Microprice as a Centered Poisson Corrector

Paper I: move-time structure, centering, and horizon effects

Lev Petersen *April 17, 2026*

Abstract. Classical microprice propagates expected future price changes through an absorbing top-of-book Markov chain. We reduce the event-time model to price-move time: for no-move kernel Q, move kernel T, one-step drift r, and fundamental matrix $N = (I - Q)^{-1}$, the pair $G_1 := Nr$, $B := NT$ give the one-move drift and the move-time kernel, and the n-move predictor is $G^{(n)} = \sum_{j=0}^{n-1} B^j G_1$. If B is ergodic with stationary π, mean move drift $\mu = \pi^\top G_1$, and projector $\Pi_B := \mathbf{1}\pi^\top$, then $G^{(n)} = n\mu\,\mathbf{1} + g - B^n g$ where $(I - B)g = G_1 - \mu\mathbf{1}$, $\pi^\top g = 0$, and $G^* := (I - B + \Pi_B)^{-1} G_1 = g + \mu\mathbf{1}$. Classical microprice is a centered Poisson corrector on the move-time chain, and the transient $-B^n g$ governs the horizon effect. The move-time resolvent $(I - \beta B)^{-1} G_1$ regularises G^* by the constant mode. Empirically, on three public panels (Stoikov BAC/CVX, WSELOB, one-day IEX), the correction performs as predicted: short-horizon errors are transient-dominated, the centered limit wins once the horizon is long enough and the move-time martingale diagnostic $\widehat{t}_M$ stays below 2, and a large $\widehat{t}_M$ is a reliable warning that discounted extrapolation will be driven by residual drift rather than by the bounded state correction.

1 Introduction

Stoikov's microprice [1] propagates expected future price changes through an absorbing top-of-book Markov chain. This paper identifies the long-run object computed by that construction and explains why that object need not be the best short-horizon forecast. Throughout we use the dictionary

$$r(x) := \mathbb{E}[p_{n+1} - p_n \mid X_n = x], \quad N := (I - Q)^{-1},$$

$$G_1 := Nr, \qquad B := NT,$$

and, when B is ergodic with stationary distribution π, set $\mu := \pi^\top G_1$, $\Pi_B := \mathbf{1}\pi^\top$, and

$$(I - B)g = G_1 - \mu\mathbf{1}, \qquad \pi^\top g = 0,$$

$$G^* := (I - B + \Pi_B)^{-1} G_1 = g + \mu\mathbf{1}.$$

Here g is the centered move-time relative value and G^* the uncentered long-run level; when the chain is centered, $\mu = 0$ and $G^* = g$. A companion note [2] uses the same primitives to cast microprice as a reinforcement-learning value function and add event-time Bellman objectives.

Remark 1.1 (Notation). X_n is the event-time chain, p_n its mid-price, $P = Q + T$ the event-time kernel split into no-move and move parts, τ_k the k-th price move time, and $Y_k := X_{\tau_k}$ the embedded move-time chain. The symbols (r, N, G_1, B) are the move-time primitives; $(\pi, \mu, \Pi_B, g, G^*)$ their ergodic accompaniments. β is reserved for the move-time resolvent $(I - \beta B)^{-1} G_1$; γ for the event-time resolvent $(I - \gamma P)^{-1} r$ of the companion note.

2 From event time to move time

Let X_n be a finite-state Markov approximation of the top of book at event time. Decompose the transition kernel as $P = Q + T$, where Q_{xy} is the no-move probability and T_{xy} the move probability, and let $r(x) := \mathbb{E}[p_{n+1} - p_n \mid X_n = x]$. Assume $\rho(Q) < 1$, so $N = (I - Q)^{-1} = \sum_{k \geq 0} Q^k$ exists. Set $G_1 := Nr$ and $B := NT$. Figure 1 summarises the reduction.

