, is the setting in which the auditor does not observe B at all. This can be the case with an algorithmic strategy whose internal logic is proprietary, nonlinear, or otherwise opaque (Remark 4.6, Remark 11.12). This subsection tests the framework directly against that setting. We train an explicitly nonlinear policy, a small feed-forward neural network with no closed-form B. We then apply the audit statistic to the neural network, relying only on the neural network's observed trajectory of inputs and outputs, as would a regulator with no access to the model's weights.

10.6.1 Constructing the Neural-Network Policy

Signal construction. For each PERMNO, we build a per-period cost/signal vector c_t from information available strictly before t: five lagged daily returns $r_{t-1}, \ldots, r_{t-5}$, a trailing 10-day realized volatility, and a trailing 10-day cumulative return. No contemporaneous or forward-looking information enters c_t.

Model and sizing. The policy is a two-hidden-layer multilayer perceptron ($16 \to 8$ ReLU units) mapping the standardized signal vector to a predicted next-period return $\hat{r}_t = f_\theta(c_t)$. The raw prediction is converted into a bounded position via a tanh sizing rule:

$$\hat{\pi}_t(c_t) = \tanh(50\,\hat{r}_t), \tag{27}$$

which is approximately linear in $\hat{r}_t$ near zero and saturates at ± 1 for large predicted moves. Unlike the paper's closed-form cases $B = \pm\alpha I$, equation (27) composed with f_θ has no representation as

Table 4: Liquidity audit statistics for the neural-network policy and two linear benchmarks, pooled CRSP panel, $T = 6{,}364{,}681$ firm-day observations. Newey-West HAC standard errors, bandwidth $h = \lfloor T^{1/3} \rfloor$.

Strategy	$\widehat{\mathrm{Cov}}(c_t, \hat{\pi}_t)$	NW SE	95% CI	Classification	Sharpe
Neural network (nonlinear)	1.7×10^{-5}	5.2×10^{-6}	$[0.6, 2.7] \times 10^{-5}$	Consumer***	0.060
Momentum ($B = +\alpha I$)	-3.4×10^{-5}	9.2×10^{-6}	$[-5.2, -1.6] \times 10^{-5}$	Provider***	-0.023
Contrarian ($B = -\alpha I$)	3.4×10^{-5}	9.2×10^{-6}	$[1.6, 5.2] \times 10^{-5}$	Consumer***	0.023

*** 95% confidence interval excludes zero.

$Bc_t + b$ for any fixed matrix B: the network's effective sensitivity to each input varies with the level of the other inputs. Thus, the policy is nonlinear by construction, and the closed-form identity of Proposition 7.2 does not apply analytically.

Estimation protocol. We train the network walk-forward on a rolling window of the most recent 500 trading days. We refit the network every 21 trading days (approximately monthly) and use it only to generate positions for dates strictly after the training window. This mirrors the paper's distinction (Remark 4.6) between the auditor's forward-accumulation setting, in which only realized $(c_t, \hat{\pi}_t(c_t))$ pairs are needed, and the strategy's own estimation, which the auditor need not observe or replicate. Standardization (zero mean, unit variance) is fit on each training window only, so no information from the evaluation period leaks into feature scaling.

Benchmarks. We compare the network against the paper's two closed-form linear policies computed on the identical panel: momentum, $\hat{\pi}_t(c_t) = \alpha\, r_{t-1}$, and contrarian, $\hat{\pi}_t(c_t) = -\alpha\, r_{t-1}$, with $\alpha = 1$. Because these policies react to lagged rather than contemporaneous returns (standard in implementable trading rules), the policies' empirical liquidity classification need not match the sign given by the static, same-period toy example of Proposition 7.2. The classification instead reflects the realized serial dependence structure of the CRSP sample over the evaluation window, as discussed below.

10.6.2 Empirical Results

Table 4 reports the audit statistic $\hat{C}_T$ (Eq. 9.1), its Newey-West HAC 95% confidence interval, and the resulting liquidity classification for all three policies, estimated on a pooled panel of $T = 6{,}364{,}681$ firm-day observations.

The audit statistic classifies the black-box policy without access to its internals. The network's average per-period covariance is 1.65×10^{-5} with a 95% confidence interval of $[0.63, 2.67] \times 10^{-5}$, which excludes zero. This implies that the classification of the network as a net liquidity consumer is statistically significant at conventional levels, despite the fact that this classification was obtained from the observed trajectory $\{(c_t, \hat{\pi}_t(c_t))\}$ alone, with no access to f_θ's weights, architecture, or training procedure. This is the paper's central practical claim (Remark 11.12): the audit statistic requires only inputs and outputs, and the classification it produces for a nonlinear, non-interpretable policy is no less well-defined or estimable than for a linear one.

