

The True Cost of Volatility Targeting: Decomposing the Insurance Premium into Opportunity and Transaction Components

Yi-Hao Lai

Department of Finance, Da-Yeh University

168 University Road, Dacun, Changhua 515006, Taiwan

Corresponding author. Email: yhlai@mail.dyu.edu.tw

September 2026

Abstract

Volatility targeting (VT) strategies scale equity exposure inversely to expected volatility, producing well-documented tail-risk reductions but persistent return shortfalls relative to buy-and-hold. We decompose this return gap—the “insurance premium”—into *opportunity cost* (returns foregone by reducing equity participation) and *direct cost* (transaction expenses from rebalancing). Using daily S&P 500 data conditioned on the CBOE VVIX index over 2012–2024 (3,262 trading days, $N = 3,261$ returns), we find that opportunity cost accounts for 90% of the total insurance premium under unconditional VT (4.05% vs. 0.43% per annum). The share is imprecisely estimated—its stationary-bootstrap 95% interval is [37%, 95%], which does not exclude direct cost exceeding opportunity cost—so what survives resampling is the level of the premium ([0.66%, 8.01%] per annum), not the split. VVIX-conditional targeting—activating VT only when the volatility-of-volatility z -score exceeds 1.0—reduces total cost by 74% in the full sample (from 4.48% to 1.18%), primarily by removing hedging in calm markets. This reduction, however,

does not survive uncertainty quantification: the conditional strategy's premium is not bounded away from zero under resampling, and cross-OOS analysis over six non-overlapping windows shows it beats buy-and-hold on Sharpe in only 2. We therefore treat VVIX conditioning as a hypothesis-generating in-sample accounting result rather than a validated trading rule. The 50/50 SPY/GLD benchmark earns a 54 basis-point average rebalancing premium (negative in two of four sub-periods), raising the bar for VT to clear. These findings redirect the VT design problem from minimizing transaction costs to measuring when opportunity cost dominates the insurance premium.

Keywords: volatility targeting, insurance premium, opportunity cost, transaction cost, volatility-of-volatility, rebalancing premium

JEL Classification: G11, G12, G17

1 Introduction

Volatility targeting (VT) scales equity exposure inversely to conditional volatility (Moreira and Muir, 2017): more when markets are calm, less when turbulence rises. Harvey et al. (2018) document reduced tail risk, and Hocquard et al. (2013) thinner left tails under constant-volatility designs; Fleming et al. (2001) establish the economic value of volatility timing.

Yet VT strategies underperform buy-and-hold in cumulative returns, a shortfall commonly attributed to trading friction; Cederburg et al. (2020) document poor out-of-sample performance across 103 volatility-managed strategies. This attribution conflates two mechanisms: the *opportunity cost* of holding cash through elevated-volatility periods that never materialize into losses, and the *direct cost* of turnover—fees, spreads and market impact. Despite a substantial literature on VT mechanics (Moreira and Muir, 2017; Barroso and Santa-Clara, 2015; Cederburg et al., 2020; Bongaerts et al., 2020; Liu et al., 2019), this decomposition has received little empirical attention.

While Bongaerts et al. (2020) condition volatility targeting on volatility states to optimize Sharpe ratios, we condition on volatility-of-volatility and ask what the investor is actually paying for. The answer reorients strategy design: if the premium is foregone upside rather than trading friction, cost reduction must target participation, not turnover.

Our primary contribution is an accounting decomposition of the VT insurance premium into opportunity and direct cost components (Section 2). The decomposition is arithmetic; its magnitude—90% opportunity cost at the point estimate—is not, and it challenges the prevailing focus on transaction cost minimization. Two exercises follow: the 50/50 benchmark’s 54 basis-point rebalancing premium (negative in two of four sub-periods), which sets the bar VT must clear, and an exploratory VVIX-conditioning exercise that cuts the premium by 74% in sample but wins only 2 of 6 out-of-sample (OOS) windows.

