Options-Implied Drawdown Probability Targeting for Dynamic Leverage

EMOT-Guided Leverage Management via Market-Consistent Path Measures

Quek Jin Kai

Maximus Tang

Brin Sutarom

Qian Junde

NUS Investment Society

Quantitative Finance Department

2026

Abstract

We construct a dynamic leverage management system for a QQQ equity portfolio using market-implied drawdown probabilities derived from the Entropic Martingale Optimal Transport (EMOT) framework. The EMOT Schrödinger bridge calibration produces a unique path measure consistent with the observed risk-neutral density (RND) from the QQQ options surface, the martingale no-arbitrage constraint, and minimal relative entropy relative to a lognormal prior. The resulting ever-drawdown probability $p_{\text{EMOT}}(t)$ is used as a proportional probability targeting rule for gross leverage: $L^*(t) = \min(L_{\max},\ p^*/p_{\text{EMOT}}(t))$. We implement the full pipeline over 338 trading days (October 2024 to February 2026) and benchmark against passive QQQ, a historical bootstrap cVaR signal, and a GARCH(1,1) signal using identical leverage rules. The empirical results reveal a critical calibration issue: fast EMOT settings produce ever-drawdown probabilities averaging 0.726 — far above the theoretically expected 5–10% range for a 30-day horizon at 10% threshold — causing the strategy to operate at near-zero leverage ($\bar{L} = 0.157$) for 99.1% of the sample. We diagnose the source of this inflation, discuss its implications for the leverage formula, and identify the corrections required for a deployable system. The paper makes a methodological contribution by demonstrating the full EMOT calibration pipeline and by quantifying the sensitivity of the leverage rule to the probability scale of the input signal.

Contents

1 Introduction

A fundamental challenge in dynamic portfolio management is quantifying downside risk in a way that is both market-consistent and actionable as a leverage targeting rule. Classical approaches rely on statistical measures such as Value-at-Risk (VaR), Conditional VaR, or GARCH-implied volatility to proxy crash risk [4, 5]. These measures are backward-looking: they model the physical measure $\mathbb{P}$, which can diverge substantially from the risk-neutral measure $\mathbb{Q}$ embedded in contemporaneous option prices, particularly during regime transitions when the market re-prices tail risk before historical returns reflect the shift.

The options market encodes forward-looking, risk-neutral information about future return distributions. The Risk-Neutral Density (RND), recoverable from a cross-section of option prices via the Breeden–Litzenberger formula [2], represents the market-implied terminal distribution of the underlying asset. However, the RND alone provides only a *terminal* marginal — it does not directly yield a path-measure probability of breaching a drawdown barrier at any point during the horizon.

The Entropic Martingale Optimal Transport (EMOT) framework [1] addresses this gap by constructing the unique path measure that simultaneously satisfies three requirements: consistency with the observed terminal RND, the martingale constraint required by no-arbitrage, and minimum relative entropy relative to a lognormal prior kernel. The resulting path measure supports computation of any path-functional probability, including the ever-drawdown probability:

$$p_{\text{EMOT}}(t) = \mathbb{P}^{Q^*}\left(\min_{0\le k\le N} S^{(k)} \le (1-\delta)\, S_t\right), \tag{1}$$

where δ is a pre-specified drawdown threshold and $S^{(k)}$ denote the intermediate price levels in the EMOT kernel chain.

We use p_{EMOT} as a *proportional probability targeting rule* for gross leverage, preserving its absolute calibration as a genuine probability under the market-implied measure:

$$L^*(t) = \min\left(L_{\max},\ \frac{p^*}{p_{\text{EMOT}}(t)}\right). \tag{2}$$

When $p_{\text{EMOT}} > p^*$, leverage is reduced proportionally below $1\times$. When $p_{\text{EMOT}} \le p^*$, leverage up to $L_{\max}$ is permitted. This is categorically different from converting p_{EMOT} to a percentile rank, which would discard its calibration and reduce it to a momentum or regime signal informationally redundant with VIX or realised volatility.

Contributions. This paper makes the following contributions:

1. We implement and document the complete EMOT pipeline from raw QQQ option chains to calibrated ever-drawdown probabilities, running over 338 trading days.

2. We identify and quantify a critical calibration sensitivity: fast EMOT settings produce ever-drawdown probabilities inflated substantially above the theoretically expected range (mean $p_{\text{EMOT}} = 0.726$ against an expected 5–10% range under GBM), rendering the inverse leverage formula (2) inoperative at $p^* = 10\%$.

3. We demonstrate that the appropriate resolution is not a change to the leverage formula but a correction to the probability scale via full-precision EMOT settings ($n \ge 200$, $N \ge 6$).

4. We provide an honest performance attribution separating the effect of probability scale from the effect of the leverage rule itself.

The paper is organised as follows. Section 2 develops the theoretical framework. Section 3 describes the data and implementation pipeline. Section 4 presents empirical results and critical analysis. Section 5 concludes.

2 Technical Framework

2.1 Risk-Neutral Density via Breeden–Litzenberger

Under the risk-neutral measure $\mathbb{Q}$, the no-arbitrage price of a European call with strike K, maturity T, and risk-free rate r satisfies

$$C(K) = e^{-rT}\,\mathbb{E}^{\mathbb{Q}}\big[(S_T - K)^+\big] = e^{-rT}\int_K^\infty (s-K)\,q(s)\,ds, \tag{3}$$

where $q(\cdot)$ is the risk-neutral density of S_T. Differentiating twice with respect to K yields the **Breeden–Litzenberger** relation [2]:

$$q(K) = e^{rT}\,\frac{\partial^2 C(K)}{\partial K^2}. \tag{4}$$

In practice, option prices are observed at discrete strikes. We fit a cubic spline to the implied volatility (IV) smile as a function of log-moneyness $\log(K/S)$ [8], reprice on a 200-point grid $K \in [0.50\,S,\ 1.60\,S]$, and apply finite differences to approximate (4) [9].