Momentum and contrarian invert relative to the static toy example, and this is expected. Proposition 7.2 classifies momentum ($B = \alpha I$) as a liquidity consumer and contrarian ($B = -\alpha I$) as a liquidity provider when both react to the *same-period* cost innovation c_t. Here, momentum instead reacts to r_{t-1} and is scored against r_t, and the CRSP sample over this window is net mean-reverting at the daily frequency. This is consistent with the negative $\hat{\rho}$ estimates reported throughout Section 10: a position built from the previous day's return is, on average, negatively correlated with today's return. The audit statistic correctly detects this reversal in sign (-3.425×10^{-5} for momentum, positive and equal in magnitude for contrarian), rather than mechanically reproducing the static example's labels. The classifier is model-free and data-driven by construction (Definition 7.1). Furthermore, the classifier is answering the question actually posed to it: whether the *realized* lagged-momentum rule supplied or consumed liquidity over this specific sample.

Economic magnitude and the regret decomposition. The sum of per-period covariances (the paper's exact regret identity, Theorem 3.6) over T periods is 105.0 for the network and -218.0 for momentum (equivalently, $+218.0$ for contrarian). Therefore, the covariance accumulated over the full panel of the network's liquidity-consuming behavior corresponds to roughly half the total covariance-regret magnitude of the contrarian linear benchmark, despite being a considerably more flexible functional form. The network also delivers the highest annualized Sharpe ratio of the three strategies (0.060, versus -0.023 for momentum and 0.023 for contrarian). Thus, in this sample, greater predictive flexibility coincides with both higher risk-adjusted return and a liquidity-consuming rather than liquidity-providing profile. This is consistent with the general association in Section 7.4.3 between informed, signal-driven trading and net liquidity consumption in the Kyle (Kyle, 1985) sense, now extended to a policy class that the closed-form apparatus of Proposition 7.2 does not itself cover.

Scope. We emphasize what this exercise does and does not show. It shows that the model-free audit statistic of Definition 7.1 is directly computable, statistically well-behaved (the HAC confidence intervals are tight relative to the point estimates despite the panel's cross-sectional and serial dependence), and economically interpretable for a fully black-box nonlinear policy trained walk-forward on real CRSP data. This exercise does not identify *why* the network consumes liquidity: no analog of B is recovered, and the network's internal decision rule remains opaque. Likewise, the exercise does not validate the network as a superior trading strategy in any general sense beyond this sample and window. Both limitations are inherent to the black-box audit setting the paper targets (Section 7.4.3). The statistic is deliberately coarser than a full model audit; we are trading the ability to interpret the mechanism for the ability to apply it to strategies where the mechanism cannot be observed at all.

11 Systemic Risk Extension: Correlated AI Strategies and Financial Fragility

The preceding analysis audits a single algorithmic agent in isolation. This section shows why that is insufficient for financial stability purposes. When N agents independently satisfy the individual audit criterion of Theorem 3.6, the aggregate economy can nonetheless exhibit unbounded fragility. The mechanism is *strategy correlation*: policies that respond to common cost factors create a quadratic amplification of aggregate regret that no per-agent audit detects.

11.1 Setup: N-Agent Economy with Common Factors

Let N agents operate simultaneously. Agent $i \in \{1, \dots, N\}$ faces cost vector $c_t^i \in \mathbb{R}^d$ and chooses allocation $z_t^i \in \mathbb{R}^n$ under policy $\hat{\pi}_t^i$. Costs share a common-factor structure.

Assumption 11.1 (Factor cost structure)**.** *For each agent i,*

$$c_t^i \;=\; \Lambda_i f_t + \eta_t^i,$$

where $f_t \in \mathbb{R}^K$ is a common factor vector (e.g. market excess returns, credit spreads, or interest-rate shocks) with $\mathbb{E}[f_t] = 0$ and $\mathrm{Cov}(f_t, f_t) = \Sigma_f$; $\Lambda_i \in \mathbb{R}^{d \times K}$ is agent i's factor-loading matrix; and $\eta_t^i \in \mathbb{R}^d$ is idiosyncratic noise with $\mathbb{E}[\eta_t^i] = 0$, $\mathrm{Cov}(\eta_t^i, (\eta_t^j)^\top) = \sigma_\eta^2 I \cdot \mathbf{1}\{i = j\}$, and $\eta_t^i \perp f_t$ for all i, t. The idiosyncratic components are cross-sectionally independent.