2 Methodology and Data

2.1 Insurance Premium Decomposition

Let r_t^{BH} and r_t^{VT} denote the daily returns of the buy-and-hold equity benchmark and the volatility-targeted strategy. VT allocates weight $w_t = \min(\sigma^*/\hat{\sigma}_t, 1)$ to equities, where σ^* is the target volatility and $\hat{\sigma}_t$ the estimate available at $t - 1$. Following the target-volatility convention studied by [Perchet et al. \(2016\)](#), we implement a 12% annual target using lagged VIX: $w_t = \min(12/\text{VIX}_{t-1}, 1)$.

Define the *total insurance premium* as the annualized return difference, where $\bar{r}$ denotes the annualized mean daily *log* return, here and in Eq. (2):

$$\text{IP}_{\text{total}} = \bar{r}^{\text{BH}} - \bar{r}^{\text{VT}}. \quad (1)$$

Let $r_t^{\text{VT, gross}}$ denote the VT return *before* transaction costs. The *opportunity cost* captures the return gap due to reduced equity participation:

$$\text{IP}_{\text{opp}} = \bar{r}^{\text{BH}} - \bar{r}^{\text{VT, gross}}. \quad (2)$$

The *direct cost* captures the cumulative transaction expense:

$$\text{IP}_{\text{direct}} = \left(\sum_t c \cdot |w_t - w_{t-1}| \right) / Y, \quad (3)$$

where $c = 5$ basis points per unit weight change and $Y = T/252$ is the sample length in years ($T = 3,261$ return observations, $Y \approx 12.94$). Both IP_{opp} and $\text{IP}_{\text{direct}}$ are accumulated over the 3,201 regime-known days (the first 60 observations are the warm-up of the z -score defined below), while Y spans the full panel; reported premia therefore understate their regime-known-domain counterparts by 1.9% of the reported value. By construction, $\text{IP}_{\text{total}} = \text{IP}_{\text{opp}} + \text{IP}_{\text{direct}}$.

Two features deserve emphasis. First, IP_{opp} is a *net* quantity: as a full-sample return gap it already nets the returns VT avoids in drawdowns—the insurance payout—against

the upside it forgoes. Second, the share $s = \text{IP}_{\text{opp}}/\text{IP}_{\text{total}}$ is a ratio of two sample means and carries sampling uncertainty; we quantify it with a stationary bootstrap (Politis and Romano, 1994), reporting a bounded percentile interval for s only when the resampled total premium stays away from zero (Supplementary Material S.1).

2.2 VoV-Conditional Targeting

Rather than applying VT unconditionally, we condition on the volatility-of-volatility (VoV) regime, computing an expanding-window z -score (minimum 60 observations) of the CBOE VVIX index (CBOE, 2012):

$$z_t^{\text{VoV}} = \frac{\text{VVIX}_t - \overline{\text{VVIX}_{1:t}}}{\text{sd}(\text{VVIX}_{1:t})}. \quad (4)$$

We define two conditional strategies:

- **S2 (VoV-Conditional):** Apply 12/VIX weight only when $z_t^{\text{VoV}} > 1.0$ and VIX is rising (5-day change > 0); otherwise hold 100% equity. The 5-day window is a judgement, discussed but not swept in Supplementary Material S.2; sensitivity to the threshold z -score is reported in Section 3.
- **S3 (Smooth VoV):** Interpolate continuously between full equity and full VT as a linear function of the clipped VoV z -score: $w_t = 1 - \text{clip}(z_t^{\text{VoV}}, 0, 1)(1 - w_t^{\text{VT}})$, so $z_t^{\text{VoV}} \leq 0$ holds 100% equity and $z_t^{\text{VoV}} \geq 1$ applies full VT. The equivalent algebraic form is in Supplementary Material S.2.