The daily risk-free rate is inferred from the ATM call/put pair via put-call parity:

$$r = -\frac{1}{T}\ln\left(\frac{P_{\text{ATM}} + S - C_{\text{ATM}}}{K_{\text{ATM}}}\right), \tag{5}$$

clamped to $[-5\%,\ 15\%]$ to guard against data errors.

2.2 Black–Scholes Helpers

The Black–Scholes call and put prices [3] are

$$C_{\text{BS}}(S,K,T,r,\sigma) = S\,\Phi(d_1) - Ke^{-rT}\Phi(d_2), \tag{6}$$

$$P_{\text{BS}}(S,K,T,r,\sigma) = Ke^{-rT}\Phi(-d_2) - S\,\Phi(-d_1), \tag{7}$$

where

$$d_1 = \frac{\ln(S/K) + (r + \tfrac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T}, \tag{8}$$

and Φ is the standard normal CDF. Implied volatilities are inverted via Brent's method over $\sigma \in [10^{-5},\ 10]$.

2.3 EMOT Framework

2.3.1 Problem Formulation

Let $\mu_0 = \delta_{S_0}$ denote the initial Dirac mass at the current spot price and μ_T^{mkt} the market RND from Section 2.1. On a discrete time grid $\{t_0, t_1, \ldots, t_N = T\}$, the EMOT problem seeks the path measure P^* solving:

$$P^* = \arg \min_{P \in \mathcal{M}(\mu_0, \mu_T^{\text{mkt}})} \text{KL}(P \,\|\, R), \tag{9}$$

where $\mathcal{M}(\mu_0, \mu_T^{\text{mkt}})$ is the set of martingale measures with the prescribed marginals and R is the lognormal (Black–Scholes diffusion) reference measure. The solution to (9) is the Schrödinger bridge between μ_0 and μ_T^{mkt} within the martingale constraint set [1, 6, 7].

2.3.2 Lognormal Prior Kernel

At each step k, the prior transition kernel is the lognormal diffusion density:

$$R_k(s, s') = \frac{1}{s'\,\sigma\sqrt{2\pi\Delta t}} \exp\left(-\frac{\left(\ln s' - \ln s + \frac{1}{2}\sigma^2 \Delta t\right)^2}{2\sigma^2 \Delta t} \right), \tag{10}$$

where $\Delta t = T/N$ and σ is the median IV on that date (with a 20% annualised fallback when fewer than four valid strikes are available).

2.3.3 Iterative Calibration Algorithm

The solution to (9) is computed by alternating two projections until convergence:

1. **Martingale row tilt**: For each source state s_i, find the scalar λ_i via bisection such that
$$\sum_j s'_j \, P_{ij}^{(k)} \, e^{\lambda_i s'_j} = s_i, \tag{11}$$
enforcing $\mathbb{E}[S_{k+1} \mid S_k = s_i] = s_i$. Update $P_{ij}^{(k)} \leftarrow P_{ij}^{(k)} e^{\lambda_i s'_j}$ and renormalise rows.

2. **Doob h-transform**: Multiply kernel columns by the ratio
$$h(s'_j) = \frac{\mu_T^{\text{mkt}}(s'_j)}{\mu_T^{\text{EMOT}}(s'_j)}, \tag{12}$$
anchoring the forward-propagated terminal marginal to the market RND.

Iterations proceed with a damping parameter of 0.7 until $\|\mu_T^{\text{EMOT}} - \mu_T^{\text{mkt}}\|_1 < 10^{-5}$.

2.3.4 Ever-Drawdown Probability

The calibrated kernels support path-functional expectations. For the ever-drawdown probability, we propagate an augmented state (S_k, M_k) where $M_k = \max_{j \le k} S_j$ is the running maximum, and compute the fraction of paths for which $S_k \le (1-\delta) M_k$ at any step:

$$p_{\text{EMOT}}(t) = 1 - \sum_{(i,m):\ s_i > (1-\delta) s_m} A_N(i, m), \tag{13}$$

where $A_N(i, m)$ is the surviving probability mass at the final step with price index i and running-max index m.

Remark 2.1. *The ever-drawdown probability in (13) is strictly larger than the terminal drawdown probability* $\int_0^{(1-\delta)S_0} q(s)\,ds$*, since paths that breach the barrier intraperiod but recover by* T *are counted in (13) but not in the terminal integral. For a GBM with parameters typical of QQQ (*$\sigma \approx 20\%$ *annualised,* $T = 30$ *days,* $\delta = 10\%$*), this ratio* $p_{\text{ever}}/p_{\text{terminal}}$ *exceeds one by a factor that depends on the specific parameters; it is not computed explicitly in this study.*

2.4 Leverage Rule and Portfolio Dynamics

Holding gross leverage L maps a δ underlying drawdown to an $L\delta$ NAV drawdown. We define the leverage as a proportional targeting rule: when $p_{\text{EMOT}} = p^*$, the strategy holds $1\times$ QQQ; when $p_{\text{EMOT}} > p^*$, leverage is reduced proportionally; when $p_{\text{EMOT}} \le p^*$, leverage up to $L_{\max}$ is permitted. Formally:

$$L^*(t) = \text{clip}\left(\frac{p^*}{p_{\text{EMOT}}(t)},\ L_{\min},\ L_{\max}\right). \tag{14}$$