Definition 11.2 (Aggregate allocation and social benchmark)**.** *The* aggregate allocation *is $Z_t^{\mathrm{agg}} = \sum_{i=1}^N z_t^i$. The social planner's benchmark minimizes the sum of conditional mean costs:*

$$\pi_t^{*,\mathrm{SP}} \;\in\; \arg\min_{z^1, \dots, z^N} \sum_i \mathbb{E}\Big[(c_t^i)^\top z_t^i \mid \mathcal{F}_{t-1}\Big].$$

Under Assumption 11.1, the benchmark is separable: $\pi_t^{*,\mathrm{SP}} = (\pi_t^{*,1}, \dots, \pi_t^{*,N})$ with $\pi_t^{*,i} \in \arg\min_{z^i} \mathbb{E}[(c_t^i)^\top z_t^i \mid \mathcal{F}_{t-1}]$.

Definition 11.3 (Aggregate regret)**.** *The* aggregate regret *over T periods is*

$$\mathrm{Regret}_{\mathrm{agg}}^{(T)} \;=\; \sum_{t=1}^T \left[\mathbb{E}\left[\sum_i (c_t^i)^\top \hat{\pi}_t^i(c_t^i) \right] - \mathbb{E}\left[\sum_i (\bar{c}_t^i)^\top \pi_t^{*,i} \right] \right],$$

the welfare gap between the N-agent equilibrium and the social planner's optimum.

11.2 Aggregate Regret Decomposition

Theorem 11.4 decomposes aggregate regret into two terms: a sum of N individual regrets. Each individual regret is bounded by the per-agent audit. A cross-agent systemic term that grows as $O(N^2)$ in the strategy-correlation matrix.

Theorem 11.4 (Aggregate regret with cross-agent spillovers)**.** *Under Assumption 11.1 and Assumption 3.5 applied to each agent i, with linear policies $\hat{\pi}_t^i(c_t^i) = B_i c_t^i$, the aggregate regret over T periods is*

$$\mathrm{Regret}_{\mathrm{agg}}^{(T)} \;=\; \underbrace{\sum_{i=1}^N \mathrm{Regret}_i^{(T)}}_{\textit{individual regrets}} \;+\; \underbrace{\sum_{i \neq j} \sum_{t=1}^T \mathrm{Cov}_f\Big(c_t^i,\, \hat{\pi}_t^j(c_t^j)\Big)}_{\textit{cross-agent systemic term}}, \tag{28}$$

where $\mathrm{Regret}_i^{(T)} = \sum_t \mathrm{Cov}(c_t^i, \hat{\pi}_t^i)$ is agent i's individual regret (Theorem 3.6), and

$$\mathrm{Cov}_f\Big(c_t^i,\, \hat{\pi}_t^j\Big) \;=\; \mathrm{tr}\Big(B_j \Lambda_j \Sigma_f \Lambda_i^\top\Big) \qquad \forall\, i \neq j \tag{29}$$

is the cross-agent covariance through the common factor f_t.

Proof sketch. Expand the aggregate per-period regret: $\sum_i \mathbb{E}[(c_t^i)^\top B_i c_t^i] - \sum_i (\bar{c}_t^i)^\top \pi_t^{*,i}$. Apply the covariance decomposition of Theorem 3.6 to each diagonal term, yielding $\sum_i \mathrm{Cov}(c_t^i, \hat{\pi}_t^i)$ as the individual component. For $i \neq j$, the cross term $\mathbb{E}[(c_t^i)^\top B_j c_t^j]$ does not appear in any single-agent regret. By Assumption 11.1, $\mathrm{Cov}(c_t^i, c_t^j) = \Lambda_i \Sigma_f \Lambda_j^\top$ (idiosyncratic components are cross-sectionally independent). Substituting $B_j c_t^j$ for the policy and applying the trace-cyclic identity gives (29). Summing over t yields (28). □

Remark 11.5 (Individual compliance ≠ systemic safety)**.** *Even if every agent individually satisfies the audit criterion* $\mathrm{Cov}(c_t^i, \hat{\pi}_t^i) \leq \varepsilon$*, the cross-agent term in* (28) *can be of order* $N(N-1) \cdot \mathrm{tr}(B_j \Lambda_j \Sigma_f \Lambda_i^\top)$*, which grows without bound as* $N \to \infty$*. A regulator auditing of firms one at a time misses this entirely. This is the* N*-agent analogue of the SR 11-7 deficiency identified in Section 1: individual model validation provides no systemic guarantee.*

11.3 The Fragility Multiplier and Phase Transition

The cross-agent term in (28) depends on how similar agents' strategies are. We characterize this via a scalar *fragility multiplier*.

