# Multi-Model Monte Carlo Portfolio Optimization Project

Derivative Securities, Spring 2025

Beatrice Lee, Jessie Yao, Malcolm Hall, Rajath Sathyanarayana, Tim Phan

April 30, 2025

## 1 Introduction

Building a resilient portfolio requires more than just historical mean-variance analysis, as we have learned throughout the course. It demands a framework that captures the full variety of possible price behaviors and interest-rate environments. In this project, we calibrate four distinct price-dynamics models (Geometric Brownian Motion, Merton jump-diffusion, Constant Elasticity of Variance, and Heston stochastic volatility) to one-year historical return data. By simulating thousands of Monte Carlo paths under each model and generating stochastic rate scenarios (Vasicek, CIR, Ho-Lee), we create a comprehensive set of forward-looking return distributions and discount factors. Using these portfolios, we optimize portfolio allocations to maximize the average Sharpe ratio, first for an equity-only scenario and then extended to include plain-vanilla options. Finally, we examine how alternative interest-rate dynamics reshape optimal weights under both Sharpe-ratio and classic mean-variance criteria.

## 2 Data

We used a dataset composed of daily stock price data for a portfolio of ten publicly traded large-cap companies from 5 different sectors: technology, healthcare, finance, energy, and consumer goods. The portfolio includes Apple (AAPL), Nvidia (NVDA), UnitedHealth Group (UNH), Eli Lilly (LLY), JPMorgan Chase (JPM), Blackstone (BX), Costco Wholesale (COST), PepsiCo (PEP), ExxonMobil (XOM), and Enphase Energy (ENPH). We obtained historical closing prices for each stock from Yahoo Finance over a one-year period using the yfinance API. After gathering the raw price data, we calculated daily returns. This was the basis for estimating key statistical measures such as the annualized mean returns and the annualized covariance matrix of returns where we assumed 252 trading days per year. These measures are the foundational inputs for portfolio construction and optimization. The historical data and model parameters feed directly into the portfolio optimization framework. Our framework is modular overall, and in the future this same framework could be used to perform the same analysis on any year long historical return data. For the basis of the model, we used a risk free rate of 0.035 found from online sources to be approximately the risk free rate for the past year. To advance this project further, we could have derived this risk free rate from data. Given more time to work on the project, we could consider a larger portfolio or work on integrating ETFs into our work.

# 3 Techniques

## 3.1 Optimizing Our Initial Portfolio

We performed constrained optimization to select weights that maximize the Sharpe ratio using our previously computed annualized return vector and covariance matrix. Since most solvers minimize, we defined the objective function as the negative Sharpe ratio. We enforced two natural constraints that weights sum to one (full investment) and each weight lies between zero and one (no short sales) and then provided an equally weighted vector as our starting guess. Finally, we called scipy.optimize.minimize with Sequential Least Squares Programming (SLSQP) to find the optimal allocation, and then inverted the sign of the objective value to report the maximum Sharpe ratio achieved.

## 3.2 Simulating the Future Portfolio

After collecting our data as described in the data section, we incorporated select model-specific parameters to better capture the underlying dynamics of each stock. Before being able to simulate the Merton jump-diffusion model, we identified jumps as days where log returns deviated by more than two standard deviations from the mean. We then estimated the jump intensity $\lambda$ as the frequency of those events per year, and computed the average jump size $\mu_j$ and its volatility $\sigma_j$. For the Constant Elasticity of Variance (CEV) model, we regressed log(squared returns) against log(prices) in a one-month rolling window to estimate beta $\beta$, which is the elasticity exponent showing how volatility scales with price. For the Heston model, we set the long-run variance $\theta$ to the annualized variance of log returns, computed the volatility of variance $\xi$ as the standard deviation of a 21-day rolling variance (annualized), and estimated the return-variance correlation $\rho$ using the Pearson correlation between returns and changes in rolling variance. We kept the mean-reversion speed $\kappa$ fixed at 2.0 for stability. These model-specific calibrations allowed for more realistic simulations of future price dynamics.

After computing these parameters, we began to build our simulation framework consisting of four Monte Carlo simulator functions: Geometric Brownian Motion, Merton, CEV, and Heston. Each of these simulations accepts an initial price $S_0$, drift $\mu$, volatility $\sigma$, time horizon $T$, time-step $\Delta t$, number of steps $n$, and the model-specific parameters. Using these inputs, each function computes the number of steps $N$ and creates a NumPy array of shape $(n, N+1)$, setting the first column to $S_0$. Then each function updates the price differently at each step.

The Geometric Brownian Motion (GBM) simulator applies the classical Black–Scholes diffusion. At each time increment, it draws a standard normal shock for each price evolution according to:

$$S_{t+\Delta t} = S_t \exp\left((\mu - \tfrac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}Z\right)$$

This ensures log-normally distributed outcomes and constant proportional volatility.

The Merton jump-diffusion function builds off of GBM by adding random jumps. After drawing the same Gaussian shock $Z$, it simulates a Poisson count $N_j \sim \text{Pois}(\lambda\Delta t)$ of jumps in each interval. If $N_j > 0$, it generates $N_j$ independent Normal$(\mu_j, \sigma_j)$ jump sizes and sums them to obtain a total jump $J$. A compensator term $\kappa = \exp(\mu_j + \frac{1}{2}\sigma_j^2) - 1$ is subtracted from the drift so that the process remains martingale under the risk-neutral measure. The price now evolves according to:

$$S_{t+\Delta t} = S_t \exp\left((\mu - \tfrac{1}{2}\sigma^2 - \lambda\kappa)\Delta t + \sigma\sqrt{\Delta t}Z + J\right)$$

The CEV (Constant Elasticity of Variance) function departs from pure exponentiation and instead uses an additive Euler-type scheme where volatility scales as $S_t^\beta$. After drawing $Z$, it floors any previous price below a tiny $\varepsilon$ to avoid zero-divisions, computes vol_term $= \max(\sigma S_t^{\beta}, 0.05)$ to prevent extreme volatility, and then updates:

$$S_{t+\Delta t} = S_t + \mu S_t \Delta t + \text{vol\_term}\sqrt{\Delta t} Z$$

The Heston function tracks both the asset price and its instantaneous variance $V_t$. At each step it draws two correlated normals $Z_1, Z_2$ with correlation $\rho$. Variance is advanced via a full-function Euler discretization:

$$V_{t+\Delta t} = V_t + \kappa(\theta - V_t)\Delta t + \xi\sqrt{\max(V_t, 0)}\sqrt{\Delta t} Z_1$$

which ensures nonnegative variance. The price then follows:

$$S_{t+\Delta t} = S_t \exp\left((\mu - \tfrac{1}{2}V_t)\Delta t + \sqrt{\max(V_t, 0)}\sqrt{\Delta t} Z_2\right)$$

These four functions provide different levels of model complexity from the simplest constant-volatility diffusion to full stochastic-volatility plus jumps which enable and constrain the dynamics with different assumptions about price dynamics affected at the microstructure level.

