М І Ф

Математика, Інформатика, Фізика: Наука та Освіта, Том 3, № 2 (2026), с. 152–163.
Mathematics, Informatics, Physics: Science and Education, Volume 3, No. 2 (2026), pp. 152–163.
Journal homepage: https://intranet.vspu.edu.ua/miph

UDC 519.217.2:519.86

Relaxation-time bound on the value of regime information in ergodic control of concentrated liquidity positions

Illia Kasadzhyk[1], Oleksandr Vasyliev[2]

[1]Odesa I. I. Mechnikov National University,
Department of Optimal Control and Economic Cybernetics, Odesa, Ukraine
illia.kasadzhyk@stud.onu.edu.ua
https://orcid.org/0009-0006-6935-0393

[2]Odesa I. I. Mechnikov National University,
Department of Optimal Control and Economic Cybernetics, Odesa, Ukraine
av5111955@gmail.com
https://orcid.org/0000-0002-3826-4883

Abstract. We consider the problem of choosing the price range width of a liquidity provider position in an automated market maker with concentrated liquidity (Uniswap V3) in the case where the market volatility regime is a hidden state of a Markov chain observed only through the posterior distribution of an online filter. For symmetric boundary-reset range policies, a closed-form ergodic growth rate of log-capital with a logarithmic loss per cycle is derived and checked against Monte Carlo simulation of the underlying random walk and against an event-driven backtest on a five-year one-minute series of the USDC/WETH 0.05 % pool. Regime-conditional moments estimated from the same observation that produced the posterior are shown not to be a forecast: the volatility contrast between regimes is inflated roughly fourfold relative to the one-day predictive estimate, and the "optimal" widths that follow are an artefact. Predictive moments over a common horizon are proposed, with a self-consistent calibration of that horizon through the expected cycle length. The value of regime information for the width choice is proved nonnegative and bounded above by a quantity proportional to the ratio of the relaxation time of the chain to the decision horizon; for the fitted one-minute model the relaxation time is 19 minutes, whereas a position lives for weeks or longer. An empirical check on non-overlapping windows with a permutation test agrees with the theorem: the gain of the regime policy with the causal posterior does not exceed 0.75 percentage points per year (1.2 with the oracle state), and none of the values is statistically different from zero.

Keywords: concentrated liquidity, hidden Markov models, ergodic control, relaxation time, value of information, capital growth rate.

e-ISSN 3041-1955 DOI: https://doi.org/10.31652/3041-1955-2026-03-02-05
2020 Mathematics Subject Classification: 93E20, 60J10, 91G80.
© 2026 Kasadzhyk I., Vasyliev O., Creative Commons Attribution 4.0 International Licence.

1. Introduction

An automated market maker with concentrated liquidity [1] allows a liquidity provider to place capital in a price range $[p_a, p_b]$ and to earn fees only while the spot price stays inside the range. Narrowing the range multiplies both the fee income and the losses from adverse price moves, so the choice of the width is a control problem. The main source of losses is the loss relative to a rebalancing portfolio (loss-versus-rebalancing, LVR), derived for a stochastic price model in [2]; the class of boundary-reset range policies was introduced in [3], the width choice was posed as a stochastic control problem in [4], impermanent loss and fee asymptotics were studied in [5], and the empirical returns of Uniswap V3 liquidity providers in [6].

In practice, regime-based control adds a hidden Markov model of volatility to this scheme [7], [8]: the market is in one of K regimes, the regime is not observed, and an online filter [9] provides the posterior distribution φ_n. The natural next step is to estimate the regime-conditional moments of returns and to choose the range width according to the current posterior. To the authors' knowledge, two questions remain unanswered here: over what horizon the conditional moments should be estimated if the position lives much longer than one observation step, and what regime information is worth for a decision that is fixed for the whole lifetime of the position.

This paper answers both questions. First, conditional moments estimated from the same observation that produced the posterior are not a forecast and inflate the volatility contrast between regimes roughly fourfold (10.6 versus 2.6 at one day); the correct estimate uses a predictive window with one horizon for the fee yield and the variance, and the horizon is determined self-consistently through the cycle length of the position. Second, for any ergodic chain with a diagonalizable transition matrix, the value of regime information is bounded above by a quantity of the order of the ratio of the relaxation time to the decision horizon. Both results are checked on a five-year one-minute series of a real pool.

2. Statement of the problem

The aim of the paper is to establish the correct horizon for estimating regime-conditional moments in the problem of choosing the width of a concentrated liquidity range and to obtain an upper bound on the value of regime information for the class of symmetric boundary-reset policies. The novelty lies in three results: a closed-form ergodic growth rate with a logarithmic loss per cycle; a definition of predictive conditional moments over a common horizon, together with a way to choose this horizon self-consistently; and an upper bound on the value of regime information in terms of the ratio of the relaxation time of the chain to the decision horizon, a bound that is new for liquidity management problems.

Price and regimes. Let S_t be the spot price of the pool (the market maker for one pair of assets), $Y_t = \ln S_t$, and let $X_n = Y_{t_n} - Y_{t_{n-1}}$ be the log-price increments on a regular one-minute grid (t_n), with unconditional second moment $s^2 = \mathbb{E}[X_n^2]$; N is the length of the series in steps. The hidden regime $Z_n \in \{1, \dots, K\}$ is a stationary ergodic Markov chain with transition matrix A, stationary distribution ν and second largest eigenvalue in modulus λ_2. Given (Z_n), the increments are independent, and their distribution is a mixture of an atom at zero (the price does not move between swaps) and a Gaussian component with scale σ_{Z_n}; the conditional second moment is denoted by $\mathbb{E}[X^2 \mid Z = j]$. The observer sees only $X_{1:n}$ and the posterior $\varphi_n(j) = \mathbb{P}(Z_n = j \mid X_{1:n})$ computed by the online filter.

Position and fees. A position of half-width $d > 0$ on the logarithmic price scale occupies the range $[p\,e^{-d},\, p\,e^{d}]$ around the entry price p. Its liquidity per unit of capital

equals

$$c(d) = \frac{L\sqrt{p}}{V} = \frac{1}{2\big(1 - e^{-d/2}\big)}, \qquad c(\infty) = \tfrac{1}{2}, \tag{1}$$

where $c(d)$ is the liquidity concentration; L is the liquidity of the position; V is its capital. The second equality in (1) follows from the value [1] $V = L\big(2\sqrt{p} - \sqrt{p_a} - p/\sqrt{p_b}\big)$ of a position with range $[p_a, p_b]$ at a price $p \in [p_a, p_b]$, taken at $p_a = pe^{-d}$, $p_b = pe^{d}$; the concentration relative to the full range equals $2c(d)$. The fee yield of the pool per unit of liquidity per minute is $\gamma\kappa_n$, where $\kappa_n = \mathrm{vol}_n\sqrt{p_n}/\mathrm{Liq}_n$ is determined by the volume vol_n of swaps (trades with the pool) in units of the base asset, the price p_n and the active liquidity Liq_n, the sum of L over the positions whose range contains p_n (fees are shared among them in proportion to L); $I = \mathbb{E}[\kappa_n]$ is the mean of κ_n, γ is the fee rate of the pool, and G is the network fee (gas) for reopening the position around the new price (a rebalance).

