

Single-Asset Adaptive Leveraged Volatility Control

Nikhil Devanathan^{*1}, Dylan Rueter¹, Stephen Boyd¹,
Emmanuel Candès¹, Trevor Hastie¹, Mykel J. Kochenderfer¹,
Arpit Aporv², David Soronow², and Igor Zamkovsky²

¹BlackRock AI Lab

²BlackRock Index Services

March 30, 2026

Abstract

This paper introduces a methodology for constructing a market index composed of a liquid risky asset and a liquid risk-free asset that achieves a fixed target volatility. Existing volatility-targeting strategies typically scale portfolio exposure inversely with a variance forecast, but such open-loop approaches suffer from high turnover, leverage spikes, and sensitivity to estimation error—issues that limit practical adoption in index construction. We propose a proportional-control approach for setting the index weights that explicitly corrects tracking error through feedback. The method requires only a few interpretable parameters, making it transparent and practical for index construction. We demonstrate in simulation that this approach is more effective at consistently achieving the target volatility than the open-loop alternative.

1 Introduction

We address the problem of constructing a financial index—a rule-based portfolio—from two assets: a liquid risky asset (e.g., an equity fund) and a liquid risk-free asset (e.g., cash or short-term government bonds). The objective is for the index to realize a pre-specified target volatility, where volatility refers to the standard deviation of returns. We operate under two constraints: the risky asset cannot be shorted, and there is an upper bound on leverage, i.e., the degree to which capital is borrowed to amplify exposure to the risky asset.

From a control-theoretic perspective, this is a setpoint tracking problem with input saturation: the “plant” is the stochastic return process of the risky asset, the control input is the portfolio weight, and the regulated output is the

^{*}email: nikhil.devanathan@blackrock.com

realized volatility of the index. The saturation arises from the leverage limit, which imposes a hard upper bound on the control authority.

Volatility-controlled indices are widely used in institutional investment, structured products, and risk-managed strategies. By maintaining a stable risk profile, such indices allow investors to separate asset selection decisions from risk-level decisions. From an engineering standpoint, the problem is appealing because it requires regulating a noisy, non-stationary process using simple, interpretable rules—a setting where feedback control has clear advantages over open-loop strategies.

The core difficulty is that asset volatility is latent and time-varying. Unlike a physical system where sensor noise corrupts a measurable state, here the “true” volatility is never directly observed—only noisy return realizations are available. Volatility estimates derived from historical returns are inherently lagged and uncertain. Moreover, financial volatility exhibits well-documented stylized facts: it clusters (periods of high volatility tend to persist), it responds asymmetrically to positive and negative returns, and it occasionally undergoes abrupt regime shifts during market crises [9, 16]. These properties cause open-loop strategies that rely solely on volatility forecasts to systematically lag behind the true volatility, resulting in persistent tracking errors.

The standard approach to volatility targeting, formalized by Moreira and Muir [14], scales portfolio exposure inversely with a forecast of variance. Large empirical studies confirm that such open-loop rules can improve Sharpe ratios and reduce tail risk across many asset classes [7]. However, these same studies highlight practical issues: high portfolio turnover, leverage spikes during volatility transitions, and sensitivity to the choice of volatility estimator.

Several refinements have been proposed. Conditional schemes restrict scaling to extreme volatility regimes [2], while smoothing techniques based on exponentially weighted moving averages (EWMA) or Bayesian updating dampen weight oscillations [1]. Despite these improvements, empirical work has shown that out-of-sample benefits can vanish once implementation frictions such as transaction costs and rebalancing delays are considered [3, 5]. Crucially, all of these approaches remain fundamentally open-loop: they set portfolio weights based on a volatility forecast without incorporating feedback from the realized volatility of the portfolio itself.

The control-systems literature has applied feedback principles to related financial problems—for example, prior work on stock trading has used proportional and PI controllers to dynamically adjust position sizes [13, 12], and earlier work has addressed volatility tracking in stock markets [4]—but a gap remains: there is no simple, interpretable closed-loop controller specifically designed for volatility targeting under leverage constraints. Existing financial methods are open-loop, and existing control-theoretic work addresses different objectives (e.g., maximizing growth, minimizing drawdown) rather than tracking a volatility setpoint.