Definition 11.6 (Strategy correlation and fragility multiplier)**.** *For symmetric agents (*$B_i = B$*,* $\Lambda_i = \Lambda$ *for all* i*), define the* systematic fraction of policy variance

$$\rho_{\mathrm{sys}} \;=\; \frac{\mathrm{tr}(B\Lambda\Sigma_f\Lambda^\top)}{\mathrm{tr}(B\Sigma_c)} \;\in\; [0,1],$$

where $\Sigma_c = \Lambda\Sigma_f\Lambda^\top + \sigma_\eta^2 I$ *is the total per-agent cost covariance. The* fragility multiplier *is*

$$M(N, \rho_{\mathrm{sys}}) \;=\; 1 + (N-1)\cdot\rho_{\mathrm{sys}}.$$

Corollary 11.7 (N^2 scaling of aggregate regret)**.** *Under the symmetric agent setting of Definition 11.6, the aggregate regret is*

$$\mathrm{Regret}_{\mathrm{agg}}^{(T)} \;=\; N\cdot T\cdot \mathrm{tr}(B\Sigma_c)\cdot\big[1+(N-1)\cdot\rho_{\mathrm{sys}}\big] \;=\; N\cdot \mathrm{Regret}_{\mathrm{individual}}^{(T)}\cdot M(N,\rho_{\mathrm{sys}}). \tag{30}$$

Three regimes follow immediately from (30)*:*

(i) ***Idiosyncratic limit*** *(*$\rho_{\mathrm{sys}} \to 0$*):* $M \to 1$*. Aggregate regret =* $N\times$ *individual regret. Diversification holds; individual audits are sufficient.*

(ii) ***Partial correlation*** *(*$0 < \rho_{\mathrm{sys}} < 1$*):* $M \in (1, N)$*. A fragility premium of* $(N-1)\cdot\rho_{\mathrm{sys}}\cdot T\cdot \mathrm{tr}(B\Sigma_c)$ *accrues beyond the sum of individual regrets. For* $N = 100$ *institutions each with* $\rho_{\mathrm{sys}} = 0.3$*, this premium is* 2,970% *of one firm's regret.*

(iii) ***Perfect correlation*** *(*$\rho_{\mathrm{sys}} \to 1$*):* $M \to N$*. Aggregate regret =* $N^2\times$ *individual regret. The system behaves as a single agent of size* N*: diversification disappears entirely and every agent bears the full systemic loss.*

The transition from regime (i) to (iii) is continuous in ρ_{sys}. We identify a natural critical threshold using random matrix theory.

Proposition 11.8 (Phase transition at the Marchenko–Pastur boundary)**.** *Consider the $N \times N$ pairwise-strategy covariance matrix Ψ_t, with the (i,j)-th entry $\psi_{ij} = \mathrm{Cov}(\hat{\pi}_t^i, \hat{\pi}_t^j)$. For large N, under the null hypothesis that all strategies are idiosyncratic ($\rho_{\mathrm{sys}} = 0$), the largest eigenvalue of Ψ_t converges to the Marchenko–Pastur upper edge*

$$\lambda_+ \;=\; \sigma_\pi^2\Big(1+\sqrt{N/T}\Big)^2 \;\approx\; \sigma_\pi^2 \cdot N/T \quad \textit{for } N/T \to c > 0$$

Following (Bai and Silverstein, 2010b), we can declare systemic fragility event *whenever $\lambda_{\max}(\hat{\Psi}_t) > \lambda_+$, signaling that common factor co-movement exceeds the level attributable to the finite-sample noise alone.*

Remark 11.9 (Connection to fire sales)**.** *In regime (iii), all agents hold proportional positions simultaneously. A liquidity shock that forces one agent to liquidate triggers liquidation by all others through two channels: (a)* direct contagion*, as falling prices cause margin calls and redemptions; and (b) the* regret channel *of this paper, since rising costs c_t force all agents to reduce $\hat{\pi}_t^i(c_t^i) = Bc_t^i$ in the same direction, compressing aggregate demand precisely when it is most needed. Equation (30), therefore, provides a quantitative bridge between the standard fire-sale narrative and the regret-based audit framework.*

11.4 An Observable Systemic Fragility Index

The regulator does not observe ρ_{sys} directly, but observes $\{(c_t^i, \hat{\pi}_t^i(c_t^i))\}$ for all i and t. We construct a consistent estimator of the fragility multiplier and derive a confidence interval for aggregate regret.

Definition 11.10 (Systemic fragility index)**.** *Given observed trajectories for all N agents over T periods, the* systemic fragility index *is*

$$\hat{F}_T \;=\; \frac{\|\hat{\Psi}_T\|_F^2}{N^2},$$

where $\hat{\Psi}_T$ is the $N \times N$ sample strategy-covariance matrix with (i,j)-th entry

$$\hat{\Psi}_{ij} \;=\; \frac{1}{T}\sum_{t=1}^{T}\big(\hat{\pi}_t^i - \bar{\pi}^i\big)^\top\big(\hat{\pi}_t^j - \bar{\pi}^j\big), \qquad \bar{\pi}^i = \frac{1}{T}\sum_{t=1}^{T}\hat{\pi}_t^i.$$

$\hat{F}_T \in [0,1]$: a value of 0 indicates perfectly idiosyncratic strategies; a value approaching 1 indicates near-identical strategies across all agents.