Policy class and criterion. We consider symmetric boundary-reset policies (τ-reset in the terminology of [3]): the position is opened around the current price, held until the price first leaves the range, and immediately reopened around the new price with half-width d, which may depend on the posterior at the moment of entry. The trajectory splits into entry–exit–rebalance cycles. The criterion is the long-run growth rate of log-capital (the Kelly criterion [10]), so the problem is one of ergodic control; the general impulse-control formulation reduces to a quasi-variational inequality [11] and has no closed-form solution; hence the policy class is restricted to a one-parameter family.

Two questions. How should the regime-conditional moments be estimated for substitution into the criterion for a position that lives H steps? How much does the policy $d(\varphi_n)$ gain over the best single width, and how does this gain depend on H?

Data. The numerical results are obtained from the full history of the USDC/WETH pool with a 0.05 % fee on the Ethereum network from 5 May 2021 to 14 September 2026: 11.6 million swap events aggregated into $N = 2\,818\,340$ one-minute increments; $s = 10.55$ basis points (bp), $I = 4.05 \cdot 10^{-4}$, $\gamma = 0.05\,\%$, $G = 5$ USD, $V = 10^6$ USD. A three-state hidden Markov model is fitted by the online EM algorithm [9]: the scales of the Gaussian components are $\sigma_j = 3.56,\ 7.72,\ 24.3$ bp with zero-increment probabilities $\pi_j = 0.56,\ 0.18,\ 0.02$ (hence $\sqrt{\mathbb{E}[X^2 \mid Z = j]} = 2.35,\ 6.99,\ 24.0$ bp), $a_{jj} = 0.980,\ 0.942,\ 0.849$, $\nu = (0.623,\ 0.332,\ 0.045)$, $|\lambda_2| = 0.9474$. We refer to the regimes as consolidation, moderate activity and turbulence. The fitted parameters reproduce the bulk of the distribution but not its tail: the square root of the unconditional second moment of the model $\sum_j \nu_j \mathbb{E}[X^2 \mid Z = j]$ is 6.745 bp against the sample 10.55 bp; hence all quantitative calculations below use posterior-weighted sample moments of the series, and the chain parameters enter only through φ_n and the mixing rate of A.

3. Main results

3.1. Ergodic growth rate and its verification. By the renewal-reward theorem [12], the growth rate of log-capital equals the ratio of the expected increment per cycle to the expected cycle length. If Y is a driftless Brownian motion with variance s^2 per step and $\kappa_n \equiv I$, the mean first exit time from $[-d, d]$ equals d^2/s^2, and

$$g(d) = \gamma\, I\, c(d) - \frac{\Lambda(d) + \gamma\,|\Delta f| + G/V}{\mathbb{E}[T(d)]}, \qquad \mathbb{E}[T(d)] = \frac{d^2}{s^2}, \tag{2}$$

where $g(d)$ is the growth rate of log-capital per step; $\Lambda(d) = -\frac{1}{2}\big(\ln R(+d) + \ln R(-d)\big)$ is the expected logarithmic loss per cycle with equally likely exits through either boundary; $R(\pm d)$ is the ratio of the value of the position at the exit through the upper or the

lower boundary to its value at entry; $|\Delta f| = \frac{1}{2}$ is the fraction of capital swapped at the boundary to restore equal asset values (at the boundary the position consists of a single asset); $\mathbb{E}[T(d)]$ is the expected cycle length in steps. For a range symmetric in log-price, the piecewise value function gives $R(+d) = (1 + e^{d/2})/2$ and $R(-d) = (e^{-d/2} + e^{-d})/2$, whence the closed form

$$\Lambda(d) = -\ln \frac{1 + e^{-d/2}}{2} \geq 0, \qquad \Lambda(d) = \frac{d}{4} + O(d^2), \qquad \Lambda(d) \to \ln 2 \ (d \to \infty). \tag{3}$$

The logarithm is essential here: replacing $\mathbb{E} \ln R$ with $\mathbb{E}R - 1$ introduces a Jensen gap of 6.5 to 7.6 percentage points (pp) per year for $1 \leq d \leq 3$ with the parameters of the pool.

Since $c(d) \approx 1/d$ and $\Lambda(d) \approx d/4$ as $d \to 0$, the term $\Lambda(d)/\mathbb{E}[T(d)]$ in (2) equals, to leading order in d, $c(d)\, s^2/4$, i.e. the full-range LVR $s^2/8$ multiplied by the relative concentration $2c(d)$; the fixed rebalancing costs $(\gamma|\Delta f| + G/V)\, s^2/d^2$ grow faster, as $1/d^2$. The ratio of fees to LVR

$$\frac{\gamma I c(d)}{c(d)\, s^2/4} = \frac{4\gamma I}{s^2} \tag{4}$$

does not depend on the width to leading order: concentration multiplies income and LVR by the same factor. On the five-year series, $4\gamma I/s^2 = 0.73$ (by regime: 2.68, 1.29, 0.29 along the Viterbi path, the most probable regime sequence given the whole series, and 2.58, 1.15, 0.34 with the contemporaneous estimates of Section 3.2). On average, therefore, the fees of the pool do not cover LVR, and it is this deficit, together with the rebalancing costs, that ruins narrow ranges. As $d \to \infty$ the loss term vanishes, since Λ is bounded while $\mathbb{E}[T] \to \infty$; formula (2) is valid provided $\mathbb{E}[T(d)] \ll N$, and on the available series this condition is violated for $d \geq 1$.