## 3.3 Optimizing the Future Portfolio

We started the optimization process by initializing our four modeling scenarios (GBM, Merton jump-diffusion, CEV, and Heston) over a one-year horizon with daily steps and 5,000 Monte Carlo paths. For each ticker and each model, we call the corresponding simulator to generate price paths and record the terminal returns into a dictionary. We then set up an optimization to maximize the average Sharpe ratio across all four models. Our objective function first imposes a heavy penalty if any weight is negative, then, for each model, computes the portfolio return series as the dot product of simulated returns and the weight vector, calculates that model's Sharpe ratio, and finally returns the negative of the mean Sharpe across models. We constrain the weights to sum to one (full investment) and lie between zero and one (no shorting), and we start from an equal-weight guess. Using SciPy's SLSQP solver, we find the optimal allocation that balances performance under all price dynamics.

## 3.4 Stochastic Interest Rate Models

To compare the uncertainty in future discount rates, we simulated three classic short-rate processes over the same year using a daily-step grid ($\Delta t = 1/252$):

- **Vasicek:** Allowing rates to mean-revert but briefly go negative.
- **CIR:** Enforcing nonnegativity via the root diffusion term.
- **Ho-Lee:** A Gaussian model with constant drift alpha.

We set $r_0$ equal to 3%, $\kappa$ equal to 5, $\theta$ equal to 4%, $\sigma$ equal to 10%, and $\alpha$ equal to 1% (for Ho-Lee). For each path, we compute a one-year discount factor:

$$D_t = \exp\left(-\sum_{i=1}^{252} r_i^{(t)}\Delta t\right)$$

Then invert it to obtain the pathwise continuously compounded rate $r_t = -\ln(D_t)$. These rate scenarios were then used for pricing our options and for setting the risk-free anchor in our portfolio optimizations.

### 3.5 Optimizing the Future Portfolio with Stochastic Interest Rates

To optimize our future portfolio with stochastic interest rates, we re-optimized our portfolio using each rate model to see how the bond environment alters the optimal allocation. In doing this, we chose to also compare our current Sharpe-ratio-based model with a Markowitz mean-variance minimization strategy.

In our Sharpe maximization strategy, we treated each model's vector as the pathwise risk free rates and maximized the average Sharpe ratio:

$$\max_{w \in \Delta_n} \frac{1}{5000} \sum_{i=1}^{5000} \frac{R_{p,i} - r_i}{\sigma_{p,i}}$$

where $R_{p,i}$ is the simulated portfolio return on path $i$, and $\sigma_{p,i}$ is its overall volatility. This produced three portfolios with one per rate-model in order to view how a steeper mean-reversion (CIR) versus unbounded negatives (Vasicek) shifts weight toward or away from duration sensitive stocks and hedges.

In our classical Markowitz mean-variance strategy, we distilled each model's $r_i$ into a constant r-bar equal to the expected value of $r_i$, subtracted this from our average equity return vector, and then solved:

$$\min_{w \in \Delta_n} w^\top \Sigma w \quad \text{s.t.} \quad w^\top \mu \geq \bar{R}_{target}$$

This yielded three portfolios that balance mean excess return against variance under each rate-model's implied discount.

### 3.6 Adding Options to Improve Sharpe Ratio

To further improve upon Sharpe ratio for our current equity-only portfolio, we introduced plain-vanilla at-the-money European calls and puts for each underlying stock. Each option was given an expiry of one year and priced using Black-Scholes formula by plugging in the average implied short-rate from our simulations as the constant, the current stock price as the strike, and each stock's annualized volatility estimated from historical returns. For each Monte Carlo path $i$ and stock $j$, we then computed the option's pathwise return:

$$R_{i,j}^{\text{opt}} = \frac{\text{Payoff}_{i,j} - \text{Premium}_j}{\text{Premium}_j}$$

where the payoff is $\max(K - S_T, 0)$ for puts and $\max(S_T - K, 0)$ for calls.

This broadens our portfolio's possible outcomes by tripling the available financial instruments. Upon re-running our Sharpe maximization over this set, we saw an increase across the three different stochastic interest rate models by using hedging puts to further decrease risk without sacrificing upside.

# 4 Results

## 4.1 Optimizing Our Initial Portfolio

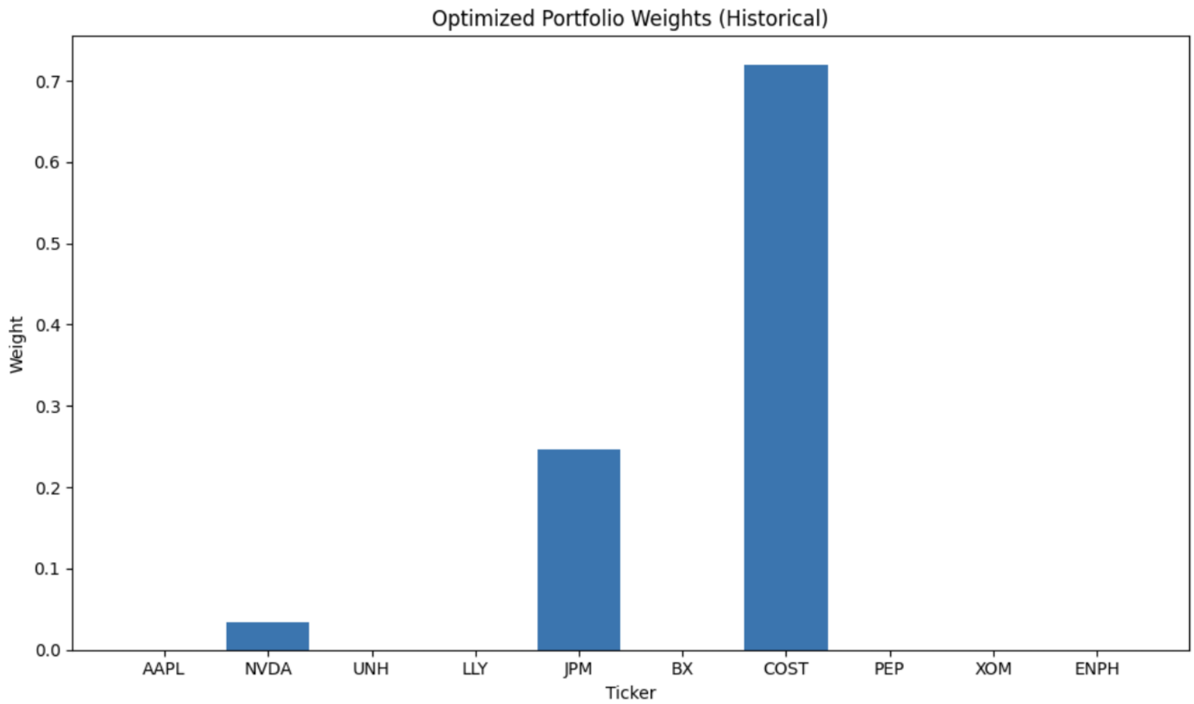


Optimization drives the portfolio to a very concentrated allocation with roughly 71.9% of capital placed in Costco (COST), 24.7% in JPMorgan (JPM), and a small 3.4% in Nvidia (NVDA), with all other positions set to zero. This reflects those three names' combination of high expected return and relatively low covariance with the rest of the stocks over the past year. The bar chart shows this allocation with Costco's bar above the others, followed by JPM and NVDA. The portfolio achieves an annualized Sharpe ratio of approximately 1.47, indicating an efficient risk-adjusted return by historical standards.

