The Kyle's view on crowded factor strategies

Junhwi Kim and Hyeonjun Kim

Department of Economics, Sungkyunkwan University

May 2026

Factor backtests almost always treat the investor as a price-taker — they execute at the closing print, and the only friction subtracted is a commission. In live trading the investor is anything but a price-taker: she is one of N imperfectly informed traders submitting correlated orders to a market maker who reads the flow. Recasting the factor investor as a noisy informed trader in a Kyle (1985) auction gives a single, closed-form summary of what is happening — an *erosion ratio* $\mathcal{E} = (1+(N-1)\phi)/(2+(N-1)\phi)$ that depends only on signal informativeness ϕ and the number of competing investors N. The same expression, fit to ten-year exponentially weighted factor returns over 1950–2025 (data from Jensen et al., 2023), sorts the eight canonical equity factors cleanly — and out-of-sample, calibrating only on pre-2010 data, beats the unconditional historical mean on six of eight factors and a linear time-trend forecast on six of eight, with average RMSE ratios of 0.88 and 0.77 respectively.

KEYWORDS Kyle model, price impact, factor investing, crowding, execution cost

1 What backtests quietly assume

A factor backtest produces gross returns. The footnote noting this is usually honest. What the footnote does not say — and what the reader is left to translate on her own — is that the gap between the gross number and what a desk actually realises is neither a fixed haircut nor an exogenous friction. It is the equilibrium consequence of trading against a market maker who reads the order flow.

The standard treatment of execution cost in the empirical factor literature is to estimate spread and impact from historical data, subtract them from gross returns, and report the difference (Novy-Marx and Velikov, 2016; Frazzini et al., 2014; Korajczyk and Sadka, 2004). That work is careful and useful. What it does not do is ask why those costs are what they are, or whether they change with the number of investors pursuing the same signal. They are read off the tape; they are not derived.

This is the gap we want to close. Once you grant that a factor investor is just an imperfectly informed trader in the sense of Kyle (1985), her execution cost is no longer a parameter to be estimated — it is an equilibrium object that depends on how many other people share her signal and how informative that signal is. A single ratio, derivable in a page of algebra, summarises the result.

2 A one-equation reading

2.1 *The setup*

Take a single risky asset with fundamental value $v \sim \mathcal{N}(p_0, \sigma_v^2)$. There are $N \geq 1$ factor investors, each observing a private signal $s_i = v + \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ independent across investors. Call the signal informativeness $\phi = \sigma_v^2/(\sigma_v^2 + \sigma_\varepsilon^2) \in (0, 1]$. Noise traders submit $u \sim \mathcal{N}(0, \sigma_u^2)$. A competitive market maker observes only the aggregate flow $y = \sum_i x_i + u$ and prices the asset at $p = \mathbb{E}[v \mid y] = p_0 + \lambda y$.

This is the multi-trader Kyle model of Kyle (1989), with the wrinkle that signals are noisy. We look for a symmetric linear equilibrium in which each investor uses a strategy of the form $x_i = \beta(s_i - p_0)$. The derivation is standard — take the first-order condition on the investor's expected profit, then impose Bayesian consistency of λ with the aggregate order flow. The fixed point yields

$$\lambda = \frac{\sigma_v \sqrt{\phi N}}{\sigma_u \Gamma}, \qquad \beta = \frac{\phi}{\lambda \Gamma}, \qquad \Gamma \equiv 2 + (N-1)\phi. \tag{1}$$

2.2 The erosion ratio

Define gross expected alpha for investor i as the revenue she would earn if she submitted her equilibrium order but the market maker did not move the price: $\Pi_i^{\text{gross}} = \beta\,\sigma_v^2$. The realised expected profit, with price impact in place, is

$$\mathbb{E}[\pi_i] \;=\; \frac{\sigma_v \sigma_u \sqrt{\phi}}{\sqrt{N}\,\Gamma}. \tag{2}$$

