Article

Not peer-reviewed version

Tail-Aware Portfolio Optimization for Listed Real-Estate Securities Under Downside Risk

Dilmi C. W. Hettiachchi-Halpe-Kankanamalage * , Abootaleb Shirvani , Nicholas Appiah , Svetlozar T. Rachev , W. Brent Lindquist , Frank J. Fabozzi

Posted Date: 30 March 2026

doi: 10.20944/preprints202603.2360.v1

Keywords: listed real-estate securities; portfolio optimization; downside risk; conditional value-at-risk (CVaR); mean–variance optimization; extreme value theory (EVT); generalized Pareto distribution (GPD); Hill tail index; robust regression; volatility clustering; FIGARCH; ARFIMA; normal double inverse Gaussian (NDIG); option pricing; implied volatility surface

 doi:10.20944/preprints202603.2360.v1

Article

Tail-Aware Portfolio Optimization for Listed Real-Estate Securities Under Downside Risk

Dilmi C.W. Hettiachchi-Halpe-Kankanamalage [1,*], Abootaleb Shirvani [2], Nicholas Appiah [1], Svetlozar T. Rachev [1], W. Brent Lindquist [1] and Frank J. Fabozzi [3]

[1] Department of Mathematics & Statistics, Texas Tech University, Lubbock, TX, USA
[2] Department of Mathematical Science, Kean University, Union, NJ, USA
[3] Carey Business School, Johns Hopkins University, Baltimore, MD, USA
* Correspondence: diwickra@ttu.edu

Abstract

This paper has presented a combined empirical framework for measuring the risk-return profile of listed real-estate securities in a non-Gaussian market situation. By leveraging daily data for 30 U.S. and international listed real estate securities from 2021 to 2024, we probe how portfolio outcomes vary according to the optimization criterion and distributional aspects of returns obscured by conventional mean-variance summaries. We build long-only and long-short portfolios using classical Markowitz mean–variance optimization methods and conditional value-at-risk (CVaR) optimization techniques, and compare their realized dynamics cumulative growth and efficient frontiers under alternative risk-free benchmarks. By applying extreme value theory, we quantify extreme-risk exposure via generalized Pareto modeling and Hill tail-index estimation, which is compared to the broader behavior of the equity market. We analyze portfolio stability and reward efficiency using volatility, Sharpe, Sortino, and Rachev ratios, as well as maximum drawdown and information ratio. In addition, robust single-factor regressions are estimated on a sector benchmark, and residual diagnostics are analyzed to define common factor dependence while minimizing the impact of outliers. To introduce a forward-looking dimension, we calibrate the double-subordinated normal inverse Gaussian specification and extract NDIG model-implied option prices and volatility surfaces. We also investigate volatility persistence through ARFIMA-GARCH modeling to assess whether listed real-estate security returns exhibit long-memory features beyond standard volatility clustering. Results indicate that listed real-estate security returns exhibit heavy tails, pronounced downside sensitivity, and persistent volatility, supporting the use of tail-aware optimization, robust estimation, and long-memory-consistent volatility diagnostics beyond standard Gaussian benchmarks.

Keywords: listed real-estate securities; portfolio optimization; downside risk; conditional value-at-risk (CVaR); mean–variance optimization; extreme value theory (EVT); generalized Pareto distribution (GPD); Hill tail index; robust regression; volatility clustering; FIGARCH; ARFIMA; normal double inverse Gaussian (NDIG); option pricing; implied volatility surface

1. Introduction

The real estate industry occupies a unique place in the world financial ecosystem as a physical form of commodity with real value-added potential, and a method of income production and saving against inflation [1]. Real estate investment trusts (REITs) and real estate exchange-traded funds (ETFs) are increasingly significant vehicles in enabling investors to gain diversified exposure to real estate markets without sacrificing liquidity. The popularity of publicly traded real estate securities during the past two decades has turned an illiquid market into a dynamic part of global portfolios by granting investors access to a range of sub-sectors (residential, commercial, industrial, infrastructure, and healthcare).

 doi:10.20944/preprints202603.2360.v1

Listed real-estate securities are strategically attractive owing to the improved diversification and yield stability they can provide. Unlike property investments, where investment is directly made, listed real estate securities can be traded intraday, provide transparent pricing, and enjoy index-oriented construction. Although these benefits come with more sensitivity to interest-rate fluctuations, sector segmentation, and liquidity shocks. In terms of empirical evidence, returns for the listed real-estate securities are characteristically tinged with massive tails and volatility and cross-market dependence—features that complicate the classic mean–variance optimization equation and inspire tail-aware risk modelling [2–6]. Moreover, the globalization of the flow of property capital has made the link between property securities and general equity markets stronger in turn, limiting the diversification advantage previously gained from the property class.

In the literature, an associated set of studies has also confirmed the distinctiveness of asset classes of real estate along with its "equity-like" integration, which are diversifying advantages with time and regimes. Early research emphasizes that listed real estate serves as a portfolio diversifier even if its co-movement with overall capital markets is economically significant and not fixed [7–9]. Later evidence suggests that REIT and securitized real estate returns are subject to the same systematic risk factors in common with equities and fixed income, including interest-rate sensitivity, term structure effects, and varying risk premia from time to time [10,11]. Recent research also provides increasing support of the view that real estate securities can be increasingly part of the equity markets in stress episodes in which cross asset correlations increase, and tail dependence increases [12–14]. These results challenge portfolios to consider non-normality and tail risk as a first-order feature and not as a side effect of the Gaussian assumption.

From a portfolio-design point of view, focusing on an exclusively long-only allocation may be excessively restrictive when implementable long-short overlays exist. Constrained long–short strategies are employed in contemporary quantitative finance to introduce feasible hedges, factor tilts, and risk controls in a manner with explicit limits on leverage, margin, and drawdowns [15–18]. The classical "shorting" objection, of unbounded losses, is solved by establishing long–short exposures with binding constraints designed to prevent portfolio value from becoming negative and that directly control tail exposure and capital erosion. In this way, under realistic conditions, flexibility on the short side can also take the form of a practical layer upon which to spread risk budgets and enhance downside sensitivity.

Listed real estate is also a barometer of bigger macroeconomic and monetary tides. Escalating interest rates worldwide, inflationary pressures, and shifts in post-pandemic developments in urban dynamics have all been re-aligning valuations for both commercial and residential subsectors. Meanwhile, institutional investors seek ETF designs with an eye to delivering tactical sentiment about property markets or as a hedge for inflation. However, the same characteristics that grant exposure to and flexibility over derivatives (high leverage, derivatives usage, and concentrated holdings) increase vulnerability to rapid price fluctuations and liquidity disruptions. The microstructure of ETF trading and creation–redemption mechanism can also exacerbate price pressures and transfer liquidity shock across linked portfolios during market stress [19,20]. The simultaneous relationship between risk and return of listed real-estate security portfolios is thus central for asset managers and policy makers.

Portfolio insurance, closely associated, should be thought of as part of the allocation problem and not an add-on post hoc. The portfolio choice and the overlay/insurance decision are treated as an integrated policy that, in theory and practice, would respect downside constraints while still meeting decision-relevant performance parameters [21–25]. A sequential approach optimizing for one criterion and then ensuring the resulting portfolio for the other can incur internally inconsistent exposures and less efficient use of the risk budget. In this way, optimization frameworks motivate us to embed downside protection directly through objectives and constraints so that portfolio outputs are both efficient and coherently insured.

Both this (the former) and recent literature focus on the need to include downside and tail-risk variables in a portfolio, particularly for structurally and liquidly exposed sectors. Mean–variance

 doi:10.20944/preprints202603.2360.v1

frameworks, despite their importance as classical models, do not accurately represent the likelihood of severe losses and asymmetric return distributions [2–5]. A more structured risk indicator, the conditional value-at-risk (CVaR) measure, introduced by Rockafellar and Uryasev [26], captures tail risk more closely. CVaR-based formulations have been established to provide practical optimization under tail constraints, and are now employed extensively in empirical portfolio applications [27,28]. Moreover, the Hill estimator and generalized Pareto distribution (GPD) approaches that are based on extreme-value theory allow for empirical quantification of heavy-tailed behavior [3,29–33]. At the same time, volatility models like ARFIMA–GARCH do a good job of accounting for persistence and volatility clustering on financial time series [34,35]. By embedding the above frameworks into a unified environment, the risk dynamics presented in the listed real-estate securities are more closely approximated.

Besides backward-looking risk estimates, forward-looking metrics from option markets offer useful information about investors' expectations for volatility and tail response. The normal inverse Gaussian (NIG) Lévy process [36] provides a relatively flexible method for measuring skewness and kurtosis that is not restricted to Gaussian. Using the Carr–Madan fast Fourier transform (FFT) option-pricing method [37], it facilitates calculating model-derived volatility surfaces efficiently. In general, the implied volatility surface is generally regarded as a summary of risk-neutral uncertainty based on a market model and can be used to encode non-Gaussian noise features such as skew and tail risk [38,39]. Related evidence also indicates that implied volatility behaviour may vary between ETF options and index options, which may be due to structure and microstructure effects that are applicable when model-implied option signals are used to interpret risk [40]. By combining these models with portfolio analysis, the gap between historical return assessment and forward-looking risk assessment is filled.

Based on this framework, this analysis outlines a holistic empirical approach to historical portfolio optimization, tail-risk diagnostics, model-implied option analysis, and long-range dependence in volatility for publicly traded real estate securities. Using a dataset of 30 U.S. and international listed real-estate securities, we examine how various optimization approaches do differently, from maximizing long-only to constraints applied to the long-short formulation under Markowitz and CVaR conditions. We further assess tail behavior through Hill tail-index estimation and GPD modeling, and we employ robust regression to capture inter-ETF co-movements while reducing the influence of outliers [41–43]. To incorporate a forward-looking dimension, we then calibrate a Normal Double-Inverse Gaussian (NDIG)-based Lévy option-pricing model and extract model-implied volatility smiles and term structures. Finally, we investigate long-range dependence in volatility through ARFIMA–GARCH modeling, allowing us to assess whether persistence in listed real-estate security returns extends beyond short-run volatility clustering [34,35].

The main contribution to this study is that of a threesome contribution. First it establishes an end-to-end analytical pipeline connecting ex-post portfolio construction to ex-ante model-implied uncertainty. Second, it expands the use of tail-aware and robust methodologies to the listed real estate security domain, where traditional Gaussian assumptions don't work. Third, it treats constraint-aware long–short overlays as implementable risk controls, as well as incorporating downside protection directly into the allocation problem, thereby bringing portfolio optimization into harmony with portfolio-insurance objectives. Collectively, these contributions together offer a consolidated framework for explaining risk-return trade-offs in listed real-estate securities, with implications for asset allocation, risk budgeting, and policy oversight.

The rest of this paper is organized as follows. Section 2 describes the dataset, security selection, and summary statistics. Section 3 develops the historical portfolio-optimization framework and reports the efficient frontier and portfolio-performance results. Section 4 presents the extreme-value tail-risk diagnostics, while Section 5 compares the portfolios using risk-adjusted performance measures. Section 6 examines robust benchmark regressions and residual behavior. Section 7 introduces the NDIG-based option-pricing analysis and the corresponding model-implied volatility surfaces. Section 8

 doi:10.20944/preprints202603.2360.v1

studies long-range dependence in volatility. Finally, Section 9 concludes with implications for investors and directions for future research.

2. Data and Descriptive Statistics

2.1. Data Description

The empirical investigation is grounded in a cross-section of 30 listed actively traded real estate securities, from different sectors and geographic regions. The sample consists of domestic, worldwide, and regional funds in residential, commercial, industrial, healthcare, and infrastructure sectors, which include businesses and corporate assets, and was used to conduct the study. Daily adjusted closing prices and market cap values in the financial sector were extracted from Bloomberg Professional Services[1] during the time range January 4, 2021 to January 1, 2025. This 4-year timeframe follows the recovery from the pandemic, interest-rate increase in 2022, and new valuations in the U.S. and Global Property markets, providing a complete picture of an expansionary and contractionary phase.

Alongside security prices, we obtain Dow Jones Industrial Average (DJIA) benchmark data to indicate a broad market trend and risk sentiment from Bloomberg Professional Services[2]. To obtain the 10-year and 3-month yield and treasury bill rates, Yahoo Finance[3] proxies long- and short-term risk-free rates, respectively. In order to be consistent with daily asset returns, the risk-free rate is expressed as:

$$R_{f,\text{daily}} = \left(1 + \frac{r_f}{100}\right)^{1/365} - 1. \tag{1}$$

Table A1 in appendix A summarizes key features, including the symbol ticker, security name, inception date, and current market capitalization, of selected instruments classified under a sector. The dataset includes large-cap and niche exposures spanning globally diversified products like *VNQI*, *REET*, or *RWO*, and narrow categories of securities such as *SRVR* (digital infrastructure), *ARE* (life sciences real estate).

Across the paper, arithmetic returns are used to optimize portfolios, construct efficient frontiers, and evaluate performance, including the buy-and-hold benchmark returns. This choice keeps the empirical analysis aligned with the linear evolution of portfolio wealth and ensures consistency across portfolio construction, risk measurement, and comparative performance assessment.

Figures 1–2 present the cumulative value of a $100 invest-and-hold position in each listed real-estate security over the sample period. The figures show that the 2022 tightening cycle affected most funds, with subsequent performance differing across residential, commercial, global, mortgage, infrastructure, and specialized real estate segments. For completeness, A1 in the Appendix B provides plots of the cumulative arithmetic return for the price plots shown in Figures 1–2.

[1] Accessed on January 30, 2025.
[2] Accessed on April 10, 2025.
[3] Accessed on February 3, 2025.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 30 March 2026 doi:10.20944/preprints202603.2360.v1

Figure 1. Cumulative price growth for selected listed real-estate securities, normalised to $100 on January 4, 2021.