High VoV signals regime uncertainty, when insurance is most valuable; low VoV signals stable regimes where VT hedging is unnecessary and costly.¹

2.3 Data

We use daily prices for SPY, GLD, VIX, and VVIX from Yahoo Finance (via `yfinance`) for 2012-01-03 to 2024-12-31 ($N = 3,262$ trading days); the start date is our own re-

¹Supplementary Material S.3 attributes the cost reduction between the z -score threshold and the VIX-rising filter; the split is directional and not validated out of sample.

striction: CBOE first disseminated VVIX in March 2012, so earlier values are backfilled rather than observable in real time (CBOE, 2012). Our own first weeks are backfilled too; we report no sensitivity to either choice (S.5). Equity and gold legs use dividend-adjusted closes; VIX and VVIX are index levels and admit no such adjustment. The weight not held in equity earns cash at the lagged effective federal funds rate (FRED series EFFR), so every strategy is self-financing and the idle weight is priced at its actual opportunity cost rather than at zero. The 3,262 price days yield 3,261 returns (2012-01-04 to 2024-12-31), the estimation sample throughout. All signals are lagged one trading day (`signal.shift(1)`) to prevent look-ahead bias. Transaction costs are 5 bps per unit of weight change, charged once on $|w_t - w_{t-1}|$ —an assumed upper bound on SPY’s effective cost in the sense of Hasbrouck (2009)—justified in Supplementary Material S.4, which also reports a 1 bps sensitivity. Generative AI assistants (Anthropic Claude, OpenAI Codex) were used to implement and review the replication code, whose gate checks thirteen headline claims against this manuscript.

3 Results

3.1 Strategy Performance

Table 1 reports performance for five strategies over 2012–2024. Buy-and-hold (S0) achieves a compound annual growth rate (CAGR) of 14.58% and a Sharpe ratio of 0.74. Unconditional VT (S1, Always 12/VIX) sacrifices 5.34 percentage points of CAGR for a 58% smaller, partly mechanical, maximum drawdown (MDD, -14.18% vs. -33.72%). S2 retains most of the BH return (CAGR 13.23%) while reducing MDD by 36% and attaining the highest Sharpe ratio (0.81).

Table 1: Strategy Performance: Full Sample (2012–2024)

	S0: BH	S1: Always	S2: VoV-	S3: Smooth	S4: 50/50
Metric	SPY	12/VIX	Cond.	VoV	SPY/GLD
CAGR (%)	14.58	9.24	13.23	10.81	9.52
Ann. Vol (%)	16.48	9.26	13.63	11.51	11.43
Sharpe	0.74	0.80	0.81	0.77	0.67
MDD (%)	−33.72	−14.18	−21.69	−22.20	−20.32
Calmar	0.43	0.65	0.61	0.49	0.47
Sortino	0.68	0.74	0.78	0.72	0.62
Mean CRRA $\gamma=5$	0.220	0.199	0.216	0.213	0.169
Avg. Weight (%)	100.0	74.1	93.4	86.0	—
Ann. Turnover (%)	0.0	872.1	1044.7	912.1	—

Notes: S0 is 100% SPY buy-and-hold. S1 sets $w = \min(12/\text{VIX}_{t-1}, 1)$ daily. S2 applies S1 only when VoV z -score > 1.0 and VIX is rising; otherwise $w = 1$. S3 continuously interpolates S1 toward full equity as the clipped VoV z -score falls below 1.0 (linear in $\text{clip}(z, 0, 1)$). S4 is a 50/50 SPY/GLD portfolio rebalanced daily to constant weights at no cost; its Sharpe (0.67) is a daily constant-weight benchmark and is distinct from the separately computed monthly-rebalancing premium (2006–2024) analysed in Section 3. All signals lagged one day. TX cost: 5 bps per weight change. Sharpe and Sortino are scored against the same lagged effective federal funds rate the portfolios accrue on their idle weight, with the Ann. Vol row above as the Sharpe denominator. MDD and volatility are *not* exposure-adjusted: average equity weight differs by up to 25.9 percentage points across the columns (Avg. Weight row), so part of the drawdown and volatility reduction is mechanical rather than an effect of timing. CRRA $\gamma = 5$ reports mean CRRA utility of the wealth path (not certainty-equivalent); a formal CE analysis requires computing $\text{CE} = u^{-1}(E[u(W)])$. Equity legs use dividend-adjusted closes and the idle weight earns the lagged effective federal funds rate; see Section 2.3. Data source: yfinance and FRED; replication package.