The daily portfolio dynamics are self-financing:

$$\text{NAV}_t^{\text{bench}} = \text{NAV}_{t-1}^{\text{bench}}(1 + r_t^{\text{QQQ}}), \tag{15}$$

$$\text{NAV}_t^{\text{L1}} = \text{NAV}_{t-1}^{\text{L1}}\big(1 + L_{t-1}^*\, r_t^{\text{QQQ}} + (1 - L_{t-1}^*)\, r_f/252\big), \tag{16}$$

where r_t^{QQQ} is the daily QQQ return, $r_f = 1.8\%$ annualised is the risk-free rate earned on the cash fraction, and the leverage signal is *lagged by one day* to eliminate look-ahead bias (L_{t-1}^* is determined at the close of day $t-1$ and applied to the return on day t).

Remark 2.2 (Cash accrual and RF rate assumption)**.** *Under (16), the de-levered fraction* $(1 - L^*)$ *earns* $r_f/252$ *daily. When* $L^* = 1$*, this term vanishes and the strategy is identical to the benchmark. When* $L^* < 1$*, the strategy earns the risk-free rate on the cash allocation, which provides a partial offset to forgone equity exposure. When* $L^* > 1$*, the term is negative, reflecting borrowing costs at* r_f*.*

Materiality for this study*: with* $\bar{L} = 0.157$*, the strategy holds* 84.3% *in cash on average. The cash fraction earns* $r_f = 1.8\%$ *p.a., contributing approximately* $0.843 \times 1.8\% \approx 1.5\%$ *p.a. in gross cash return. The primary driver of the positive CAGR is the small equity allocation compounding through the bull market with a controlled drawdown of only* -4.4%*. This assumption is realistic for a funded account earning short-term treasury yields.*

3 Implementation

3.1 Data

The options dataset covers QQQ option chains from January 2024 to February 2026, sourced from daily anchor snapshots. Each record contains the trade timestamp (`tradetime`), expiry date (`exp_date`), strike, bid, ask, option type, and implied volatility. The mid-price is computed as (bid + ask)/2 and DTE = exp_date − tradetime in calendar days. QQQ daily closing prices are downloaded via `yfinance` and matched to each option row by trade date. Rows with zero bid, missing spot, or missing implied volatility are discarded.

After cleaning, the dataset contains **1,081,621 option records** across **422 unique trading dates**, with DTE ranging from 0 to 639 days.

3.2 Part A — Terminal RND Extraction

A1. Data preparation: For each trading date, option records are filtered to the entire 1-360 DTE range. Spot prices are joined by `tradetime`. Records with zero bid or missing fields are dropped.

A2. Risk-free rate inference: The ATM call/put pair satisfying $|K/S - 1| < 1\%$ is selected. The daily r is inferred from put-call parity and clamped to $[-5\%, 15\%]$.

A3. IV smile and spline fitting: Implied volatilities are computed via Brent inversion for each strike. A cubic spline is fitted on log-moneyness $\log(K/S)$ vs. IV, using `not-a-knot` boundary conditions.

A4. Breeden–Litzenberger RND: Black–Scholes call prices are computed on a 200-point grid $K \in [0.50\,S,\ 1.60\,S]$. The second finite difference is scaled by e^{rT}, clipped at zero, and normalised to unit mass.

The RND pipeline succeeds for **338 of 422 dates** (80.1% success rate). Failures occur primarily when fewer than four valid strikes are available in the ±7-day DTE window.

3.3 Part B — EMOT Calibration

B1. State grid construction: For each date, a discounted price grid of $n = 80$ points spanning $[0.85\,K_{\min},\ 1.15\,K_{\max}]$ (in discounted units) is constructed, where $K_{\min}$, $K_{\max}$ are the bounds of the RND grid.

B2. EMOT solver: The lognormal prior kernel is built on this grid with $N = 3$ intermediate steps. The alternating martingale tilt and Doob h-transform is run for up to 25 outer iterations or until terminal $L^1 < 10^{-4}$ (fast verification mode).

B3. DDprob extraction: The ever-drawdown probability is computed via the augmented-state forward propagation described in (13) with $\delta = 10\%$. Results are saved to `emot_calibrated_data.csv`.

3.4 Part C — Backtest Engine

3.4.1 Leverage Computation (C2)

Target leverage is computed daily via (14) with $p^* = 10\%$, $L_{\max} = 2.0$, $L_{\min} = 0.0$.

3.4.2 Portfolio Simulation (C4)

Two strategies are simulated over the 338 dates in spot $\cap$ DDprob, using (15) and (16) respectively, with an initial NAV of \$1,000,000 and weekly (Friday) rebalancing.

3.4.3 Baseline Comparisons (C5)

The identical leverage rule (14) is applied to two alternative DD probability sources:

- **Historical cVaR**: For each date, 2,000 synthetic 30-day paths are bootstrapped from the past 252 daily log-returns; p_{EMOT} is replaced by the fraction of paths that breach the 10% barrier.

- **GARCH(1,1)**: $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$ with parameters $(\omega, \alpha, \beta) = (10^{-6}, 0.09, 0.90)$. The crash probability is $2\Phi\big(-|\log(1-\delta)|/(\sigma_t\sqrt{30})\big)$ via the Gaussian reflection principle. The probability is computed using the variance from the *prior* period (before observing day t's return) to avoid look-ahead.