Verification protocol. To separate the error of the ergodic limit from the error of the random walk model itself, we evaluate the same policy in three independent ways (Fig. 1; results in Table 1 and Fig. 2):

(a) formula (2) with the parameters I, s^2 of the series;
(b) direct Monte Carlo (MC) simulation of the same model (Gaussian increments with variance s^2, constant I, exact piecewise value function) over the horizon N with 96 trajectories, closing the cycle still open at the end of the horizon and realizing its loss;
(c) an exact event-driven backtest, i.e. a replay of the policy on the real series: the endpoints $pe^{\pm d}$ are not rounded to the price grid of the pool, the exit is recorded at a node of the one-minute grid, fees accrue as $\gamma L \sum_n \sqrt{p_n}\, \kappa_n$ over the minutes inside the range, the swap at the boundary costs $\gamma|\Delta f|V$, the gas G is charged for every rebalance, and ruin is recorded when the capital is insufficient to pay for the gas of the next rebalance.

TABLE 1. Capital growth rate, % per year (annualized as $100\big(e^{525\,600\,g} - 1\big)$): formula (2), Monte Carlo in the same model (s.e. is the standard error) and backtest on the real series ($V = 10^6$ USD, $\gamma = 0.05\,\%$, $G = 5$ USD)

d	$\mathbb{E}[T]/N$	Formula	Monte Carlo (± 1 s.e.)	Backtest	Rebalances
0.03	0.000	-76.6	-71.9 ± 0.5	ruin	1570
0.07	0.002	-42.6	-41.5 ± 1.1	-45.1	563
0.20	0.013	-14.6	-16.0 ± 1.5	-29.0	90
0.50	0.080	-3.4	-4.1 ± 1.7	-4.1	11
1.00	0.319	$+0.7$	-0.8 ± 1.8	$+3.0$	2
1.50	0.717	$+2.1$	$+0.8 \pm 1.8$	$+5.1$	0

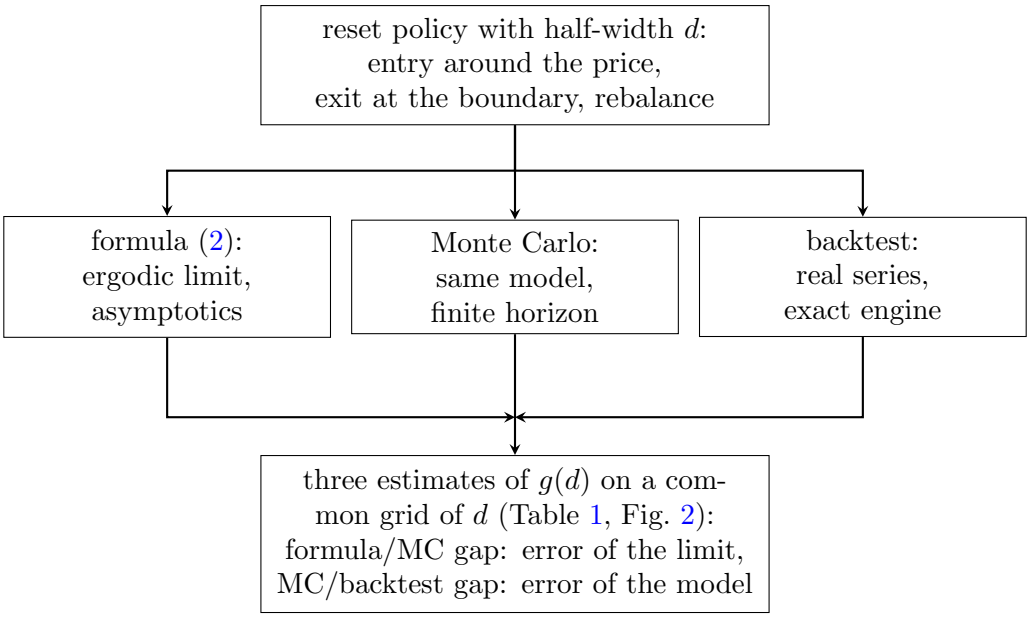

Fig. 1. Verification protocol for the ergodic formula: three independent estimates of the same policy separate the error of the ergodic limit from the error of the random walk model

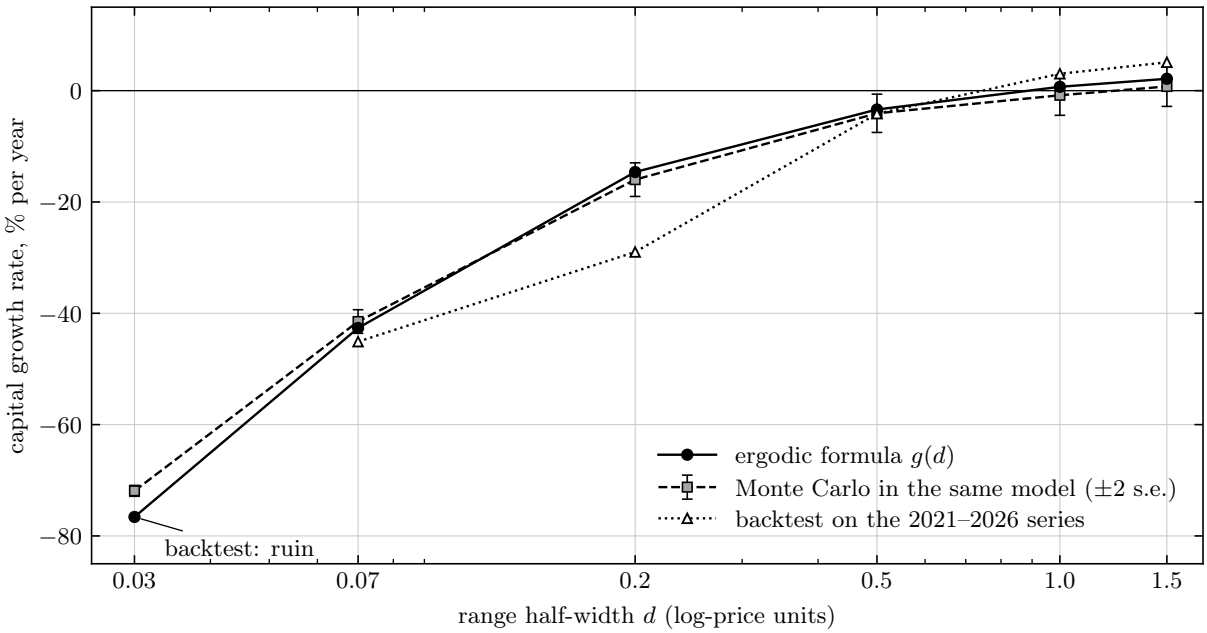

Fig. 2. Three estimates of the growth rate as functions of the half-width d: the formula with the logarithmic loss (3), Monte Carlo in the same model and the backtest on the 2021–2026 series

The formula and Monte Carlo agree within one or two standard errors for $d \geq 0.07$. At $d = 0.03$ the formula is lower by 4.7 pp: the mean exit time of the discrete walk exceeds d^2/s^2 by about 4 % at $d/s = 28$, which reduces the loss term by the same factor. For $d \geq 0.2$ the formula is systematically higher than Monte Carlo by 0.7–1.5 pp. This comes from the linearization of fees, since (2) uses $F = \gamma I c(d) \mathbb{E}[T]$ in place of $\ln(1 + F)$, and F grows from 0.04 at $d = 0.2$ to 0.39 at $d = 1.5$. With $\ln(1 + F)$ in place of F the formula gives -15.5, -4.5, -0.7, $+0.55$ versus -16.0, -4.1, -0.8, $+0.8$ from Monte

Carlo, within half a standard error, and after this correction the contribution of the unclosed cycle at $d \geq 1$ cannot be distinguished from zero.