Figure 2. Cumulative growth of the remaining listed real-estate securities, normalised to $100 on January 4, 2021.

 doi:10.20944/preprints202603.2360.v1

3. Historical Portfolio Optimization

One of the basic principles of modern investment analysis, with an integrated quantitative model that provides insight into risk vs return balance under uncertainty, is portfolio optimization. Especially in such a sectoral, liquidity-shock, and interest-rate-dependent industry as listed real estate securities, efficient risk-adjusted investments need to be found. In this section, we formulate a full benchmark of historical portfolio strategies under the classical mean–variance model of [2] and the tail-risk–sensitivity CVaR model [26], which then uses the data generated for such investments to develop a comparison between the historical portfolio strategies. Applying constraints, leverage, and risk preferences, we assess portfolio performance, diversification advantages, and downside protection after both long-only and long-short configurations. The study also generates and illustrates efficient frontiers under alternative risk-free benchmarks in order to demonstrate how short- and long-term yield environments contribute to the level of return attainable.

3.1. Efficient Frontier Analysis

We first examine the efficient frontiers generated by the mean–variance and CVaR optimization frameworks. Each strategy was implemented using one mean–variance specification and two CVaR specifications, where the CVaR optimization was based on confidence levels of 95% and 99%. These approaches provide complementary perspectives on portfolio construction. While mean–variance optimization seeks to control overall dispersion around the mean for a given target return, CVaR optimization focuses on limiting expected losses in the tail of the return distribution, offering stronger protection under adverse market conditions. This distinction is especially relevant for listed real-estate securities, whose returns often display asymmetry and heavy-tailed behavior.

Figure 3 reports the mean–variance and CVaR(95%) efficient frontiers as of 31 December 2024. In each panel, the black curve represents the efficient frontier and the red line denotes the corresponding capital market line. The tangency portfolio is given by the point at which the capital market line touches the frontier. In both cases, the tangency allocation lies near *ITB*, indicating that the most attractive reward-to-risk trade-off depends on whether total risk or downside tail risk is used as the relevant portfolio criterion.

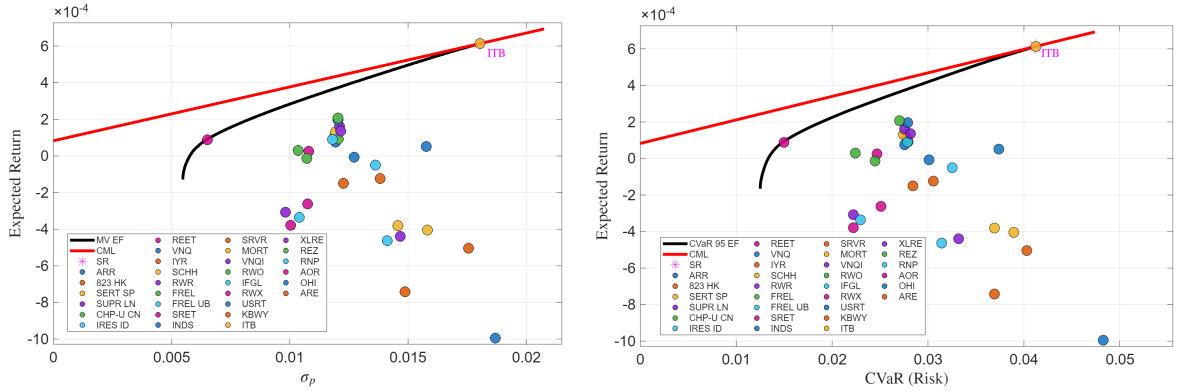

Figure 3. Mean-variance (left) and CVaR(95%) (right) efficient frontiers as of 31 December 2024. The red line is the capital market line computed using the 3-month U.S. Treasury bill yield as the risk-free rate.

Figure 4 shows the corresponding efficient frontiers for 29 April 2022. Relative to the end-of-sample frontiers, the April 2022 results reflect market conditions during the monetary tightening cycle, when rising interest rates and repricing pressures weighed heavily on the listed real-estate sector. The differences between the two dates show that the attainable risk-return trade-off is not fixed, but instead varies with the prevailing market environment.

 doi:10.20944/preprints202603.2360.v1

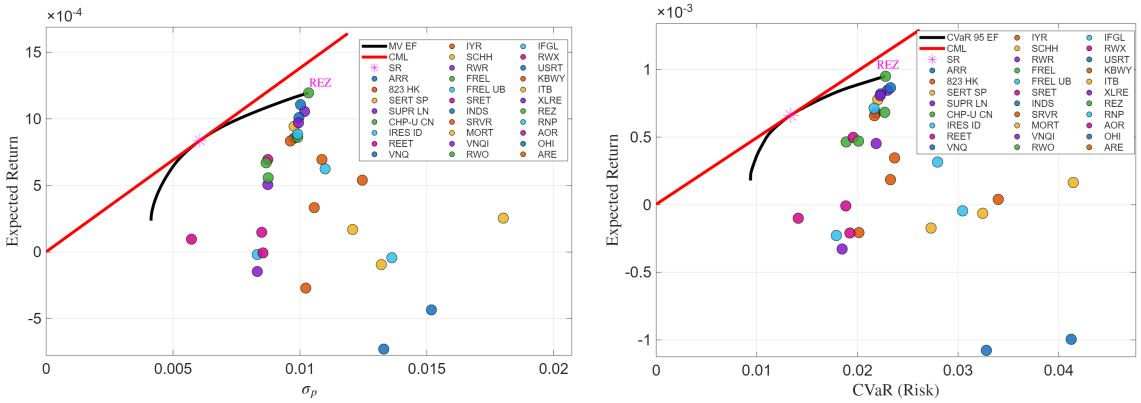

Figure 4. Mean–variance (left) and CVaR(95%) (right) efficient frontiers as of 29 April 2022. The red line indicates the capital market line calculated using the 3-month U.S. Treasury bill yield as the risk-free rate.

3.2. Long-Only and Long–Short Portfolio Performance

We consider two investment strategies: long-only and long–short. The long-only strategy imposes the constraints

$$0 \le w_i(t_k) \le 1, \qquad \sum_i w_i(t_k) = 1,$$

where $w_i(t_k)$ denotes the weight assigned to asset i over the holding period $[t_k, t_{k+1})$. These constraints ensure full investment without leverage or short exposure. The long-short strategy relaxes the bounds to

$$-\frac{1}{30} \le w_i(t_k) \le 1 + \frac{1}{30}, \qquad \sum_i w_i(t_k) = 1,$$

allowing limited short positions and modest leverage while preserving the full-investment constraint. In both cases, rebalancing is subject to a turnover restriction so that portfolio weights remain implementable from one rebalancing date to the next. The long-only strategy, therefore, provides more conservative exposure, whereas the long-short strategy permits additional flexibility at the cost of greater sensitivity to market movements.

We consider two efficient frontiers for each optimization regime: the minimum-risk portfolio and the tangency portfolio. To distinguish among the portfolios, we use the following notation:

LO MVP long-only and mean–variance; the least feasible solution;

LO TVP long-only, mean–variance, tangency; solution;

LO C95 long only, CVaR(95%), no-risk;

LO TC95 long only, CVaR(95%), tangency solution;

LO C99 long-only, CVaR(99%), minimum-risk solution;

LO TC99 long only, CVaR(99%), tangency solution.

The same notation, with LO replaced by LS, is used for the corresponding long-short portfolios.

We analyze realized performance by comparing the cumulative ($100) investment in each optimized portfolio over the non-sample period. Based on this concept, we model a passive benchmark in a buy-and-hold portfolio (BHP) initialized with a similar weight over the 30 securities as the first example, and then the portfolio is held, i.e., without rebalancing. This construction allocates the capital to both assets at its very first formation date and is thus equally distributed, yet the portfolio weights will, by this construction, begin to drift endogenously with relative price movements. Therefore, at the time of inception, BHP is kept equally weighted with respect to its first constituents; its weight is therefore even more likely to shift to the better-performing constituents over time in a process resembling the concentration effect described in Forsyth [44]. In a stylized setting, if n_i denotes the number of units purchased in asset i at time 0, then initial wealth is $W(0) = \sum_{i=1}^{N} n_i S_i(0)$, and under equal initial allocation with no subsequent trading, terminal wealth becomes $W(T) = \sum_{i=1}^{N} n_i S_i(T)$ [44]. This

 doi:10.20944/preprints202603.2360.v1

benchmark, therefore, provides a neutral passive reference against which the value added by dynamic historical optimization can be assessed.

Figure 5 estimates the compounded value of $100 investments under each strategy-optimization combination based on the evaluation period. For long only, the minimum-risk portfolios are fairly close to the buy-and-hold benchmark, and the tangency portfolios yield much higher cumulative growth. Tangency portfolios again exhibit the best cumulative performance when pursuing the long-short strategy, which is also characterized by the better flexibility it affords due to the reduction of exposure and leverage. Conversely, superior performance comes at higher portfolio turnover and even higher implementation costs.

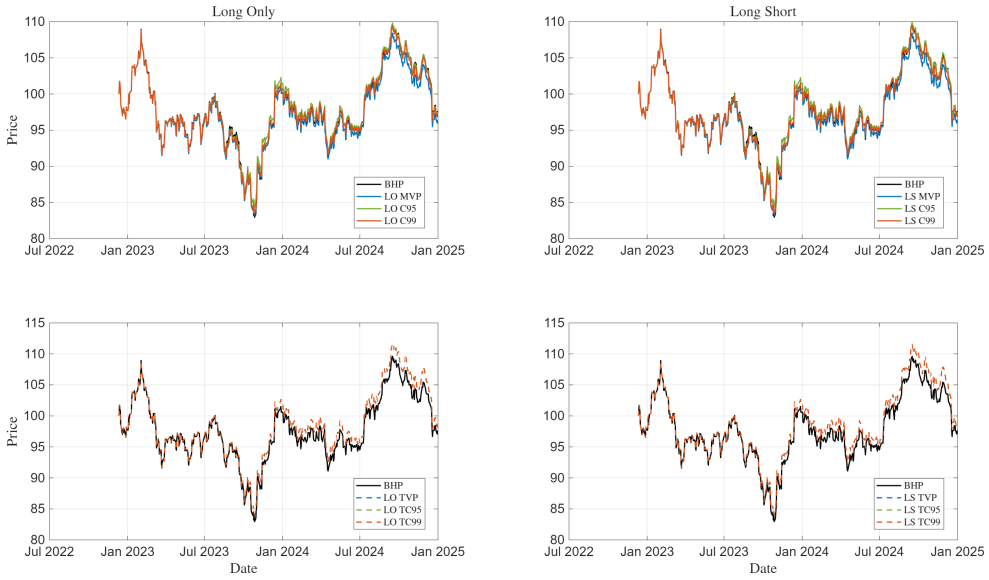

Figure 5. Cumulative value (price) of $100 invested in the portfolio under (left) long-only and (right) long–short strategies under different combinations of optimization and EF solution. (BHP serves as the Benchmark)

Overall, we can say that the historical portfolio performance results indicate that both the optimization criterion and the portfolio constraints have an important impact on realized performance in the listed real-estate security world. Minimum-risk portfolios provide more conservative exposure, and tangency portfolios are higher return-seeking ones. In terms of this, again the disparity between long–only and long–short allocations shows that reduced short exposure may result in relative long-term performance gains, but this enhancement has to be evaluated in relation to the additional trading and risk management related costs of such strategies.

4. Tail Risk Estimation

The portfolio results in Section 3 show that listed real-estate security allocations differ materially in their realized risk–return profiles under mean–variance and CVaR objectives. Those comparisons, however, do not directly describe the frequency or persistence of extreme downside realizations. To address this issue, this section studies the behavior of the lower tail of portfolio returns using tools from extreme value theory (EVT). The analysis is carried out first for benchmark portfolio constructions and then for the optimized synthetic portfolios generated under long-only and long–short allocation rules.

Because EVT is naturally formulated for upper tails, we work with portfolio losses rather than negative returns. Let R_t denote the daily return of a portfolio at time t, and define the corresponding loss variable by

$$L_t = -R_t. \tag{2}$$

 doi:10.20944/preprints202603.2360.v1

Studying large negative returns is therefore equivalent to studying the upper tail of L_t. Throughout this section, threshold exceedances and tail-index diagnostics are applied to the upper tail of the loss distribution. Unless otherwise stated, the threshold is taken at the 95th percentile of losses, corresponding to the worst 5% of return realizations.

4.1. Modeling Tails with the Generalized Pareto Distribution

We begin with a peaks-over-threshold analysis of portfolio losses. Under standard EVT conditions, the conditional distribution of exceedances above a sufficiently high threshold is well approximated by a generalized Pareto distribution (GPD) [3,5]. Its cumulative distribution function is

$$F_{\mathrm{GPD}}(x;\sigma,\xi) = \begin{cases} 1-\left(1+\dfrac{\xi x}{\sigma}\right)^{-1/\xi}, & \xi \neq 0, \\ 1-\exp(-x/\sigma), & \xi = 0, \end{cases} \tag{3}$$

where $\sigma > 0$ is the scale parameter and ξ is the shape parameter. The sign of ξ determines the effective tail class: $\xi > 0$ indicates a heavy tail with Pareto-type decay, $\xi = 0$ corresponds to an exponential tail, and $\xi < 0$ implies bounded support. When $\xi > 0$, the associated tail index is $\alpha = 1/\xi$, so that smaller values of α correspond to heavier tails.

Following McNeil et al. [5], we fit the GPD to exceedances above the 95th percentile of daily losses using maximum likelihood estimation. Figure 6 reports the empirical exceedance distributions and fitted GPD curves for the buy-and-hold portfolio at depths $d = 20$, $d = 30$, and $d = 50$. The fitted curves track the empirical exceedance behavior reasonably well, and the fit becomes smoother as the number of exceedances increases.