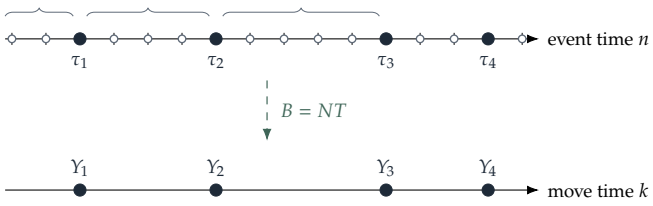

Figure 1. The embedded move-time chain. Event-time transitions (top) alternate no-move steps (open circles, absorbed into $N = (I - Q)^{-1}$) and price-move steps (filled). The move times τ_k define the embedded chain $Y_k = X_{\tau_k}$ (bottom) with kernel $B = NT$.

Proposition 2.1 (Embedded move-time chain)**.** *With $\tau_0 := 0$ and $\tau_{k+1} := \inf\{n > \tau_k : p_n \neq p_{n-1}\}$, $Y_k := X_{\tau_k}$ is a Markov chain with transition kernel B, and $G_1(x) = \mathbb{E}_x[p_{\tau_1} - p_{\tau_0}]$ is the expected price change at the next move given the current state.*

Proof. The first price move occurs after n no-move steps followed by one move with kernel $Q^n T$; summing over $n \geq 0$ gives $\mathbb{P}_x(Y_1 = y) = ((I - Q)^{-1} T)_{xy} = B_{xy}$. Similarly $\mathbb{E}_x[p_{\tau_1} - p_{\tau_0}] = ((I - Q)^{-1} r)_x = G_1(x)$. The Markov property of (Y_k) is the strong Markov property at (τ_k). □

Remark 2.2 (Centering). When B is ergodic, $B^n \to \Pi_B$. Stoikov's convergence condition $B^n G_1 \to 0$ is therefore equivalent to $\pi^\top G_1 = 0$: only the invariant constant mode can obstruct convergence.

3 The centered Poisson corrector

Define the n-move price adjustment

$$G^{(n)} := \sum_{j=0}^{n-1} B^j G_1, \quad \mathbb{E}[p_{\tau_n} - p_{\tau_0} \mid Y_0 = x] = G^{(n)}(x).$$

Proposition 3.1 (Finite-horizon decomposition)**.** *If B is ergodic with stationary π and $h := G_1 - \mu\mathbf{1}$, the unique centered solution of $(I - B)g = h$, $\pi^\top g = 0$, satisfies*

$$G^{(n)} = n\mu\,\mathbf{1} + g - B^n g \quad \textit{for every } n \geq 1.$$

With $G^* := (I - B + \Pi_B)^{-1} G_1$, $G^* = g + \mu\mathbf{1}$; when $\mu = 0$, $G^* = g$ and $G^{(n)} \to G^*$.

Proof. Write $G_1 = \mu\mathbf{1} + h$ with $\pi^\top h = 0$. Since $h = (I - B)g$, $G^{(n)} = \sum_{j=0}^{n-1} B^j G_1 = n\mu\,\mathbf{1} + \sum_{j=0}^{n-1} B^j (I-B) g = n\mu\,\mathbf{1} + g - B^n g$. Also $(I - B + \Pi_B)(g + \mu\mathbf{1}) = h + \mu\mathbf{1} = G_1$, so $G^* = g + \mu\mathbf{1}$. If $\mu = 0$, $B^n g \to \Pi_B g = \mathbf{1}\pi^\top g = 0$, hence $G^{(n)} \to G^*$. □

Figure 2 illustrates the decomposition: a linear drift $n\mu$, a bounded state correction g, and a transient $-B^n g$ whose decay produces the empirical horizon effect.

Corollary 3.2 (Move-time martingale correction)**.** *Under the assumptions of Proposition 3.1, $M_k := p_{\tau_k} - k\mu + g(Y_k)$ is a move-time martingale. When $\mu = 0$, the corrected price $p_{\tau_k} + G^*(Y_k)$ is itself a move-time martingale.*

Proof. From $(I - B)g = G_1 - \mu\mathbf{1}$, $\mathbb{E}[g(Y_{k+1}) - g(Y_k) \mid Y_k] = -(G_1(Y_k) - \mu)$, while $\mathbb{E}[p_{\tau_{k+1}} - p_{\tau_k} \mid Y_k] = G_1(Y_k)$. Adding gives $\mathbb{E}[M_{k+1} - M_k \mid Y_k] = 0$. □