3.2 Insurance Premium Decomposition

Table 2 presents the central result: the decomposition of the insurance premium into opportunity and direct cost components.

Table 2: Insurance Premium Decomposition (%/year)

Strategy	Opportunity Cost	Direct Cost	Total Premium	Opp. Share (%)	Δ vs. S1 (%)
S1: Always 12/VIX	4.05	0.43	4.48	90.4	—
S2: VoV-Conditional	0.66	0.52	1.18	55.7	−73.7
S3: Smooth VoV	2.88	0.46	3.34	86.3	−25.4

Notes: Opportunity cost = $\bar{r}^{\text{BH}} - \bar{r}^{\text{VT,gross}}$ (return gap before TX). Direct cost = $\sum c|w_t - w_{t-1}|$ annualized (TX friction). Total = Opp. + Direct. Δ vs. S1 shows percentage reduction in total premium relative to unconditional VT. S1’s 4.48% and the 5.34 pp CAGR gap in Table 1 are one quantity under two conventions: continuously compounded, that gap is $\ln(1.1458) - \ln(1.0924) = 4.77$ pp, and the residual 0.29 pp is the 60 warm-up days entering the annualisation denominator but not the numerator. Data: replication package; 2012–2024, $N = 3,261$ return observations.

Three findings emerge. First, for unconditional VT (S1), opportunity cost dominates at 90% of the total premium (4.05% vs. 0.43%). The cost is foregone equity participation in elevated-VIX periods followed by positive returns.

Second, VoV conditioning (S2) reduces the total premium by 74%, from 4.48% to 1.18% per annum. The reduction operates almost entirely through the opportunity cost channel: S2’s opportunity cost falls from 4.05% to 0.66% (an 84% reduction) while its direct cost *increases* from 0.43% to 0.52%. S2 engages VT only when the VoV z -score exceeds 1.0 and VIX is rising, removing the calm-period hedging behind most of S1’s opportunity cost.

Third, smooth blending (S3) reduces total cost by only 25.4% (Table 2), so partial hedging in low-VoV regimes still incurs substantial opportunity cost without proportionate risk reduction.

Sampling uncertainty. The point estimates above are imprecise, and the two strategies are imprecise in different ways. For unconditional VT (S1), the stationary bootstrap puts the opportunity-cost share at 90.4% with a 95% interval of [36.5%, 94.7%]; the total premium itself is bounded away from zero ([0.66%, 8.01%] per annum). The claim that survives resampling is therefore the *level* of the premium, not its composition: the lower bound of the share sits well below parity, so the resampled distribution admits paths on which direct cost is the larger component. The 90% figure should be read as a point estimate drawn from a strongly left-skewed distribution, not as an identified share, and the ordering it implies is a property of the realised sample rather than one established by resampling. For VoV-conditional VT (S2), *no* bounded interval for the share exists: the resampled total premium straddles zero (2,392 of 10,000 replications produce a total premium ≤ 0 ; interval $[-2.76\%, 3.97\%]$ per annum), so the share’s denominator is not signed and the ratio has no bounded connected interval. This is a substantive result, not a numerical nuisance: once sampling uncertainty is admitted, S2’s premium cannot be distinguished from zero, which independently corroborates the cross-OOS evidence in §3.5 and is the reason we treat VoV conditioning as hypothesis-generating rather than as a validated cost-reduction rule.