The share of gross alpha lost to price impact is then

$$\boxed{\mathcal{E} \;\equiv\; \frac{\Pi_i^{\text{gross}} - \mathbb{E}[\pi_i]}{\Pi_i^{\text{gross}}} \;=\; \frac{1+(N-1)\phi}{2+(N-1)\phi}} \tag{3}$$

The expression depends on only two things: signal informativeness ϕ and the number of competing investors N. Four observations:

1. **Kyle benchmark.** When $N = 1$, $\mathcal{E} = 1/2$ regardless of ϕ. A monopolistic informed trader always surrenders exactly half her gross alpha — the standard Kyle result.
2. **Crowding monotonicity.** For fixed $\phi > 0$, $\mathcal{E}$ is increasing in N and approaches one as $N \to \infty$. Competition among informed traders drives net alpha to zero.
3. **Signal-quality paradox.** For fixed $N > 1$, $\mathcal{E}$ is increasing in ϕ. Better signals make order flow more informative, which makes it easier for the market maker to price it. In a crowded strategy, refining the signal raises gross alpha but raises the erosion rate faster.
4. **Concave externality.** The marginal erosion from the N-th entrant is ϕ/Γ^2, strictly decreasing in N. The first few imitators inflict most of the damage; the rest is rounding error.

What we want to flag here is not a new equilibrium — the symmetric linear case is contained in Kyle (1989) and Foster and Viswanathan (1996). What is new is the recasting: $\mathcal{E}$ is a two-parameter summary that maps directly onto questions a desk actually asks. How crowded is this trade? How much edge do we think we have? Plug in and read off the share you can expect to keep.

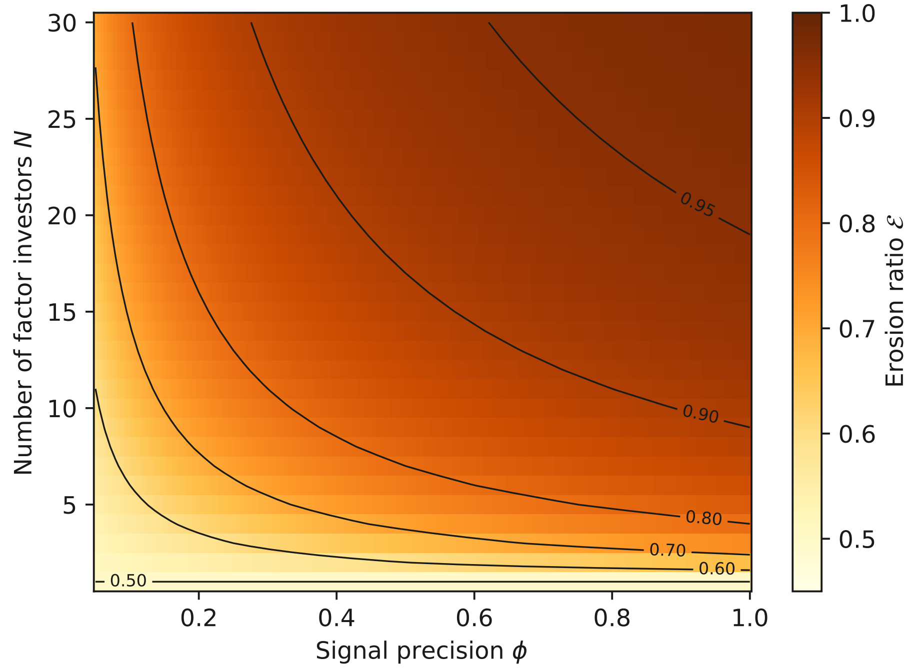

Figure 1. Erosion ratio $\mathcal{E}(\phi, N)$. The lower-left corner — modest signal precision, few competitors — is where alpha actually survives. The upper-right — well-known factors pursued at scale — is where it does not. Note how compressed the "alive" region is.

3 What the data say

The model predicts that observed factor returns should decline over time as N rises: in our notation, $r(t) \propto 1/\Gamma(t) = 1/(2 + \kappa(t))$, where $\kappa(t) \equiv (N(t)-1)\phi$ is what we will call the *crowding intensity*. The premium is constant before crowding bites and decays toward zero afterwards. Conveniently, κ is also the only thing the erosion ratio depends on: $\mathcal{E} = (1+\kappa)/(2+\kappa)$. We do not need to separately identify N and ϕ from return data — the data can speak only to their product.