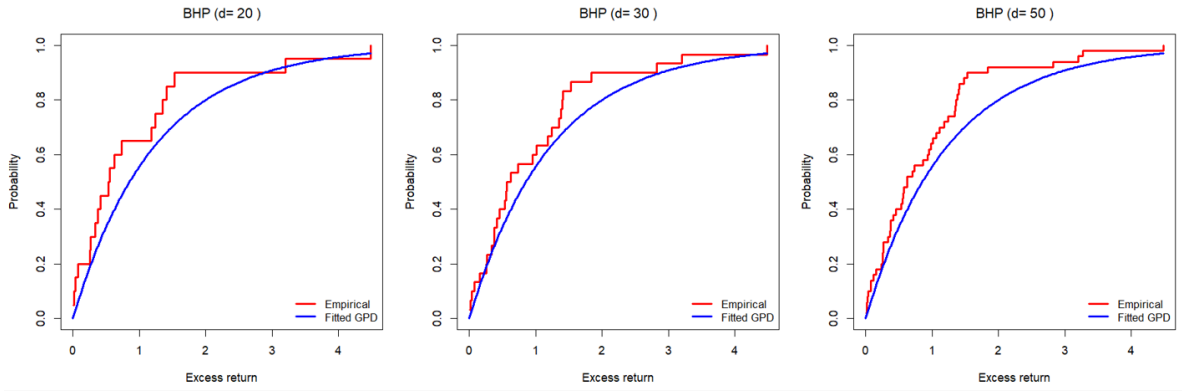

Figure 6. Empirical exceedance CDFs (red) and fitted GPDs (blue) for the BHP portfolio at depths $d = 20$, 30, and 50.

Across the benchmark specifications, the estimated shape parameters are positive, indicating that the tail of the loss distribution decays more slowly than it would under a light-tailed benchmark. This is consistent with the broader empirical literature on financial returns and supports the use of tail-sensitive risk diagnostics in the listed real-estate security setting [3,4].

4.2. Hill Estimator for Tail Index Inference

To complement the parametric GPD analysis, we estimate tail thickness nonparametrically using the Hill estimator [29]. Let $X_1, \ldots, X_n$ denote a sample of losses arranged in descending order,

$$X_{(1)} \geq X_{(2)} \geq \cdots \geq X_{(n)}.$$

 doi:10.20944/preprints202603.2360.v1

For the largest k observations, the Hill estimator of the Pareto tail index α is

$$\hat{\alpha}_{k,n}^{(H)} = \left(\frac{1}{k} \sum_{i=1}^{k} \left(\log X_{(i)} - \log X_{(k+1)} \right) \right)^{-1}. \tag{4}$$

The choice of k is central: very small values use only the most extreme observations and therefore tend to be highly variable, whereas large values reduce variance but may introduce substantial bias if the selected observations are no longer in the genuine tail region. Following De Haan and Peng [32] and McNeil et al. [5], we focus on ranges in which $k/n \in [0.01, 0.05]$.

For visual assessment, we also report Wald-type confidence bands based on Haeusler and Segers [33]. The corresponding interval is

$$CI_n^{\text{Wald}}(\theta, k) = \left[\left(1 + \frac{z_{\sqrt{\theta}/2}}{\sqrt{k}} \right)^{-1} \hat{\alpha}_{k,n}^{(H)}, \left(1 - \frac{z_{\sqrt{\theta}/2}}{\sqrt{k}} \right)^{-1} \hat{\alpha}_{k,n}^{(H)} \right], \tag{5}$$

where z_p denotes the p-th quantile of the standard normal distribution. In this setting, Hill plots are used primarily as diagnostics: a relatively flat segment indicates a locally stable tail estimate, whereas sharp oscillation or abrupt trend changes indicate threshold sensitivity or departure from a simple Pareto regime.

Figures 7 and 8 report left-tail Hill plots for the buy-and-hold and market-cap-weighted real estate portfolios, each benchmarked against the DJIA. Under the loss transformation $L_t = -R_t$, these correspond to upper-tail diagnostics for extreme negative returns.

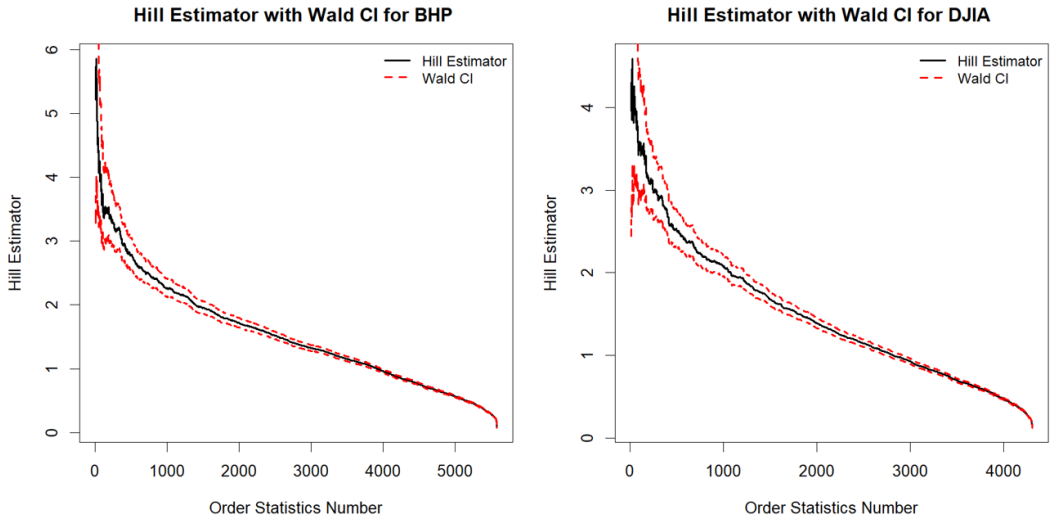

Figure 7. Left-tail Hill estimator with 95% Wald confidence intervals for (left) the BHP and (right) the DJIA benchmark.

For the buy-and-hold benchmark, the Hill curve begins at a higher level than the corresponding DJIA curve, indicating a heavier downside tail. As k increases, the curve moves into a more regular region, suggesting that extreme losses are not isolated anomalies but reflect persistent heavy-tailed behavior relative to the broad equity benchmark.

The market-cap-weighted benchmark exhibits somewhat smoother Hill behavior and narrower confidence bands than the buy-and-hold portfolio, suggesting that capitalization weighting dampens the contribution of the most extreme downside realizations. Even so, both real estate benchmark portfolios remain materially heavier-tailed than the DJIA. Taken together, the GPD and Hill results indicate that downside tail risk in this sector is both statistically and economically relevant.

 doi:10.20944/preprints202603.2360.v1

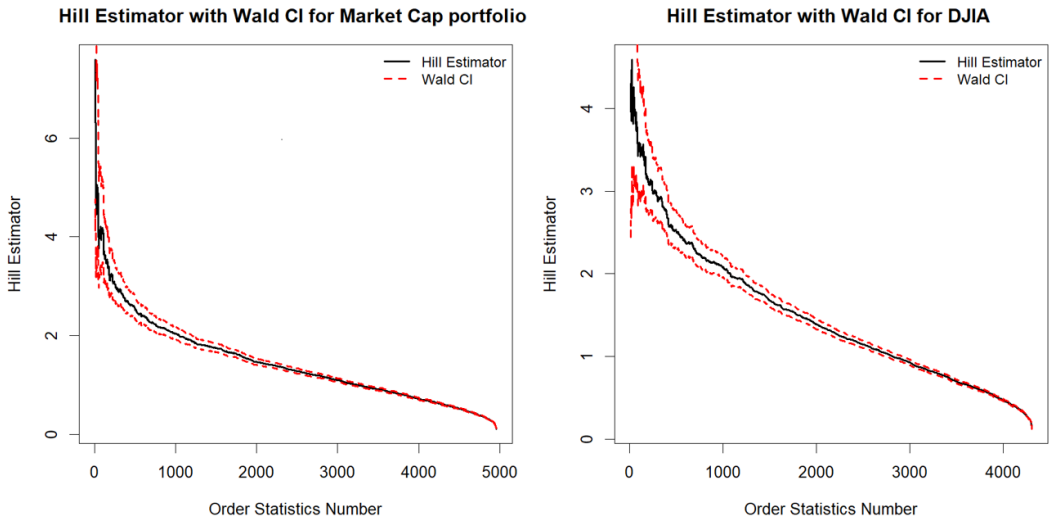

Figure 8. Left-tail Hill estimator with 95% Wald confidence intervals for (left) the market-cap-weighted portfolio and (right) the DJIA benchmark.

4.3. Tail Diagnostics for Optimized Synthetic Portfolios

We next extend the EVT analysis to the optimized synthetic portfolios constructed under the LO and LS allocation rules. The purpose is twofold. First, we compare the empirical tail behavior induced by different optimization criteria. Second, we assess whether the corresponding tail region is sufficiently regular to support meaningful threshold-based inference.

The optimized synthetic portfolio return series is generated using a two-year rolling-window historical optimization framework. For each window, portfolio weights are estimated from past return observations and then used to form the corresponding synthetic portfolio return series for tail-risk analysis. This procedure follows Lindquist et al. [45]; full construction details are reported in Appendix C.

We begin with grouped kernel density estimates of the synthetic return distributions. For the long-only portfolios, Figure 9 reports the grouped densities for the first rolling window, separated into minimum-risk and tangency groups, with the 7th percentile marked as a common empirical left-tail cutoff. The LO return distributions are relatively concentrated, and the differences across MVP, C95, C99, TVP, TC95, and TC99 are visible but not erratic. In particular, the left tails display moderate variation in spread and asymmetry, suggesting that a Hill-based comparison of the associated loss series is feasible.

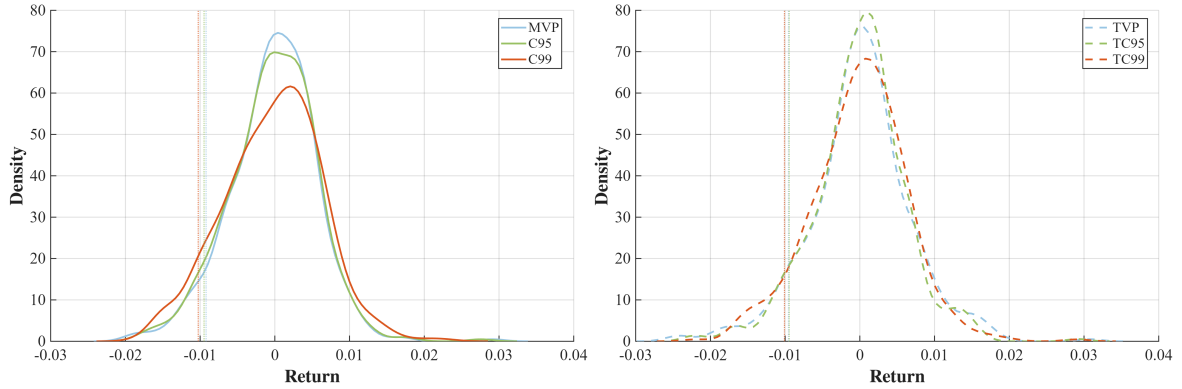

Figure 9. Grouped kernel density estimates (left) the minimum-risk group and (right) tangency group for long-only synthetic portfolios in the first rolling window.

To examine tail thickness more directly, we consider raw Hill plots using the Hill estimator (4) for the long-only loss series. The central practical issue is the selection of the number k of upper-order statistics used in the tail fit. Very small values of k typically produce highly variable estimates

 doi:10.20944/preprints202603.2360.v1

because they rely on only a few extreme observations, whereas large values of k can induce bias by incorporating observations that no longer belong to the genuine tail region.

To determine k in a systematic way, we use the asymptotic mean squared error (AMSE) criterion described by de Sousa and Michailidis [46]. Under the exponential regression approximation for upper-order log spacings,

$$i\big(\log X_{(i)} - \log X_{(i+1)}\big) \approx \left(\alpha^{-1} + b_{n,k}\left(\frac{i}{k+1}\right)^{-\rho}\right)e_i, \qquad i = 1, \ldots, k, \tag{6}$$

where the e_i are independent standard exponential random variables, α is the tail index, and $b_{n,k} = b\left(\frac{n+1}{k+1}\right)$ for a positive function $b(\cdot)$ satisfying $b(x) \to 0$ as $x \to \infty$, the AMSE of the inverse Hill estimator can be written as

$$\mathrm{AMSE}_{H_{k,n}^{-1}} = \left(\frac{b_{n,k}}{1-\rho}\right)^2 + \frac{\alpha^{-2}}{k}. \tag{7}$$

The threshold is then selected by

$$k_{\mathrm{opt}} = \arg\min_{1<k<n} \mathrm{AMSE}_{H_{k,n}^{-1}}. \tag{8}$$

This provides a data-driven alternative to purely visual threshold selection and balances the reduction in variance from using more observations against the bias introduced by moving too far away from the extreme tail.

The resulting raw Hill plots for the long-only portfolios are reported in Figure 10. The AMSE-selected thresholds, indicated by circles, generally fall within regions where the curves are comparatively stable. This does not imply that the LO tails are exactly Pareto, but it does indicate that local tail-index diagnostics remain interpretable. Some additional variability appears for the portfolios associated with tighter downside constraints, especially at very small thresholds, yet the overall behavior remains regular enough for the Hill–AMSE procedure to provide useful local information.