3.3 Structural Advantages of the 50/50 Benchmark

The 50/50 SPY/GLD benchmark attains a Sharpe ratio of 0.67 against 0.77–0.81 for the VT strategies (Table 1)—a gap far narrower than their drawdown advantage, and obtained with no forecast. Two structural advantages account for it.

First, SPY and GLD exhibit near-zero correlation ($\rho = 0.057$) over the extended 2006–2024 sample period.² Over the 2012–2024 period in Table 1, this produces a 50/50 portfolio volatility of 11.43%—30.7% below SPY’s 16.48%, computed before rounding. Second, monthly rebalancing of the 50/50 portfolio generates a full-sample average premium of 54 basis points per annum (CAGR of 10.023% vs. 9.487% for buy-and-hold)—though the empirical premium is negative in two of four sub-periods (2006–09: –95 bps; 2010–14:

²The correlation and rebalancing premium estimates use the 2006–2024 sample rather than the 2012–2024 estimation period.

−56 bps), so it is a long-run average rather than a per-period guarantee—consistent with Booth and Fama (1992) and with our weight sweep from 30/70 to 70/30.³ The premium arises mechanically from buy-low/sell-high behaviour under fixed-weight rebalancing, requiring no signal, transaction rule or parameter estimate. We make no regime-conditional claim for gold: no archived result in this package reports gold returns by VoV regime.

The 54 bps empirical figure and the 81 bps Booth and Fama (1992) approximation are two estimates of one quantity, not a premium plus a separate diversification benefit. S2’s Sharpe advantage over 50/50 reflects different risk exposures rather than better skill: S2 is predominantly equity with state-contingent hedging, whereas 50/50 earns its risk-adjusted return from cross-asset diversification and mechanical rebalancing.

3.4 Regime Analysis

The insurance premium varies sharply across VoV regimes. In “HighVoV_Falling” periods (7.7% of the sample) S1’s opportunity cost spikes to 20.87%—the post-crisis recoveries where VT misses the rebound—against 2.30% in “LowVoV_Falling” periods (45.5% of the sample). S2 removes this asymmetry by construction: it holds 100% equity in low-VoV or falling-VIX regimes, incurring zero opportunity cost on the 76% of days when VoV is below the threshold.

Figure 1 shows the accounting as a wealth path: the S1–S0 gap widens through the long unshaded stretches and narrows at crash troughs. The premium is paid continuously in forgone participation, not when the hedge is needed.

³The 54 bps estimate uses dividend-adjusted prices (`yfinance auto_adjust=True`); the same window yields approximately 63 bps on the raw-Close convention (`auto_adjust=False`), which the replication package also reports (both net of 1 bps). Supplementary Material S.5 reconciles the two conventions.




Figure 1: Cumulative wealth (log scale) of buy-and-hold SPY (S0), unconditional VT (S1), VoV-conditional VT (S2), and the 50/50 SPY/GLD benchmark (S4), 2012–2024, with VoV-regime shading. Strategy rules, price convention (dividend-adjusted closes with an EFFR cash leg on the idle weight), and the CAGR convention are those of the replication package; the plotting script asserts each curve’s CAGR against Table 1 before writing the figure. Source: replication package.

3.5 Cross-OOS Robustness

We evaluate robustness on six non-overlapping two-year windows spanning 2013–2024, comparing per-window Sharpe ratios (replication package). S2 outperforms BH SPY on a Sharpe-ratio basis in only 2 of 6 windows (2017–18 and 2019–20), while S3 also wins in 2 of 6 and the 50/50 SPY/GLD benchmark wins in 3 of 6. Unconditional VT (S1) wins the same 2 of 6, and in the same two windows, so this record cannot separate state-contingent conditioning from always-on targeting. Diebold-Mariano (DM) tests⁴ confirm that no strategy difference clears the conservative $|t| > 3.0$ heuristic adapted from [Harvey et al. \(2016\)](#) (S1 vs. S0: $t = 2.34$; S2 vs. S0: $t = 0.71$). Both wins fall in the 2018 Volmageddon and 2020 COVID windows.