3.1 Procedure

We take monthly value-weighted long-short returns from the Jensen et al. (2023) global factor dataset — 153 U.S. equity factors over 1926–2025 — and pull out eight canonical ones: value (book-to-market), momentum (12–1), size, quality (QMJ), profitability (gross profit/assets), investment (asset growth), low beta, and short-term reversal. For each factor we smooth the monthly return with an exponentially weighted moving average

(EWM, halflife 120 months, annualised) and fit

$$r(t) \;=\; \frac{2\,r_0}{2+\kappa(t)} + \text{noise} \tag{4}$$

by nonlinear least squares under three parametric shapes for $\kappa(t)$: linear-onset, squared linear-onset ("quadratic"), and sigmoid. The best fit per factor is chosen by BIC. We use EWM rather than a flat 120-month window because the latter weights a return from 1985 the same as one from 2024; the EWM downweights distant history smoothly while preserving the same effective 10-year window, which matters most when the data we care about — the post-2000 fade — is at the right edge of the sample.

A note on what *linear, quadratic, sigmoid* refer to. These labels describe the shape of $\kappa(t)$, not $r(t)$. Since $r(t) = 2r_0/(2+\kappa(t))$, a κ that grows linearly in time produces an $r(t)$ that decays as a hyperbola, not as a straight line. This is why every fitted curve in Figure 2 below looks like a smooth decreasing arc, including for factors where the BIC selects the linear-κ family. There is no bug here — it is exactly what $1/(2+\text{linear})$ looks like.

A separate note on the quadratic shape. The intended form is $\kappa(t) = b\,(\max\{0,\, t-t_0\})^2$ — zero until t_0, then growing as a squared linear-onset. A first cut at the code used $\max(0,\, b(t-t_0)^2)$, which is the same thing inverted: because $(t-t_0)^2$ is already non-negative, the $\max$ does nothing, and the resulting κ is a symmetric parabola centred at t_0. That makes $r(t)$ a bump — peak at t_0, decay on *both* sides — which is not what we want. The fits below use the corrected form.

3.2 Results

Three tiers emerge from Table 1. Value and low beta sit at the top — $\mathcal{E}_T \approx 0.99$–$1.00$, their pre-crowding premia compressed almost entirely. Onset for both lands in the late 1980s to early 1990s for low beta and post-2010 for value (under the quadratic-κ fit), the decades in which index funds, smart-beta wrappers, and systematic allocators explicitly targeted these characteristics. Momentum, size, investment, and reversal sit in a middle band ($\mathcal{E}_T \approx 0.78$–$0.83$): visible decay but still nonzero residual premia. Quality and profitability come out lowest ($\mathcal{E}_T = 0.67$ and 0.73 respectively) — the factors closest to the Kyle benchmark and so the ones the model treats as still paying. The EWM smoother weights the post-2000 fade more heavily than a flat ten-year window does, which mainly affects size and quality: size moves from "essentially dead" under flat rolling means to "decayed but alive" under EWM, and quality drifts off the Kyle floor.

Table 1. Time-series calibration of crowding intensity for eight U.S. equity factors, 1926–2025 (target: EWM monthly returns, halflife 120 months). $\hat{r}_0$ is the pre-crowding annual premium implied by the fit; $\hat{\kappa}_T$ is the crowding intensity as of 2025; "Onset" is t_0 (linear-κ, quadratic-κ) or t_{mid} (sigmoid-κ); $\mathcal{E}_T$ is the implied 2025 erosion ratio. Best κ-family chosen per factor by BIC.