We cast volatility targeting as a proportional feedback-control problem. The controller observes the realized volatility of the index (estimated through a EWMA of squared returns), computes the tracking error relative to the target,

and adjusts the portfolio weight accordingly. The EWMA serves as a simple state observer, analogous to a low-pass filter that extracts the volatility signal from noisy return data. We use clipping to enforce the leverage constraint, a technique with well-established properties for constrained control [11, 6], and define the tracking error in logarithmic space for scale invariance. The controller includes a smoothing term that acts as a lag filter on the control signal, preventing excessive reactivity to transient fluctuations.

A key feature of our approach is its deliberate simplicity. The volatility controller requires only four interpretable parameters: a proportional gain, upper and lower bounds on the control signal, and a smoothing factor. This parsimony is a design choice: in index construction, transparency and auditability are paramount, and complex black-box methods face significant barriers to adoption. We validate the method through parameter sensitivity analysis showing robustness across a wide range of gain and smoothing values, and through temporally out-of-sample testing on a diverse cohort of ETFs demonstrating generalization beyond the tuning asset.

2 Methodology

Our goal is to construct an index from a single risky asset and a single risk-free asset (for illustration, we use cash) which achieves a specified level of volatility $\sigma^{\text{tar}} > 0$ without exceeding a limit $L > 0$ on the weight allocated to the risky asset. We call σ^{tar} the target volatility and L the leverage limit. We denote by t_1, t_2, \dots the time sequence at which we observe the returns of the risky asset, realize interest on cash balances, and have the opportunity to rebalance the index. We define the trailing returns of the risky asset and risk-free asset at time t_k as $r_k, r_k^{\text{rf}} \in \mathbb{R}$ respectively.

Index weights and returns An index is a pair of weights (w_k, c_k) at time t_k such that $w_k \in \mathbb{R}$ is the weight of the risky asset and $c_k \in \mathbb{R}$ that of cash. The weights must satisfy the conditions

$$w_k + c_k = 1, \quad 0 \leq w_k \leq L.$$

The weights represent the fraction of the index invested in the risky asset and cash, respectively. A negative cash weight corresponds to leverage, or borrowing cash to invest more in the risky asset, which we can take in a limited capacity. When taking leverage, we pay interest on the borrowed cash rather than accruing positive interest on the cash. Between two consecutive times t_{k-1} and t_k , the index return r_k^{ind} is given by

$$r_k^{\text{ind}} = r_k w_{k-1} + r_k^{\text{rf}} c_{k-1}.$$

Realized volatility The realized volatility $\hat{\sigma}_k$ of the risky asset at time t_k is defined as the square root of an EWMA of the squared risky asset returns

$r_1, \dots, r_k,$

$$(\hat{\sigma}_k)^2 = \frac{1 - \beta}{1 - \beta^k} \sum_{j=1}^k \beta^{k-j} (r_j)^2, \quad (1)$$

with a halflife $h > 0$ and decay factor $\beta = \exp(-\log(2)/h)$ [10]. (This is an EWMA estimate of the second moment, but we use the fact that daily returns have means that are much smaller in magnitude than the returns.) Note that $\hat{\sigma}_k$ is known at the time that w_k is decided.

Realized index volatility Analogously, we define the realized index volatility $\hat{\sigma}_k^{\text{ind}}$ at time t_k as the square root of the EWMA of the squared index returns $r_1^{\text{ind}}, \dots, r_k^{\text{ind}},$

$$(\hat{\sigma}_k^{\text{ind}})^2 = \frac{1 - \beta}{1 - \beta^k} \sum_{j=1}^k \beta^{k-j} (r_j^{\text{ind}})^2. \quad (2)$$

Note that $\hat{\sigma}_k^{\text{ind}}$ is known at the time that w_k is decided.

Tracking error We define our relative instantaneous volatility tracking error e_k at time t_k as

$$e_k = \log(\hat{\sigma}_k^{\text{ind}}/\sigma^{\text{tar}}). \quad (3)$$

The instantaneous tracking error e_k is approximately the percentage difference between the realized index volatility $\hat{\sigma}_k^{\text{ind}}$ and the target volatility σ^{tar} . Over a period $t_j, t_{j+1}, \dots, t_k,$ the average absolute volatility tracking error is

$$\frac{1}{k - j + 1} \sum_{i=j}^k |e_i|.$$

The volatility control problem The problem is to determine a policy that chooses w_k and c_k at time t_k based on information known at time t_k (such as $\hat{\sigma}_k, \hat{\sigma}_k^{\text{ind}}, e_k$) in such a way that the average absolute volatility tracking error is minimized.