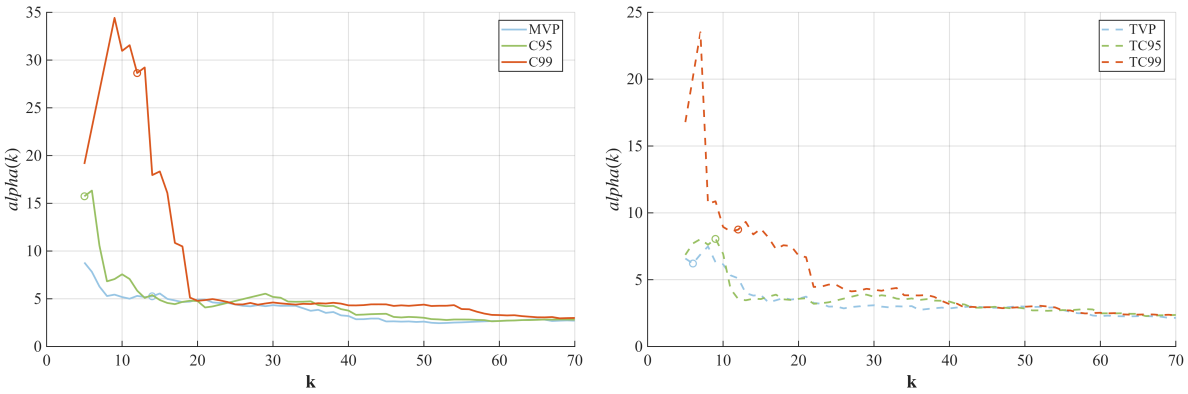

Figure 10. Raw Hill plots (left) the minimum-risk group and (right) the tangency group for the long-only synthetic portfolios in the first rolling window. Circles indicate the AMSE-selected threshold locations.

The long-short portfolios present a different picture. In the raw LS Hill plots, the extreme left tail is markedly less regular, particularly for the portfolios associated with the 99% CVaR criterion. The Hill curves are more sensitive to the threshold choice, and the local behavior in the far tail is often too erratic to support a stable interpretation. For this reason, we treat the LS case more cautiously and avoid reporting a single-tailed index estimate based on the raw Hill graphs alone.

To investigate whether the observed instability is partly driven by local roughness in the empirical tail, we smooth the LS loss distributions using kernel density estimation with alternative bandwidth multipliers. Even after this additional smoothing, however, the far-left tail remains visually unstable, especially for the most tail-focused portfolios. The LS diagnostics therefore, suggest that threshold sensitivity is not merely a plotting artifact but reflects a genuinely irregular extreme-tail structure.

 doi:10.20944/preprints202603.2360.v1

In the LS case, the diagnostic is based on the survival function rather than on a density fit. Let $\widehat{F}(x)$ denote the KDE-based estimator of the loss CDF. The corresponding survival function is

$$S(x) = 1 - \widehat{F}(x). \tag{9}$$

If the tail is locally Pareto, then

$$S(x) = P(X > x) = \left(\frac{x_m}{x}\right)^{\alpha}, \qquad x \geq x_m, \tag{10}$$

which implies the log–linear relation

$$\log S(x) = \alpha \log x_m - \alpha \log x. \tag{11}$$

Thus, a plot of $\log S(x)$ against $\log x$ should be approximately linear over any region in which a Pareto approximation is plausible, with slope $-\alpha$.

Following this logic, we examine KDE-smoothed log–log survival plots for the LS loss series. Figure 11 reports representative examples for the first rolling window using bandwidth multiplier $m = 2.5$. The purpose of this step is diagnostic rather than inferential. If a stable Pareto regime were present, one would expect an approximately linear segment over the relevant tail region. Instead, the fitted relationships remain sensitive to both smoothing choice and the portion of the tail under consideration.

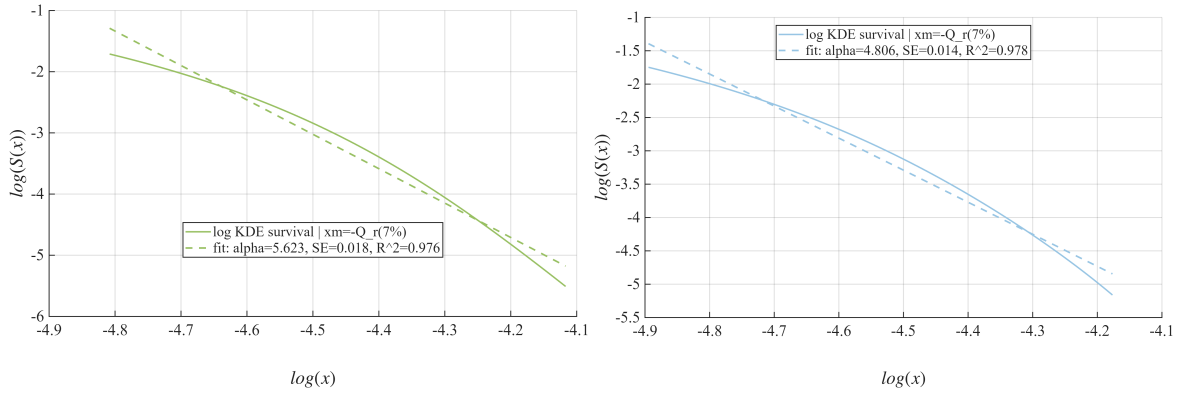

Figure 11. Log–log survival diagnostics (left) C95 and (right) MVP for LS synthetic portfolios in the first rolling window, based on KDE-smoothed loss distributions with bandwidth multiplier $m = 2.5$.

In the present application, the LS log–log survival plots do not display a sufficiently stable linear regime to justify reporting a reliable Pareto tail index. Although smoothing reduces some of the local roughness, the fitted relationships remain sensitive to the bandwidth and to the exact tail window being examined. For that reason, the LS diagnostics are interpreted conservatively: they indicate that the downside tail is more irregular and less stable than in the long-only case, but they do not support a robust claim about a unique value of α.

Overall, the synthetic-portfolio diagnostics confirm that downside-tail behavior depends materially on both the optimization criterion and the portfolio constraint set. Long-only portfolios exhibit heavy but comparatively more stable tails, for which Hill plots with AMSE-selected thresholds provide useful local diagnostics. Long–short portfolios, by contrast, display substantially greater irregularity in the far tail, especially under tighter CVaR constraints, and therefore call for a more cautious interpretation of EVT-based tail summaries.

5. Risk Metrics

To thoroughly assess portfolio performance, we use a set of historical risk-adjusted metrics that include both volatility, downside risk, and tail asymmetry combined. Specifically, we consider the

 doi:10.20944/preprints202603.2360.v1

volatility, *Sharpe ratio*, *Sortino ratio*, *Rachev ratio*, *maximum drawdown*, and *information ratio* of each portfolio configuration. The criteria selected due to their convergent view into properties of the distribution of returns are: volatility measures the overall spread; Sharpe and Sortino ratios are used to estimate risk-adjusted returns based on symmetric and downside volatility; and the Rachev ratio of tails also demonstrates reward-to-risk asymmetry at tails. The maximum drawdown is an indicator of the measured degree of actual loss, and the information ratio gives a measure of the return of active returns against benchmarks. Collectively, these offer a multi-dimensional measure of stability and reward efficiency, in addition to the ordinary mean–variance analysis.

Comparative results between the LO, LS, and benchmark BHP are shown in Figure 12. The results prove that CVaR-constrained LS (LS C95 and LS C99) provides the best risk-adjusted performance, especially under tail-sensitive measures, Rachev and Sortino.

Thus, on *Panel (a)*, which includes the information ratio, a single strategy with a fixed value of LS C95 and LS C99 presents higher performance relative to other strategies (i.e., more excess returns per unit of tracking error, while LS C95 and LS C99 had the same tracking error). BHP and variance-driven, long-only portfolios, similar to their more defensive structures and less responsive to regime shifts, have almost zero information ratios.

Panel (b) shows maximum drawdown, which means the long–short portfolios, and in particular LS MVP, LS C95, and LS C99, have a smaller peak-to-trough drop than the BHP or long-only portfolios. The results support the enhanced downside strength and capital preservation offered by short-side exposure to CVaR-based constraints. By contrast, BHP shows the biggest drawdowns, because it has less hedged exposure to market corrections.

The results of the Rachev ratio further suggest *Panel (c)* the above conclusion. Both LS C95 and LS C99 exhibit the highest Rachev ratios in both configurations, which is evidence that optimization performed on a CVaR-based basis provides optimal risk-reward trade-off in the tails due to reducing the most extreme losses while keeping upside potential intact. Long-only portfolios exhibit lower and more uniform Rachev ratios, further verifying their capacity to exploit asymmetry in the distribution of returns.

Panel (d) shows Sortino ratio, which reduces downside risk only. Although all ratios are negative – characteristic of difficult listed real estate securities during 2021-2024 long-short portfolios do again less negatively, indicating a smoother downward trend which is better managed for unfavorable returns.

The same trend is also found in *Panel (e)*, where the Sharpe ratio is added. While all are negative, the LS portfolios, especially LS C95 and LS C99 give the least negative Sharpe ratios, in line with their capacity to give relatively higher mean returns with lower volatility. The BHP and LO MVP, however, have the weakest outcomes, indicating higher volatility and less compensation for returns received.

Lastly, *Panel (f)* highlights the overall Volatility. Long–short strategies show lower standard deviations than long-only portfolios consistently, particularly LS MVP and LS C95, highlighting the stabilizing function of short positions in reducing global portfolio risk.

In general, the comparative evidence presented in Figure 12 indicates the conclusion of this study. In a cross-curve performance analysis, CVaR-constrained long–short portfolios (LS C95, LS C99) exhibit significantly better risk and tail-adjusted performance. In addition to lowering volatility and drawdowns, they improve reward asymmetry and active efficiency measured by Rachev and information ratios. The buy-and-hold benchmark, though easy to understand and widely diversified continues to be highly vulnerable to the whims of downturns—demonstrating that tail-sensitive optimization of a portfolio is important when designing a listed real-estate security portfolio.

 doi:10.20944/preprints202603.2360.v1

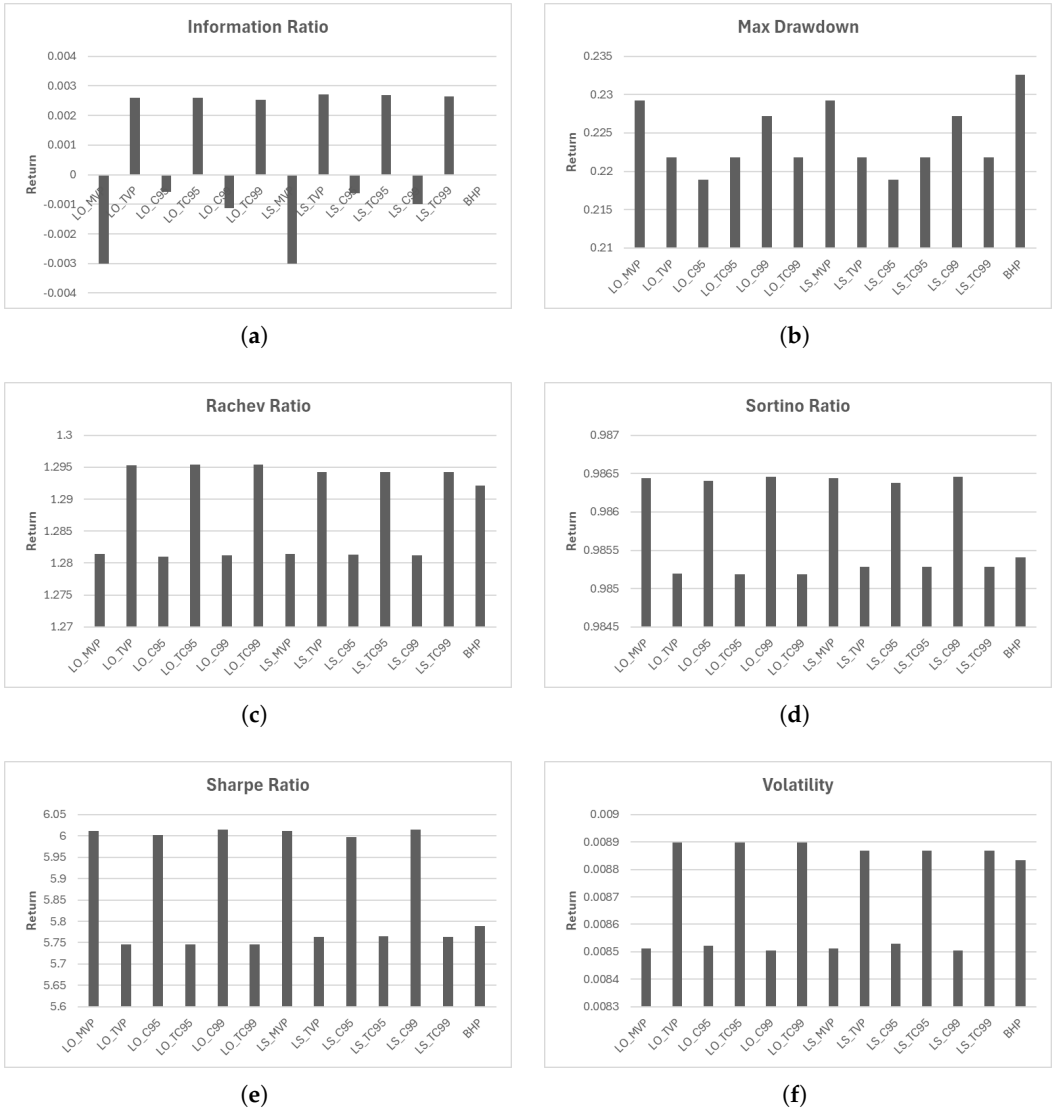

Figure 12. Portfolio risk vs. performance metrics comparison.(a) information ratio, (b) maximum drawdown, (c) Rachev ratio, (d) Sortino ratio, (e) Sharpe ratio, and (f) volatility.