The discrepancy with the real series is of a different nature, and its sign depends on the width. In the middle region ($d \approx 0.2$) the series is 13–14 pp worse than the model: a symmetric walk does not capture trends or heavy tails, under which the exit happens mostly in the direction of the move and every rebalance locks in a loss. In the wide region ($d \geq 1$) the series is 4 pp better than the model: the fee yield and the volatility are positively correlated; the turbulent regime (by $\arg\max_j \varphi_n(j)$) occupies 4.7 % of the time but generates 28 % of the fees, a fee density six times the average, and a wide position that a spike does not push out of the range collects them in full. In the narrow region all three estimates agree: a narrow position is ruined.

Thus the formula describes its own model accurately, while the model deviates from the series by up to 14 pp depending on the width; structural properties are therefore derived below from the formula, and quantitative checks use the exact backtest engine.

3.2. Predictive moments and a self-consistent horizon. Regime-based control in this class chooses a width d_j for each regime and applies d_j with probability $\varphi_n(j)$ at entry, substituting the conditional moments (I_j, s_j^2) into (2). The natural estimate of the second moment, the φ-weighted average $\sum_n \varphi_n(j) X_n^2 / \sum_n \varphi_n(j)$ (contemporaneous, since it is taken at the same step as φ_n) is not a forecast: φ_n is the posterior *after* the observation X_n, so the weighted average partly reproduces X_n^2 itself. Formula (2), in contrast, requires the variance of the increments over the lifetime of a position opened at time n. The correct moments are the predictive moments over the horizon H:

$$s_j^2(H) = \mathbb{E}\Big[\frac{1}{H}\sum_{m=1}^{H} X_{n+m}^2 \,\Big|\, Z_n = j\Big], \qquad I_j(H) = \mathbb{E}\Big[\frac{1}{H}\sum_{m=1}^{H} \kappa_{n+m} \,\Big|\, Z_n = j\Big], \tag{5}$$

where H is the estimation horizon in steps and Z_n is the regime at the moment the position is opened; both moments are averaged over the same window $n+1, \ldots, n+H$.

Estimation protocol (Fig. 3). Write $\theta_j(H) = (I_j(H), s_j^2(H))$ and, for $\xi_{n+m} \in \{\kappa_{n+m}, X_{n+m}^2\}$, let $\bar{\xi}_n(H)$ be the mean of ξ_{n+m} over $m = 1, \ldots, H$; each coordinate of $\theta_j(H)$ is estimated by the ratio of $\sum_n \varphi_n(j)\, \bar{\xi}_n(H)$ to $\sum_n \varphi_n(j)$, where φ_n uses only $X_{1:n}$ (the online filter runs causally) and the sums run over the whole series. These estimators are consistent for the posterior-weighted means $\mathbb{E}[\varphi_n(j)\, \bar{\xi}_n(H)]/\mathbb{E}[\varphi_n(j)]$, the quantities used by a policy acting on the observed φ_n; they are mixtures of the conditional means given the hidden Z_n and coincide with them for an exact filter, $\varphi_n = e_{Z_n}$. A one-step-ahead estimate of I_j combined with s_j^2 over a long window produces a spurious combination of "high fee yield at low volatility", for which the formula promises hundreds of percent per year at $d \approx 0.01$; therefore, a common horizon is mandatory.

Table 2 shows the scale of the distortion. The contemporaneous estimate gives a volatility contrast s_3/s_1 of 10.6 between the turbulent regime and the consolidation regime. The correct predictive estimate gives a contrast of 5.0 one step ahead, 2.6 at one day and 1.45 at one month. The widths $d^* = (0.0013, 0.014, 1.5)$ that follow from the contemporaneous estimate are an artefact: with predictive moments, from the daily horizon on, the optimum in all three regimes lies at the upper bound $d = 1.5$, and the splitting by regime survives only up to and including the hourly horizon, where $d^* = 0.075$ for consolidation. This boundary optimum lies in the region $\mathbb{E}[T(d)] \sim N$, where the ergodic limit (2) no longer describes the finite series; the formula is used here only to compare the regimes, and the quantitative conclusions of Section 3.4 rest on the exact backtest engine, which does not require $\mathbb{E}[T] \ll N$. Even the hourly optimum $d^* = 0.075$ for consolidation is not sharp: $g(d)$ stays within 10 % of its maximum for d from 0.05 to 0.17.

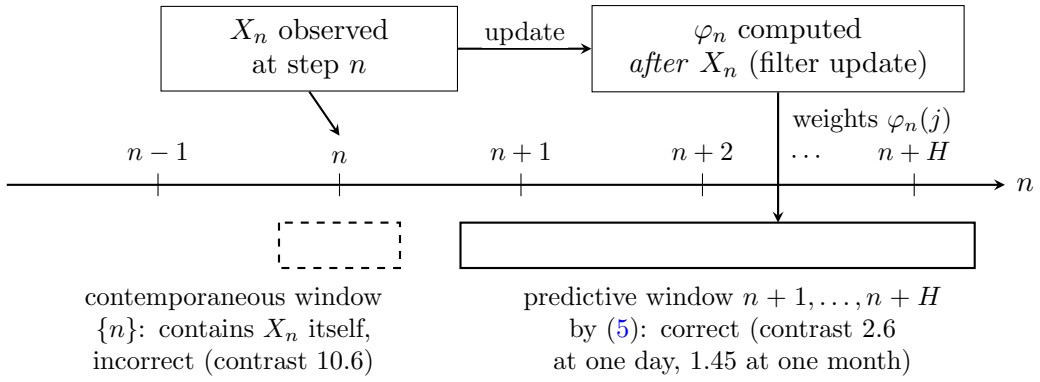

FIG. 3. Contemporaneous and predictive estimates of the conditional moments on the time axis in grid steps: the posterior φ_n is computed after X_n, so the correct estimation window starts at step $n+1$