## 4.2 Optimizing the Future Portfolio using GBM, Merton, CEV and Heston

### Geometric Brownian Motion

Under GBM, the portfolio becomes concentrated in six stocks, with 38.8% of capital allocated to Costco (COST), 21.3% to JPMorgan (JPM), 15.4% to Apple (AAPL), 12.6% to Eli Lilly (LLY), 6.7% to Blackstone (BX), and 5.2% to Nvidia (NVDA). The remaining four stocks (UNH, PEP, XOM, ENPH) receive zero weight. This mix reflects each asset's trade-off of expected return versus volatility under the assumed price dynamics, with Costco's combination of stable cash flows and low correlation driving its dominant allocation. The resulting portfolio achieves an annualized Sharpe ratio of approximately 1.70. Below the two figures show the optimized weights for the portfolio and the return distribution after running 5,000 trials.

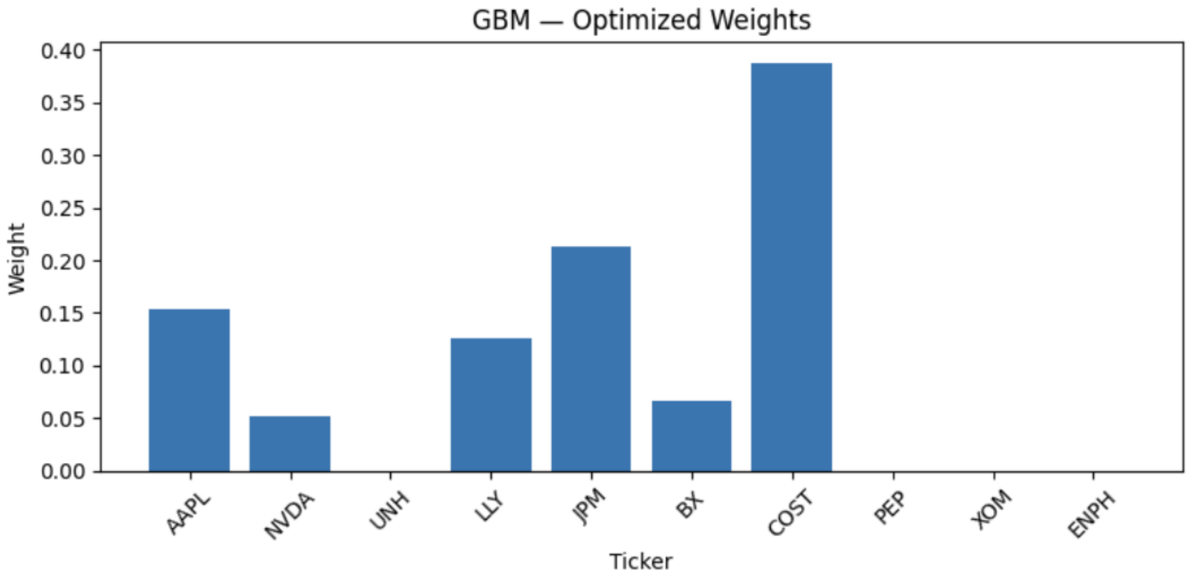


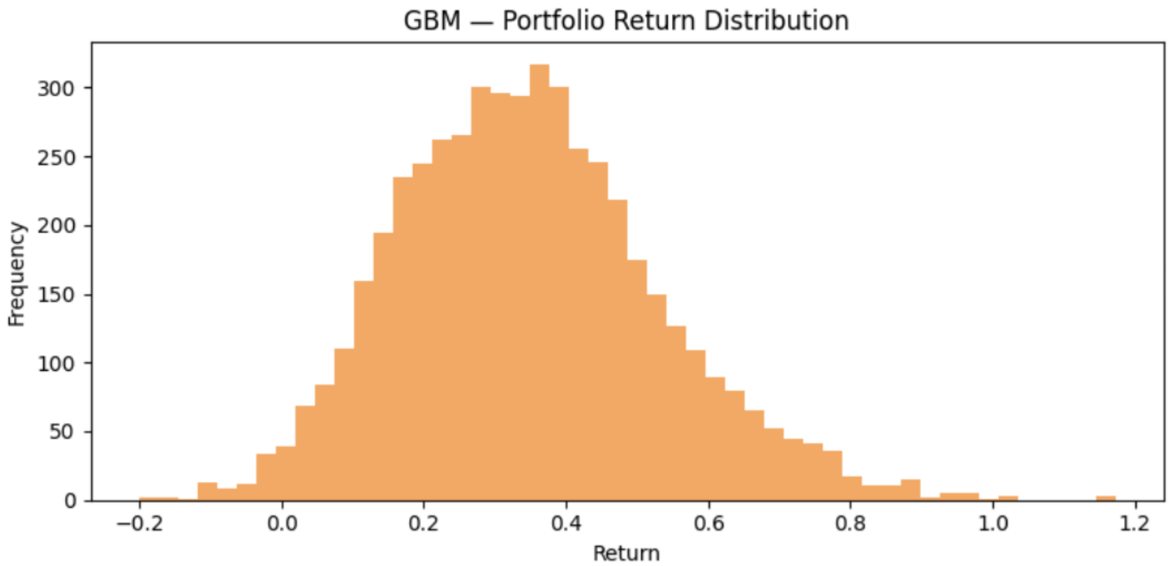


### Merton

Under the Merton jump-diffusion assumptions, the optimizer again concentrates the portfolio in five stocks, with 52.8% of capital allocated to Costco (COST), 14.6% to Apple (AAPL), 13.6% to Eli Lilly (LLY), 10.6% to JPMorgan (JPM), and 8.5% to Nvidia (NVDA), with all remaining stocks receiving zero weight. The achievable annualized Sharpe ratio drops to approximately 0.75 under the Merton dynamics. This reflects the increased downside risk from jump events and the heavier tail behavior of asset returns relative to simpler models. Below the two figures show the optimized weights for the portfolio and the return distribution after running 5,000 trials.

