

The Ornstein-Uhlenbeck Process for Liquidity

Session 4 · Intuition only — full math in Session 25 (Track 2)

Samir Asaf, PhD, CFA, CMA, CTP, CM&AA

Senior Partner, Regent Financial LLC, New York

Former Finance Instructor, Stanford University

Primary Text: Liquidity Illusion (Forthcoming, 2026)

Graduate Finance Course · Spring 2027 · Session 4 of 32

What we'll cover today

1

Why we need a process model

Description \neq prediction

2

Three properties our model must have

Mean reversion • stochastic • bounded variance

3

The OU process

SDE form and verbal translation

4

What each parameter does

$\kappa \cdot \bar{L} \cdot \sigma$ — calibrated values

5

Self-diagnostic for Track 2

If today felt easy, Track 2 may suit you

Recap: Session 3 in five points

Where we are after the previous lecture:

Liquidity illusion defined

Gap between DCF marks and secondary clearing prices

Empirical evidence shown

4 regimes, 5× mean variation, real \$150B market

Failed patches identified

Three industry attempts to fix DCF internally — all failed

Tension visible

\$100M NAV vs \$65M clearing price on the same fund stake

Today's gap

We have description — we need a process model for prediction

From description (S3) to model (S4).

From data to model: what we have, what we need

WHAT WE HAVE (from Session 3)

- ▶ Time series of secondary discounts 2003–2024
- ▶ Implied liquidity state $L(t)$ at quarterly frequency
- ▶ Four regimes with distinct statistical character
- ▶ Mean / std / autocorrelation by regime

Description is rich. Prediction requires more.

WHAT WE NEED (to value)

- ▶ A MODEL of how $L(t)$ evolves through time
- ▶ Closed-form properties: mean, variance, autocorrelation
- ▶ Conditional distribution: $L(t+h) \mid L(t) = \ell$
- ▶ Stationary distribution: long-run L behavior
- ▶ Forecast-ability for asset cash flow valuation

Without a process model, we describe the past but cannot value the future.

Three properties our model must have

Property 1

Mean reversion

Discounts widen in crises but return toward a long-run average. The unconditional mean exists and is finite. Rules out random walk.

→ Vasicek, CIR, OU all satisfy this

Property 2

Stochastic shocks

Period-to-period changes are partly random — not deterministic. Cannot be predicted from $L(t)$ alone with certainty.

→ Brownian noise dW captures this

Property 3

Bounded variance

Long-run variance is finite. $L(t)$ does not drift indefinitely. Unconditional distribution well-defined.

→ Mean reversion + Brownian noise = bounded variance

These three together rule out simple Brownian motion (unbounded variance) and rule in the OU family.

Why not just use Brownian motion?

Brownian motion is simpler. Why is it wrong for liquidity?

BM unbounded variance

Var grows linearly: $\text{Var}(W_t) = t \rightarrow \infty$. Liquidity is bounded — discounts never exceed 100%.

BM no mean reversion

Once shocked away from a level, stays there. Liquidity clearly reverts: 2009 \rightarrow 2011 in <2 yrs.

BM no persistence pattern

Liquidity persistence varies: 1.5-yr half-life. BM has no such timescale.

BM doesn't fit autocorrelation

Empirical autocorr decays exponentially. BM's would not decay.

OU fixes all four

Mean reversion + Brownian noise + bounded variance + tunable half-life

OU = the simplest extension of BM that captures observed liquidity dynamics.

The Ornstein-Uhlenbeck process

The SDE (Stochastic Differential Equation):

$$dL_t = \kappa(\bar{L} - L_t) dt + \sigma dW_t$$

Verbal translation:

$L(t)$ tends to revert toward \bar{L} at speed κ , with random shocks of magnitude σ .