TABLE 2. Regime-conditional moments depending on the estimation horizon: the contemporaneous estimate versus the predictive estimates (5), $s_j = \sqrt{s_j^2(H)}$; d_j^* is the maximizer of (2) over $0.001 \le d \le 1.5$

Horizon H	s_1, bp	s_2, bp	s_3, bp	s_3/s_1	I_3/I_1	d_1^*, d_2^*, d_3^*
contemporaneous	3.33	9.88	35.14	10.57	14.65	0.0013, 0.014, 1.5
1 min (ahead)	5.82	10.98	29.15	5.01	8.85	0.012, 1.5, 1.5
1 hour	6.79	11.24	26.19	3.86	5.78	0.075, 1.5, 1.5
1 day	7.99	11.63	21.01	2.63	2.88	1.5, 1.5, 1.5
1 week	8.95	11.30	17.80	1.99	1.81	1.5, 1.5, 1.5
1 month	9.23	10.51	13.40	1.45	1.34	1.5, 1.5, 1.5
1 year	9.31	9.39	9.96	1.07	1.11	1.5, 1.5, 1.5

Self-consistent calibration. The horizon H in (5) is determined by the problem itself: a position of half-width d_j lives on average $\mathbb{E}[T_j] = d_j^2/s_j^2$ steps, so the moments should be estimated over the very horizon they determine. The map

$$H \;\mapsto\; \big(I_j(H), s_j^2(H)\big) \;\mapsto\; d_j^* = \arg\max_d g\big(d; I_j(H), s_j^2(H)\big) \;\mapsto\; d_j^{*2}/s_j^2(H)$$

is iterated for each j on a logarithmic scale in H with damping $H \leftarrow \sqrt{H\,\mathbb{E}[T_j]}$ (Fig. 4). The fixed point lies on the bound $H \le N/4 \approx 489$ days imposed by the length of the series; the estimate $\mathbb{E}[T_j] = d_j^{*2}/s_j^2$ without this bound gives 1600–1900 days, of the order of the length of the whole sample (1957 days). At the 489-day horizon, the s_j equal 9.11, 9.24, 9.96 bp (contrast 1.09), the I_j differ by 12 %, and the d_j^* coincide. Both bounds, $H \le N/4$ and $d \le 1.5$, are computational constraints; inside them, the self-consistent problem pushes the decision horizon far beyond the memory of the regime.

3.3. Relaxation-time bound on the value of regime information. The observation of Section 3.2 reflects a structural property of the problem rather than a feature of the particular series.

Theorem 1. *Let (Z_n) be a stationary ergodic Markov chain with a diagonalizable transition matrix $A = P\,\mathrm{diag}(1, \lambda_2, \dots, \lambda_K)\,P^{-1}$, stationary distribution $\nu = (\nu_j)$ and second largest eigenvalue in modulus λ_2, $|\lambda_2| < 1$; given the trajectory (Z_n), let the pairs (κ_n, X_n) be conditionally independent over n, and let the conditional distribution of a pair depend only on Z_n and have finite conditional means $\bar\theta_i = \big(\mathbb{E}[\kappa \mid Z = i],\ \mathbb{E}[X^2 \mid Z = i]\big)$; let*

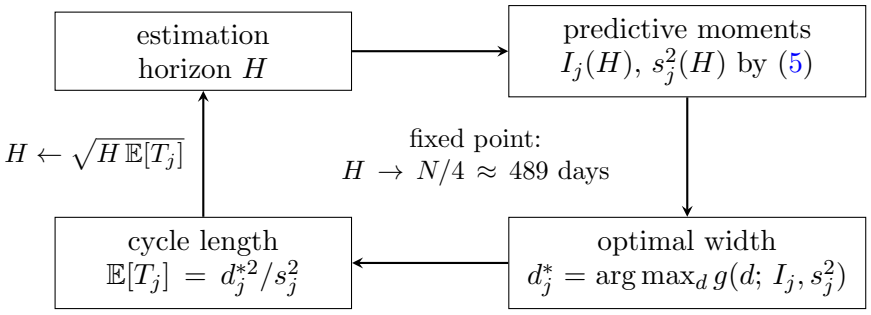

FIG. 4. Self-consistent calibration: the horizon for estimating the moments equals the cycle length that these moments determine

$\theta_j(H)$ be as in Section 3.2, let $D = [d_{\min}, d_{\max}] \subset (0, \infty)$, and let

$$\mathcal{V}(H) = \sum_j \nu_j \max_{d \in D} g\big(d; \theta_j(H)\big) - \max_{d \in D} \sum_j \nu_j\, g\big(d; \theta_j(H)\big) \tag{6}$$

be the value of regime information, the gain of the best regime policy over the best single width. Then

$$0 \le \mathcal{V}(H) \le \frac{C\, t_{\mathrm{rel}}}{H} \quad \forall H \ge 1, \qquad t_{\mathrm{rel}} = \frac{1}{1 - |\lambda_2|}, \tag{7}$$

where $C = 2K\, C_A\big(a_D \max_i \mathbb{E}[\kappa \mid Z = i] + b_D \max_i \mathbb{E}[X^2 \mid Z = i]\big)$ with $a_D = \gamma \max_{d \in D} c(d)$, $b_D = \max_{d \in D}\big(\Lambda(d) + \gamma|\Delta f| + G/V\big)/d^2$ and $C_A = \|P\|\,\|P^{-1}\|$ (spectral norms, the columns of P normalized to unit length); t_{rel} is the relaxation time of the chain, the inverse of its absolute spectral gap. The statement applies to policies that fix d according to the state at entry.