⁴DM tests use the negative daily return, $L_t = -r_t$, with Bartlett-kernel Newey–West heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; the specification and the archived implementation are in Supplementary Material S.5.

The threshold enters S2 only: S1 carries no threshold, so the sweep bears on the conditional strategy and not on the headline decomposition. At $z^* = 0.5, 1.0$ and 1.5 the S2 opportunity-cost share is 63%, 56% and 64%, and the premium reduction relative to unconditional VT is 62%, 74% and 76% on 20.9%, 14.3% and 8.3% of days. The share is non-monotone, so the 56% at our operating point is a local minimum rather than a stable property, and the reduction rises as engagement falls—consistent with the reduction coming from disengaging rather than from conditioning well. We do not offer that reduction as robustness evidence: it is the quantity the resampling above leaves unidentified.

4 Discussion

Our decomposition yields three implications. First, the conventional view that VT is costly because of high turnover (S1’s annualized turnover is 872%) is misleading: at 5 bps per trade that turnover generates only 43 bps of annual direct cost against 405 bps of foregone equity returns. Reducing trading frequency or using futures therefore addresses about a tenth of the premium at the point estimate. It is not a bound: the resampled opportunity share falls to 36.5% at its lower limit, so the direct-cost share can be as large as 63%.

Second, the VVIX-conditioning exercise suggests a mechanism but, on a 2-of-6 cross-OOS record, does not validate a deployable rule. In low-VoV regimes the forecast may be accurate yet economically unnecessary: staying fully invested avoids the opportunity cost that drives most of the unconditional premium, consistent with [Liu et al. \(2019\)](#), who find gains only in the financial crisis. It carries the volatility-of-volatility risk channel of [Huang et al. \(2019\)](#) into portfolio construction rather than asset pricing; [Bollerslev et al. \(2009\)](#), [Todorov \(2010\)](#) and [Bekaert and Hoerova \(2014\)](#) price variance risk, a different quantity.

Third, the 54 bps rebalancing premium is the long-run bar VT must clear to deliver net positive alpha.

Three limitations warrant emphasis. The 2012 start excludes the 2008 crisis. The 2-of-6 cross-OOS record means the conditioning rule cannot stand as an independent trading-rule contribution. And the rule itself has not been validated by nested out-of-sample testing; walk-forward optimization and extension beyond SPY are left to future work.

5 Conclusion

We decompose the volatility targeting insurance premium into opportunity and direct cost components. The dominant component is opportunity cost, not transaction cost—90% at the point estimate, though the bootstrap 95% interval of [37%, 95%] does not by itself exclude the reverse ordering. The path to cheaper VT therefore runs through identifying when reduced equity exposure is worth its opportunity cost, not through lower-friction execution. VVIX-conditional targeting reduces cost by 74% in the full sample, but that reduction survives neither resampling nor the cross-OOS record, and no DM comparison is Harvey-significant. We therefore keep the decomposition as the robust contribution and treat VoV conditioning as hypothesis-generating.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used generative AI assistants (Anthropic Claude, OpenAI Codex) for drafting and editing prose and for literature triage. Their use in the research process itself is described in Section 2.3. All empirical results were produced by the author-controlled code in the replication package; no result, table, or figure was generated by an AI model. The author reviewed and edited all AI-assisted content and takes full responsibility for the content of the published article.

Data availability

All data used in this study are from public sources: daily dividend-adjusted closes for SPY and GLD from Yahoo Finance, the VIX and VVIX indices published by the Chicago Board Options Exchange, and the effective federal funds rate from FRED. A self-contained replication package accompanies this submission and is available to the editor and referees on request: the bundled price series, together with a script that checks the paper’s thirteen headline claims (the cost decomposition and the two rebalancing premia). It will be deposited in a public repository with a citable DOI on acceptance.