Factor	κ-family	$\hat{r}_0$ (%)	$\hat{\kappa}_T$	Onset	$\mathcal{E}_T$
Value	Quadratic	3.0	89.7	2013	0.99
Momentum	Linear	10.7	3.1	1961	0.80
Size	Sigmoid	3.3	4.0	1990	0.83
Quality	Sigmoid	3.8	1.0	1975	0.67
Profitability	Sigmoid	4.8	1.7	1974	0.73
Investment	Quadratic	5.8	2.5	1984	0.78
Low Beta	Quadratic	1.1	579.8	1992	1.00
Reversal	Quadratic	5.1	3.8	1974	0.83

3.3 Out-of-sample evaluation

The in-sample exercise above fits the model on the entire 1926–2025 window, which is informative about *whether* the functional form can describe the data but is silent on *whether it forecasts*. The genuine test is to hold out the recent period and ask whether a model calibrated only on history would have anticipated what came next. We split each factor's smoothed series at January 2010 — a natural boundary, post-quants-crisis and post-GFC, and a date by which the academic crowding story was already well established (McLean and Pontiff, 2016; Chordia et al., 2014) — fit on the train portion only, and forecast the post-2010 EWM target.

We compare three forecasters:

1. **Kyle-tax** — our $r(t) = 2r_0/(2+\kappa(t))$ with the train-set BIC-best κ-family.
2. **Historical mean** — a constant equal to the pre-2010 EWM average, the natural null for return-forecasting exercises in this tradition.
3. **Linear time-decay** — $r(t) = a + b\,t$ fit by OLS on the train portion; the simplest non-constant alternative.

Performance is summarised by the test-period RMSE for each forecaster and by Campbell–Thompson out-of-sample R^2, $R^2_{\text{OS}} = 1 - \mathrm{RMSE}^2_{\text{Kyle}}/\mathrm{RMSE}^2_{\text{bench}}$, against each benchmark.

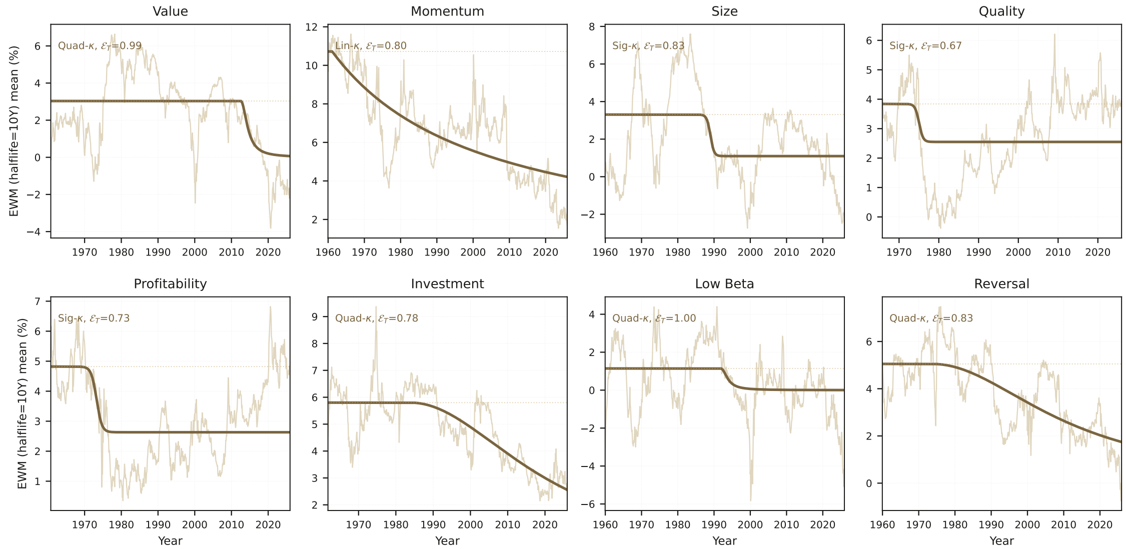

Figure 2. EWM (halflife=10Y) factor returns (light) and fitted crowding-decay curve (dark). Each panel reports the best-BIC κ-family (Lin-κ, Quad-κ, Sig-κ) and the implied terminal erosion ratio. The horizontal dotted line marks $\hat{r}_0$. Every fitted curve is a smooth arc because $r(t) = 2r_0/(2+\kappa(t))$ inverts κ; the panel annotation refers to the shape of κ, not of the curve drawn.