6. Robust Regression of Securities on the Benchmark Portfolio

The preceding sections documented the distributional features of listed real-estate security returns—heavy tails, volatility clustering, heterogeneous risk profiles and analyzed their implications for mean–risk efficiency. We now examine how individual securities load on the sector factor and quantify the residual risk that remains after controlling for common variation. Since the tail behavior identified earlier generates high-leverage observations and local bursts of volatility, classical ordinary least squares estimation is not well-suited for stable inference on factor exposures. Robust procedures with bounded influence provide substantially greater stability under such error structures.

We consider a single-factor market model for security i, $R_{i,t} = \alpha_i + \beta_i R_t^{\mathrm{BHP}} + \varepsilon_{i,t}$, where R_t^{BHP} denotes the BHP returns. The parameter β_i measures the security's sensitivity to the sector factor, while α_i captures the component of expected return orthogonal to the benchmark at the daily frequency. Given the empirical evidence on heavy tails (Section 4), we estimate (α_i, β_i) via an MM procedure [41–43]. The estimator begins with a high-breakdown S-estimator that is consistent under broad contamination models and then refines the fit through an M-step using a bounded ψ-function (Huber/Tukey) solved via iteratively reweighted least squares. This yields regression coefficients with finite influence and high asymptotic efficiency under light tails while remaining stable when extreme observations occur.

 doi:10.20944/preprints202603.2360.v1

Table 1 reports robust estimates for representative securities, including sample size (n), slope $\hat{\beta}_i$ and its 95% confidence interval, intercept $\hat{\alpha}_i$, and pseudo-R^2. The reported value $n = 530$ corresponds to a two-year daily estimation window after alignment of benchmark and security returns. The broad, diversified funds; *REET*, *VNQ*, *IYR*, and all exhibit positive factor loadings that are materially below one, with comparatively tight confidence intervals and the highest pseudo-R^2 values in the table. This indicates that these securities comove positively with the sector benchmark, but their day-to-day fluctuations are less than one-for-one with the buy-and-hold benchmark. In contrast, subsector exposures display more heterogeneous loading patterns. *ITB* (home construction) remains positively exposed but with lower fit, consistent with the greater cyclicality documented earlier. *OHI* (healthcare REITs) has the smallest slope and the weakest fit, reflecting looser integration with the aggregate sector factor. The mortgage REIT *ARR* lies between these cases, with materially larger residual dispersion. Across all specifications, intercepts remain economically small once high-influence days are down-weighted, reinforcing that cross-sectional differences arise mainly from factor exposure and idiosyncratic volatility rather than persistent abnormal performance.

Table 1. Estimated regression coefficients (β) with 95% confidence intervals from robust regressions of selected securities on the BHP.

Security	n	β	95% CI(β)	SE(β)	α	R^2
REET	530	0.336	[0.303 , 0.369]	0.0169	6.76×10^{-5}	0.2534
VNQ	530	0.351	[0.313 , 0.389]	0.0193	6.78×10^{-5}	0.2409
IYR	530	0.358	[0.320 , 0.395]	0.0190	1.22×10^{-4}	0.2466
ARR	530	0.341	[0.279 , 0.403]	0.0314	-1.22×10^{-4}	0.1511
ITB	530	0.343	[0.286 , 0.400]	0.0290	4.78×10^{-4}	0.1816
OHI	530	0.192	[0.144 , 0.241]	0.0248	1.84×10^{-4}	0.0907

6.1. Residual Diagnostics

To assess model adequacy, we examine residuals standardized by a robust scale estimator. For security i, fitted values are $\widehat{R}_{i,t} = \widehat{\alpha}_i + \widehat{\beta}_i R_t^{\mathrm{BHP}}$ and residuals $\widehat{\varepsilon}_{i,t} = R_{i,t} - \widehat{R}_{i,t}$. Standardization is based on the median absolute deviation (MAD), adjusted for Gaussian consistency: $\widehat{\sigma}_i = \mathrm{MAD}(\widehat{\varepsilon}_{i,\cdot})/0.67449$, and $\widehat{u}_{i,t} = \widehat{\varepsilon}_{i,t}/\widehat{\sigma}_i$.

Diagnostics proceed on two dimensions. First, plots of $\widehat{u}_{i,t}$ against fitted values—using point opacity proportional to the final MM robustness weights—test for remaining curvature or scale effects. Second, Ljung–Box statistics applied to $\widehat{u}_{i,t}$ and $\widehat{u}_{i,t}^2$ assess the presence of autocorrelation or additional volatility structure. Figure 13 shows that after robust fitting, no systematic nonlinearities remain, and weight shading indicates that large innovations have limited impact on the slope. The magnitude and dispersion of standardized residuals align with the coefficient evidence: broad aggregates (*REET*, *VNQ*, *IYR*) exhibit concentrated residual distributions, while subsector securities (*ITB*, *OHI*) display broader residual patterns consistent with their higher idiosyncratic volatility. *ARR* occupies an intermediate position. Ljung–Box statistics indicate weak residual serial dependence, suggesting that the robust market model captures the dominant component of daily co-movement.

 doi:10.20944/preprints202603.2360.v1

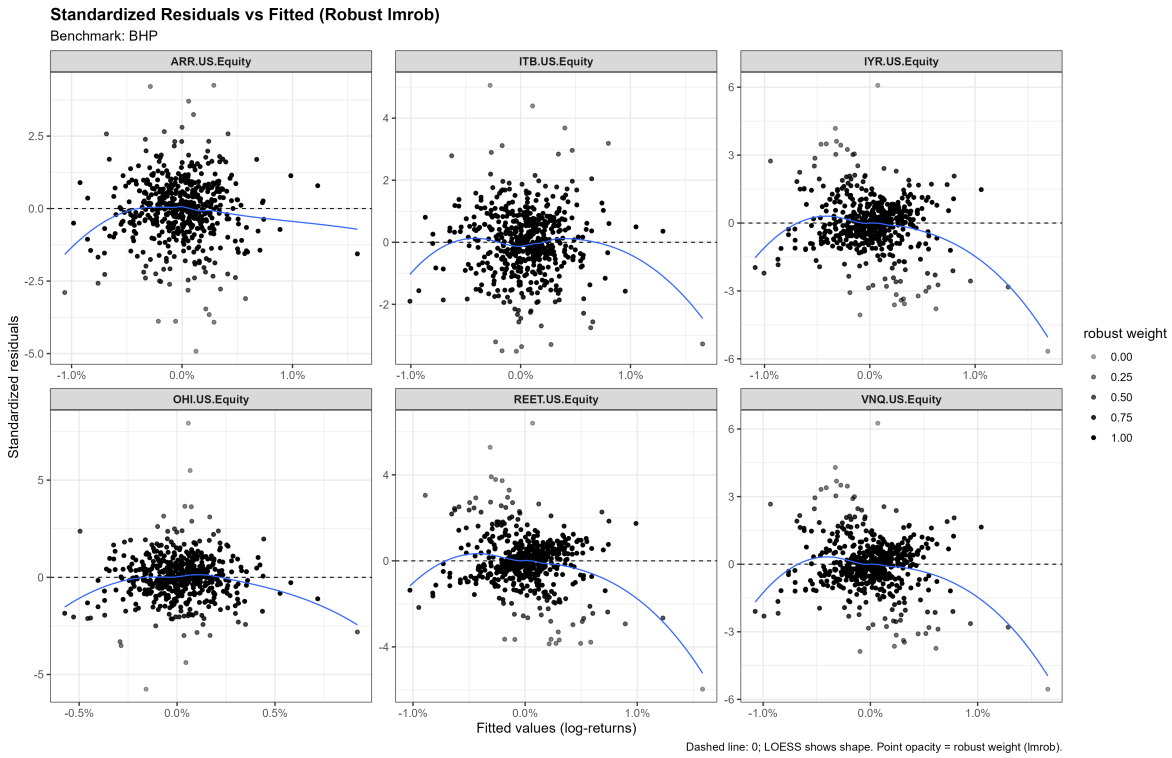

Figure 13. Standardized residuals versus fitted values from robust regressions of six securities on the BHP. Dashed line: zero. LOESS smooths highlight any curvature. Point opacity reflects robust weights (lighter = down-weighted).

Overall, the robust regressions delineate two segments of the listed real-estate security space. Diversified aggregates—characterized by high factor loadings, low residual dispersion, and limited sensitivity to heavy-tailed shocks—form the core of sector exposure and thus anchor the Markowitz and CVaR efficient frontiers. Subsector securities exhibit differentiated and more volatile loading patterns, reflecting exposure to narrower economic drivers. These differences in factor sensitivity and residual risk correspond directly to the wider dispersion, heavier tails, and more varied efficiency positions documented earlier, thereby linking the regression evidence to the portfolio-level results in Sections 3.1 and 5.

7. Option Pricing with Lévy-Based Dynamics: A Double-Subordinated Normal Inverse Gaussian Approach

To model the heavy-tailed and asymmetric dynamics of listed real-estate security portfolios, we adopt the NDIG process introduced by Shirvani et al. [47]. This framework extends the classical NIG model through *double subordination*, capturing both the information-time deformation of market activity and the stochastic scaling of volatility. Traditional models such as Black–Scholes assume constant variance and Gaussian returns, which fail to describe the observed kurtosis, skewness, and volatility clustering in sectoral real-estate security data. The NDIG process, by contrast, provides a structurally consistent method for modeling these features through two independent Lévy subordinators that drive the evolution of price and volatility over intrinsic time.

Following Shirvani et al. [47], we define the log-price dynamics as:

$$S_t = e^{X_t}, \quad t \in [0, \tau], \tag{12}$$

$$X_t = X_0 + \mu t + \gamma U(t) + \rho T(U(t)) + \sigma B_{T(U(t))}, \quad t \geq 0, \tag{13}$$

where B_t is a standard Brownian motion, and $U(t)$, $T(t)$ are Lévy subordinators representing (i) the intrinsic market activity time and (ii) the stochastic volatility-time process, respectively. The drift term

 doi:10.20944/preprints202603.2360.v1

μ captures the expected rate of return, while σ denotes instantaneous volatility. The nested structure $T(U(t))$ induces time deformation and volatility clustering, allowing return increments to exhibit semi-heavy tails and self-similar scaling. This double subordination ensures that the price process follows a martingale under the risk-neutral measure $\mathbb{Q}$, maintaining consistency with no-arbitrage conditions.

Under $\mathbb{Q}$, the characteristic function of the NDIG log-price process, $\varphi_{\ln S_\tau^{(\mathbb{Q})}}(v)$, is used to value European options through the Fourier-based Carr–Madan approach [37]:

$$C(S_0, r, k, \tau) = \frac{e^{-r\tau - ak}}{\pi} \int_0^\infty e^{-ivk} \frac{\varphi_{\ln S_\tau^{(\mathbb{Q})}}(v - i(a+1))}{a^2 + a - v^2 + i(2a+2)v}\, dv, \tag{14}$$

where $a > 0$ ensures numerical convergence, r is the risk-free rate, and $k = \ln K$ denotes the strike in log space. The FFT method is employed for efficient computation, enabling arbitrage-free pricing across the strike–maturity grid. Model parameters are estimated via moment-matching, aligning the NDIG-implied cumulants with those empirically obtained from security returns. Because the calibration is performed to return moments rather than observed option prices, the surfaces reported below should be interpreted as NDIG model-implied surfaces rather than market-implied surfaces extracted from traded options.

This NDIG setting unifies volatility clustering, tail asymmetry, and market activity heterogeneity within one coherent pricing model. Its flexibility is particularly suitable for listed real-estate securities, where liquidity constraints and macro-financial shocks generate prolonged volatility persistence and asymmetric return distributions.

Figures 14–17 illustrate the empirical outcomes of the NDIG option-pricing model for the listed real-estate security portfolio. These results combine the model's theoretical structure—double subordination of market and volatility time—with its empirical calibration to security return data, highlighting how intrinsic time deformation and stochastic volatility jointly shape the pricing kernel.

Figure 14 presents the NDIG-implied call and put price surfaces computed using the Carr–Madan FFT approach. Both surfaces exhibit the expected *no-arbitrage geometry*: call prices increase with maturity and decline with strike, while put prices rise with strike and maturity. The steep gradients observed near short maturities and low strike regions emphasize the high marginal value of near-term downside protection. Unlike the log-normal benchmark of Black–Scholes, the NDIG-generated surfaces are asymmetric, reflecting the negative skewness and excess kurtosis inherent in security-return distributions. This asymmetry originates from the interaction of the two subordinators—$U(t)$ and $T(U(t))$—which accelerate market time during volatile episodes and amplify volatility clustering. These results closely parallel the stylized behavior reported in Shirvani et al. [47], confirming that NDIG pricing preserves realistic jump asymmetry across asset classes.

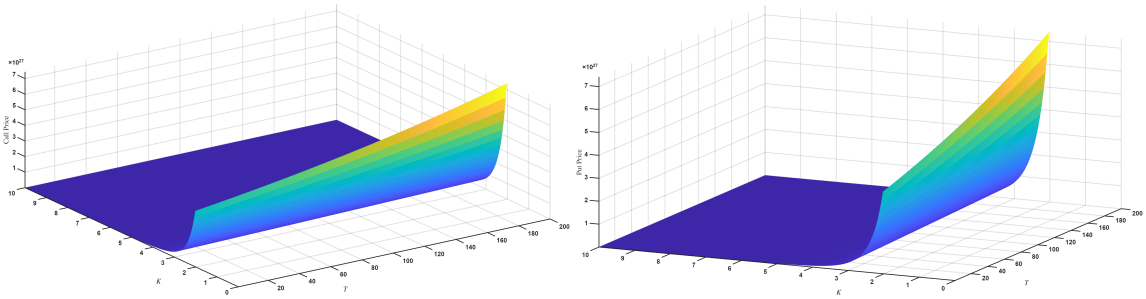

Figure 14. NDIG-implied (left) call and (right) put option price surfaces across strike (K) and maturity (T) dimensions. The convex structure and asymmetry reflect semi-heavy tails and downside jump sensitivity in returns.