Theorem 11.11 (Consistent systemic audit estimator)**.** *Under Assumption 11.1, Assumption 3.2(c) for each agent, and the L^2-mixingale condition of Theorem 9.1 applied jointly across agents, as $T \to \infty$:*

$$\frac{1}{\sqrt{T}}\left(\widehat{\mathrm{Regret}}_{\mathrm{agg}}^{(T)} - \mathrm{Regret}_{\mathrm{agg}}^{(T)}\right) \;\xrightarrow{d}\; \mathcal{N}\big(0, \sigma_{\mathrm{SYS}}^2\big)\,, \tag{31}$$

where σ_{SYS}^2 is the long-run variance of the aggregate cross-product $\sum_i (c_t^i)^\top \hat{\pi}_t^i + \sum_{i\neq j}(c_t^i)^\top B_j c_t^j$, consistently estimated by a HAC estimator with bandwidth $h = \lfloor T^{1/3} \rfloor$. The estimator

$$\widehat{\mathrm{Regret}}_{\mathrm{agg}}^{(T)} \;=\; N \cdot \hat{C}_T \cdot \hat{M}, \qquad \hat{M} = 1 + (N-1)\hat{\rho}_{\mathrm{sys}},$$

where $\hat{C}_T$ is the per-agent trajectory estimator (9.1) and $\hat{M}$ is the estimated fragility multiplier, is computable in $O(T \cdot N^2 \cdot nd)$ time.

Remark 11.12 (Regulatory implementation)**.** *A macro-prudential regulator implementing Theorem 11.11 collects the N-vector of weight allocations $(\hat{\pi}_t^1, \ldots, \hat{\pi}_t^N)$ and factor returns c_t daily. The matrix $\hat{\Psi}_T$ can be updated online in $O(N^2 \cdot n)$ per day via Welford's algorithm (Remark 9.3). The fragility index $\hat{F}_T$ and its Marchenko–Pastur threshold (Proposition 11.8) are then reported alongside the per-firm covariance sums at each regulatory cycle. This architecture is compatible with existing SR 11-7 / SR-26-2 reporting and requires no access to any agent's internal optimization engine.*

11.5 Empirical Calibration: CRSP 2016–2025

To calibrate the fragility multiplier for realistic parameter values, we use the AR(1) estimates of Table 2 to bound ρ_{sys}. The momentum policy $\hat{\pi}_t(c_t) = Bc_t$ with $B = \alpha I$ and the common Fama–French market factor $f_t = r_{m,t} - r_{f,t}$ imply

$$\rho_{\text{sys}} \;=\; \frac{\alpha \cdot \text{Var}(\beta r_{m,t})}{\alpha \cdot \text{Var}(\beta r_{m,t}) + \alpha \cdot \sigma_\eta^2} \;=\; R^2_{\text{mkt}}, \tag{32}$$

where R^2_{mkt} is the cross-sectional R^2 of the market-factor regression. Using the CRSP 2016–2025 data of Section 10, R^2_{mkt} averages 0.27 in normal years but rises to 0.61 during the COVID crisis period ($\hat{\rho} = -0.144$), reflecting the sharp co-movement of individual stocks with the aggregate market during stress. Table 5 reports the implied fragility multipliers.

Table 5: Fragility multiplier $M(N, \rho_{\text{sys}}) = 1 + (N-1) \cdot \rho_{\text{sys}}$ for $N \in \{50, 100\}$ AI-driven allocators. ρ_{sys} approximated by R^2_{mkt} from CRSP data; crisis-period figures are upper bounds (see robustness notes in Section 10).

Period	$\hat{\rho}_{\text{sys}}$ (R^2_{mkt})	$M(N{=}50)$	$M(N{=}100)$
Normal (avg. 2016–2019, 2021)	0.27	14.2×	27.7×
COVID contraction (Jan–Mar 2020)	0.47	24.1×	47.5×
COVID crisis (Apr–Jun 2020)	0.61	30.8×	60.0×
2022 (near-zero autocorr.)	0.19	10.2×	19.8×

Two conclusions are immediate. First, even under normal conditions, aggregate regret for 100 correlated AI strategies is approximately 27 times the individual regret. Second, during crisis episodes, ρ_{sys} nearly doubles, pushing M to 60 for $N = 100$ institutions. This represents a form of *liquidity spiral*: stress raises correlation, which raises the fragility multiplier, which raises aggregate welfare loss, which deepens stress. The systemic audit framework of Theorem 11.11 is designed to detect the onset of this spiral before it fully develops.

11.6 Policy Implications for Systemic AI Regulation

Three policy implications follow from Corollary 11.7 and Theorem 11.11.