The Kyle-tax model beats the historical-mean null on six of eight factors and the linear-decay alternative on six of eight, with cross-factor mean RMSE ratios of 0.88 and 0.77 respectively. Where the model wins it wins decisively: size ($R^2_{\text{OS}} = 0.79$), momentum (0.34), reversal (0.52), value (0.29). The two factors where the Kyle model *loses* to the historical mean — quality and profitability — are precisely the two the in-sample fit flagged as not-yet-decayed; the model is forecasting continued fade and the data is not delivering one, so the constant null wins by default. We do not consider this a bug. A model that says "erosion has bitten this factor" should predict worse performance going forward, and on quality and profitability that prediction is being falsified in real time — which is exactly what the in-sample $\mathcal{E}_T$ values for those two factors were already telling us.

Two further notes on interpretation. First, the comparison is genuinely out-of-sample: the EWM smoother at each post-2010 month uses only data up to that month, but the parameters being fit (r_0, κ-family, onset, rate) are frozen at the 2009-12 boundary. Second, the linear-decay benchmark is harder to beat than it might look — it is a flexible two-parameter alternative that, on data that genuinely fades, can mimic the Kyle-tax forecast over short horizons. That the Kyle-tax form still wins on six factors and by 23% in average RMSE is, we think, the stronger of the two results.

Table 2. Out-of-sample evaluation. Train: 1926–2009 (EWM target). Test: 2010–2025. RMSE in percentage points of annualised return. R^2_{OS} is Campbell–Thompson, positive meaning the Kyle-tax model beats the benchmark. Bold marks positive R^2_{OS}.

Factor	κ-family	$\text{RMSE}_{\text{Kyle}}$	$\text{RMSE}_{\bar{r}}$	RMSE_{lin}	R^2_{OS} vs $\bar{r}$	R^2_{OS} vs lin
Value	Sigmoid	2.82%	3.36%	4.03%	**0.29**	**0.51**
Momentum	Linear	2.55%	3.14%	3.85%	**0.34**	**0.56**
Size	Sigmoid	1.75%	3.80%	2.25%	**0.79**	**0.40**
Quality	Sigmoid	2.09%	1.68%	1.48%	-0.54	-0.99
Profitability	Sigmoid	2.37%	1.69%	3.40%	-0.98	**0.51**
Investment	Sigmoid	1.71%	2.39%	1.53%	**0.49**	-0.25
Low Beta	Quadratic	1.52%	1.67%	4.75%	**0.18**	**0.90**
Reversal	Sigmoid	1.64%	2.38%	3.40%	**0.52**	**0.77**
cross-factor mean					0.13	0.30

Two caveats are worth stating plainly. First, κ is a composite. We cannot separate "many investors with weak signals" from "few investors with strong signals" using returns alone — only their product matters. Second, factor returns decline for reasons other than informed-trader entry: data-mining corrections, structural changes, macro regime shifts. The model is silent on these, and the fits attribute all systematic decay to crowding. That is a strong assumption. What the exercise establishes is not the size of the crowding effect in isolation, but the fact that a remarkably parsimonious functional form — one parameter for the timing of decay, one for its rate, plus the level — is enough to sort the canonical factors into the tiers most practitioners already feel they are in.

4 What to take away

A few things, stated plainly:

1. The relevant "cost" of running a factor strategy is not a commission. It is the share of gross alpha the market maker extracts because she can read aggregate order flow. That share has a closed-form expression in terms of two quantities a practitioner already has rough numbers for.
2. Improving signal quality is not unambiguously good. In a crowded trade, a more precise