The decomposition of model-implied variance is visualized in Figures 15 and 16, which plot the integrands $Q_1(K,t)$ and $Q_2(K,t)$ derived from the Fourier-based variance representation of the NDIG model. These integrands capture how different strike–maturity regions contribute to the overall model-implied variance.

 doi:10.20944/preprints202603.2360.v1

Figure 15 shows $Q_1(K,t)$, the integrand primarily associated with right-tail (upside) contributions. The surface exhibits moderate elevation at long maturities and low strikes, suggesting that while upside shocks are priced, their contribution to overall variance remains limited. In contrast, Figure 16 depicts $Q_2(K,t)$, the integrand capturing downside (left-tail) variance. Here, variance concentration is far more pronounced, particularly for short-dated maturities and lower strike values. These peaks correspond to periods of high downside risk premia—consistent with the post-2020 tightening cycle, when volatility persistence and asymmetric tail risk dominated security dynamics. The shift of variance mass from short-dated to longer maturities through 2023–2024 indicates the gradual normalization of risk, mirroring the diffusion–jump transition typical in NDIG dynamics. This behavior confirms that most of the model-implied variance in the listed real-estate security sector originates from left-tail insurance demand, in agreement with the double-subordinated Lévy evidence in Shirvani et al. [47].

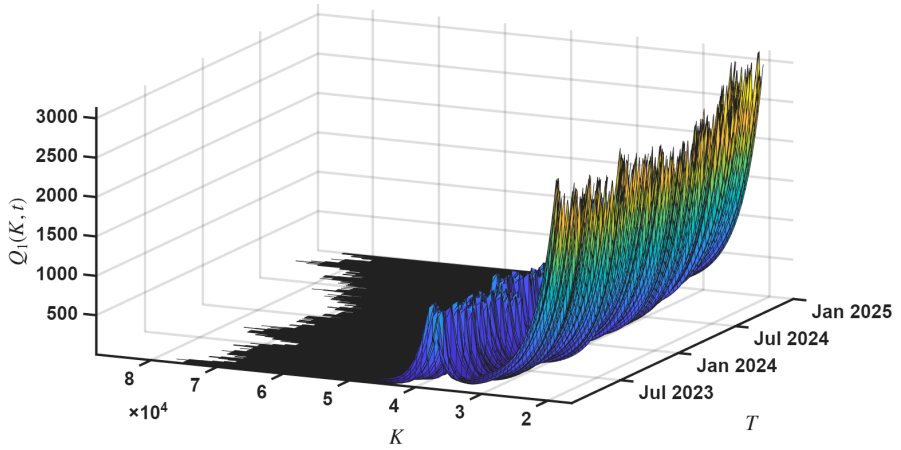

Figure 15. Variance integrand $Q_1(K,t)$ highlighting upside risk contributions over strike (K) and maturity (T). The smoother curvature indicates limited sensitivity of call-side variance to volatility clustering.

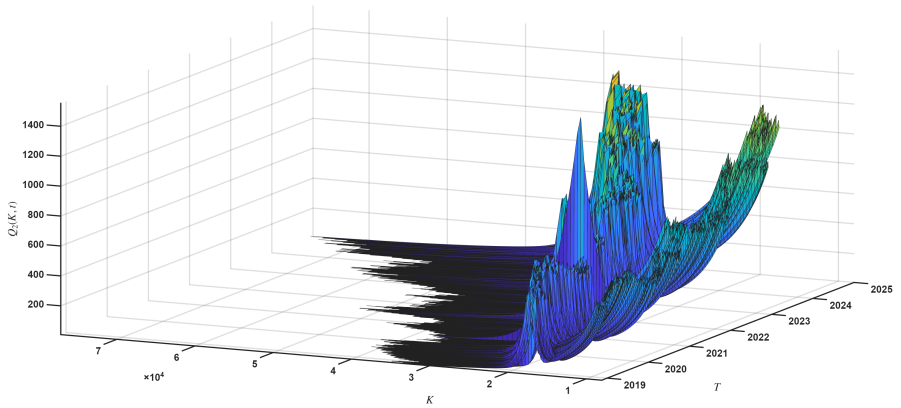

Figure 16. Variance integrand $Q_2(K,t)$ emphasizing left-tail risk concentration in short maturities and low-strike regions. Peaks correspond to periods of heightened downside risk and elevated volatility persistence.

Figure 17 displays the resulting implied-volatility surface derived from the NDIG-calibrated call prices. The surface reveals a well-defined *volatility smile* at short maturities and gradual flattening as time to maturity increases. This curvature captures two key empirical phenomena: (i) investors' strong demand for short-term protection against jumps, reflected in elevated out-of-the-money put volatility, and (ii) the progressive dominance of diffusive variance at longer horizons, yielding a smoother term structure. The persistence of the smile across moderate maturities demonstrates the ability of the NDIG process to reproduce both instantaneous and persistent volatility effects within a single modeling framework. This dynamic behavior, combining stochastic time deformation and volatility clustering, reinforces the NDIG model's explanatory power relative to simpler Lévy or diffusion-based counterparts.

 doi:10.20944/preprints202603.2360.v1

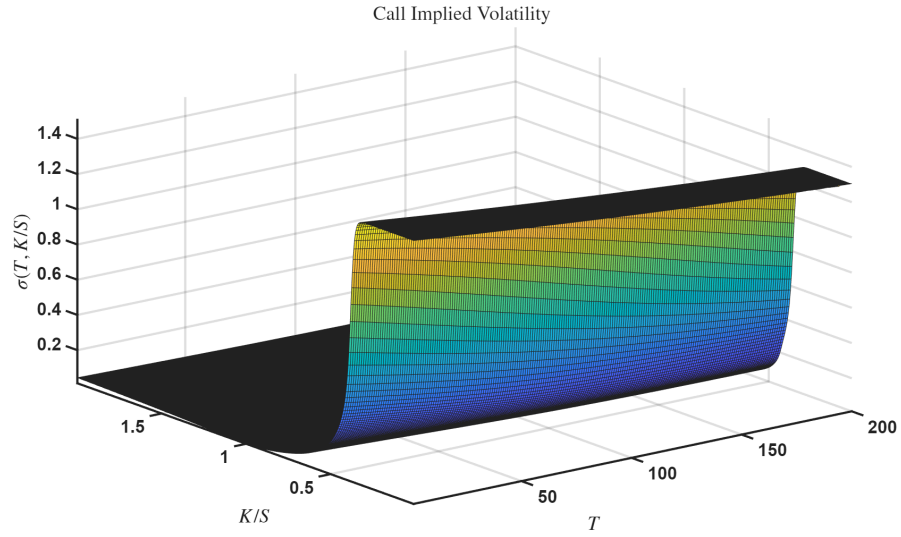

Figure 17. NDIG-implied volatility surface showing pronounced short-term smile and gradual term-structure flattening. Heavy tails and volatility clustering drive curvature at low moneyness and short maturities.

Overall, the NDIG option-pricing results support the model's ability to replicate key stylized facts of option markets linked to the listed real-estate security sector. The asymmetry in call–put surfaces, the left-tail dominance of variance integrands, and the short-maturity smile all indicate that the underlying security dynamics are driven by persistent tail risk and state-dependent volatility—properties well captured by the double-subordinated structure. These empirical outcomes extend the findings of Shirvani et al. [47] to a new asset class, demonstrating that NDIG dynamics provide a unified foundation for modeling both digital-asset and real-estate-based portfolios under uncertainty.

8. Long-Range Dependence in Volatility

The preceding sections document heavy tails, downside asymmetry, and option-implied evidence of non-Gaussian risk in listed real-estate securities. We now turn to the temporal dimension of this behavior by examining whether volatility shocks dissipate quickly or remain persistent over time. In financial markets, long-range dependence (LRD) refers to the slow, hyperbolic decay of autocorrelations in conditional variance, implying that the effects of volatility shocks can persist over extended horizons. Identifying such persistence is important for risk measurement, volatility forecasting, and dynamic portfolio allocation. To assess whether the BHP exhibits long-memory behavior in volatility, we estimate an AR(1)-FIGARCH(1,d_σ,1) model following Baillie [34] and Bollerslev [35].

The FIGARCH model expands the GARCH specification by providing a fractional differencing operator in the conditional variance equation. In [0,1], we specify the degree of long memory using d_σ at $d_\sigma = 0$, the model reduces to a standard GARCH process, which has short memory. When d approaches 1, it implies strong persistence and hyperbolic decay of the volatility autocorrelations. Hence, this framework allows us to distinguish between short-lived volatility clustering and more sustained persistence in conditional variance.

Figure 18 shows the rolling estimates of the fractional differencing parameter d_σ based on the AR(1)–FIGARCH(1,d_σ,1) model. Estimates are nearly homogenous throughout most of the sample, suggesting large volatility persistence. The small temporary declines in d around mid-2023 and early 2024 illustrate temporary stabilization of volatility and are likely to be followed by a market recovery. On the other hand, in general, we can see how listed real-estate security volatility is a memory phenomenon at work, where shocks decay slowly as opposed to dissipating quickly as is the case with short-memory models.

 doi:10.20944/preprints202603.2360.v1

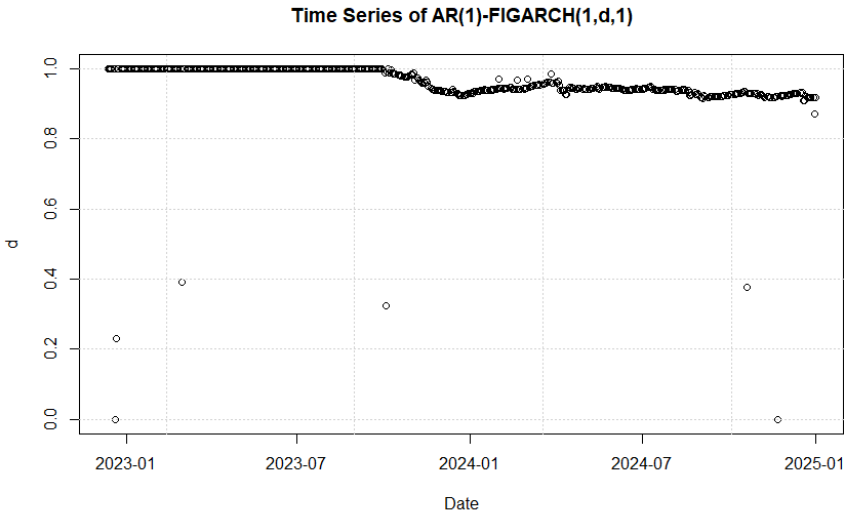

Figure 18. Rolling estimates of the fractional differencing parameter d_σ from an AR(1)–FIGARCH(1,d_σ,1) model fit to BHP returns over two years, rolling windows. Strong persistence and long-memory behavior in volatility are both indicated by large d_σ values.

To analyse persistence in the conditional mean and variance of the BHP return series, we also fitted an ARFIMA(1,d,1)–GARCH(1,1) model. The estimated fractional differencing parameter was approximately zero, indicating little evidence of long-memory in the mean, and while the variance dynamics remained robustly strong, full parameter estimates are included in appendix Table A2.

These results provide evidence that volatility in the buy-and-hold listed real-estate security portfolio is long-range, which is consistent with its persistence in volatility clustering rather than mean-level dependence. In practice, this suggests that volatility shocks may carry longer-term impact on future risks as well, and thus models allowing long time variance–like FIGARCH–may be suitable for modeling to or forecasting portfolio risk over a period of time. However, for the fact that it is largely based on variance rather than mean returns, short memory frameworks such as GARCH or CVaR optimization are also enough to carry out downside risk control when volatility persistence is taken into account.

9. Conclusions

Real-estate securities traded on exchanges offer the liquidity and accessibility of equities, but their risk behavior is shaped by forces that are often more complex than those captured by conventional portfolio models. Using daily data for a broad cross-section of listed real-estate securities over 2021-2024, this study shows that the risks in this market are not merely episodic or incidental. Instead, they are embedded in the return-generating process itself, through heavy tails, pronounced downside sensitivity, and persistent volatility. In such an environment, Gaussian assumptions and variance-based summaries alone are often too narrow to describe the full extent of downside exposure, especially during periods of market stress.

The portfolio results highlight this point clearly. Tail-aware optimization produces allocations that are more closely aligned with the observed risk characteristics of listed real-estate security returns, while mean–variance optimization remains a useful benchmark for illustrating diversification trade-offs. The broader set of performance and risk diagnostics—including volatility, Sharpe and Sortino ratios, the Rachev ratio, maximum drawdown, and information ratio—further shows that meaningful differences in stability and reward efficiency emerge across strategies once downside risk is taken seriously. Extreme-value analysis, through generalized Pareto modeling and Hill tail-index estimation, reinforces this picture by showing that tail losses decay slowly and that adverse outcomes are more persistent than standard models would suggest.

 doi:10.20944/preprints202603.2360.v1

The robust regression analysis adds another layer of insight by showing that dependence within the listed real-estate universe is not uniform. Broad and diversified securities are influenced more strongly by sector-wide movements, whereas specialized funds display greater idiosyncratic behavior and wider residual variation. This distinction helps clarify why benchmark-based explanations are informative in some cases but incomplete in others. The option-implied analysis complements the historical evidence by bringing in a forward-looking perspective. The Lévy-based implied price and volatility surfaces display asymmetry, curvature, and short-maturity effects consistent with the same nonlinear risk patterns observed in realized returns, suggesting that market prices themselves reflect persistent tail risk and state-dependent volatility.