2.1 Open-loop approach

In the ideal case, the returns of the risk-free asset are known constants and the returns of the risky asset are independently and identically distributed according to a normal distribution with mean μ and standard deviation $\sigma > 0$ at each time t_j . In this case, the index returns at time t_{k+1} are distributed according to a normal distribution with mean $w_k \mu + c_k r_{k+1}^{\text{rf}}$ and standard deviation $w_k \sigma$. If we assume that $|\mu| \ll \sigma$ and $r_k^{\text{rf}} \ll \sigma$, then the expected value of $(r_{k+1})^2$ is approximately σ^2 and the expected value of $(r_{k+1}^{\text{ind}})^2$ is approximately $w_k^2 \sigma^2$. It follows by the linearity of expectation that the estimated variance $(\hat{\sigma}_k)^2$ of the risky asset should be approximately $(\sigma)^2$.

In this case, a natural approach to choosing w_k and c_k such that the expected value of $(r_{k+1}^{\text{ind}})^2$ is equal to σ^{tar} is to set

$$w_k = \min\{\sigma^{\text{tar}}/\hat{\sigma}_k, L\}, \quad c_k = 1 - w_k. \quad (4)$$

When $\sigma^{\text{tar}}/\hat{\sigma}_k \leq L$, we set the risky asset weight w_k so that the expected value of the index variance is equal to the target variance. When $\sigma^{\text{tar}}/\hat{\sigma}_k > L$, we set $w_k = L$, accepting that this is the best we can do given the leverage limit. We call this approach the “open-loop” approach because the weights w_k are set without any regard to the previous index weights w_1, \dots, w_{k-1} .

This methodology is summarized in Algorithm 1.

Algorithm 1 Open-loop solution to index construction

- 1: **Input:** σ^{tar}, L
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: Observe risky asset volatility $\hat{\sigma}_k$
 - 4: Set risky asset weight $w_k = \min\{\sigma^{\text{tar}}/\hat{\sigma}_k, L\}$
 - 5: Set cash weight $c_k = 1 - w_k$
 - 6: **end for**
-

In practice, although this method may sometimes perform reasonably, it fares poorly when there are changes in the prevailing volatility of the risky asset. This is illustrated in §3.

2.2 Volatility control

We now modify the methodology from §2.1 to better react to shifts in the volatility of the risky asset. To achieve this, we add a *control parameter* $\kappa_k \in \mathbb{R}$ to our policy (4) and set w_k and c_k such that

$$\begin{aligned} w_k &= \min\{\exp(\kappa_k)\sigma^{\text{tar}}/\hat{\sigma}_k, L\} \\ c_k &= 1 - w_k. \end{aligned} \quad (5)$$

The $\exp(\kappa_k)$ scales the risk target and thus allows us to adjust for persistent overestimation or underestimation of the volatility of the risky asset. Operating in logarithmic space is motivated by the empirical observation that volatility varies multiplicatively: for an asset with a mean volatility of $X\%$, periods of $2X\%$ and $\frac{1}{2}X\%$ volatility are approximately equally likely. This choice also aligns with prior work on log-optimal portfolios, which has established the benefits of logarithmic formulations for scale-invariant control [15, 8].

It remains to discuss how to dynamically set the control parameter κ_k , and we describe a simple methodology, which depends on only four parameters. The four parameters are a proportional gain $g > 0$, a control parameter lower bound $\kappa_{\min} < 0$, an upper bound $\kappa_{\max} > 0$, and a smoothing factor $\theta \in (0, 1)$. After setting these parameters, we compute κ_k via

$$\kappa_k = (1 - \theta)\text{clip}(-ge_k; I) + \theta\kappa_{k-1} \quad (6)$$

in which

$$I = [\kappa_{\min}, \kappa_{\max}], \quad \text{clip}(t; [a, b]) = \begin{cases} a, & t < a, \\ t, & a \leq t \leq b, \\ b & t > b. \end{cases} \quad (7)$$

This controller is designed to adjust κ_k based on the discrepancy between the realized volatility of the index $\hat{\sigma}_k^{\text{ind}}$ and the target volatility σ^{tar} . To simplify, assume $\theta = 0$ (no smoothing) and that the volatility of the index has recently been lower than the target. Then the error e_k will be negative, and the term $-ge_k$ will be positive, meaning that we will target a higher volatility in the next iteration. Conversely, if the volatility of the index has recently been higher than the target, then e_k will be positive, and $-ge_k$ will be negative, meaning that we will target a lower volatility in the next iteration.