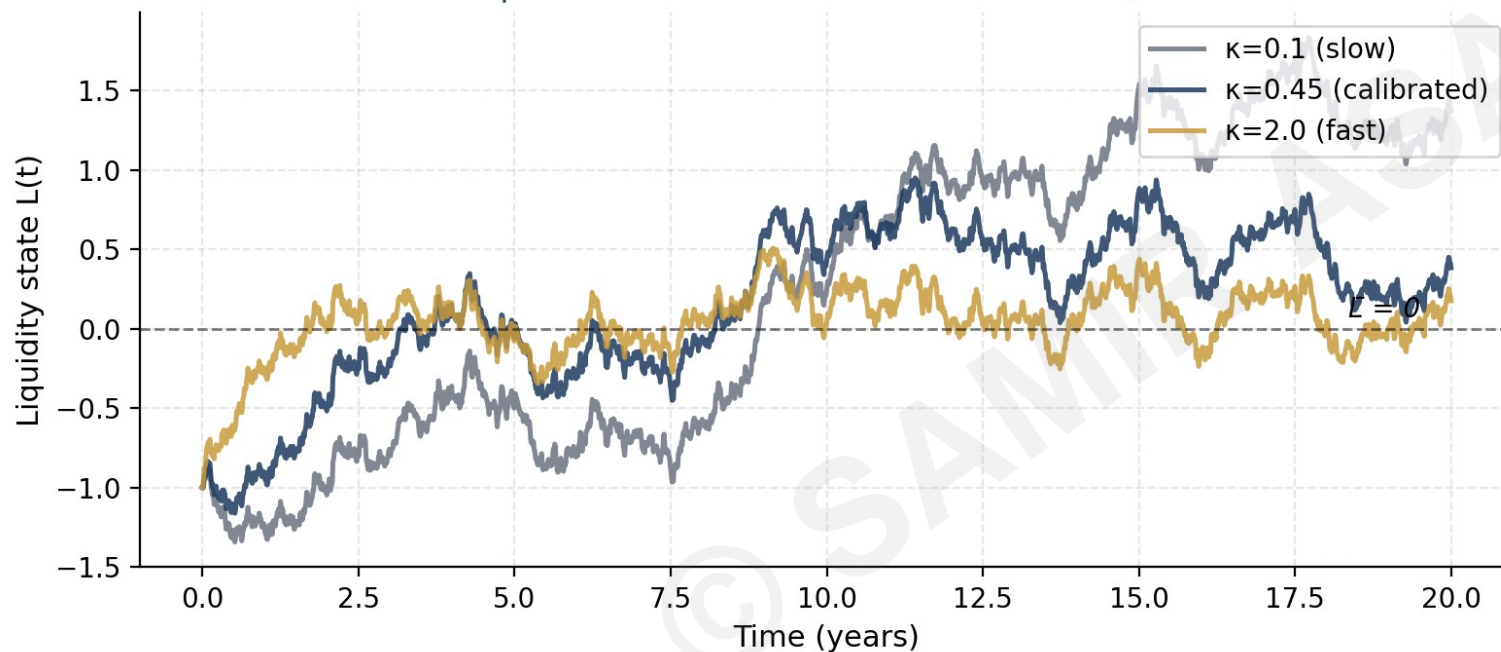
History

Originally derived by Leonard Ornstein and George Uhlenbeck in 1930 to model the velocity of a particle in a viscous fluid.

Same mathematics describes any mean-reverting stochastic variable: rates (Vasicek), volatility, and our liquidity state.

κ controls mean reversion speed

Three OU paths with same initial state and noise, different κ



Half-life formula

$$\text{Half-life} = \frac{\ln 2}{\kappa}$$

Time to revert halfway to \bar{L} .

κ	Half-life
$\kappa = 0.1$	6.9 years
$\kappa = 0.45$	1.5 years
$\kappa = 2.0$	0.35 years

Calibrated $\kappa=0.45$ (1.5-yr half-life) matches secondary market data.

Analogy: radioactive decay uses the same half-life math.

\bar{L} (long-run mean) and σ (volatility)

Calibrated GE-LAV parameters: $\kappa=0.45/\text{yr}$ · $\bar{L}=1.0$ · $\sigma=0.32$ · stationary std ≈ 0.337

\bar{L} economic role

State liquidity gravitates toward in the long run, absent persistent shocks

\bar{L} calibration

$\bar{L}=1.0$ corresponds to historical average premium near 3.5%

σ economic role

Magnitude of period-to-period random shocks to the liquidity state

σ calibration

$\sigma=0.32$ produces stationary std of ~ 0.337 , matching regime-to-regime variation