Taken together, the findings suggest that listed real-estate security portfolios are best evaluated through a framework that combines tail-sensitive optimization, robust dependence analysis, and model-implied market information. Such an approach is especially valuable when macro-financial conditions amplify downside asymmetry and volatility persistence. Future research can build on this framework by allowing portfolio weights to adjust dynamically through rolling or expanding estimation windows and by linking portfolio exposure more explicitly to changing tail risk, volatility regimes, and dependence structures. Extensions based on regime-switching models and option-implied signals may also help improve portfolio adjustments as market conditions evolve.

Appendix A. Real Estate Security Descriptions

Table A1. Listed real-estate securities included in the study, categorized by sector and size as of 30 January 2025. The table reflects the empirical universe used in the analysis, including one duplicated Bloomberg venue code for FREL.

Ticker (BBG)	Name	Inception Date	Current Market Cap (million)	Exchange/ Venue (Trading)
	Residential			
ARR	ARMOUR Residential REIT, Inc.	5/2/2008	1,374.2 USD	NYSE
IRES ID	Irish Residential Properties REIT plc	16/4/2014	565.3 EUR	Euronext Dublin
ITB	iShares U.S. Home Construction ETF	5/5/2006	2,140 USD	NYSE Arca
REZ	iShares Residential and Multisector Real Estate ETF	5/4/2007	795.2 USD	NASDAQ
	Commercial			
823 HK	Link Real Estate Investment Trust (Link REIT)	9/6/2005	108,460.7 HKD	HKEX
SERT SP	Sasseur Real Estate Investment Trust	30/10/2017	886 EUR	SGX
SUPR LN	Supermarket Income REIT plc	21/7/2017	1,040.6 GBP	LSE
VNQ	Vanguard Real Estate ETF	9/29/2004	33,890 USD	NYSE Arca
IYR	iShares U.S. Real Estate ETF	6/19/2000	3,760 USD	NYSE Arca
SCHH	Schwab U.S. REIT ETF	1/13/2011	7,990 USD	NYSE Arca
RWR	SPDR Dow Jones REIT ETF	4/27/2001	2,030 USD	NYSE Arca
FREL US	Fidelity MSCI Real Estate Index ETF	2/5/2015	1,060 USD	US Composite (listed NYSE Arca)
SRET	Global X SuperDividend REIT ETF	3/16/2015	183.4 USD	NASDAQ
FREL UB	Fidelity MSCI Real Estate Index ETF	2/5/2015	1,060 USD	Nasdaq BX venue (BBG: UB)
	Healthcare / Senior Living			
CHP-U CN	Chartwell Retirement Residences	14/11/2003	10,460.8 CAD	TSX
OHI	Omega Healthcare Investors, Inc.	31/3/1992	10,713.4 USD	NYSE
	Global / International			
VNQI	Vanguard Global ex-U.S. Real Estate ETF	11/1/2010	3,320 USD	NASDAQ
RWO	SPDR Dow Jones Global Real Estate ETF	5/13/2008	1,120 USD	NYSE Arca
IFGL	iShares International Developed Real Estate ETF	11/19/2007	91.7 USD	NASDAQ
RWX	SPDR Dow Jones International Real Estate ETF	12/15/2006	381.7 USD	NYSE Arca

Continued on next page

 doi:10.20944/preprints202603.2360.v1

Table A1. *Cont.*

Ticker (BBG)	Name	Inception Date	Current Market Cap (million)	Exchange/ Venue (Trading)
USRT	iShares Core U.S. REIT ETF	5/4/2007	2,990 USD	NYSE Arca
REET	iShares Global REIT ETF	7/10/2014	3,950 USD	NYSE Arca
	Mortgage / Other			
MORT	VanEck Mortgage REIT Income ETF	8/17/2011	302.1 USD	NYSE Arca
KBWY	Invesco KBW Premium Yield Equity REIT ETF	12/2/2010	223.6 USD	NASDAQ
RNP	Cohen & Steers REIT & Preferred and Income Fund	6/27/2003	1,040 USD	NYSE
	Industrial / Infrastructure			
INDS	Pacer Industrial Real Estate ETF	5/15/2018	141.3 USD	NYSE Arca
SRVR	Pacer Data & Infrastructure Real Estate ETF	5/16/2018	444.9 USD	NYSE Arca
	Diversified			
AOR	iShares Core 60/40 Balanced Allocation ETF	11/11/2008	2,410 USD	NYSE Arca
XLRE	Real Estate Select Sector SPDR Fund	10/8/2015	7,490 USD	NYSE Arca
	Specialized			
ARE	Alexandria Real Estate Equities, Inc.	27/3/1997	12,661.1 USD	NYSE

Appendix B. Arithmetic and Logarithmic Returns

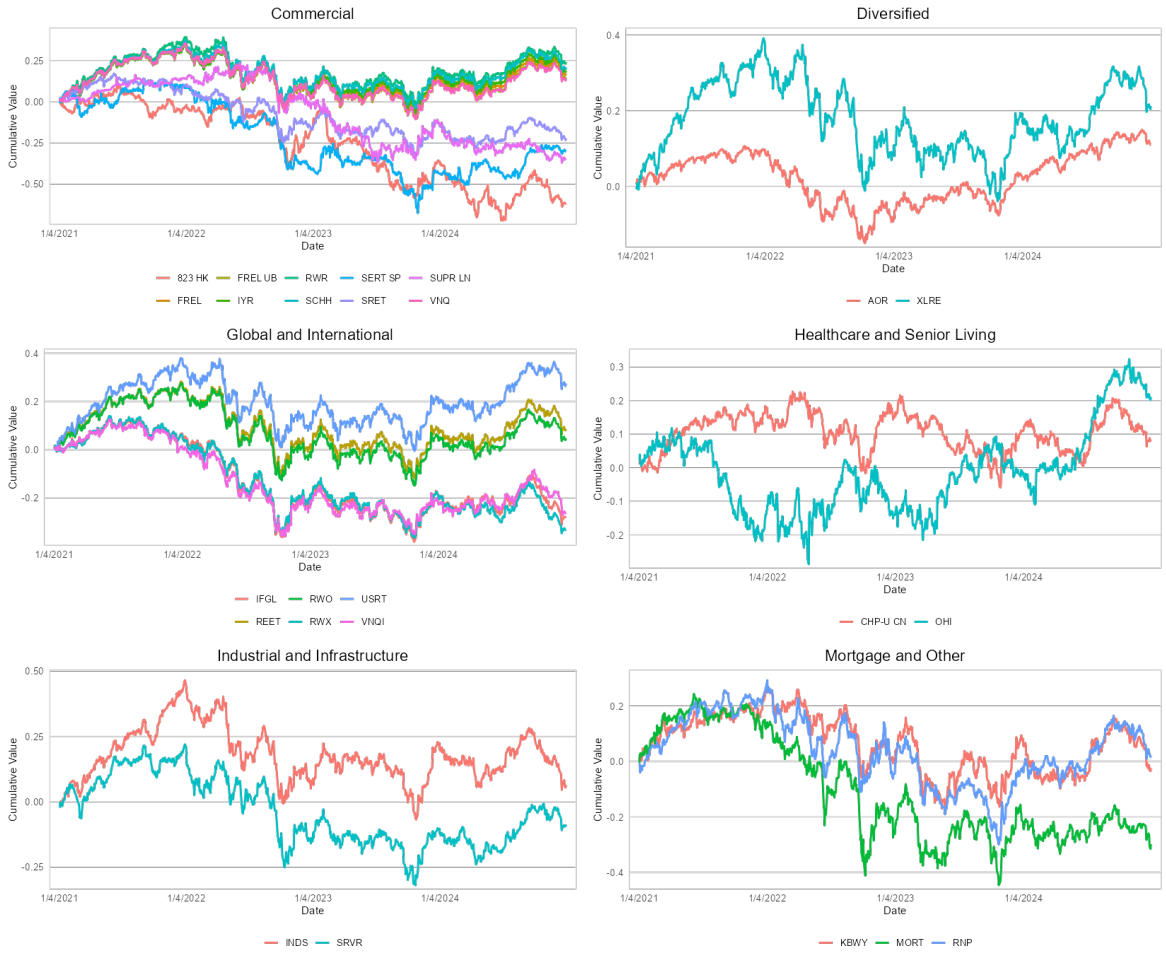

Figure A1. *Cont.*

 doi:10.20944/preprints202603.2360.v1

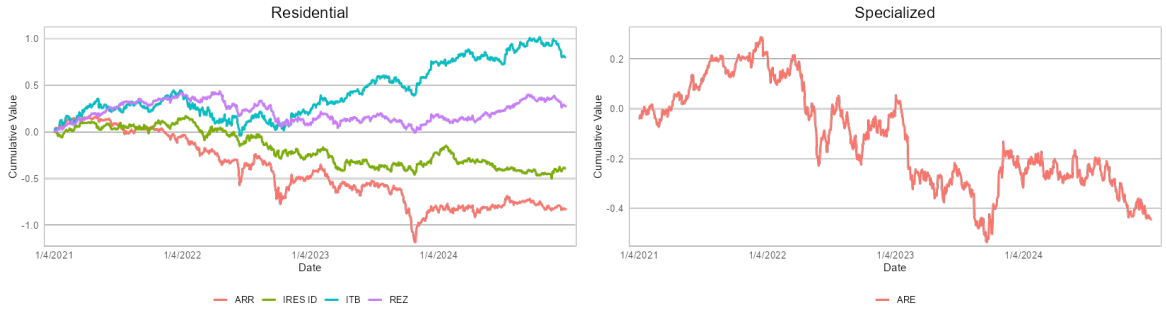

Figure A1. Cumulative arithmetic returns of each listed real-estate security under a buy-and-hold strategy, normalised to $100 on January 4, 2021.

To assess potential differences in return definitions, we compute the absolute divergence between arithmetic and logarithmic returns, defined as

$$|D_t| = |R_t - r_t|, \quad t = 1, 2, \ldots, T, \tag{A1}$$

where $R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$ and $r_t = \ln(P_t / P_{t-1})$. Figures A2 and A3 illustrate the time-series behavior of $|D_t|$ across the thirty listed real-estate securities over the period 2021-2024. The differences are generally negligible during stable market phases but widen notably during episodes of heightened volatility—particularly during the 2022 monetary tightening cycle and the mid-2023 repricing phase. This divergence underscores how nonlinear compounding effects become more pronounced when price fluctuations are large.

The comparison serves a diagnostic purpose and is not used directly in subsequent analyses. Throughout the study, arithmetic returns are employed for portfolio optimization and efficient-frontier construction, whereas logarithmic returns are used for stochastic modeling tasks such as volatility estimation and Lévy-based option pricing. The $|D_t|$ metric provides supporting evidence for the non-Gaussian and leptokurtic nature of listed real-estate security returns, motivating the use of CVaR-constrained optimization, Hill tail-index estimation, and ARFIMA–GARCH volatility modeling in the main body of the paper.

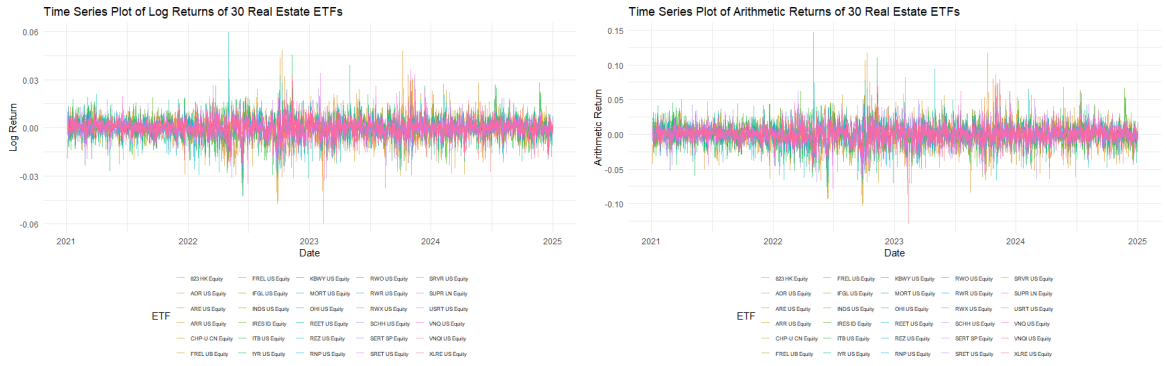

Figure A2. Time-series plots of (a) log returns and (b) arithmetic returns for each of the 30 listed real-estate securities over 2021-2024.

 doi:10.20944/preprints202603.2360.v1

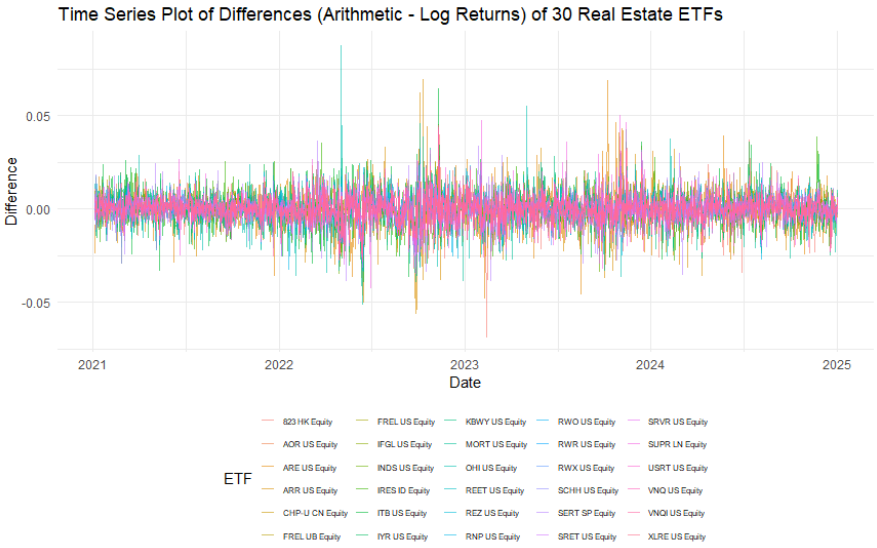

Figure A3. Time-series plot of the absolute difference between arithmetic and log returns.