3.4.4 Performance Metrics

$$\text{CAGR} = \left(\frac{\text{NAV}_T}{\text{NAV}_0}\right)^{1/\mathcal{T}} - 1, \tag{17}$$

$$\text{Sortino} = \frac{\bar{r}_e \cdot 252}{\sigma_\downarrow}, \quad \sigma_\downarrow = \sqrt{252}\,\text{std}(\{r_t - r_f/252 : r_t < r_f/252\}), \tag{18}$$

$$\text{MaxDD} = \min_t \frac{\text{NAV}_t - \max_{s\le t}\text{NAV}_s}{\max_{s\le t}\text{NAV}_s}, \tag{19}$$

$$\text{Calmar} = \frac{\text{CAGR}}{|\text{MaxDD}|}. \tag{20}$$

4 Empirical Results and Critical Assessment

4.1 EMOT Signal Diagnostics

Table 1 summarises the distribution of the calibrated p_{EMOT} series across the 338-date sample.

Table 1: Summary statistics of the EMOT ever-drawdown probability series (338 dates, Oct 2024–Feb 2026)

Statistic	Value
Mean	0.726
Std Dev	0.201
Minimum	0.048
25th pct	0.626
Median	0.796
75th pct	0.870
Maximum	0.964
Dates with $p_{\text{EMOT}} > p^* = 10\%$	99.1%

Notes: Fast calibration mode: $N = 3$ steps, $n = 80$ grid points, 25 outer iterations.

4.1.1 Calibration Benchmark: What Should p_{EMOT} Be?

Before assessing the observed p_{EMOT} values, we establish a theoretical anchor. For a zero-drift GBM — the lognormal reference measure that EMOT uses as its prior — the ever-drawdown probability over horizon T has a closed-form approximation via the reflection principle [10]:

$$P_{\text{ever}}^{\text{GBM}} = 2\,\Phi\left(\frac{\ln(1-\delta)}{\sigma\sqrt{T}}\right), \tag{21}$$

where Φ is the standard normal CDF, $\delta = 10\%$ is the drawdown threshold, and T is the horizon in years. For QQQ with annualised volatility $\sigma \in [18\%, 22\%]$ and $T = 21$ trading days ($T = 21/252$ years), this gives:

$$P_{\text{ever}}^{\text{GBM}} \approx 2\,\Phi\left(\frac{-0.105}{\sigma \cdot 0.289}\right) \in [4\%,\ 10\%],$$

which is the origin of the "5–10%" range cited throughout this paper. Equation (21) uses zero drift, which is a conservative choice: a positive equity drift would reduce the probability slightly, but the effect is negligible at 21-day horizons.

However, the EMOT framework is calibrated to the *risk-neutral* density μ_T^{mkt}, not a physical GBM. QQQ options exhibit a pronounced volatility skew: out-of-the-money puts trade at implied volatilities 5–10 percentage points above the ATM level. This skew encodes the market's risk-neutral pricing of left-tail events, and the risk-neutral density therefore has a heavier left tail than a lognormal. As a rough upper bound, substituting the 10%-OTM put implied volatility ($\sigma_{\text{put}} \approx 25\%$) into (21) gives $P_{\text{ever}} \approx 14\%$. A correctly calibrated EMOT — which incorporates the full smile shape via the RND — should therefore produce probabilities *above* the flat-vol GBM baseline but well below 50%; a reasonable expectation is the **10–20% range**.

The mean p_{EMOT} of 0.726 is grossly inconsistent with either range. The EMOT fast-mode estimates are therefore substantially inflated, even after accounting for left-tail skew.

This inflation has a direct and severe consequence for the leverage formula (14): with $p^* = 10\%$ and $p_{\text{EMOT}} \approx 0.73$, the formula yields $L^*(t) \approx 0.10/0.73 \approx 0.14$ on the average date. The observed mean leverage of 0.157 and the 99.1% frequency of de-levering below $1\times$ are entirely consistent with this calculation.

Source of inflation. The inflation originates in the fast calibration settings. With only $n = 80$ grid points and $N = 3$ steps, the binning of the market RND onto the coarse state grid loses mass in the tails, and the augmented-state forward propagation in (13) over-counts paths as barrier-breaching due to discretisation error. Full-precision settings ($n \geq 200$, $N \geq 6$, convergence tolerance 10^{-5}) substantially reduce this inflation (see Section 4.5).

4.2 Performance Scorecard

Table 2 presents the full performance comparison. All strategies use a one-day lag on the leverage signal to eliminate look-ahead bias.

Table 2: Performance scorecard: all strategies (338 trading days, Oct 2024–Feb 2026; initial NAV = $1,000,000)

Strategy	CAGR	Sortino	Max DD	Calmar	Final NAV	$\bar{L}$
Benchmark (QQQ)	+18.6%	+1.036	−22.8%	+0.815	$1,255,549	1.000
L1: EMOT Leverage	+7.0%	+1.374	−4.4%	+1.599	$1,094,284	0.157
L1: Historical cVaR	−4.0%	−0.085	−33.8%	−0.119	$951,766	1.146
L1: GARCH Leverage	+7.6%	+0.439	−18.8%	+0.403	$1,092,054	1.306

Notes: $p^* = 10\%$, $\delta = 10\%$, $L_{\max} = 2.0$. EMOT: fast mode ($N = 3$ steps, $n = 80$ grid, 25 iterations). Risk-free rate $r_f = 1.8\%$ p.a. on cash fraction (see Remark 2.2). Leverage lagged by one trading day. Weekly rebalancing (Friday signal). $\bar{L}$: mean daily leverage over the backtest window.