The gain parameter g influences how reactive κ_k is to discrepancies between the realized and target volatilities. The clipping imposes reasonable limits on κ_k to prevent it from becoming too large or too small. The smoothing factor θ determines how quickly κ_k can change; a larger value of θ means that κ_k will change more slowly, while a smaller value implies κ_k will change more quickly.

Since the volatility control methodology assumes our index has some history (for volatility estimation), we recommend either running or simulating the open-loop method until some time t_{start} before switching to the volatility control methodology.

The volatility control methodology is summarized in Algorithm 2 and its performance is demonstrated in §3.

Algorithm 2 Volatility control for index weight selection

- 1: **Input:** $\sigma^{\text{tar}}, L, g, \kappa_{\min}, \kappa_{\max}, \theta$
 - 2: Set $\kappa_0 = 0$
 - 3: **for** $t_k = t_{\text{start}}, \dots$ **do**
 - 4: Observe risky asset volatility $\hat{\sigma}_k$
 - 5: Produce an index volatility estimate $\hat{\sigma}_k^{\text{ind}}$ from historical index returns
 - 6: Compute volatility control error as in (3)
 - 7: Compute volatility control parameter as in (6)
 - 8: Set risky asset weight $w_k = \min\{\exp(\kappa_k)\sigma^{\text{tar}}/\hat{\sigma}_k, L\}$
 - 9: Set cash weight $c_k = 1 - w_k$
 - 10: **end for**
-

3 Results

In this section, we present results about the performance of Algorithms 1 and 2. We first compare the methods on a single asset, then analyze parameter sensitivity, and finally evaluate generalization across a broad cohort of ETFs.

3.1 Baseline Comparison

Simulation details Our simulation spans the trading days from June 8, 2000 to December 31, 2024, inclusive. The risky asset is IVV, the iShares Core S&P 500 exchange traded fund (ETF), which provides broad exposure to large-cap U.S. equities. We use close-to-close returns for IVV and assume that returns are observed, new weights are computed, and the index is rebalanced at market close. We model transaction costs as a 5 basis point bid-ask spread applied to each rebalancing trade. The return on cash is the daily Federal Funds Effective Rate. We compute the risky asset volatility estimate $\hat{\sigma}_k$ as the square root of an EWMA of the squared risky asset returns r_1, \dots, r_k with a half-life of 126 trading days. We set the leverage limit to $L = 1.5$ and the target volatility to $\sigma^{\text{tar}} = 0.15/\sqrt{252}$, corresponding to an annualized volatility of 15%.

In Algorithm 2, we set $g = 55$, $\kappa_{\min} = -1$, $\kappa_{\max} = 1$, and $\theta = 0.6$; the selection of these parameters is discussed in §3.2. We run the method for the first 10 trading days before engaging control. We compute the index volatility estimate $\hat{\sigma}_k^{\text{ind}}$ as the square root of an EWMA of the squared index returns $r_1^{\text{ind}}, \dots, r_k^{\text{ind}}$ with a half-life of 126 trading days.

Discussion Figure 1 shows the cumulative returns of IVV, the open-loop method, and the volatility control method. Figure 2 shows the running estimated volatility of IVV and that of the two methods over the same period. The Monte Carlo methodology described in Appendix A was used to compute 10th and 90th percentiles for the running volatility estimate of an asset with a true volatility of 15%.

Table 1 summarizes the performance of the three methods over the simulation period. The volatility tracking error is computed as the mean absolute error between the running volatility estimate of the method and the target volatility σ^{tar} , scaled by $\sqrt{252}$ to give an annualized value. The Sharpe ratio is the annualized return in excess of the Fed Funds Effective Rate divided by annualized volatility, the Kalmar ratio is the annualized return divided by maximum drawdown, and turnover is the sum of absolute weight changes over the period. Although IVV does not track a volatility target, we still provide a volatility tracking error for IVV for reference.

As shown in the figures and table, the volatility control method outperforms the open-loop method in terms of return, volatility tracking error, and maximum drawdown.




Figure 1: Cumulative returns of three indices.




Figure 2: Running annualized estimated volatility of three indices. 15% target volatility is shown as a solid black line. A 90% confidence interval for the running volatility estimate of an asset with a true volatility of 15% is shown as a shaded region.