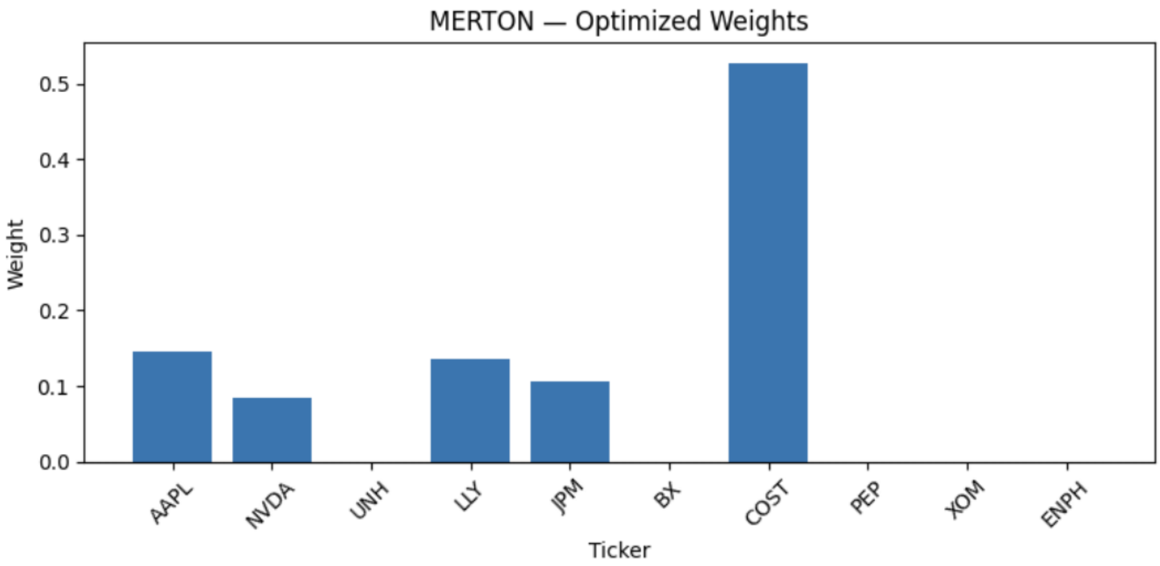


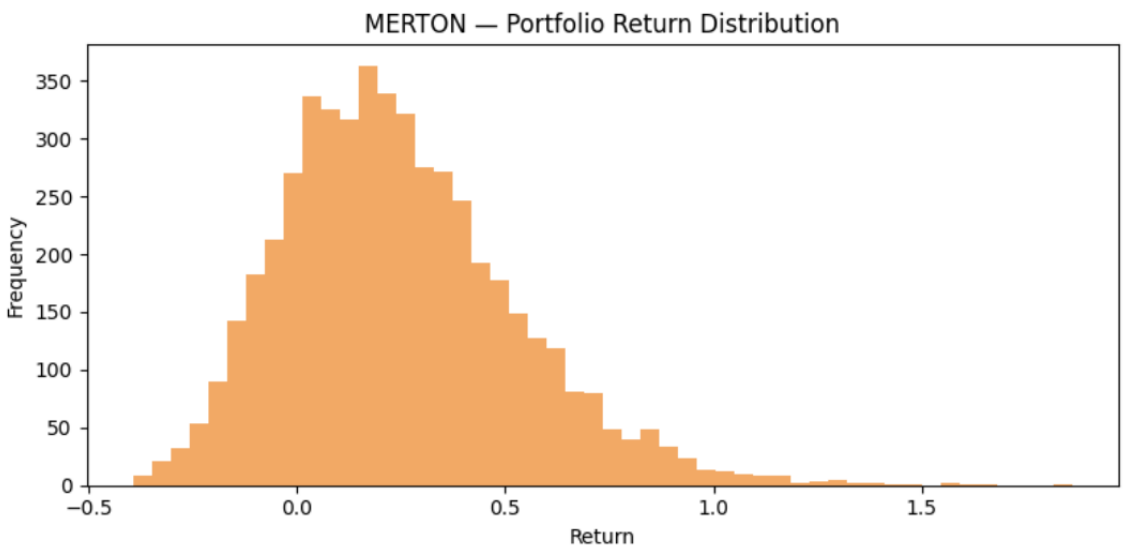


### Constant Elasticity of Variance (CEV)

Under the CEV dynamics, the optimizer places 98.6% of the portfolio in Costco (COST), with only negligible stakes in Apple, Nvidia, Eli Lilly, and JPMorgan, and zero allocation to the remaining names. This extreme concentration reflects Costco's status as the lowest-risk asset under the CEV specification, where volatility scales with the stock price raised to a power $\beta$. Because our implementation also imposes a minimum volatility "floor," the optimizer interprets Costco as having nearly deterministic returns. As a result, the annualized Sharpe ratio skyrockets to an implausible value of approximately 272.6, signaling that the model's numerical settings have artificially crushed perceived risk and created a degenerate optimization outcome.

This behavior highlights a clear flaw in our CEV implementation. Under the calibrated model, volatility follows $\sigma_{\text{CEV}}(S_t) = \sigma S_t^{\beta}$, which—when $\beta$ is compressed due to limited data and noisy regressions—causes the volatility of high-priced stocks like Costco to collapse toward zero. The optimizer then allocates nearly the entire portfolio to this "risk-free" asset, inflating the Sharpe ratio unrealistically. Although we enforced a small volatility floor, it was insufficient to restore realistic risk levels for high-priced assets.

In future work, our team should consider estimating $\beta$ over longer time windows, applying tighter bounds or priors on its value, or exploring alternative discretization schemes like log-CEV to better preserve randomness. Additionally, imposing a maximum weight constraint could prevent the optimizer from over-concentrating in any single asset. Below, the two figures show the optimized weights for the portfolio and the return distribution after running 5,000 trials.