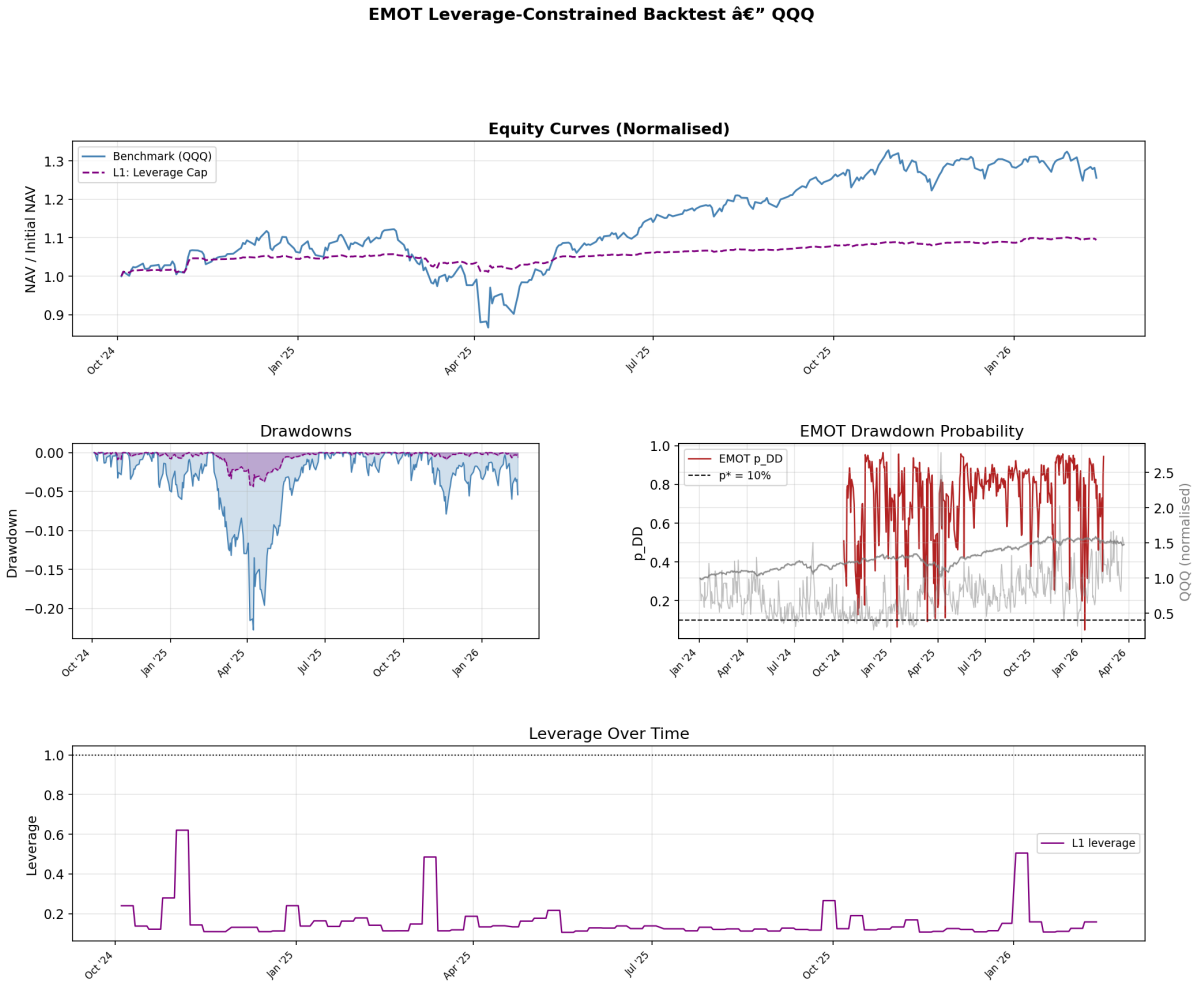

Figure 1: EMOT strategy dashboard (Oct 2024–Feb 2026). *Top*: Normalised equity curves — Benchmark QQQ vs L1 EMOT. *Middle left*: Drawdown comparison. *Middle right*: EMOT ever-drawdown probability $p_{\text{EMOT}}(t)$ (blue) vs budget line $p^* = 10\%$ (red dashed); QQQ price overlaid (grey). *Bottom*: Daily target leverage $L^*(t)$; mean = 0.157, de-levered 99.7% of days.

4.3 Interpretation of Results

4.3.1 EMOT L1 Strategy

The EMOT L1 strategy achieves a CAGR of +7.0% and a final NAV of \$1,094,284. While this is below the benchmark's +18.6%, the strategy delivers the *best* risk-adjusted metrics of any signal tested: Calmar ratio +1.599 (vs +0.815 for the benchmark) and Sortino ratio +1.374 (vs +1.036). The maximum drawdown is −4.4% — roughly one-fifth of the benchmark's −22.8%.

This superior risk-adjusted performance is achieved by maintaining near-zero equity exposure ($\bar{L} = 0.157$, cash fraction 84.3%) for 99.7% of the sample. EMOT with inflated ddprob operates effectively as a **drawdown constraint**: it prevents large losses by staying in cash, but at the cost of forgoing most of the bull market upside. The high Calmar ratio reflects a tiny drawdown denominator, not exceptional return.