Table 1: Summary of performance metrics for three indices.

	IVV	Open-Loop	Volatility Control
Volatility Tracking Error	5.2%	2.3%	0.4%
Annualized Return	7.8%	6.8%	8.2%
Annualized Volatility	19.1%	14.9%	14.9%
Sharpe Ratio	0.31	0.33	0.42
Kalmar Ratio	0.14	0.18	0.22
Maximum Drawdown	55.3%	38.6%	37.1%
Turnover	0%	93%	1105%

3.2 Parameter Sensitivity

The volatility controller has two key tunable parameters: the proportional gain g and the smoothing factor θ . We fix $\kappa_{\min} = -1$ and $\kappa_{\max} = 1$ throughout, as these bounds primarily serve as safety limits. To understand sensitivity to g and θ , we conduct a grid search over $g \in \{0, e^0, e^{0.5}, e^1, \dots, e^5\}$ (12 values) and $\theta \in \{0, 0.1, \dots, 0.9\}$ (10 values). For each parameter combination, we simulate the volatility controller on IVV over the period June 8, 2000 to December 31, 2009 and compute three metrics: volatility tracking error, change in Kalmar ratio relative to IVV, and portfolio turnover.

Figure 3 shows heatmaps of these metrics over the (g, θ) grid. Tracking error decreases monotonically in g and is largely insensitive to θ . Change in Kalmar ratio is positive for nearly all parameter combinations, with a ridge of particularly favorable values between $(g, \theta) = (148, 0.8)$ and $(55, 0.6)$. Turnover increases monotonically in g and decreases monotonically in θ .

These surfaces suggest that many parameter combinations achieve good performance. Higher gain reduces tracking error but increases turnover; the smoothing factor θ effectively mitigates turnover while preserving the low tracking error achieved by high gain. Based on this analysis, reasonable choices include $(g, \theta) \in \{(55, 0.6), (90, 0.8)\}$. For both the baseline comparison in §3.1 and out-of-sample testing, we select $(g, \theta) = (55, 0.6)$ from within this robust region.




Figure 3: Parameter sensitivity of the volatility controller on IVV (June 2000–December 2009). Left: volatility tracking error (0.29%–2.52%). Center: change in Kalmar ratio relative to the underlying asset (−0.056 to +0.066). Right: portfolio turnover (108%–3558%). Lighter colors indicate higher values. All three metrics vary smoothly over the parameter space, indicating robustness to parameter choice.

3.3 Multi-Asset Evaluation

The preceding analysis establishes that the volatility controller is robust to parameter choice on a single asset. We now assess whether the method generalizes to assets not used for tuning. We evaluate both methods on a broad cohort of ETFs using a temporally out-of-sample methodology: parameters are selected using IVV data from June 2000 to December 2009 (the period analyzed in §3.2), and we then test on a diverse set of ETFs over the disjoint period January 1, 2010 to December 31, 2024. We use the parameters $(g, \theta) = (55, 0.6)$ identified above, applied uniformly across all assets without per-asset tuning. As in the baseline comparison, we include a 5 basis point bid-ask spread.

The ETF cohort consists of all ETFs in our dataset with positive price data throughout the test period and annualized returns exceeding 5%. This filter yields a diverse set of 44 ETFs spanning U.S. and international equities, commodities, and sector funds: ACWI, EFA, EWD, EWJ, EWK, EWL, EWN, EWT, IAI, IAK, IAT, IAU, IBB, IDU, IEO, IGM, IGV, IHF, IHI, IJH, IJR, IOO, ITA, ITB, IVV, IWB, IWD, IWF, IWM, IWV, IWO, IXN, IYC, IYF, IYH, IYJ, IYK, IYM, IYR, IYW, OEF, SOXX, and THD.

Figures 4 and 5 show histograms of three metrics across the ETF cohort for the volatility controller and open-loop method, respectively. For the volatility controller, tracking error is predominantly below 0.4% and entirely below 0.6%, demonstrating that the controller achieves its core objective across diverse assets. The change in Kalmar ratio is positive for the majority of ETFs, with values as high as 0.25, though a minority exhibit negative values—none below −0.15. Turnover ranges from 800% to 1800%.

For the open-loop method, tracking error ranges from 1% to 3%, consistent with the IVV baseline. The change in Kalmar ratio is similar in shape but

smaller in magnitude (-0.1 to 0.2), with most of the mass slightly below zero. Turnover is substantially lower (70%–100%), reflecting the method’s less reactive nature.