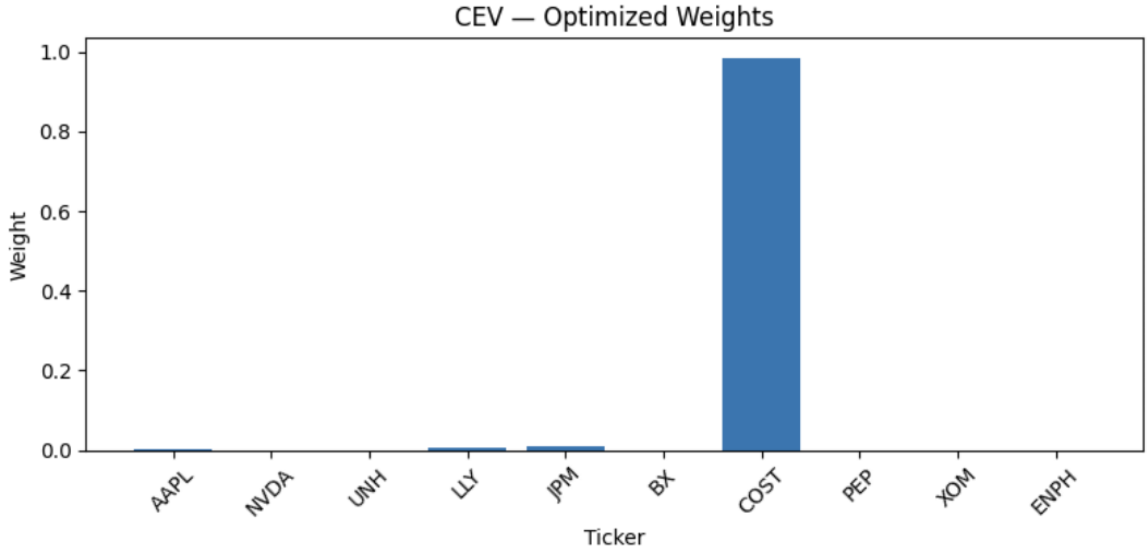


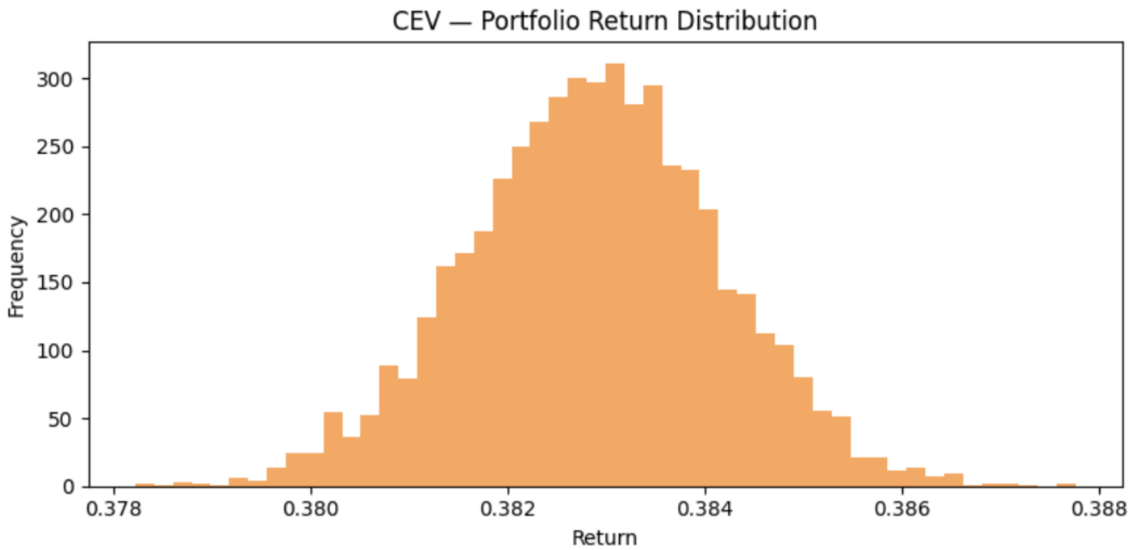


### Heston

Under the Heston stochastic-volatility specification, the optimized portfolio again has a large percent in Costco at 38.5 %, followed by JPMorgan at 21.3 %, Apple at 15.6 %, Eli Lilly at 11.6 %, Blackstone at 7.8 %, and a small 5.2 % allocation to Nvidia. The other four stocks are excluded. This blend reflects how the Heston model's mean-reverting variance and correlation structure value relatively stable, low-covariance equities while still capturing upside from growth names. The resulting annualized Sharpe ratio of approximately 1.745 indicates that, even once volatility itself is random and path-dependent, this mix delivers strong risk-adjusted returns in our simulated scenarios. Below the two figures show the optimized weights for the portfolio and the return distribution after running 5,000 trials.

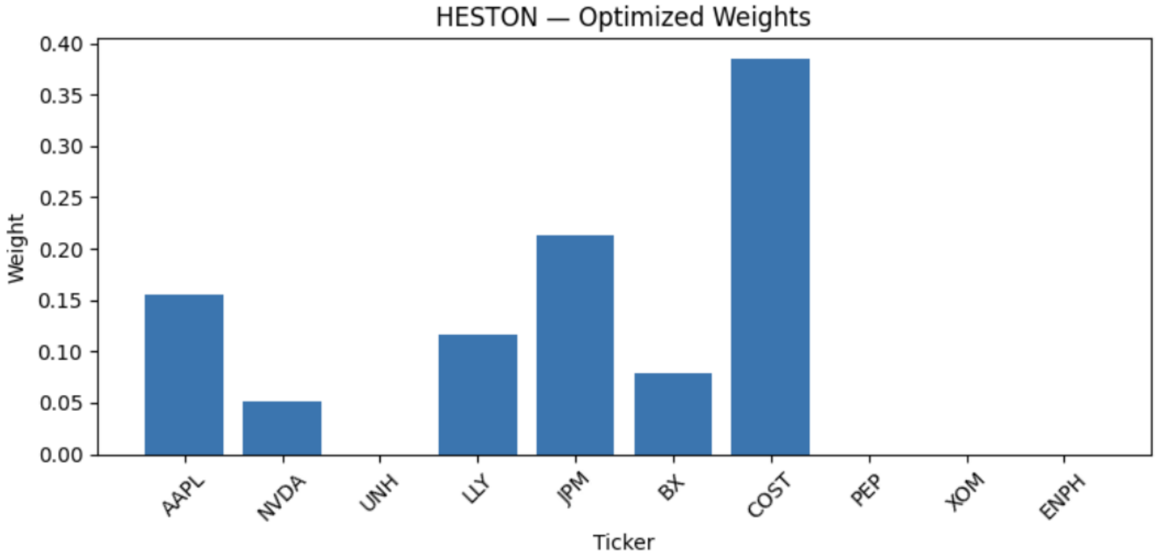


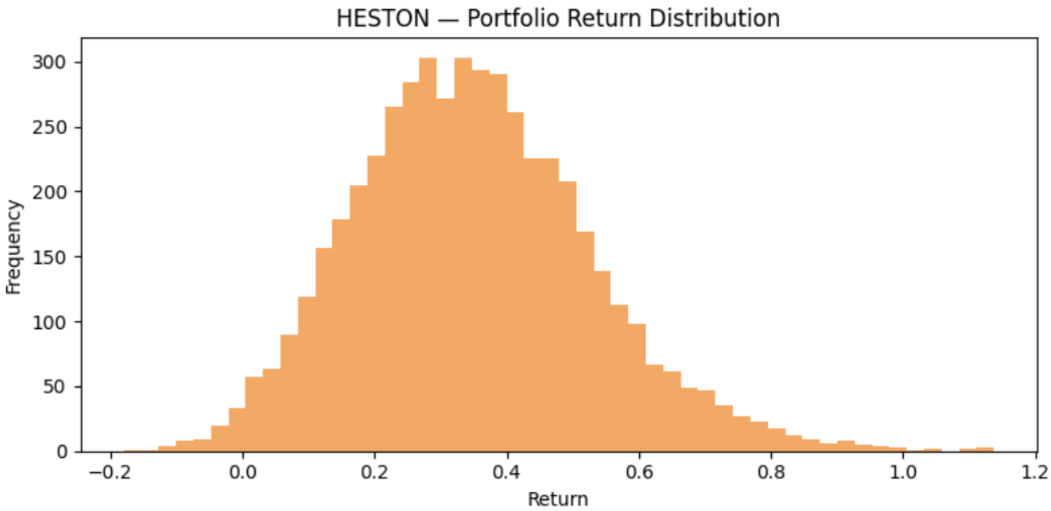


## 4.3 Portfolio Optimization Across GBM, Merton, and Heston Models (excluding CEV)

From these initial four models, we see that they all indicate selecting similar assets in the portfolio. Simple diffusions (GBM, Heston) deliver comparable Sharpe ratios above 1.7, jump-diffusion penalizes tail risk, and the naïve CEV setup collapses to a degenerate solution. After seeing the poor results with CEV we focused on GBM, Merton jump-diffusion, and Heston models for further optimization. After minimizing the negative of the average Sharpe across these three scenarios, we obtained the following allocations: Costco (COST) 41.3%, JPMorgan (JPM) 19.7%, Apple (AAPL) 15.5%, Eli Lilly (LLY) 11.9%, Blackstone (BX) 5.9%, with small residual stakes in Nvidia (NVDA) 5.6% and zeros elsewhere. When we plug these weights back into each model's simulations, the resulting annualized Sharpe ratios are 1.693 under GBM, 0.708 under Merton (reflecting the drag from jump risk), and 1.720 under Heston's stochastic-volatility paths. Below the two figures show the optimized weights for the portfolio and the return distribution for the three models after running 5,000 trials.