Proof. Nonnegativity follows from the fact that the maximum of a sum does not exceed the sum of the maxima:

$$\max_{d \in D} \sum_j \nu_j\, g(d; \theta_j) \le \sum_j \nu_j \max_{d \in D} g(d; \theta_j).$$

By the Markov property, $\theta_j(H) = \frac{1}{H} \sum_{m=1}^{H} \sum_i (A^m)_{ji}\, \bar\theta_i$. Diagonalizability gives $A^m - \mathbf{1}\nu^\top = P\,\mathrm{diag}(0, \lambda_2^m, \dots, \lambda_K^m)\, P^{-1}$, whence $|(A^m)_{ji} - \nu_i| \le C_A |\lambda_2|^m$ (convergence to equilibrium for an ergodic chain, see [13]). Therefore, for each coordinate $r = 1, 2$ of $\theta = (\theta^{(1)}, \theta^{(2)}) = (I, s^2)$ and the unconditional moments $\bar\theta = \sum_i \nu_i \bar\theta_i$,

$$\big|\theta_j^{(r)}(H) - \bar\theta^{(r)}\big| \le \frac{1}{H} \sum_{m=1}^{H} K\, C_A |\lambda_2|^m \max_i \big|\bar\theta_i^{(r)}\big| \le \frac{K\, C_A\, t_{\mathrm{rel}}}{H} \max_i \big|\bar\theta_i^{(r)}\big|.$$

The function $g(d; \theta)$ in (2) is affine in $\theta = (I, s^2)$ with the coefficients $\gamma c(d)$ and $-\big(\Lambda(d) + \gamma|\Delta f| + G/V\big)/d^2$, bounded in absolute value on D by a_D and b_D; hence $|g(d;\theta) - g(d;\theta')| \le a_D |I - I'| + b_D |s^2 - s'^2|$ for all $d \in D$. Consequently, $\big|\max_d g(d; \theta_j(H)) - \max_d g(d; \bar\theta)\big| \le a_D \big|I_j(H) - \bar\theta^{(1)}\big| + b_D \big|s_j^2(H) - \bar\theta^{(2)}\big|$ for every j, and the same bound holds for the maximum of the weighted sum. The difference of the two terms in (6) does not exceed twice this bound, which gives (7) with the stated C. □

Remark 2. Bound (7) need not be sharp; it holds with the same C if the $\theta_j(H)$ are replaced by convex combinations of them and ν by any probability vector (in particular, by the distribution of the state at reset times), hence for the posterior-weighted moments of Section 3.2. For a reversible chain the mixing time satisfies $t_{\mathrm{mix}}(\varepsilon) \le \big\lceil t_{\mathrm{rel}} \ln\big(1/(\varepsilon \min_j \nu_j)\big)\big\rceil$ [14]; without reversibility the diagonalization gives $t_{\mathrm{mix}}(\varepsilon) \le \big\lceil t_{\mathrm{rel}} \ln\big(KC_A/(2\varepsilon)\big)\big\rceil$. For the fitted model, $|\lambda_2| = 0.9474$, $t_{\mathrm{rel}} = 19.0$ min and $C_A = 2.67$.

The chain is not reversible; direct computation from A^h gives $t_{\mathrm{mix}}(1/4) = 24$ min, and the rows of A^h coincide with the stationary distribution to within $6{\cdot}10^{-5}$ by $h = 180$ min. This fast mixing is a property of the one-minute model as fitted: the residual contrast of 2.6 at the daily horizon in Table 2 is produced by volatility memory longer than the one-minute model implies; φ_n tracks this memory like any estimate of the current variance.

3.4. Empirical check. The quantity $\mathcal{V}(H)$ is computed in two independent ways (Fig. 5). The analytical approach evaluates (6) by formula (2) with predictive moments and with the time averages of $\varphi_n(j)$ in place of ν_j, and gives 0.50, 0.00, 0.00, 0.00 pp per year for the one-minute posterior at the horizons of one hour, one day, one week and one month; for the daily posterior (a three-state model of the same form fitted to daily increments, mean dwell 2 to 29 days) it gives 0.08, 0.27, 0.00, 0.00.

The empirical approach does not use the formula. We split the series into non-overlapping windows $W_k = \{n_k, \dots, n_k + H - 1\}$, $n_k = n_0 + kH$, $k = 0, \dots, M-1$ ($M = 1943$ daily, 277 weekly, 64 monthly windows; $n_0 = 20\,000$ steps of filter burn-in); for each of the nine widths $d \in \{0.03, \dots, 1.5\}$ we run the exact backtest engine over W_k to get the log-growth rate $g_k(d)$. Window k contributes to regime j with the weight $w_k(j)$, equal to $\varphi_{n_k-1}(j)$ for the one-minute posterior, to the daily posterior at n_k, or to the indicator of the oracle Viterbi state. Then $\hat{\mathcal{V}}(H)$ is (6) with $\hat{g}_j(d) = \sum_k w_k(j)\, g_k(d) \big/ \sum_k w_k(j)$ in place of $g(d; \theta_j(H))$ and $\hat{\nu}_j = \frac{1}{M}\sum_k w_k(j)$ in place of ν_j; confidence intervals are obtained by a bootstrap over windows. Since $\hat{\mathcal{V}}(H) \geq 0$ by construction, the bootstrap interval cannot test $\mathcal{V}(H) = 0$, so we run a permutation test: we permute the vectors w_k across windows 1000 times, and the p-value is the fraction of permutations with $\hat{\mathcal{V}}(H)$ at least the observed value.

The best single width at all horizons is the cap $d = 1.5$, so a regime can gain only by narrowing. The empirical value at one day is zero for all three regime signals; at one week it is 0.33 pp for the one-minute posterior ($p = 0.13$), zero for the daily posterior and 0.47 for the oracle ($p = 0.27$); at one month it is 0.75 ($p = 0.14$; the turbulent regime has a total weight of only 4 windows out of 64), zero and 1.16 ($p = 0.18$; 5 windows), respectively. No value is significant at the 0.10 level. Where the value is nonzero, the best regime policy is $(d_1, d_2, d_3) = (1.5, 1.5, 0.5)$: entry in the turbulent regime calls for a narrower range, contrary to the rule "widen the range as volatility rises", because the turbulent minutes carry a disproportionate share of the fees.

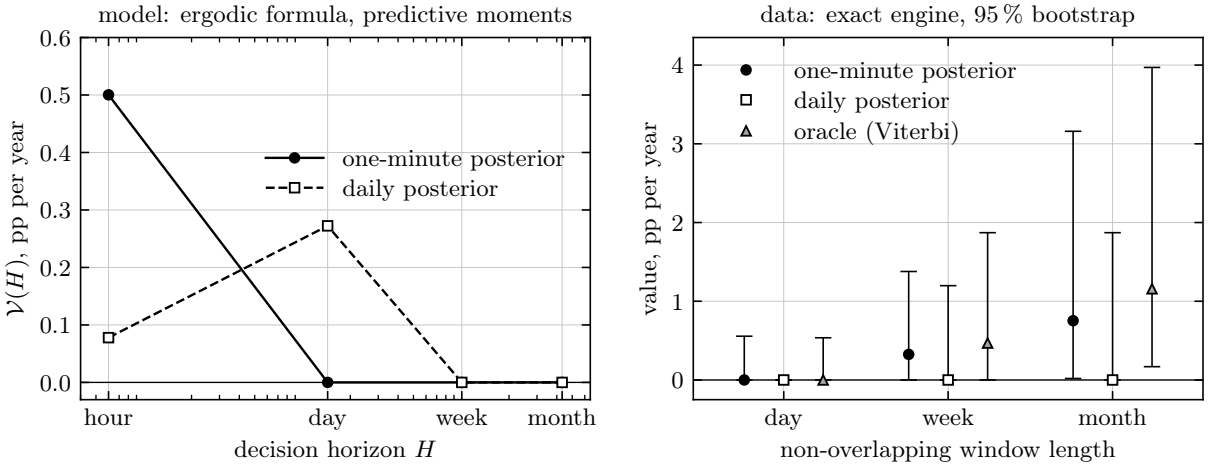

FIG. 5. Value of regime information $\mathcal{V}(H)$: analytically by (6) (left) and empirically with the exact backtest engine and 95 % bootstrap intervals (right)