Aggregate position limits. If $\hat{F}_T$ exceeds the Marchenko–Pastur threshold of Proposition 11.8, a regulator can cap individual position sizes at

$$z_{\max} \;=\; \frac{\varepsilon}{N \cdot \rho_{\text{sys}} \cdot \|\bar{c}\|}$$

to maintain aggregate regret below a target ε. This is the systemic analog of single-firm leverage limits, derived here from first principles via the regret decomposition.

Strategy diversity requirements. Minimizing aggregate regret subject to fixed total capital is equivalent to minimizing $M(N, \rho_{\text{sys}})$, which is achieved by minimizing ρ_{sys}. This provides a formal welfare justification for regulatory requirements that AI-driven strategies use distinct factor-loading matrices Λ_i. These are analogous to capital diversification rules, but applied to algorithmic model design rather than balance-sheet composition.

Timing of intervention. Because $\hat{F}_T$ is updated online in $O(N^2 \cdot n)$ per day (Remark 11.12), the regulator observes the fragility multiplier in near-real time. A rising $\hat{F}_T$ preceding a volatility spike is an early warning of imminent systemic stress, complementing standard VaR-based indicators and satisfying the "forward-looking" mandate of SR 11-7.

11.7 Connection to Coordination Failures and Runs

The fragility multiplier provides a formal link between the regret framework and classic models of bank runs and coordination failures. Under perfect correlation ($\rho_{\text{sys}} = 1$), equation (30) collapses to

$$\text{Regret}_{\text{agg}}^{(T)} \;=\; N^2 \cdot T \cdot \text{tr}(B\Sigma_c). \tag{33}$$

This is the payoff structure of a coordination game: each agent individually minimizes regret, but collectively they produce quadratic aggregate loss, as in a prisoner's dilemma. The "run equilibrium" arises when a common cost shock causes all agents to simultaneously de-risk. De-risking becomes a policy that individually reduces each agent's $\text{Cov}(c_t^i, \hat{\pi}_t^i)$ but, in aggregate, drains market liquidity and raises c_t^i for all agents through price impact, precisely the mechanism of Corollary 6.4. Equation (33) therefore quantifies the welfare loss from the run equilibrium relative to the planner's optimum ($\rho_{\text{sys}} = 0$) in a closed form that is empirically estimable from observable data alone.

This connection is more than analogical. In the two-period version of the model, the game played by N agents is a linear-quadratic game with price-impact externality, and the Nash equilibrium regret is exactly $M(N, \rho_{\text{sys}})$ times the social optimum. The full game-theoretic mixed-strategy equilibria arise when agents have private information about the factor loading Λ_i. Deriving these equilibria is an important direction for future work.

12 Conclusion

A strategy's liquidity role is reflected in the covariance between its trades and the costs it faces, and that covariance is recoverable from trajectory data alone. This is the paper's central claim, and three further results follow from it.

The classification is exact and model-free. For any linear policy, $\text{Cov}(c_t, \hat{\pi}_t(c_t)) = \text{tr}(B\Sigma_c)$ (Proposition 2.2): the sign of a single statistic, computed from observed costs and decisions alone, places the strategy exactly into the Kyle (1985) dichotomy of liquidity demander or supplier. No order-flow sign, quote data, or knowledge of the strategy's signal is required. Momentum strategies are liquidity consumers, and contrarian strategies are liquidity providers as a matter of algebra, not estimation; the only object that must be estimated is the policy matrix B implicit in the observed trajectory, and even that estimation is unnecessary if one only wants the sign.

The same statistic recovers the effective spread. Under an AR(1) cost process, the auto-covariance correction to the covariance sum equals αs^2, the product of strategy size and the squared Roll (1984) implied spread (Theorem 7.7). This is not an analogy or an approximation: it is an

exact closed-form identity connecting a performance-attribution statistic to a market-microstructure quantity, derived from the same trajectory data and at no additional computational cost. Empirically, the implied spread recovered this way tracks realized illiquidity throughout the sample. The implied spread roughly triples during the COVID-19 liquidity crisis and collapses to its sample minimum during the 2022 rate-shock episode, when persistent common-factor repricing dominates short-horizon return dynamics (Section 10.2). The correction is smallest precisely when markets are deep and largest precisely when microstructure breaks down, which is the behavior one wants from an illiquidity proxy.

Aggregate liquidity imbalance produces a closed-form fire sale. Extending the single-agent result to N correlated strategies under a common-factor cost structure (Section 7.6) shows that the welfare loss from liquidity-balance failure scales as N^2 in the strategy-correlation coefficient (Corollary 7.20). When agents' liquidity demand and supply fail to net to zero, prices must adjust to absorb the imbalance. This happens under perfect strategy correlation, when many funds run the same factor-driven allocation. The resulting price-impact regret compounds the covariance regret multiplicatively rather than additively. This gives the familiar fire-sale mechanism a closed form computable from the same observable trajectories used for the single-agent classification, with no model of the deleveraging process itself required.