Proposition 3.3 (Move-time discounted regularisation)**.** *For $\beta \in (0,1)$, $g_\beta := (I - \beta B)^{-1} G_1$ satisfies $g_\beta = \mu\mathbf{1}/(1-\beta) + (I - \beta B)^{-1}(G_1 - \mu\mathbf{1})$, so the only singular term as $\beta \uparrow 1$ is the constant mode. If $\mu = 0$ and $\|B^k G_1\|_\infty \le C_\rho \rho^k \|G_1\|_\infty$,*

$$\|g_\beta - G^*\|_\infty \le (1-\beta)\,\frac{C_\rho \rho}{(1-\rho)^2}\,\|G_1\|_\infty.$$

Proof. The Neumann series gives $g_\beta = \mu\mathbf{1}/(1-\beta) + (I-\beta B)^{-1} h$ with $h = G_1 - \mu\mathbf{1}$. If $\mu = 0$, $G^* - g_\beta = \sum_{k\ge 1}(1-\beta^k) B^k G_1$, and $1 - \beta^k \le k(1-\beta)$ gives the stated bound. □

Remark 3.4 (Estimating the gap). The bound rate ρ is the spectral radius of B restricted to the centered subspace $\{v : \pi^\top v = 0\}$; equivalently, the modulus of the subdominant eigenvalue of B. It is available in closed form from the estimated B and gives a quantitative prediction for the horizon crossover.

The resolvent is therefore a regularised proxy for G^*: nearly lossless when $\mu \approx 0$, but dominated by the constant mode when centering fails. Unlike the event-time resolvent $V_\gamma = (I - \gamma P)^{-1} r$ of the companion note, it discounts only after the move-time reduction.

4 Empirical results

4.1 Evaluation protocol

We score predictors against grouped, state-conditioned future-price curves. Fix a horizon h and let c index a spread/imbalance bucket. With $\bar{\Delta}_h(c)$ the average future mid-price change and $\bar{f}(c)$ the average predictor value in bucket c,

$$\mathrm{RMSE}_h(f) := \Big(|C_h|^{-1} \sum_{c \in C_h} \big(\bar{\Delta}_h(c) - \bar{f}(c)\big)^2 \Big)^{1/2}.$$

Remark 4.1 (Calibration, not forecasting). Grouped RMSE measures bucket-averaged calibration rather than pointwise forecast quality: a state-constant predictor within each bucket attains perfect calibration with zero forecasting skill. Pointwise RMSE is the relevant quantity for trading; grouped RMSE is used here because it aligns directly with the state-indexed objects $G_1(x)$, $G^*(x)$, and the weighted mid.

Horizons are clock-time counts rather than move counts. Panel-level geometric-mean event-per-move ratios (Stoikov ≈ 10, IEX ≈ 5, WSELOB ≈ 40) imply the 60 s WSELOB horizon covers only one to two price moves, while the 180 s BAC horizon covers roughly twenty – so the predicted horizon effect should be most visible on WSELOB. "Disc(.99)" denotes the move-time resolvent $(I - 0.99\,B)^{-1} G_1$.

The move-time martingale diagnostic of Corollary 3.2 is, with $\Delta P_k := p_{\tau_{k+1}} - p_{\tau_k}$,

$$\widehat{\xi}_k := \Delta P_k - \widehat{\mu} + G^*(Y_{k+1}) - G^*(Y_k),$$

$$\widehat{t}_M := \big|\overline{\widehat{\xi}} \,/\, \big(s_{\widehat{\xi}}/\sqrt{n}\big)\big|.$$

Small $\widehat{t}_M$ means the estimated correction is close to a move-time martingale. Sample sizes span 3.9k observed moves (KGHM) to 163k (CVX); IEX samples are smallest at 2.9–7.0k.

4.2 The horizon effect, at a glance

Figure 3 shows the horizon effect predicted by Proposition 3.1 directly. Each panel plots the grouped-RMSE ratio to the raw mid as the horizon h sweeps from 5 to 600 s; a ratio below 1 means the predictor beats the raw mid. On Stoikov the centered corrector G^* already starts below 1 and improves to ≈ 0.47 around 180 s before drifting back as the stationarity assumption degrades. On WSELOB, G^* begins with a ratio of 2.7 (transient-dominated), crosses parity at $h^* \approx 14$ s, and reaches ≈ 0.55 between 30 and 60 s. On IEX the centered corrector hugs unity throughout while Disc(.99) diverges above 4 at short horizons and never recovers – the residual-drift failure mode of Proposition 3.3.