Acknowledgements

All errors are my own.

References

- Barroso, P., Santa-Clara, P., 2015. Momentum has its moments. *Journal of Financial Economics* 116(1), 111–120. <https://doi.org/10.1016/j.jfineco.2014.11.010>
- Bekaert, G., Hoerova, M., 2014. The VIX, the variance premium and stock market volatility. *Journal of Econometrics* 183(2), 181–192. <https://doi.org/10.1016/j.jeconom.2014.05.008>
- Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22(11), 4463–4492. <https://doi.org/10.1093/rfs/hhp008>
- Bongaerts, D., Kang, X., van Dijk, M., 2020. Conditional volatility targeting. *Financial Analysts Journal* 76(4), 54–71. <https://doi.org/10.1080/0015198X.2020.1790853>
- Booth, D.G., Fama, E.F., 1992. Diversification returns and asset contributions. *Financial Analysts Journal* 48(3), 26–32. <https://doi.org/10.2469/faj.v48.n3.26>

- CBOE, 2012. Double the fun with CBOE's VVIX Index. White Paper, Chicago Board Options Exchange. <https://cdn.cboe.com/resources/indices/documents/vvix-termstructure.pdf>
- Cederburg, S., O'Doherty, M.S., Wang, F., Yan, X.S., 2020. On the performance of volatility-managed portfolios. *Journal of Financial Economics* 138(1), 95–117. <https://doi.org/10.1016/j.jfineco.2020.04.015>
- Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *Journal of Finance* 56(1), 329–352. <https://doi.org/10.1111/0022-1082.00327>
- Harvey, C.R., Liu, Y., Zhu, H., 2016. ...and the cross-section of expected returns. *Review of Financial Studies* 29(1), 5–68. <https://doi.org/10.1093/rfs/hhv059>
- Harvey, C.R., Hoyle, E., Korgaonkar, R., Rattray, S., Sargaison, M., Van Hemert, O., 2018. The impact of volatility targeting. *Journal of Portfolio Management* 45(1), 14–33. <https://doi.org/10.3905/jpm.2018.45.1.014>
- Hasbrouck, J., 2009. Trading costs and returns for U.S. equities: Estimating effective costs from daily data. *Journal of Finance* 64(3), 1445–1477. <https://doi.org/10.1111/j.1540-6261.2009.01469.x>
- Hocquard, A., Ng, S., Papageorgiou, N., 2013. A constant-volatility framework for managing tail risk. *Journal of Portfolio Management* 39(2), 28–40. <https://doi.org/10.3905/jpm.2013.39.2.028>
- Huang, D., Schlag, C., Shaliastovich, I., Thimme, J., 2019. Volatility-of-volatility risk. *Journal of Financial and Quantitative Analysis* 54(6), 2423–2452. <https://doi.org/10.1017/S0022109018001436>
- Liu, F., Tang, X., Zhou, G., 2019. Volatility-managed portfolio: Does it really work? *Journal of Portfolio Management* 46(1), 38–51. <https://doi.org/10.3905/jpm.2019.1.107>

- Moreira, A., Muir, T., 2017. Volatility-managed portfolios. *Journal of Finance* 72(4), 1611–1644. <https://doi.org/10.1111/jofi.12513>
- Perchet, R., de Carvalho, R.L., Heckel, T., Moulin, P., 2016. Predicting the success of volatility targeting strategies: Application to equities and other asset classes. *Journal of Alternative Investments* 18(3), 21–38. <https://doi.org/10.3905/jai.2016.18.3.021>
- Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. *Journal of the American Statistical Association* 89(428), 1303–1313. <https://doi.org/10.1080/01621459.1994.10476870>
- Todorov, V., 2010. Variance risk-premium dynamics: The role of jumps. *Review of Financial Studies* 23(1), 345–383. <https://doi.org/10.1093/rfs/hhp035>