The two approaches agree: over the lifetime of a position (one week or longer), splitting the width by regime yields at most 0.75 pp per year with either causal posterior and 1.2 pp with the oracle Viterbi state, and neither value is statistically different from zero. For the one-minute posterior this is what Theorem 1 predicts, since $t_{\mathrm{rel}} = 19$ min $\ll H$; the growth of the empirical values with H tracks the 95th percentile of the permutation null (0.05, 0.58, 1.16 pp at $M = 1943, 277, 64$), a selection effect of the maximum over widths and regimes. With $D = [0.03, 1.5]$, $K = 3$, $C_A = 2.67$ and the one-step moments (Table 2, $H = 1$) in place of $\bar{\theta}_i$, bound (7) gives 38 pp per year at the monthly horizon (2.2 pp for $D = [0.5, 1.5]$), loose because $a_D, b_D \sim 1/d_{\min}$ but of the observed order for wide ranges.

3.5. Reproducibility. The one-minute series is built by two independent, cross-checked collection pipelines from the full log of the pool's swap events, keyed by (block number, log index); the per-minute aggregates are the price of the last swap, the total volume and the mean active liquidity, with the price and liquidity carried forward in empty minutes. The backtest engine implements the piecewise value function of a Uniswap V3 position without approximations and is verified by 31 identities (value at entry, boundary continuity, the constant-product limit, monotonicity of the concentration). The online filter is implemented in C++ and verified bit for bit against an independent NumPy implementation. A streaming regime-identification pipeline, a policy simulator built on precomputed runs of the exact engine (63 starts, 11 strategies) and a panel of predictive moments for the current posterior are available in the password-protected web laboratory at https://appliedmathlab.com/chimera/ (access is provided on request).

Conclusions. The main tool is formula (2) with the logarithmic loss per cycle (3). It agrees with Monte Carlo simulation of the same model within one or two standard errors for $d \geq 0.07$, and the ratio of fees to LVR (4) does not depend on the range width to leading order. Regime-conditional moments should be estimated by (5) over a common predictive horizon determined self-consistently through the cycle length; the contemporaneous estimate inflates the regime contrast roughly fourfold and produces spurious optimal widths. Theorem 1 gives the bound Ct_{rel}/H for the value of regime information; for the fitted one-minute model, $t_{\mathrm{rel}} = 19$ min, and over the lifetime of a position the value is indistinguishable from zero both analytically and empirically.

Further work: bound (7) for impulse controls in continuous time; the cost of delayed or noisy filter information; price models with jumps.

Conflict of interest and ethics. The authors declare no conflict of interests. The authors also declare full adherence to all journal research ethics policies, namely involving the participation of human subjects anonymity and/or consent to publish. An AI tool (Claude, Anthropic) was used only to assist with technical data processing (help with the data-processing and simulation code, verified by the authors) and with translating the text from the authors' Ukrainian original; the model, the analysis, the interpretation of the results and the conclusions are the authors' own, and the authors take full responsibility for the code, the computations and the content.

Acknowledgements. The authors declare no special funding for this work.

References

1. Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D. (2021). *Uniswap v3 Core*. https://uniswap.org/whitepaper-v3.pdf
2. Milionis, J., Moallemi, C. C., Roughgarden, T., Zhang, A. L. (2022). *Automated market making and loss-versus-rebalancing,* arXiv preprint arXiv:2208.06046. https://doi.org/10.48550/arXiv.2208.06046

3. Fan, Z., Marmolejo-Cossío, F., Moroz, D. J., Neuder, M., Rao, R., Parkes, D. C. (2023). *Strategic liquidity provision in Uniswap V3,* 5th Conference on Advances in Financial Technologies (AFT 2023), LIPIcs, **282**, 25:1–25:22. https://doi.org/10.4230/LIPIcs.AFT.2023.25
4. Cartea, Á., Drissi, F., Monga, M. (2024). *Decentralized finance and automated market making: Predictable loss and optimal liquidity provision,* SIAM Journal on Financial Mathematics, **15** (3), 931–959. https://doi.org/10.1137/23M1602103
5. Echenim, M., Gobet, E., Maurice, A.-C. (2026). *Uniswap V3: Impermanent loss modeling and swap fees asymptotic analysis,* SIAM Journal on Financial Mathematics, **17** (1), 1–40. https://doi.org/10.1137/23M1606149
6. Heimbach, L., Schertenleib, E., Wattenhofer, R. (2022). *Risks and returns of Uniswap V3 liquidity providers,* Proceedings of the 4th ACM Conference on Advances in Financial Technologies (AFT '22), 89–101. https://doi.org/10.1145/3558535.3559772
7. Hamilton, J. D. (1989). *A new approach to the economic analysis of nonstationary time series and the business cycle,* Econometrica, **57** (2), 357–384. https://doi.org/10.2307/1912559
8. Rabiner, L. R. (1989). *A tutorial on hidden Markov models and selected applications in speech recognition,* Proceedings of the IEEE, **77** (2), 257–286. https://doi.org/10.1109/5.18626
9. Cappé, O. (2011). *Online EM algorithm for hidden Markov models,* Journal of Computational and Graphical Statistics, **20** (3), 728–749. https://doi.org/10.1198/jcgs.2011.09109
10. Kelly, J. L., Jr. (1956). *A new interpretation of information rate,* Bell System Technical Journal, **35** (4), 917–926. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x
11. Øksendal, B., Sulem, A. (2019). *Applied Stochastic Control of Jump Diffusions,* 3rd ed., Springer, Cham.
12. Asmussen, S. (2003). *Applied Probability and Queues,* 2nd ed., Springer, New York.
13. Norris, J. R. (1997). *Markov Chains,* Cambridge University Press, Cambridge.
14. Levin, D. A., Peres, Y. (2017). *Markov Chains and Mixing Times,* 2nd ed., American Mathematical Society, Providence, RI.