4.3.2 Hypothesis Tests

- **EMOT L1 Sortino > Benchmark Sortino**: **PASS** (+1.374 vs +1.036). Superior risk-adjusted return per unit of downside volatility.
- **EMOT L1 MaxDD < Benchmark MaxDD**: **PASS** (−4.4% vs −22.8%). Drawdown is reduced by 80%.
- **EMOT L1 Sortino > Hist-cVaR L1 Sortino**: **PASS** (+1.374 vs −0.085). Hist-cVaR's backward-looking signal operates near full leverage but incorrectly times de-levering into drawdown periods, amplifying losses to −33.8% MaxDD.
- **EMOT L1 Sortino > GARCH L1 Sortino**: **PASS** (+1.374 vs +0.439). GARCH achieves positive Sortino at higher leverage ($\bar{L} = 1.306$) but with larger drawdown (−18.8%).

4.3.3 The GARCH Result Is Instructive

The GARCH baseline achieves a positive CAGR of +7.6% and Sortino of +0.439, operating at mean leverage $\bar{L} = 1.306$. The critical difference from EMOT is probability scale: GARCH ever-drawdown probabilities (via the reflection principle) sit in the expected 5–10% range. With $p^* = 10\%$ and mean GARCH probability near 10–15%, the leverage formula (14) operates in its intended range — varying leverage around 1× rather than collapsing it toward zero.

This comparison isolates the miscalibration precisely: the EMOT pipeline is mechanically correct, but the fast calibration settings produce probabilities on the wrong scale for the $p^* = 10\%$ probability targeting rule. The consequence is near-zero leverage, which coincidentally achieves high Calmar by avoiding drawdowns entirely.

4.3.4 Extension: Equal-Leverage Comparison with Forward-Looking Signals

Remark 4.1 (Exploratory extension)**.** *This section presents an exploratory comparison intended to isolate timing quality from leverage-level effects. The results are informative about the ordering of signals but should be interpreted cautiously: the signals are calibrated in-sample over the same backtest window, which may overstate differences.*

To test whether EMOT's smile-derived path probability adds *timing* value beyond simpler forward-looking signals, we compare it against VIX rank and ATM IV rank signals — all calibrated to the same mean leverage $\bar{L} = 1.0$ over the same Oct 2024–Feb 2026 window.

The calibration uses a scalar multiplier k applied to each leverage series such that $\overline{\text{clip}(k \cdot L, 0, 2)} = 1.0$. This preserves each signal's temporal variation (its timing) while eliminating the leverage-level difference. EMOT is rescaled with $p^* = \bar{p}_{\text{EMOT}} \approx 0.726$; VIX and ATM IV rank signals use their existing 0.5/rank formula rescaled. An additional variant, *EMOT rank signal*, applies the rank formula directly to p_{EMOT} to test whether the ddprob ordering is informative.

Table 3: Equal-leverage comparison: all signals at mean $\bar{L} = 1.0$ (Oct 2024–Feb 2026; Benchmark shown for reference at $\bar{L} = 1.0$)

Strategy	CAGR	Sortino	Max DD	Calmar
Benchmark (L=1)	+18.6%	+1.036	−22.8%	+0.815
EMOT rescaled ($p^* = \bar{p}$)	+20.8%	+0.978	−25.5%	+0.813
VIX rank (rescaled)	+13.4%	+0.904	−15.1%	+0.892
ATM IV rank (rescaled)	+18.1%	+1.026	−16.5%	+1.096
EMOT rank signal	+9.6%	+0.421	−31.4%	+0.305

Notes: All strategies clipped to EMOT window (Oct 2024–Feb 2026) and calibrated to $\bar{L} = 1.0$. EMOT rank signal avg $\bar{L} = 1.128$ due to clip asymmetry over the 63-day rolling window; if rescaled to 1.0, Calmar would be further reduced.

Three findings emerge from Table 3:

(i) EMOT rescaled achieves the highest CAGR (+20.8%) but essentially the same Calmar as the benchmark (0.813 vs 0.815). Forcing EMOT to operate at full leverage eliminates its risk advantage without providing meaningful return alpha. The inflated ddprob has low temporal variation (std = 0.201, primarily clustered near 0.73), so the leverage series varies little around 1.0 when rescaled.

(ii) ATM IV rank achieves the best Calmar (+1.096) at equal leverage — beating the benchmark, VIX rank, and EMOT rescaled. The smile's median implied volatility contains genuine forward-looking information about the distribution of near-term returns, even without the full EMOT path computation.

(iii) EMOT rank signal is the worst performer (Calmar +0.305), confirming that the *ordering* of the inflated ddprob is not informative for leverage timing. Periods flagged as high-risk by the fast EMOT solver do not systematically precede worse realised returns.

The implication is that EMOT's value, when correctly calibrated, will come from its *absolute probability level* (risk budget interpretation) rather than from its rank ordering. At full-precision settings, the ddprob will vary meaningfully between 5% and 15%, making the $L = p^*/p_{\text{EMOT}}$ formula operate as intended.

4.4 Advantages of the EMOT Framework

Despite the calibration issue identified above, the EMOT approach possesses genuine conceptual advantages:

- **Market-consistent path measure**: EMOT is the unique framework that simultaneously matches the market-implied terminal density, satisfies no-arbitrage, and

remains closest to a diffusion prior via minimum relative entropy. No parametric model (GARCH, Heston) achieves all three simultaneously.

- **Absolute probability interpretation**: When correctly calibrated, p_{EMOT} is a genuine probability under the market-implied measure, directly interpretable as a risk budget. This is superior to VIX-based signals, which are $\mathbb{P}$-measure quantities and require additional assumptions to convert to $\mathbb{Q}$-measure probabilities.
- **Forward-looking skewness and fat tails**: The RND from which EMOT is calibrated reflects the full shape of the options smile, including risk-neutral skewness and excess kurtosis. GARCH and historical simulation discard this information by assuming Gaussian innovations or stationary return distributions.
- **Rapid reaction to regime shifts**: By refreshing the 30-DTE RND daily, the strategy re-prices drawdown risk at each option chain observation, reacting to term-structure shifts in implied volatility before historical returns reflect the new regime.