Both methods generalize consistently from IVV to the broader ETF cohort: the volatility controller achieves significantly lower tracking error at the cost of higher turnover, while the open-loop method offers lower turnover but substantially worse volatility tracking. The high turnover of the volatility controller is appropriate for highly liquid assets where rebalancing costs are minimal; for less liquid assets, the open-loop method may be preferable despite its weaker tracking performance.




Figure 4: Distribution of performance metrics for the volatility controller across an ETF cohort (January 2010–December 2024). Left: volatility tracking error. Center: change in Kalmar ratio relative to the underlying asset. Right: portfolio turnover.




Figure 5: Distribution of performance metrics for the open-loop method across the same ETF cohort. Left: volatility tracking error. Center: change in Kalmar ratio relative to the underlying asset. Right: portfolio turnover.

Acknowledgments

We would like to acknowledge Gabriel Maher, Rob Tibshirani, and Ludovic Breger for their work on an initial methodology we built upon. We would additionally like to thank Veronica Mai, Logan Bell, Alex Tzikas, Raphael Chinchilla, and Ron Kahn for helpful discussions and feedback.

References

- [1] Mauro Bernardi, Daniele Bianchi, and Nicolas Bianco. *Smoothing volatility targeting*. 2022. DOI: 10.48550/ARXIV.2212.07288.
- [2] Dion Bongaerts, Xiaowei Kang, and Mathijs van Dijk. “Conditional Volatility Targeting”. In: *Financial Analysts Journal* 76.4 (Sept. 2020), pp. 54–71. DOI: 10.1080/0015198x.2020.1790853.
- [3] Scott Cederburg et al. “On the performance of volatility-managed portfolios”. In: *Journal of Financial Economics* 138.1 (Oct. 2020), pp. 95–117. DOI: 10.1016/j.jfineco.2020.04.015.
- [4] J. Cvitanic, R. Liptser, and B. Rozovskii. “Tracking volatility (stock markets)”. In: *IEEE Conference on Decision and Control (CDC)*. Vol. 2. 2000, pp. 1189–1193. DOI: 10.1109/cdc.2000.912016.
- [5] T. E. Duncan, B. Pasik-Duncan, and L. Stettner. “Growth optimal portfolio under proportional transaction costs with obligatory diversification”. In: *IEEE Conference on Decision and Control (CDC)*. 2008, pp. 5582–5589. DOI: 10.1109/cdc.2008.4739164.
- [6] E.G. Gilbert and K.T. Tan. “Linear systems with state and control constraints: the theory and application of maximal output admissible sets”. In: *IEEE Transactions on Automatic Control* 36.9 (1991), pp. 1008–1020. DOI: 10.1109/9.83532.
- [7] Campbell R. Harvey et al. “The Impact of Volatility Targeting”. In: *The Journal of Portfolio Management* 45.1 (Oct. 2018), pp. 14–33. DOI: 10.3905/jpm.2018.45.1.014.
- [8] Chung-Han Hsieh. “Necessary and Sufficient Conditions for Frequency-Based Kelly Optimal Portfolio”. In: *IEEE Control Systems Letters* (2020), pp. 1–1. DOI: 10.1109/lcsys.2020.3002214.
- [9] Ruimeng Hu. “Asymptotic Optimal Portfolio in Fast Mean-Reverting Stochastic Environments”. In: *IEEE Conference on Decision and Control (CDC)*. 2018, pp. 5771–5776. DOI: 10.1109/cdc.2018.8619309.
- [10] Kasper Johansson et al. “A Simple Method for Predicting Covariance Matrices of Financial Returns”. In: *Foundations and Trends^o in Econometrics* 12.4 (2023), pp. 324–407. DOI: 10.1561/08000000047.
- [11] Michal Kvasnica and Miroslav Fikar. “Clipping-Based Complexity Reduction in Explicit MPC”. In: *IEEE Transactions on Automatic Control* 57.7 (July 2012), pp. 1878–1883. DOI: 10.1109/tac.2011.2179428.