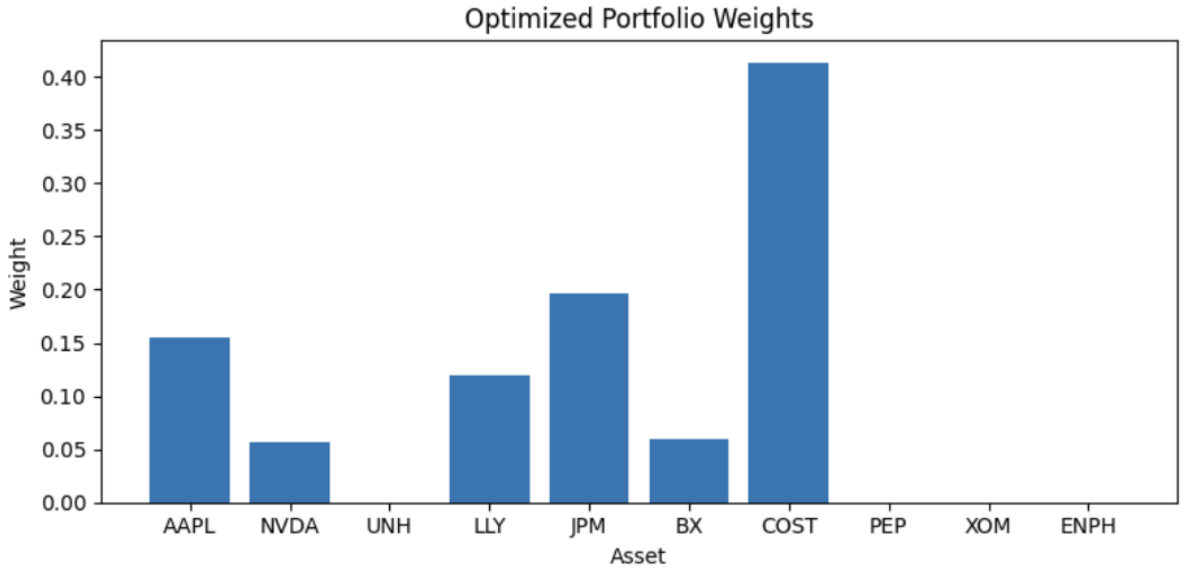
Optimized Portfolio Weights
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Weight
AAPL
NVDA
UNH
LLY
JPM
BX
COST
PEP
XOM
ENPH
Asset

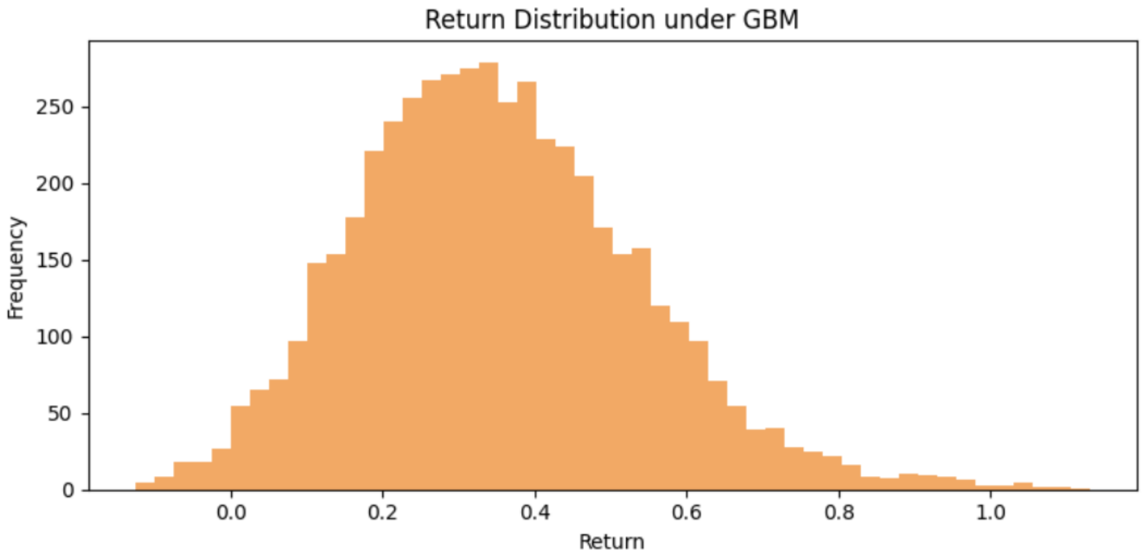
Return Distribution under GBM
250
200
150
100
50
0
Frequency
0.0
0.2
0.4
0.6
0.8
1.0
Return

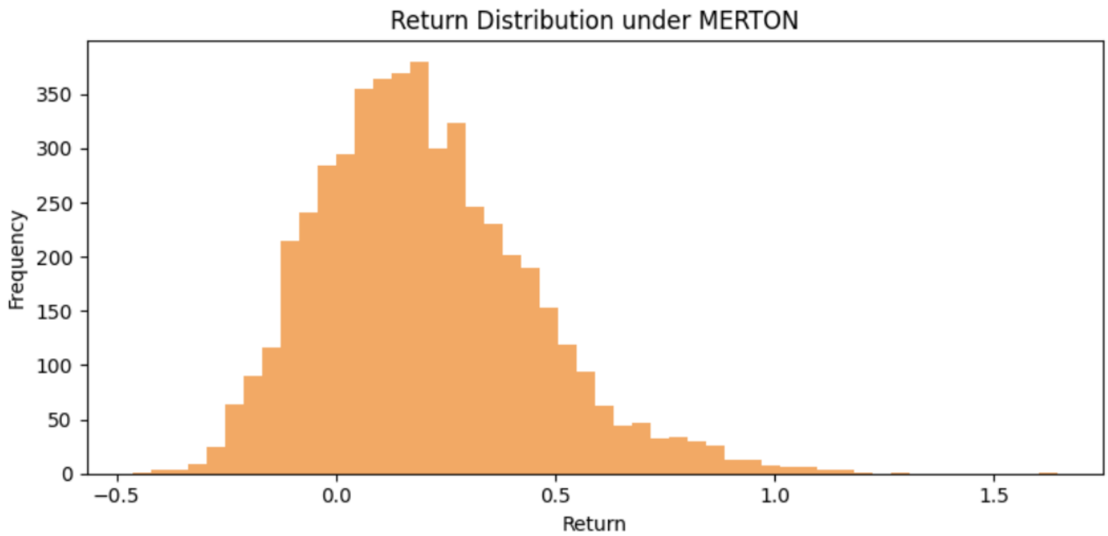
Return Distribution under MERTON
350
300
250
200
150
100
50
0
Frequency
−0.5
0.0
0.5
1.0
1.5
Return