What the classifier recovers, and what it does not. The trajectory-based classifier of this paper is deliberately coarser than order-flow-based microstructure measures. The classifier requires only realized costs (asset prices) and position sizes, not signed trades, quote revisions, or limit-order-book depth. The classifier recovers a strategy's *average* liquidity role over the estimation window rather than its behavior at any single trade. This is a feature when the internal logic of the strategy is unobservable, the setting motivating this paper. When finer-grained order-flow data is available, standard measures (Hasbrouck, 1991) or trade-classification algorithms will recover intraday dynamics that this statistic cannot. The two approaches are complements rather than substitutes. The Roll-spread identity of Theorem 6.3 gives an explicit bridge between the two approaches, expressing a daily-frequency, trajectory-based quantity in the same units as a standard microstructure spread estimator, so that one can validate it against the other on data where both are available. We view extending this bridge to intraday frequencies, and benchmarking the implied spread directly against effective and realized spread measures computed from TAQ data, as the natural next step for this line of work.

A closing remark on horizon. The discounted identity of Theorem 3.8 shows that a single-period covariance estimate, scaled by the effective horizon $1/(1-\gamma)$, measures long-run discounted regret directly. The same scaling applies to the liquidity classification: a strategy's liquidity role, once identified from a short window of trajectory data, extrapolates to its long-run discounted contribution to market liquidity without re-estimation at every horizon. This makes the single covariance statistic a practical tool not only for one-off classification but also for monitoring how a strategy's liquidity role evolves as market conditions change.

References

Aldridge, Irene, "Regret Equals Covariance: A Closed-Form Characterization for Stochastic Optimization," *arXiv:2605.14019 [econ.EM]*, 2026.

Almgren, Robert and Neil Chriss, "Optimal Execution of Portfolio Transactions," *J. Risk*, 2001, *3*, 5–39.

Bai, Zhidong and Jack W. Silverstein, *Spectral Analysis of Large Dimensional Random Matrices*, 2nd ed., New York: Springer, 2010.

_ and _ , *Spectral Analysis of Large Dimensional Random Matrices*, 2nd ed., Springer, 2010.

Dai, Wei, Mamdouh Medhat, Robert Novy-Marx, and Savina Rizova, "Reversals and the Returns to Liquidity Provision," *Fin. Analysts J.*, 2024, *80* (2), 122–151.

Elmachtoub, Adam N. and Paul Grigas, "Smart "Predict, then Optimize"," *Mgmt. Sci.*, 2022, *68* (1), 9–26.

Engle, Robert F., "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 1982, *50* (4), 987–1007.

Fama, Eugene F. and Kenneth R. French, "A Five-Factor Asset Pricing Model," *J. Fin. Econ.*, 2015, *116* (1), 1–22.

Gatheral, Jim, "No-Dynamic-Arbitrage and Market Impact," *Quant. Finance*, 2010, *10* (7), 749–759.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Hasbrouck, Joel, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 1991, *46* (1), 179–207.

Hazan, Elad, "Introduction to Online Convex Optimization," *Foundations and Trends in Optimization*, 2016, *2* (3–4), 157–325.

Jegadeesh, Narasimhan, Jiang Luo, Avanidhar Subrahmanyam, and Sheridan Titman, "Short-Term Reversals and Longer-Term Momentum around the World," *Rev. Fin. Stud.*, 2025, *38* (12), 3673–3728.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

Ledoit, Olivier and Michael Wolf, "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *J. Multivariate Anal.*, 2004, *88* (2), 365–411.

Mamais, Panagiotis, "Explaining and Predicting Momentum Performance Shifts Across Time and Sectors," *J. Forecast.*, 2025. doi:10.1002/for.3232.

McLeish, D. L., "A Maximal Inequality and Dependent Strong Laws," *Annals of Probability*, 1975, *3* (5), 829–839.

Nagel, Stefan, "Evaporating Liquidity," *Rev. Fin. Stud.*, 2012, *25* (7), 2005–2039.

Roll, Richard, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *J. Finance*, 1984, *39* (4), 1127–1139.

Shleifer, Andrei and Robert W. Vishny, "Liquidation Values and Debt Capacity: A Market Equilibrium Approach," *Journal of Finance*, 1992, *47* (4), 1343–1366.

Welford, B. P., "Note on a Method for Calculating Corrected Sums of Squares and Products," *Technometrics*, 1962, *4* (3), 419–420.

Zinkevich, Martin, "Online Convex Programming and Generalized Infinitesimal Gradient Ascent," in "Proceedings of the Twentieth International Conference on Machine Learning (ICML)" 2003, pp. 928–936.