4.3 Public BAC/CVX benchmarks

Table 1 gives the Stoikov panel numbers. BAC is mixed: G^* materially improves on the raw mid but is itself beaten by roughly 2× by the weighted mid. This is because BAC trades in a one-tick regime, so the imbalance-weighted linear interpolation already captures most of the state-conditional signal; CVX has meaningful spread variation and the weighted mid degrades. In all three cells, $\widehat{t}_M < 2$.

4.4 WSELOB: the horizon crossover

WSELOB gives the clearest horizon crossover; Figure 3 traces it as a sweep, Table 2 pins down the per-asset numbers. At 10 s, the raw mid is best on PEKAO, PKOBP, and PZU, while G_1 is best on KGHM and PKNORLEN; G^* can still be 6–83% worse than the raw mid, consistent with a material transient $B^n g$. By 60 s the ordering reverses: G^* is best on four names and Disc(.99) on the fifth, with grouped-RMSE reductions of 20–56% relative to the raw mid. The weighted mid is uncompetitive, and $\widehat{t}_M < 2$ for every name.

4.5 IEX: centering as a diagnostic

The IEX venue-local panel is deliberately harder: a single venue, a short sample, and a coarse state description leave more room for residual drift. The IEX panel of Figure 3 makes this visible – G^* hugs unity at every horizon while Disc(.99) sits above 4 at short h and never recovers. BAC remains near-centered with $\widehat{t}_M = 0.16$ and G^* modestly better than the raw and weighted mid. SPY and GLD do not: $\widehat{t}_M = 3.09$ and 2.24, and the discounted surrogate no

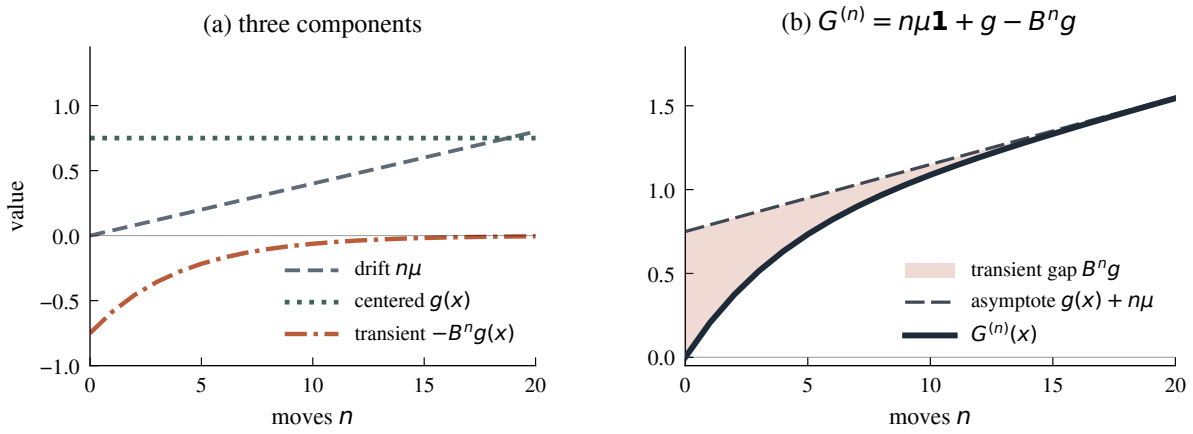

Figure 2. Conceptual plot of $G^{(n)}(x)$ for a single state x, with synthetic $\mu = 0.04$, $g(x) = 0.75$, and geometric decay $\|B^n g\| = g\,\rho^n$ with $\rho = 0.78$. *(a)* the three components $n\mu$, $g(x)$, $-B^n g(x)$. *(b)* their sum $G^{(n)}(x)$ approaches the asymptote $g(x) + n\mu$ as the shaded transient gap $B^n g$ decays.