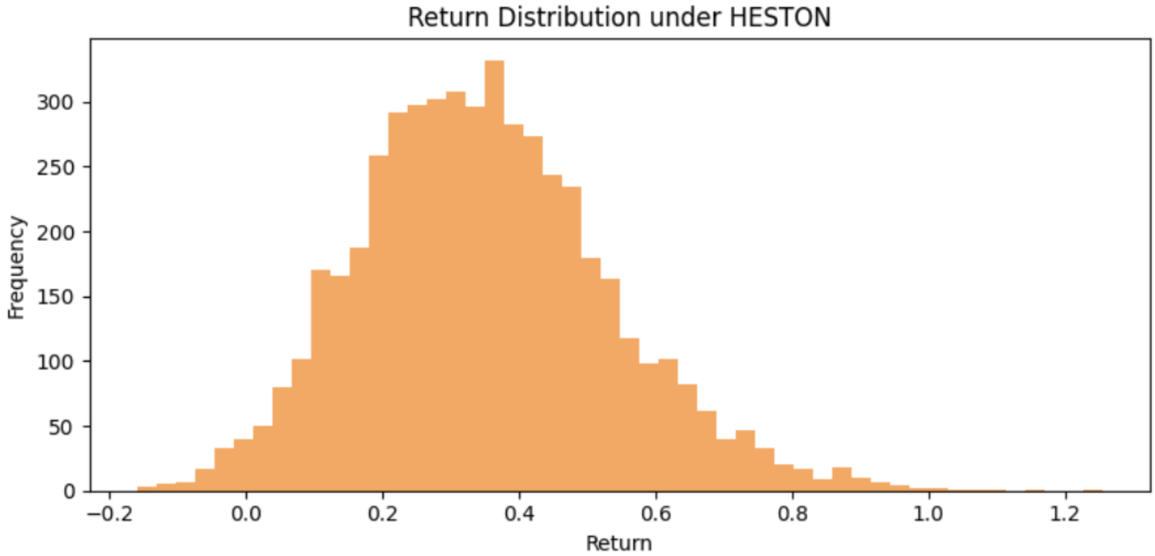


## 4.4 Optimizing the Future Portfolio with Stochastic Interest Rates

Under the Vasicek framework, which permits short-term rates to dip below zero, the portfolio tilts toward defensive stocks less sensitive to discount-rate swings. Costco commands 41.00 % of capital, followed by 21.15 % in JPMorgan, 14.33 % in Apple, 11.55 % in Eli Lilly, 6.45 % in Blackstone, and 5.52 % in Nvidia, with the remainder at zero. This allocation yields annualized Sharpe ratios of 1.679 under GBM, 0.656 under Merton jump-diffusion, and 1.678 under Heston stochastic volatility, reflecting how negative-rate scenarios modestly boost bond-like equity values while still rewarding core growth names.

When interest rates follow the CIR process which enforces nonnegative yields and stronger mean reversion, the optimizer favors slightly higher equity exposure to offset generally higher discount factors. The resulting weights are nearly identical: 40.99 % in Costco, 21.14 % in JPMorgan, 14.34 % in Apple, 11.54 % in Eli Lilly, 6.45 % in Blackstone, and 5.53 % in Nvidia. Sharpe ratios under each model were 1.676 (GBM), 0.654 (Merton), and 1.675 (Heston), showing that the different interest rate model did not alter the portfolio's core tilt.

Under the Ho–Lee arithmetic short-rate model, which imposes a constant drift on rates, the optimized allocation again centers on the same six names: 41.08 % Costco, 21.14 % JPMorgan, 14.32 % Apple, 11.53 % Eli Lilly, 6.38 % Blackstone, and 5.55 % Nvidia and zeros elsewhere. The corresponding Sharpe ratios are 1.664 for GBM, 0.645 for Merton, and 1.662 for Heston. This middle-ground rate path produces a blend that sits between the defensive bias of Vasicek and the slightly more aggressive tilt of CIR, but it still rewards the same resilient equities.

Across all three interest-rate scenarios, the optimal portfolio remains stable with a diversified mix of low-covariance, high-efficiency stocks, and the annualized Sharpe ratios shift by only a few hundredths. Rate uncertainty leads to only subtle rebalancing without completely changing the allocation. This consistency highlights how robust a multi-model approach can be, and it demonstrates that even as rates follow very different stochastic paths, the same core equity selection continues to maximize risk-adjusted returns.

After completing this baseline understanding of the interest rate models, we constructed two different portfolios to illustrate the contrast between a scenario-based Sharpe-ratio approach and the classical mean–variance optimization. In the first method, we treated each of our three calibrated price models—GBM, Merton jump-diffusion, and Heston stochastic volatility—as a separate market "scenario" and chose the weight vector that maximized the average Sharpe ratio across them. This process drove the allocation toward a handful of high-efficiency names: about 41.7 % in Costco, 19.8 % in JPMorgan, 15.1 % in Apple, 11.9 % in Eli Lilly, 5.9 % in

Nvidia, and 5.7 % in Blackstone, yielding an average Sharpe of 1.392.

However, the Markowitz (MVO) portfolio took the historical average return (approximately 7.50 %) as a binding constraint and minimized variance subject to meeting that target. The resulting allocation was far more diversified with no position in Apple or Blackstone, roughly 28.4 % in Costco, 27.4 % in PepsiCo, 15.1 % in ExxonMobil, 10.3 % in both JPMorgan and UnitedHealth, 6.4 % in Eli Lilly, and 2.1 % in Nvidia. This produced a volatility of 14.20 % and a much lower Sharpe of 0.303. When plotted side by side, we see a concentrated, high–Sharpe portfolio that bets heavily on the most resilient names across our simulations on the left and a broad, target-return portfolio that sacrifices risk-adjusted efficiency for diversification.

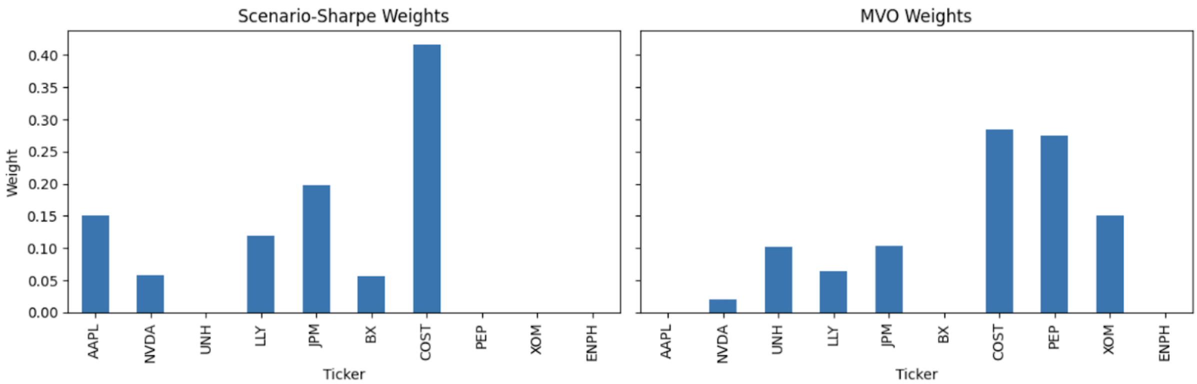


This comparison tells us that our stochastic models are doing exactly what they're designed to do. They capture dynamics like volatility clustering and mean reversion that the sample mean–variance framework overlooks. When we optimize across these models, we see portfolios that favor equities most resilient to each form of risk resulting in significantly higher theoretical Sharpe ratios. At the same time, the scenario–Sharpe allocation's concentration in just six stocks highlights a central trade-off. The richness in model structure can yield more "efficient" portfolios on paper, but it also amplifies model-risk. The classical MVO portfolio, though less efficient, naturally guards against any single model's misspecifications by spreading bets across a wider range of equities. These results suggest a balanced path forward for investors of leveraging these stochastic models to identify opportunities while acknowledging potential exposure to risk.