Appendix

A Additional corollaries

Remark A.1 (Corollaries). *(i) Stationary policy: if $\hat{\pi}_t = \hat{\pi}$ for all t, then* $\mathrm{Regret}^{(T)} = T \cdot \mathrm{Cov}(c, \hat{\pi}(c))$, *growing linearly in T, matching the rate known from online learning. Policy selection based on a single-period covariance diagnostic is, therefore, sufficient for horizon-T optimization.*

(ii) Effective horizon: Regret_γ *equals $T_{\mathrm{eff}} = (1-\gamma)^{-1}$ periods of undiscounted regret. For daily rebalancing with $\gamma = 0.99$, $T_{\mathrm{eff}} = 100$ trading days. (iii) Regime-switching: in a two-regime Markov-switching model with transition matrix P, $|\Delta_T| \le T\|\bar{c}^{(1)} - \bar{c}^{(2)}\| \cdot \rho(P)^{1/2} \cdot \sup_t \|\mathbb{E}[\hat{\pi}_t \mid \mathcal{F}_{t-1}] - \pi_t^*\|$, so fast-mixing regimes (small $\rho(P)$) permit (2) as an approximation.*

B Proof of Theorem 6.3 (AR(1) multi-period regret)

Proof. **Step 1: Per-period and cross-period costs.** By Remark 6.2, $\mathrm{Regret}^{(T)} = \sum_t \mathbb{E}[c_t^\top (Bc_t + b)]$.

Inline claim 1 (on-diagonal): $\mathbb{E}[c_t^\top B c_t] = \mathrm{tr}(B\Sigma_c)$. *Proof*: since $\mathbb{E}[c_t] = 0$, $\mathbb{E}[c_t^\top B c_t] = \mathbb{E}[\mathrm{tr}(Bc_t c_t^\top)] = \mathrm{tr}(B\,\mathbb{E}[c_t c_t^\top]) = \mathrm{tr}(B\Sigma_c)$; and $\mathbb{E}[b^\top c_t] = 0$. Summing over t gives the raw covariance sum $\sum_t \mathbb{E}[c_t^\top (Bc_t + b)] = T \cdot \mathrm{tr}(B\Sigma_c)$.

Inline claim 2 (off-diagonal): for $t \neq s$, $\mathbb{E}[c_t^\top B c_s] = \mathrm{tr}(BA^{|t-s|}\Sigma_c)$. *Proof*: by iterating the AR(1) recursion, $c_t = A^{t-s}c_s + \sum_{j=0}^{t-s-1} A^j \varepsilon_{t-j}$ for $t > s$, so $\mathrm{Cov}(c_t, c_s) = A^{t-s}\Sigma_c$ (the innovation terms are independent of c_s). Then $\mathbb{E}[c_t^\top B c_s] = \mathrm{tr}(B\,\mathbb{E}[c_s c_t^\top]) = \mathrm{tr}(B(A^{t-s}\Sigma_c)^\top)$; symmetry of Σ_c and the cyclic trace property give $\mathrm{tr}(B(A^{t-s}\Sigma_c)^\top) = \mathrm{tr}(BA^{t-s}\Sigma_c)$, establishing Inline claim 2.

Step 2: AR(1) auto-covariance correction. The raw sum $T \cdot \mathrm{tr}(B\Sigma_c)$ treats each period as independent. Because the policy weight Bc_t covaries with future costs through the AR(1) propagation $c_s = A^{s-t}c_t + \text{noise}$ $(s > t)$, we subtract the cross-period interaction:

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\sum_{s=t+1}^{T} c_s^\top B A^{s-t} c_t\right] &= \sum_{t=1}^{T}\sum_{s=t+1}^{T} \mathrm{tr}(BA^{s-t}\Sigma_c) \quad \text{(Inline claim 2)} \\ &= \sum_{t=1}^{T}(T-t)\,\mathrm{tr}(BA^t\Sigma_c),\end{aligned}$$

where the last equality reparametrizes by $k = s - t$ and uses stationarity of the autocovariance structure.

Step 3: Assemble. Combining Steps 1 and 2 yields (10). □

C Proof of Proposition (Computational complexity and end-to-end audit cost)

Proof. Mean computation costs $O(T(n+d))$ (one pass each for $\bar{c}$ and $\bar{\pi}$); forming each summand $(c_t - \bar{c})^\top(\hat{\pi}_t - \bar{\pi})$ costs $O(nd)$ and summing over T terms gives $O(T \cdot nd)$; the Newey-West estimator with $h = \lfloor T^{1/3} \rfloor$ lags costs $O(T \cdot T^{1/3}) = O(T^{4/3}) \le O(T \cdot nd)$ whenever $nd \ge T^{1/3}$, which holds in any multi-asset setting. The end-to-end bound follows since each audit component is individually $O(T \cdot nd)$ and the maximum dominates. □