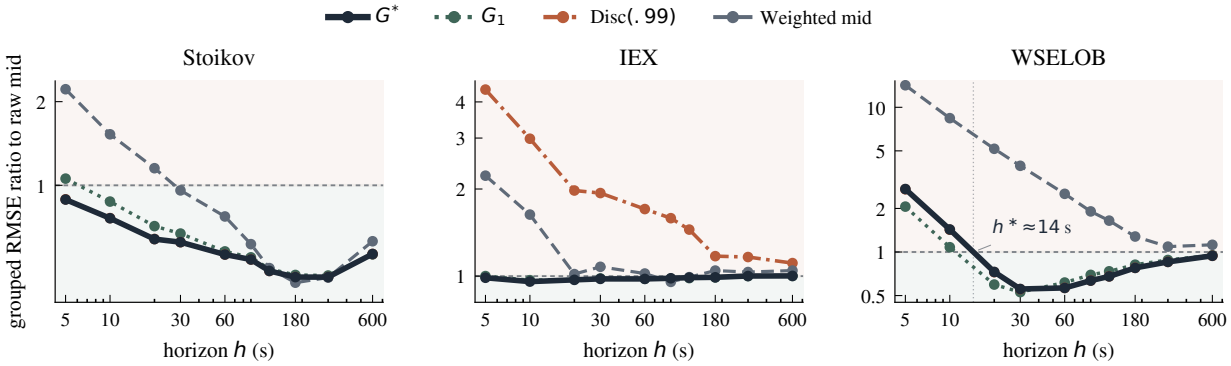

Figure 3. Horizon effect on each public panel: geometric-mean grouped-RMSE ratio to the raw mid as the forecasting horizon h sweeps from 5 to 600 s. Below the dashed $y = 1$ line, the predictor beats the raw mid; the shaded band turns clay where G^* is still transient-dominated and moss where the centered limit has taken over. WSELOB's crossover at $h^* \approx 14$ s is marked; IEX never crosses because $\widehat{t}_M$ flags residual drift.

Table 1. Public BAC/CVX grouped-RMSE benchmarks and move-time diagnostics. Best predictor per row in bold; lower is better.

		grouped RMSE					diagnostic	
Asset	h (s)	Raw	WMid	G_1	G^*	Disc(.99)	$\widehat{t}_M$	n_{mv}
BAC	60	0.001812	**0.000435**	0.000909	0.000892	0.000892	1.12	1.9k
BAC	180	0.002287	**0.000345**	0.000561	0.000547	0.000547	1.12	1.9k
CVX	10	0.001828	0.003941	0.000910	**0.000697**	0.000699	0.84	163k

Table 2. WSELOB summary by asset. Ratios >1 at 60 s favour G^* over the raw or weighted mid.

	best predictor		60 s improvement over G^*		diagnostic	
Asset	@10 s	@60 s	Raw/G^*	WMid/G^*	$\widehat{t}_M$	n_{mv}
KGHM	G_1	G^*	1.25	2.07	1.97	3.9k
PEKAO	Raw mid	Disc(.99)	1.97	6.61	1.08	12.7k
PKNORLEN	G_1	G^*	1.54	4.14	0.68	10.9k
PKOBP	Raw mid	G^*	2.03	4.94	1.85	17.2k
PZU	Raw mid	G^*	2.26	6.37	0.77	14.0k

longer shadows the centered correction. This is the failure mode of Proposition 3.3: when the constant mode is not well centered, discounted extrapolation is driven by residual drift rather than the bounded state correction. Sample sizes are small (2.9k, 5.6k, 7.0k moves), so the IEX panel is best read as a qualitative stress test of the diagnostic.

5 Conclusion

Move-time sampling gives $G^{(n)} = n\mu\mathbf{1} + g - B^n g$ (Figure 2); empirically (Figure 3), G^* dominates the raw mid past an asset-specific $h^* \approx 14$ s on WSELOB whenever $\widehat{t}_M < 2$, while IEX exhibits the residual-drift failure mode of Proposition 3.3. The companion note [2] adds event-time Bellman objectives.

References

[1] S. Stoikov, *The Micro-Price: A High Frequency Estimator of Future Prices*, Quantitative Finance, 18(12):1959–1966, 2018.

[2] L. Petersen, *Microprice, Bellman Values, and Semi-Markov Reinforcement Learning*, companion note, 2025.