- [12] Shirzad Malekpour and B. Ross Barmish. “On stock trading using a controller with delay: The Robust Positive Expectation Property”. In: *IEEE Conference on Decision and Control (CDC)*. 2016, pp. 2881–2887. DOI: 10.1109/cdc.2016.7798698.
- [13] Shirzad Malekpour, James A. Primbs, and B. Ross Barmish. “On stock trading using a PI controller in an idealized market: The robust positive expectation property”. In: *IEEE Conference on Decision and Control (CDC)*. 2013, pp. 1210–1216. DOI: 10.1109/cdc.2013.6760047.
- [14] Alan Moreira and Tyler Muir. *Volatility Managed Portfolios*. Tech. rep. National Bureau of Economic Research, Apr. 2016. DOI: 10.3386/w22208.
- [15] Anton V. Proskurnikov and B. Ross Barmish. “On the Benefit of Non-linear Control for Robust Logarithmic Growth: Coin Flipping Games as a Demonstration Case”. In: *IEEE Control Systems Letters* 7 (2023), pp. 2275–2280. DOI: 10.1109/lcsys.2023.3286312.
- [16] Xun Yu Zhou and G. Yin. “Continuous-time mean-variance portfolio selection with regime switching”. In: *IEEE Conference on Decision and Control (CDC)*. Vol. 1. 2002, pp. 383–388. DOI: 10.1109/cdc.2002.1184524.

A Uncertainty Quantification for Volatility Estimates

Suppose we have constructed an index with returns $r_1^{\text{ind}}, r_2^{\text{ind}}, \dots$ and we estimate its volatility using some method (e.g., the EWMA defined in (2)) to obtain estimates $\sigma_1^{\text{eval}}, \sigma_2^{\text{eval}}, \dots$. Even if the index had a volatility exactly equal to the target σ^{tar} , any finite-sample estimate σ_k^{eval} will almost surely differ from σ^{tar} due to sampling variability. To evaluate whether our index is successfully tracking the target, we must understand how much deviation to expect from a perfectly tracking index.

We propose the following Monte Carlo methodology to quantify this uncertainty:

1. Sample N synthetic return values $r_1^{\text{synth}}, \dots, r_N^{\text{synth}}$ independently from a normal distribution with mean 0 and standard deviation σ^{tar} .
2. Compute volatility estimates $\sigma_1^{\text{synth}}, \dots, \sigma_N^{\text{synth}}$ from the synthetic returns using the same estimation method used to compute σ_k^{eval} from the index returns.
3. Discard the first n values (a burn-in period) and compute statistics of interest from $\sigma_{n+1}^{\text{synth}}, \dots, \sigma_N^{\text{synth}}$.

Because the synthetic returns have true volatility σ^{tar} by construction, the distribution of σ_k^{synth} characterizes the sampling variability we would observe if our index were tracking the target perfectly. This distribution provides the confidence bands shown in Figure 2.

B Closed-Form Approximation for EWMA Estimators

While the Monte Carlo approach above is general, we can derive a closed-form approximation when using EWMA-based volatility estimates. We first consider the simpler case of a simple moving average (SMA) estimator, then extend to EWMA.

For an SMA estimator with window size m , the squared volatility estimate from i.i.d. normal returns with true standard deviation σ is

$$(\sigma^{\text{SMA}})^2 = \frac{1}{m} \sum_{i=1}^m (r_i)^2.$$

Since each $r_i \sim \mathcal{N}(0, \sigma^2)$, we have $(r_i/\sigma)^2 \sim \chi_1^2$, and the sum of m independent χ_1^2 variables follows a χ_m^2 distribution. Thus $(\sigma^{\text{SMA}})^2 \sim (\sigma^2/m)\chi_m^2$, or equivalently, σ^{SMA} follows a scaled chi distribution with m degrees of freedom and scale σ/\sqrt{m} .

For an EWMA estimator with decay factor $\beta = \exp(-\log(2)/h)$ corresponding to half-life h , the squared volatility estimate is a weighted sum of squared returns. In the infinite-sample limit,

$$(\sigma^{\text{EWMA}})^2 = (1 - \beta) \sum_{i=1}^{\infty} \beta^{i-1} r_i^2.$$

This is no longer an exact chi-squared, but we can approximate its distribution by moment matching. Using $\mathbb{E}[r_i^2] = \sigma^2$ and $\text{Var}[r_i^2] = 2\sigma^4$ (since $r_i^2/\sigma^2 \sim \chi_1^2$), we obtain

$$\begin{aligned} \mathbb{E}[(\sigma^{\text{EWMA}})^2] &= \sigma^2, \\ \text{Var}[(\sigma^{\text{EWMA}})^2] &= (1 - \beta)^2 \frac{2\sigma^4}{1 - \beta^2} = \frac{2\sigma^4(1 - \beta)}{1 + \beta}. \end{aligned}$$

A scaled chi-squared $(\sigma^2/\nu)\chi_\nu^2$ has mean σ^2 and variance $2\sigma^4/\nu$. Matching variance gives the effective degrees of freedom $\nu = (1 + \beta)/(1 - \beta)$. Thus σ^{EWMA} is approximately distributed as a scaled chi distribution with $(1 + \beta)/(1 - \beta)$ degrees of freedom and scale $\sigma\sqrt{(1 - \beta)/(1 + \beta)}$.