## 4.5 Optimizing the Future Portfolio with Options

To complete our project we added at-the-money European options to our core equities, then optimized weights by averaging across all six interest-rate scenarios (Vasicek, CIR, Ho–Lee) and the three price-dynamics models (GBM, Merton, Heston). This framework produced the following allocation: 31.01 % in Costco, 14.27 % in JPMorgan, 10.60 % in Apple, 9.61 % in Eli Lilly, 6.51 % in Blackstone, 3.97 % in Nvidia, combined with protective puts on UnitedHealth (3.97 %), Blackstone (1.08 %), PepsiCo (5.83 %), ExxonMobil (2.65 %), and Enphase (10.51 %). All call positions remained at zero.

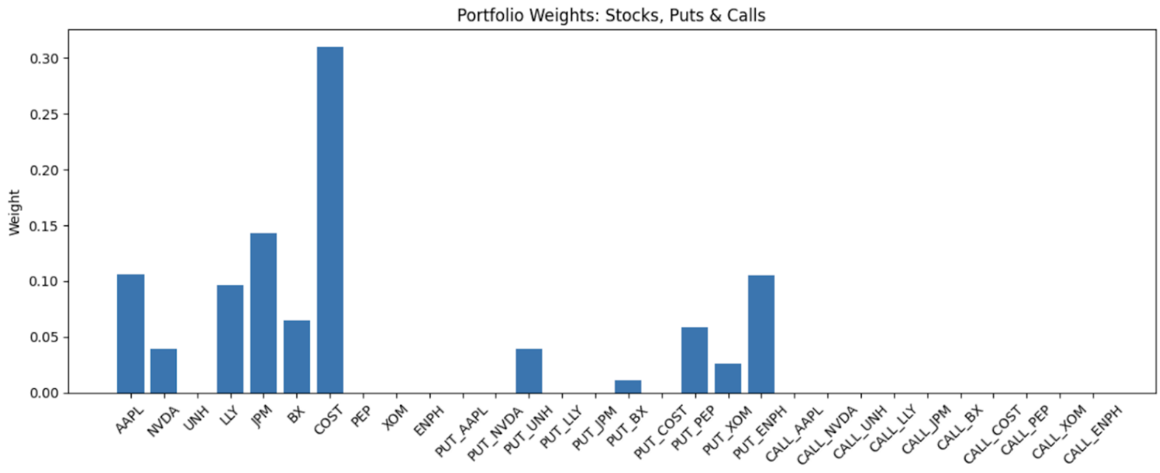


The resulting Scenario Sharpe ratio jumps to 2.3309 (This is the average over all models). Model-by-model, the portfolio achieves annualized Sharpe ratios of 2.215 under GBM, 2.566 under Merton's jump-diffusion, and 2.212 under Heston. The higher Sharpe in the Merton case shows how the put insulates against heavy-tailed shocks, while the strong performance under GBM and Heston confirms the overlay's accretive effect even in smoother markets. This shows that targeted allocation to plain-vanilla puts can substantially enhance risk-adjusted returns and bolster portfolio resilience. The portfolio's risk was more evenly spread across a mix of stocks and options, enhancing diversification. Extending this to other scenarios, since our portfolio was relatively small, we can hypothesize that calls can also add an advantage under other circumstances. In the future, we could build upon our project's framework to include a wider variety of options.

# 5 Conclusion

Throughout this project we demonstrated the value of a multi-model Monte Carlo approach to portfolio construction. When driven by simple Geometric Brownian Motion and by the Heston stochastic-volatility model, our optimizations produced similar, well-diversified allocations across Costco, JPMorgan, Apple, Eli Lilly, Blackstone, and Nvidia, each achieving an annualized Sharpe ratio around 1.7. Introducing Merton's jump-diffusion dynamics with its heavier tails and sudden shock risk drove allocations even more toward low-correlation stocks like Costco and JPMorgan, but lowered the attainable Sharpe to roughly 0.75. Our CEV implementation exposed how sensitive the results are to parameter calibration and discretization as our implementation led to an implausible Sharpe above 200.

By comparing these four price-dynamics side by side, we see that both diffusion-only and stochastic-volatility models reward steady, low-covariance equities. By optimizing our portfolio using the 3 effective models, the portfolio strikes a balance that performs well across both constant-volatility and stochastic-volatility worlds, while still accounting for jump-diffusion scenarios that will meaningfully erode risk-adjusted returns. Adding Vasicek, CIR, and Ho–Lee short-rate generators into both our discounting and objective functions revealed how rate-model choice subtly shifts optimal weights. It showed us how robustness to rate uncertainty can be integrated into portfolio design. Extending our framework to include plain-vanilla options further improved Sharpe ratios, proving intuition learned in class that utilizing options gives an investor significant advantages.

We think the Heston stochastic-volatility framework and the Merton jump-diffusion model are the two most compelling tools for capturing the dynamics we anticipate in future markets. Heston's model generates time-varying volatility with mean-reverting behavior, reflecting true market dynamics. Merton's jump-diffusion component explicitly accounts for crash risk, ensuring that rare but impactful events like policy surprises or earnings gaps are integrated into simulations and portfolio sizing rather than being systematically understated. On the rate side, the CIR short-rate model provides a realistic, positive-only, mean-reverting description of developed-market policy rates. Taken together, we think that this model best accounts for future markets by blending these approaches by modeling asset prices with stochastic volatility and jump risk while discounting cash flows under a CIR process.

We learned in this project that building an effective investment process will blend insights across these different models. Doing so, investors can move beyond one-shot mean-variance analysis and instead build portfolios that are not only optimized for past data but also adaptable to a broad spectrum of future market environments. This project also clearly displayed the added value of even the simplest options in driving value for investors.

# References


1. Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, *81*(3), 637–654.

2. Merton, R. C. (1976). Option Pricing When Underlying Stock Returns Are Discontinuous. *Journal of Financial Economics*, *3*(1–2), 125–144.

3. Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A Theory of the Term Structure of Interest Rates. *Econometrica*, *53*(2), 385–407.

4. Vasicek, O. (1977). An Equilibrium Characterization of the Term Structure. *Journal of Financial Economics*, *5*(2), 177–188.

5. Ho, T. S. Y., & Lee, S. B. (1986). Term Structure Movements and Pricing Interest Rate Contingent Claims. *Journal of Finance*, *41*(5), 1011–1029.

6. Heston, S. L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, *6*(2), 327–343.

7. Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.

8. Yahoo Finance. (n.d.). Historical market data retrieved via the *yfinance* API.