4.5 Limitations and Paths to Resolution

4.5.1 EMOT Calibration Precision

The primary limitation is the use of fast calibration settings for computational tractability in the batch run. With $n = 80$ grid points and $N = 3$ steps, discretisation error substantially inflates the ever-drawdown probability. The expected resolution is to increase grid resolution and step count: $n \geq 200$ grid points, $N \geq 6$ steps, and a convergence tolerance of 10^{-5}. The computational cost increases roughly as $O(n^2 \cdot N)$, and full-precision settings are expected to reduce the inflation, though this has not been verified empirically in the present study.

An alternative, computationally cheaper resolution is probability rescaling: scale the fast-mode output by the ratio of the GBM theoretical mean to the empirical EMOT mean:

$$\tilde{p}(t) = p_{\text{EMOT}}(t) \times \frac{\bar{p}_{\text{GBM}}}{\bar{p}_{\text{EMOT}}}, \tag{22}$$

where $\bar{p}_{\text{GBM}} \approx 0.07$ and $\bar{p}_{\text{EMOT}} \approx 0.73$, giving a rescaling factor of approximately 0.096. This preserves the temporal variation of p_{EMOT} (which contains genuine signal) while correcting the level.

4.5.2 Single Regime

The 338-day backtest window (October 2024–February 2026) represents a predominantly bullish, low-volatility equity regime. QQQ returned +18.6% over the period. Any risk-reducing strategy that de-levers will appear to underperform in such an environment, but this does not imply it would underperform over a full market cycle. Testing across crisis regimes (2008, 2020, 2022) is essential before drawing conclusions about strategy quality.

4.5.3 Frictionless Leverage

The backtest ignores transaction costs, borrowing costs for $L > 1$, and market impact. Daily rebalancing of a leveraged QQQ position incurs bid-ask spreads on each adjustment; for a strategy with $\bar{L} = 0.157$ this is less material than for a strategy maintaining $L \approx 2$, but it remains a real cost in production.

4.5.4 Fixed-δ Approximation

Equation (14) approximates the effective constraint by using a fixed barrier δ regardless of the current leverage level. A rigorous treatment would require re-calibrating EMOT at δ/L for each target leverage L, as a leveraged portfolio of L on QQQ experiences a 10% NAV drawdown from a $10\%/L$ underlying move. This introduces a fixed-point problem in the leverage computation but is solvable via bisection.

4.5.5 Implementation Note: Look-Ahead Bias

An important implementation discipline is the one-day lag on all signals in the NAV simulation. Without the lag — applying today's signal to today's return — performance metrics are spuriously inflated due to look-ahead correlation between the leverage decision and same-day returns. All results in Table 2 are computed with the correct lag L^*_{t-1} applied to r_t^{QQQ}.

5 Conclusion

We have presented an end-to-end pipeline implementing the EMOT framework for market-consistent leverage management of a QQQ equity portfolio. The pipeline correctly executes each component: RND extraction via Breeden–Litzenberger, martingale Schrödinger bridge calibration via alternating tilt and h-transform, and path-measure ever-drawdown probability via augmented-state forward propagation.

The principal empirical finding has two parts. First, fast EMOT settings produce ever-drawdown probabilities substantially inflated above the theoretically expected range (mean $p_{\text{EMOT}} = 0.726$ against an expected 5–10% range), causing the leverage formula to operate at near-zero leverage ($\bar{L} = 0.157$) for 99.7% of the sample. Paradoxically, this miscalibration produces the best risk-adjusted metrics of all strategies tested: Calmar +1.599 and Sortino +1.374, with MaxDD of only −4.4% — achieved by staying near cash throughout the period. All four hypothesis tests pass. The GARCH baseline, operating at a consistent probability scale with $\bar{L} = 1.306$, achieves +7.6% CAGR but with larger drawdown, isolating the issue to the probability level rather than the formula.

Second, in an exploratory equal-leverage comparison (all signals calibrated to $\bar{L} = 1.0$), fast-mode EMOT adds no timing value over the benchmark — its rescaled Calmar (+0.813) is essentially identical to QQQ buy-and-hold (+0.815). Strikingly, ATM IV rank achieves the best Calmar (+1.096) at equal leverage, suggesting the smile's median implied volatility contains genuine forward-looking information even without the full path model. EMOT rank (percentile rank of ddprob) is the *worst* signal at full leverage (Calmar +0.305), indicating that the ordering of the fast-mode output is not predictive of return outcomes. Whether this conclusion holds under full-precision EMOT settings remains an open question.

The EMOT framework retains its theoretical advantages: it is the unique path measure satisfying simultaneous market-RND consistency, no-arbitrage, and minimal relative entropy. Full-precision calibration settings ($n \geq 200$, $N \geq 6$) are expected to reduce the probability inflation and bring the signal closer to the intended 5–10% operating range, but this remains to be verified empirically.

Four extensions merit immediate investigation: (i) full-precision EMOT settings ($n \geq 200$, $N \geq 6$) to verify whether higher resolution recovers a correctly scaled probability

and whether the absolute ddprob level then operates as intended in the targeting rule; (ii) proper bisection-based leverage computation to resolve the fixed-δ approximation: since a leveraged-QQQ portfolio experiences a 10% NAV drawdown from a $10\%/L$ underlying move, each target leverage L requires EMOT to be re-calibrated at δ/L, introducing a fixed-point solved by bisection; (iii) extension of the backtest window to include 2020 and 2022 crisis regimes, where forward-looking signals should demonstrate their advantage over backward-looking baselines on a longer track record; and (iv) integration of the 25Δ put overlay (the selection function is implemented but not wired into the portfolio simulator) to provide convex tail protection alongside the probability targeting rule.