Joint calibration

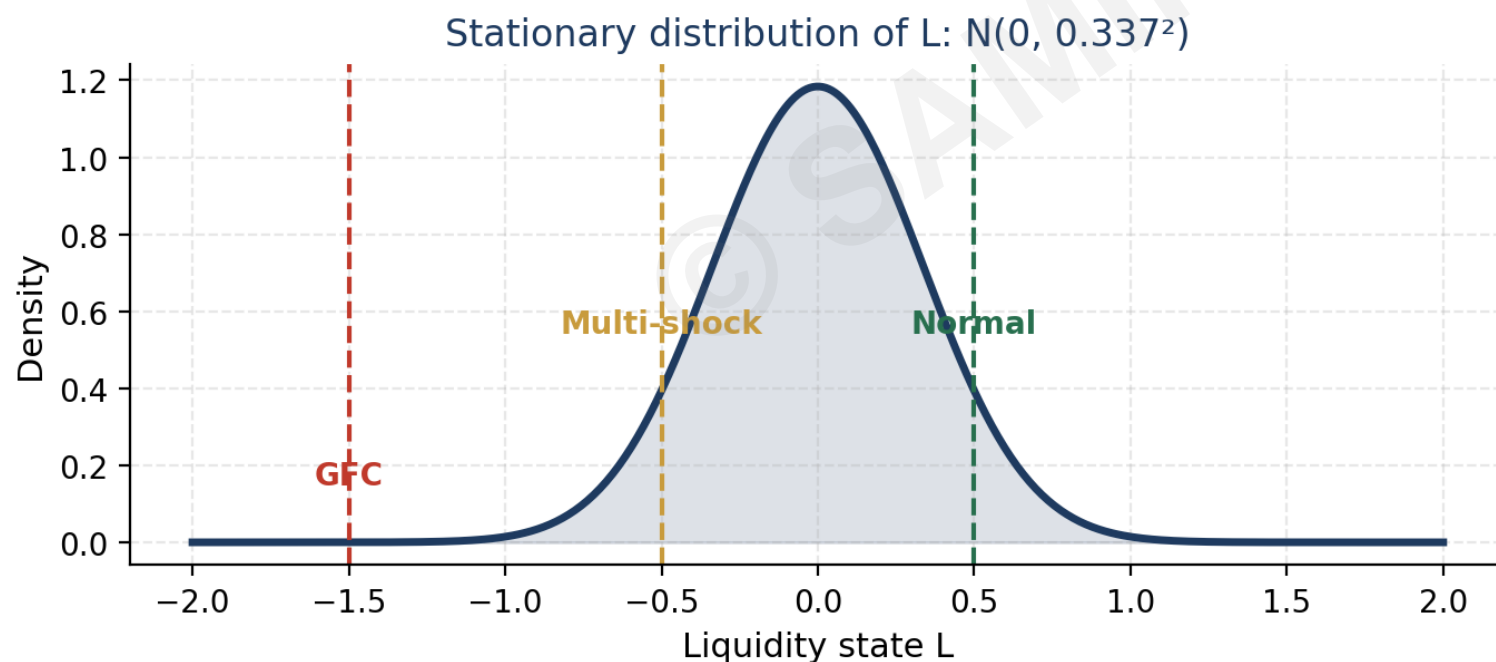
κ, \bar{L}, σ jointly fit to 2003-2024 secondary market data via MLE

See Book Chapter 7 for full calibration methodology and likelihood derivation.

Stationary distribution: where $L(t)$ lives long-run

If we run the OU process long enough, distribution of $L(t)$ converges to a fixed Gaussian.

$$L_{\infty} \sim N\left(\bar{L}, \frac{\sigma^2}{2\kappa}\right)$$



Interpretation

Long-run mean

$$E[L_{\infty}] = \bar{L} = 0$$

Long-run variance

$$\text{Var}[L_{\infty}] = \sigma^2/(2\kappa) \approx 0.114$$

Long-run std dev

$$\approx 0.337$$

GFC depth ($L = -1.5$)

$$\sim 3.7\sigma \text{ event}$$

Frequency: GFC

$$< 1\% \text{ of time}$$

Conditional distribution: $L(t+h)$ given $L(t)$

OU's most useful property: closed-form conditional distribution at any horizon h

Conditional mean

$$E[L(t+h) \mid L_t = \ell] = \bar{L} + (\ell - \bar{L}) \cdot \exp(-\kappa h) \text{ — exponential decay to } \bar{L}$$