Appendix C. Construction of Benchmark and Synthetic Portfolio Returns

Appendix C.1. Buy-and-Hold Benchmark Returns

The BHP is initialized by allocating the same fraction of wealth to each security at the portfolio formation date and then holding those positions without rebalancing. If $P_{i,0}$ denotes the initial price of asset i, and W_0 is the initial portfolio value, then equal initial allocation implies

$$n_i = \frac{W_0/N}{P_{i,0}}, \qquad i = 1, \ldots, N,$$

where n_i is the number of shares purchased in asset i. The benchmark portfolio value at time t is therefore

$$W_t^{\mathrm{BHP}} = \sum_{i=1}^{N} n_i P_{i,t}.$$

The corresponding simple return is computed as

$$R_t^{\mathrm{BHP}} = \frac{W_t^{\mathrm{BHP}} - W_{t-1}^{\mathrm{BHP}}}{W_{t-1}^{\mathrm{BHP}}}.$$

Because no rebalancing is performed after inception, the effective asset weights evolve endogenously through time as relative prices change. Hence, the BHP is equally weighted only at inception and may become increasingly concentrated in the assets that perform best over the sample period [44].

Appendix C.2. Synthetic Portfolio Return Series

For the historically optimized portfolios, return series are constructed using a rolling estimation window. Let τ denote the window length and let $r_{i,k}$ be the return on asset i at time k. For each trading day t, the optimization problem is solved using return observations from days $t-\tau, \ldots, t-1$, producing the portfolio weight vector

$$w(t) = \big(w_1(t), \ldots, w_n(t)\big)'.$$

The realized one-day portfolio return on day t is then

$$r_{p,t} = \sum_{i=1}^{n} w_i(t)\, r_{i,t}.$$

 doi:10.20944/preprints202603.2360.v1

To obtain the empirical return distribution associated with the portfolio chosen for day t, the fixed weights $w(t)$ are also applied to each historical return vector in the same estimation window, yielding the synthetic portfolio returns

$$r_{p,k}^{(t)} = \sum_{i=1}^{n} w_i(t)\, r_{i,k}, \qquad k = t-\tau, \ldots, t-1.$$

The set

$$\left\{ r_{p,k}^{(t)} : k = t-\tau, \ldots, t-1 \right\}$$

is interpreted as the in-window empirical distribution of the day-t optimized portfolio and is used in the historical tail-risk analysis. This rolling-window construction follows the historical optimization and synthetic-return methodology described in Lindquist et al. [45].

Appendix D. ARFIMA(1,d,1) Parameter Estimates for BHP Returns

To examine persistence in the conditional mean and variance of the BHP return series, we fitted an ARFIMA(1,d,1)–GARCH(1,1) model. The ARFIMA–GARCH(1,1) estimates suggest that the fractional integration parameter is effectively zero, indicating no meaningful long-memory behavior in the conditional mean of BHP returns. In the mean equation, the MA(1) coefficient is statistically significant under robust inference, while the AR(1) term is only marginally significant. In the variance equation, the large and highly significant β_1 estimate points to strong volatility persistence, whereas ω and α_1 are not statistically significant when robust standard errors are used.

Table A2. ARFIMA–GARCH(1,1) parameter estimates for BHP returns with robust standard errors.

Parameter	Estimate	Std. Error	t-value	p-value
μ	-0.000056	0.000440	-0.12645	0.899374
AR(1)	-0.312458	0.174585	-1.78972	0.073499
MA(1)	0.405633	0.151128	2.68403	0.007274
d (ARFIMA)	0.000000	0.045786	0.00000	1.000000
ω	0.000002	0.000003	0.55154	0.581261
α_1	0.028362	0.037867	0.74900	0.453859
β_1	0.946145	0.015643	60.48199	0.000000

Appendix E. Diagnostics of Option Pricing

Figure A4 compares the empirical return density (kernel estimate) with fitted parametric alternatives, including the Normal, Student-t, and NDIG specifications. The NDIG fit tracks the observed peak and tail behavior more closely than the Gaussian benchmark, while the Student-t provides an intermediate heavy-tailed approximation.

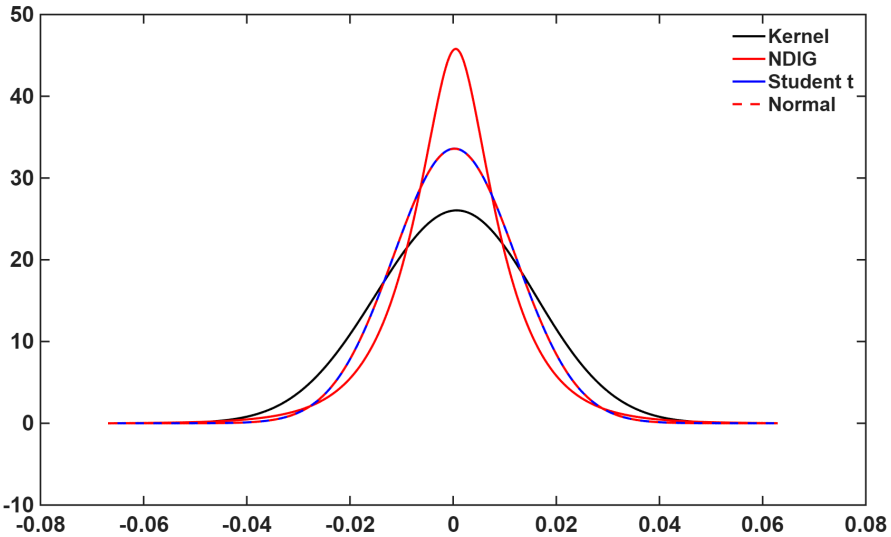

Figure A4. Kernel density of returns with fitted Normal, Student-t, and NDIG densities for comparison.

 doi:10.20944/preprints202603.2360.v1

References

1. Fabozzi, F.J., Ed. *Handbook of Finance: Financial Markets and Instruments*; Vol. 1, John Wiley & Sons, 2008.
2. Markowitz, H.M. Portfolio selection. *The Journal of Finance* **1952**, *7*, 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x.
3. Embrechts, P.; Klüppelberg, C.; Mikosch, T. *Modelling Extremal Events: For Insurance and Finance*; Springer, 2013.
4. Cont, R. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance* **2001**, *1*, 223–236. https://doi.org/10.1088/1469-7688/1/2/304.
5. McNeil, A.J.; Frey, R.; Embrechts, P. *Quantitative Risk Management: Concepts, Techniques and Tools (Revised Edition)*; Princeton University Press, 2015.
6. Rachev, S.; Ortobelli, S.; Stoyanov, S.; Fabozzi, F.J.; Biglova, A. Desirable properties of an ideal risk measure in portfolio theory. *International Journal of Theoretical and Applied Finance* **2008**, *11*, 19–54.
7. Gyourko, J.; Keim, D.B. What Does the Stock Market Tell Us About Real Estate Returns? *Real Estate Economics* **1992**, *20*, 457–485.
8. Eichholtz, P.M.A. Does International Diversification Work Better for Real Estate than for Stocks and Bonds? *Financial Analysts Journal* **1996**, *52*, 56–62.
9. Ling, D.C.; Naranjo, A. Economic Risk Factors and Commercial Real Estate Returns. *Journal of Real Estate Finance and Economics* **1997**, *14*, 283–307.
10. Ling, D.C.; Naranjo, A. The Integration of Commercial Real Estate Markets and Stock Markets. *Real Estate Economics* **1999**, *27*, 483–515.
11. Clayton, J.; MacKinnon, G. The Relative Importance of Stock, Bond and Real Estate Factors in Explaining REIT Returns. *Journal of Real Estate Finance and Economics* **2003**, *27*, 39–60.
12. Hoesli, M.; Oikarinen, E. Are REITs Real Estate? Evidence from International Sector Returns. *Journal of International Money and Finance* **2012**, *31*, 1823–1850.
13. Oikarinen, E.; Hoesli, M.; Serrano, C. The Long-Run Dynamics between Direct and Securitized Real Estate. *Journal of Real Estate Research* **2011**, *33*, 73–104.
14. Longin, F.; Solnik, B. Extreme correlation of international equity markets. *The Journal of Finance* **2001**, *56*, 649–676.
15. Asness, C.S.; Frazzini, A.; Pedersen, L.H. Leverage Aversion and Risk Parity. *Financial Analysts Journal* **2012**, *68*, 47–59.
16. Frazzini, A.; Pedersen, L.H. Betting Against Beta. *Journal of Financial Economics* **2014**, *111*, 1–25.
17. Grossman, S.J.; Vila, J.L. Optimal Dynamic Trading with Leverage Constraints. *Journal of Financial and Quantitative Analysis* **1992**, *27*, 151–168.
18. Grinold, R.C.; Kahn, R.N. *Active Portfolio Management*, 2 ed.; McGraw–Hill, 2000.
19. Madhavan, A. Exchange-Traded Funds, Market Structure, and the Flash Crash. *Financial Analysts Journal* **2012**, *68*, 20–35.
20. Ben-David, I.; Franzoni, F.; Moussawi, R. Do Exchange-Traded Funds Increase Volatility? *The Journal of Finance* **2018**, *73*, 2471–2535.
21. Leland, H.E. Who Should Buy Portfolio Insurance? *The Journal of Finance* **1980**, *35*, 581–594.
22. Leland, H.E.; Rubinstein, M. Replicating Portfolios and Investment Policies with Portfolio Insurance. *Journal of Financial Economics* **1981**, *10*, 383–423.
23. Black, F.; Jones, R. Simplifying Portfolio Insurance. *The Journal of Portfolio Management* **1987**, *13*, 48–51.
24. Black, F.; Perold, A.F. Theory of Constant Proportion Portfolio Insurance. *Journal of Economic Dynamics and Control* **1992**, *16*, 403–426.
25. Brandt, M.W. Portfolio Choice Problems. In *Handbook of Financial Econometrics*; Aït-Sahalia, Y.; Hansen, L.P., Eds.; Elsevier, 2010; Vol. 1, pp. 269–336.
26. Rockafellar, R.T.; Uryasev, S. Optimization of conditional value-at-risk. *Journal of Risk* **2000**, *2*, 21–41. https://doi.org/10.21314/JOR.2000.038.
27. Krokhmal, P.; Palmquist, J.; Uryasev, S. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk* **2002**, *4*, 11–27. https://doi.org/10.21314/JOR.2002.057.
28. Jorion, P. *Value at Risk: The New Benchmark for Managing Financial Risk*, 3rd ed.; McGraw–Hill, 2007.
29. Hill, B.M. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics* **1975**, *3*, 1163–1174.
30. Pickands, J. Statistical inference using extreme order statistics. *The Annals of Statistics* **1975**, *3*, 119–131.
31. Balkema, A.A.; De Haan, L. Limit laws for order statistics. *The Annals of Probability* **1974**, *2*, 792–804.

 doi:10.20944/preprints202603.2360.v1

32. De Haan, L.; Peng, L. Comparison of tail index estimators. *Statistica Neerlandica* **1998**, *52*, 60–70.
33. Haeusler, E.; Segers, J. Assessing confidence intervals for the tail index by Edgeworth expansions for the Hill estimator. *Statistics & Probability Letters* **2007**, *77*, 67–73.
34. Baillie, R.T. Long memory processes and fractional integration in econometrics. *Journal of Econometrics* **1996**, *73*, 5–59.
35. Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **1986**, *31*, 307–327.
36. Barndorff-Nielsen, O.E. Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics* **1997**, *24*, 1–13.
37. Carr, P.; Madan, D.B. Option Valuation Using the Fast Fourier Transform. *Journal of Computational Finance* **1999**, *2*, 61–73. https://doi.org/10.21314/JCF.1999.043.
38. Lee, R.W. The Moment Formula for Implied Volatility at Extreme Strikes. *Mathematical Finance* **2004**, *14*, 469–480.
39. Gatheral, J. *The Volatility Surface: A Practitioner's Guide*; John Wiley & Sons, 2006.
40. Ivanov, S.I. The Implied Volatility of ETF and Index Options. *Journal of Finance and Investment Analysis* **2011**, *1*, 1–20. If you cite the SSRN version, switch journal fields to note=SSRN working paper.
41. Rousseeuw, P.J.; Leroy, A.M. *Robust Regression and Outlier Detection*; John Wiley & Sons, 2003.
42. Yohai, V.J. High breakdown-point and high-efficiency robust estimates for regression. *The Annals of Statistics* **1987**, *15*, 642–656.
43. Maronna, R.A.; Martin, R.D.; Yohai, V.J.; Salibián-Barrera, M. *Robust Statistics: Theory and Methods (with R)*; John Wiley & Sons, 2019.
44. Forsyth, P.A. A Buy and Hold Portfolio Loses Diversification, 2024. White paper, University of Waterloo.
45. Lindquist, W.B.; Rachev, S.T.; Hu, Y.; Shirvani, A. *Advanced REIT Portfolio Optimization*; Springer, 2022.
46. de Sousa, B.; Michailidis, G. A Diagnostic Plot for Estimating the Tail Index of a Distribution. *Journal of Computational and Graphical Statistics* **2004**, *13*, 1–22. https://doi.org/10.1198/106186004X12335.
47. Shirvani, A.; Mittnik, S.; Lindquist, W.B.; Rachev, S. Bitcoin Volatility and Intrinsic Time Using Double-Subordinated Lévy Processes. *Risks* **2024**, *12*, 82. https://doi.org/10.3390/risks12080082.