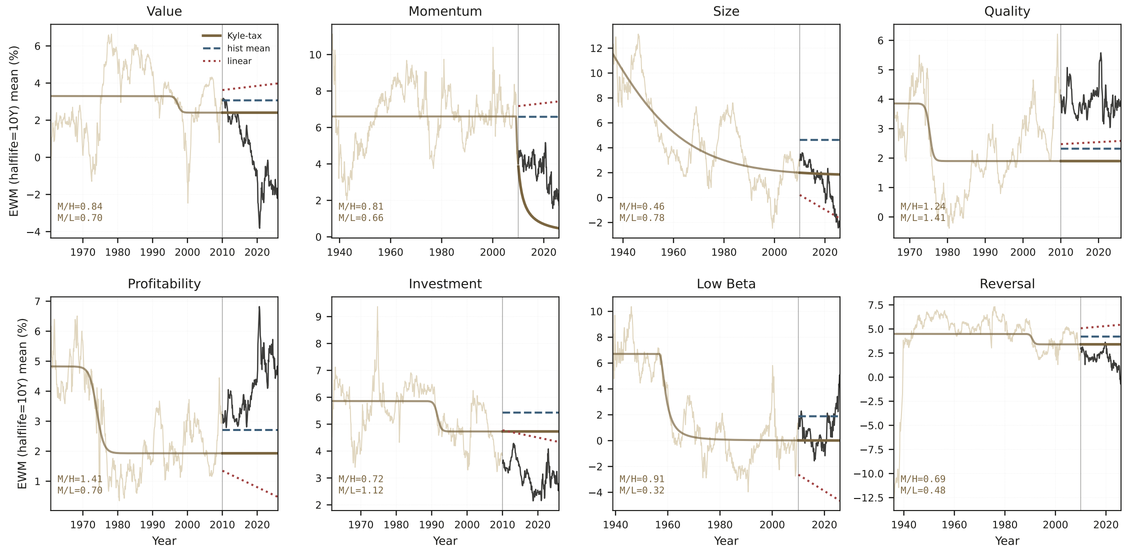

Figure 3. Out-of-sample forecasts versus realised EWM target. Light line: train target (1926–2009). Dark line: test target (2010–2025). Gold: Kyle-tax train fit (faint) and OOS forecast (solid). Blue dashed: historical-mean benchmark. Red dotted: linear time-decay. The vertical line marks the train/test split. Annotation `M/H` and `M/L` are RMSE ratios of the Kyle-tax model to the historical-mean and linear-decay benchmarks.

signal makes the order flow more legible to the market maker, and the erosion rate rises with ϕ. Past a turning point, refining the signal reduces net profit.

3. The damage from crowding is concave. The first few imitators take most of the surplus; the marginal twentieth entrant changes very little. This is why "early to a factor" matters so much and "arriving late to a popular factor" is so unrewarding — it is not a smooth gradient.
4. The cross-factor data are consistent with the mechanism. The factors investors found first, that were easiest to encode, that ETF wrappers were built around — value, size, low beta — have lost essentially all their premium. The factors that arrived later or remained harder to implement well — quality, profitability — still pay.

None of this proves crowding is the only force at work. It says that one equation, derived from a one-shot Kyle auction, is enough to organise what we see in seventy-five years of factor returns.

References

Tarun Chordia, Avanidhar Subrahmanyam, and Qing Tong. Have capital market anomalies attenuated in the recent era of high liquidity and trading activity? *Journal of Accounting and Economics*, 58(1):41–58, 2014.

F. Douglas Foster and S. Viswanathan. Strategic trading when agents forecast the forecasts of others. *Journal of Finance*, 51(4):1437–1478, 1996.

Andrea Frazzini, Ronen Israel, and Tobias J. Moskowitz. Trading costs of asset pricing anomalies. *Working Paper*, 2014.

Theis Ingerslev Jensen, Bryan Kelly, and Lasse Heje Pedersen. Is there a replication crisis in finance? *Journal of Finance*, 78(5):2465–2518, 2023.

Robert A. Korajczyk and Ronnie Sadka. Are momentum profits robust to trading costs? *Journal of Finance*, 59(3):1039–1082, 2004.

Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

Albert S. Kyle. Informed speculation with imperfect competition. *Review of Economic Studies*, 56(3):317–355, 1989.

R. David McLean and Jeffrey Pontiff. Does academic research destroy stock return predictability? *Journal of Finance*, 71(1):5–32, 2016.

Robert Novy-Marx and Mihail Velikov. A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1):104–147, 2016.