Conditional variance

$$\text{Var}[L(t+h) \mid L_t = \ell] = \sigma^2 / (2\kappa) \cdot (1 - \exp(-2\kappa h)) \text{ — grows to stationary}$$

Useful for forecasting

Given today's L , compute mean & variance of future L at any horizon

Useful for valuation

Integrate cash flows weighted by transition density

Used in Session 11

Optimal exit timing uses transition density to compute continuation value

Track 2: derivation in Karatzas-Shreve §5.6. Track 1: applied in Sessions 11-12.

Connection to models you may already know

The OU process has many cousins in finance. Same math, different domains:

Model	Author / Year	Domain	Why OU-like
Vasicek	Vasicek 1977	Short interest rates	$dr = \kappa(\bar{r} - r)dt + \sigma dW$ — same SDE
CIR	Cox-Ingersoll-Ross 1985	Interest rates (non-negative)	Square-root variant, similar mean reversion
Hull-White	Hull-White 1990	Term-structure models	Extends Vasicek with time-varying params
Heston	Heston 1993	Stochastic volatility	Variance is mean-reverting (CIR cousin)
Schwartz	Schwartz 1997	Commodity prices	Log-prices follow OU
GE-LAV	Asaf 2026	Private market liquidity	$L(t)$ follows OU

Discrete-time recursion: how to simulate OU

Euler-Maruyama discretization for hands-on simulation (problem set 1 uses this):

Step 1: Choose Δt

Small time step, e.g., $\Delta t = 1/12$ (monthly) or $1/4$ (quarterly)

Step 2: Draw $\varepsilon \sim N(0,1)$

Standard normal random variable, independently at each step

Step 3: Update L

$$L_{\{t+\Delta t\}} = L_t + \kappa(\bar{L} - L_t) \cdot \Delta t + \sigma \cdot \sqrt{\Delta t} \cdot \varepsilon$$

Step 4: Iterate

Apply repeatedly to generate a path of arbitrary length

Convergence

As $\Delta t \rightarrow 0$, discrete path converges to true OU path

Code template provided in problem set 1. About 8 lines of Python.

Self-diagnostic: are you ready for Track 2?

Track 2 (Researcher) does the full OU mathematics in Session 25 — existence, uniqueness, stationary distribution proof.

If today felt easy

Consider Track 2 (Researcher)

- ✓ You're comfortable with probability theory limits
- ✓ Stochastic calculus and Itô's lemma feel approachable
- ✓ You want to do GE-LAV proofs, not just apply them
- ✓ Career direction: PhD or quant research

If today felt hard (or right at the edge)

Track 1 (Practitioner) is for you

- ✓ You can apply OU intuition; deep proofs not required
- ✓ Career direction: LP, GP, regulator, consulting
- ✓ You want case-driven sessions in S25-31
- ✓ Math intuition lectures (S19-24) cover what you need

Decision by end of Session 5. Default = Track 1.

Why the OU process specifically

”

If the question is 'what is the simplest model that captures mean reversion, stochastic shocks, and bounded variance,' the answer is Ornstein-Uhlenbeck. Anything simpler fails one of those tests; anything more complex needs to be justified.

— Liquidity Illusion, Chapter 2, §2.7

The OU choice is principled, not arbitrary. Track 2 students will revisit this in Session 25.

Session 4 summary

What we accomplished today

- 1 The OU process models $L(t)$ as mean-reverting around \bar{L} with random shocks of magnitude σ
- 2 Three parameters: κ (speed), \bar{L} (long-run mean), σ (vol). Calibrated: $\kappa=0.45/\text{yr}$, $\bar{L}=1.0$, $\sigma=0.32$
- 3 Half-life = $\ln(2)/\kappa \approx 1.5$ years (matches data)
- 4 Stationary distribution: $\text{Normal}(\bar{L}, \sigma^2/2\kappa) \rightarrow$ GFC at $L=-1.5$ is $\sim 3.7\sigma$ event
- 5 OU has many cousins: Vasicek (rates), Heston (vol), Hull-White (term structure)

Next session

Session 5: Term structure of private capital — how the OU process pricing applies across asset lifetimes