Figure 6 shows the approximate standard deviation of σ^{EWMA} as a function of half-life h for $\sigma^{\text{tar}} = 0.15/\sqrt{252}$ (annualized volatility of 15%).

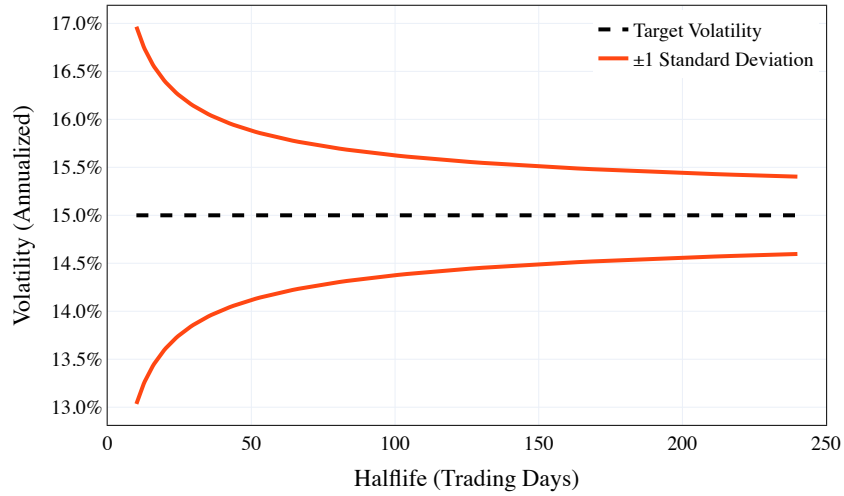


Figure 6: Approximate standard deviation of the EWMA volatility estimate as a function of half-life h for $\sigma^{\text{tar}} = 0.15/\sqrt{252}$.

Figure 7 validates this approximation by comparing the closed-form chi distribution to the empirical distribution from Monte Carlo simulation with $N = 10000$ samples.

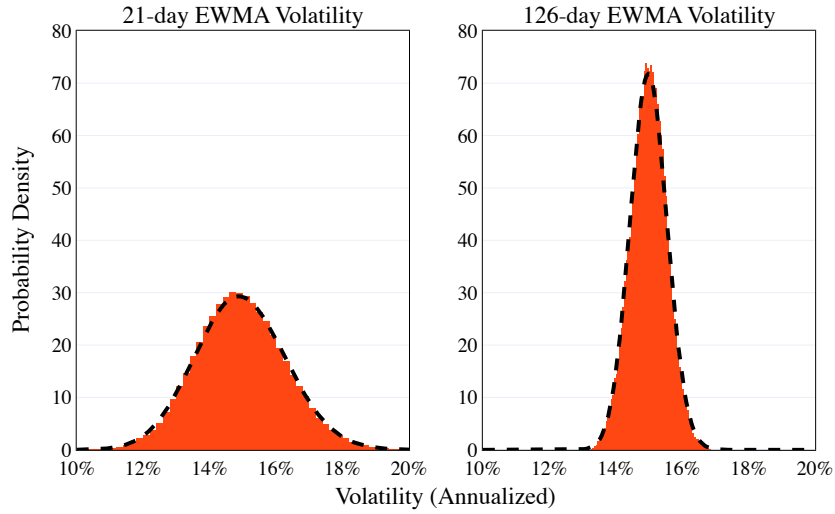


Figure 7: Histogram of EWMA volatility estimates from Monte Carlo simulation ($N = 10000$, burn-in $n = 252$) compared to the chi distribution approximation (dashed lines) for halfives $h = 21$ and $h = 126$.

Figures 6 and 7 show that EWMA volatility estimates with longer halfives have significantly less variance. As such, we recommend using halfives of at least 63 when evaluating the performance of our index construction methods.