УДК 519.217.2:519.86

Оцінка цінності режимної інформації через час релаксації в ергодичному керуванні позиціями концентрованої ліквідності

Ілля Касаджик, Олександр Васильєв

Анотація. Розглянуто задачу вибору ширини цінового діапазону позиції провайдера ліквідності в автоматизованому маркет-мейкері з концентрованою ліквідністю (Uniswap V3), коли ринковий режим волатильності є прихованим станом марковського ланцюга і спостерігається лише через апостеріорний розподіл онлайн-фільтра. У класі симетричних політик з перевідкриттям на межі отримано формулу ергодичного темпу зростання логарифму капіталу в замкненому вигляді з логарифмічною втратою за цикл; формулу звірено з моделюванням Монте-Карло відповідного випадкового блукання та з подієвим бек-тестом на п'ятирічному хвилинному ряді пулу USDC/WETH 0,05 %. Показано, що оцінювання умовних за режимом моментів за тим самим спостереженням, яке породило апостеріорний розподіл, не дає прогнозу: контраст волатильності між режимами завищується приблизно вчетверо проти прогнозної оцінки на добовому горизонті, а «оптимальні» ширини, що з такої оцінки випливають, є артефактом. Запропоновано прогнозні моменти на спільному горизонті та самоузгоджене калібрування горизонту через очікувану тривалість циклу. Доведено, що цінність режимної інформації для вибору ширини невід'ємна і обмежена зверху величиною, пропорційною відношенню часу релаксації ланцюга до горизонту рішення; для підігнаної хвилинної моделі час релаксації становить 19 хвилин, тоді як позиція живе тижні й довше. Емпірична перевірка на неперекривних вікнах з перестановковим тестом узгоджується з теоремою: виграш режимної політики за причинним апостеріорним розподілом не перевищує 0,75 відсоткового пункту річних (за оракульним станом 1,2), і жодне зі значень не відрізняється від нуля статистично.

Ключові слова: концентрована ліквідність, приховані марковські моделі, ергодичне керування, час релаксації, цінність інформації, темп зростання капіталу.

Список використаних джерел

1. Adams H., Zinsmeister N., Salem M., Keefer R., Robinson D. Uniswap v3 Core. 2021. URL: https://uniswap.org/whitepaper-v3.pdf
2. Milionis J., Moallemi C. C., Roughgarden T., Zhang A. L. Automated Market Making and Loss-Versus-Rebalancing. *arXiv preprint arXiv:2208.06046*. 2022. DOI: https://doi.org/10.48550/arXiv.2208.06046
3. Fan Z., Marmolejo-Cossío F., Moroz D. J., Neuder M., Rao R., Parkes D. C. Strategic Liquidity Provision in Uniswap V3. *5th Conference on Advances in Financial Technologies (AFT 2023). Leibniz International Proceedings in Informatics (LIPIcs)*. 2023. Vol. 282. P. 25:1–25:22. DOI: https://doi.org/10.4230/LIPIcs.AFT.2023.25
4. Cartea Á., Drissi F., Monga M. Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision. *SIAM Journal on Financial Mathematics*. 2024. Vol. 15, No. 3. P. 931–959. DOI: https://doi.org/10.1137/23M1602103
5. Echenim M., Gobet E., Maurice A.-C. Uniswap V3: Impermanent Loss Modeling and Swap Fees Asymptotic Analysis. *SIAM Journal on Financial Mathematics*. 2026. Vol. 17, No. 1. P. 1–40. DOI: https://doi.org/10.1137/23M1606149
6. Heimbach L., Schertenleib E., Wattenhofer R. Risks and Returns of Uniswap V3 Liquidity Providers. *Proceedings of the 4th ACM Conference on Advances in Financial Technologies (AFT '22)*. 2022. P. 89–101. DOI: https://doi.org/10.1145/3558535.3559772
7. Hamilton J. D. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*. 1989. Vol. 57, No. 2. P. 357–384. DOI: https://doi.org/10.2307/1912559
8. Rabiner L. R. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*. 1989. Vol. 77, No. 2. P. 257–286. DOI: https://doi.org/10.1109/5.18626
9. Cappé O. Online EM Algorithm for Hidden Markov Models. *Journal of Computational and Graphical Statistics*. 2011. Vol. 20, No. 3. P. 728–749. DOI: https://doi.org/10.1198/jcgs.2011.09109
10. Kelly J. L., Jr. A New Interpretation of Information Rate. *Bell System Technical Journal*. 1956. Vol. 35, No. 4. P. 917–926. DOI: https://doi.org/10.1002/j.1538-7305.1956.tb03809.x
11. Øksendal B., Sulem A. Applied Stochastic Control of Jump Diffusions. 3rd ed. Cham: Springer, 2019. 436 p.
12. Asmussen S. Applied Probability and Queues. 2nd ed. New York: Springer, 2003. 438 p.
13. Norris J. R. Markov Chains. Cambridge: Cambridge University Press, 1997. 237 p.
14. Levin D. A., Peres Y. Markov Chains and Mixing Times. 2nd ed. Providence, RI: American Mathematical Society, 2017. 447 p.

Про авторів / About the authors

Ілля Касаджик, здобувач другого (магістерського) рівня вищої освіти, кафедра оптимального керування та економічної кібернетики, факультет математики, фізики та інформаційних технологій, Одеський національний університет імені І. І. Мечникова, вул. Всеволода Змієнка, 2, м. Одеса, 65082, Україна;

Illia Kasadzhyk, Master's student, Department of Optimal Control and Economic Cybernetics, Faculty of Mathematics, Physics and Information Technologies, Odesa I. I. Mechnikov National University, 2 Vsevoloda Zmiienka Str., Odesa 65082, Ukraine;

Олександр Васильєв, кандидат фізико-математичних наук, доцент, кафедра оптимального керування та економічної кібернетики, факультет математики, фізики та інформаційних технологій, Одеський національний університет імені І. І. Мечникова, вул. Всеволода Змієнка, 2, м. Одеса, 65082, Україна;

Oleksandr Vasyliev, Candidate of Science in Physics and Mathematics, Associate Professor, Department of Optimal Control and Economic Cybernetics, Faculty of Mathematics, Physics and Information Technologies, Odesa I. I. Mechnikov National University, 2 Vsevoloda Zmiienka Str., Odesa 65082, Ukraine.

Отримано / Received 28.09.2026
Прийнято до друку / Accepted 26.10.2026
Опубліковано / Published 25.11.2026