A Appendix

A.1 EMOT Algorithm Summary

Algorithm 1 Multi-step EMOT Calibration (Per Date)

Require: Initial marginal $\mu_0 = \delta_{S_0}$, target terminal marginal μ_T^{mkt} (from RND), prior vol σ (median IV), steps N, price grid $\{s_1, \ldots, s_n\}$, tolerance ε, damping η

Build lognormal prior kernels $R_0, \ldots, R_{N-1}$ on s-grid
Initialise $P^{(k)} \leftarrow R_k$ for $k = 0, \ldots, N-1$
repeat
    for $k = 0$ **to** $N-1$ **do**
        for each row i **do**
            Bisect to find λ_i s.t. $\sum_j s_j P_{ij}^{(k)} e^{\lambda_i s_j} = s_i$
            $P_{ij}^{(k)} \leftarrow P_{ij}^{(k)} e^{\lambda_i s_j}$; renormalise row i
        end for
    end for
    Doob h-transform: $P_{ij}^{(N-1)} \leftarrow P_{ij}^{(N-1)} \cdot \mu_T^{\text{mkt}}(j)/\mu_T^{\text{EMOT}}(j)$
    Forward-propagate: $\mu_T^{\text{EMOT}} \leftarrow \mu_0 P^{(0)} \cdots P^{(N-1)}$
    Damped update: $P^{(k)} \leftarrow \eta P_{\text{old}}^{(k)} + (1-\eta) P_{\text{new}}^{(k)}$; renormalise
until $\|\mu_T^{\text{EMOT}} - \mu_T^{\text{mkt}}\|_1 < \varepsilon$
Extract ever-drawdown probability via augmented-state propagation
return p_{EMOT}

A.2 Configuration Parameters

Table 4: Full backtest configuration

Parameter	Symbol	Value
Backtest period	—	Oct 2024 – Feb 2026
Ticker	—	QQQ
Trading days	—	338
RND dates computed	—	338 / 422 (80.1%)
EMOT horizon	T	30 calendar days
EMOT steps (fast mode)	N	3
Price grid points (fast)	n	80
Convergence tolerance	ε	10^{-4}
Outer iterations (fast)	—	25
Damping	η	0.7
Prior vol (fallback)	σ	20% p.a.
Drawdown threshold	δ	10%
Risk budget	p^*	10%
Maximum leverage	$L_{\max}$	2.0×
Minimum leverage	$L_{\min}$	0.0×
Risk-free rate	r_f	1.8% p.a.
Initial NAV	—	\$1,000,000
Rebalance frequency	—	Weekly (Friday)
Leverage lag	—	1 trading day
Observed mean p_{EMOT}	—	0.726
Observed mean leverage	$\bar{L}$	0.157
Dates with $p_{\text{EMOT}} > p^*$	—	99.1%

References

[1] Backhoff-Veraguas, J., Bartl, D., Beiglböck, M., & Eder, M. (2020). Adapted Wasserstein distances and stability in mathematical finance. *Finance and Stochastics*, **24**(3), 601–632.

[2] Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, **51**(4), 621–651.

[3] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**(3), 637–654.

[4] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, **50**(4), 987–1007.

[5] Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, **2**(3), 21–41.

[6] Léonard, C. (2014). A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems - A*, **34**(4), 1533–1574.

[7] Chen, Y., Georgiou, T. T., & Pavon, M. (2016). Optimal transport over a linear dynamical system. *IEEE Transactions on Automatic Control*, **62**(5), 2137–2152.

[8] Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. Hoboken, NJ: Wiley Finance.

[9] Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, **53**(2), 499–547.

[10] Magdon-Ismail, M., Atiya, A. F., Pratap, A., & Abu-Mostafa, Y. S. (2004). On the maximum drawdown of a Brownian motion. *Journal of Applied Probability*, **41**(1), 147–161.

Disclaimer

This research material has been prepared by NUS Invest. NUS Invest specifically prohibits the redistribution of this material in whole or in part without the written permission of NUS Invest. The research officer(s) primarily responsible for the content of this research material, in whole or in part, certifies that their views are accurately expressed, and they will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this research material. Whilst we have taken all reasonable care to ensure that the information contained in this publication is not untrue or misleading at the time of publication, we cannot guarantee its accuracy or completeness, and you should not act on it without first independently verifying its contents. Any opinion or estimate contained in this report is subject to change without notice. We have not given any consideration to and we have not made any investigation of the investment objectives, financial situation or particular needs of the recipient or any class of persons, and accordingly, no warranty whatsoever is given and no liability whatsoever is accepted for any loss arising whether directly or indirectly as a result of the recipient or any class of persons acting on such information or opinion or estimate. You may wish to seek advice from a financial adviser regarding the suitability of the securities mentioned herein, taking into consideration your investment objectives, financial situation or particular needs, before making a commitment to invest in the securities. This report is published solely for information purposes, it does not constitute an advertisement and is not to be construed as a solicitation or an offer to buy or sell any securities or related financial instruments. No representation or warranty, either expressed or implied, is provided in relation to the accuracy, completeness or reliability of the information contained herein. The research material should not be regarded by recipients as a substitute for the exercise of their own judgement. Any opinions expressed in this research material are subject to change without notice.