Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Mathematical Framework for Dynamic Downside-Minimizing Portfolio with Exotic Option Overlay

Antarik Bansal

rikbansal02@gmail.com

BASIS Chandler https://orcid.org/0009-0007-0664-4025

Research Article

Keywords: Downside risk, semivariance optimization, portfolio optimization, exotic options, Monte Carlo simulation, dynamic rebalancing

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11613518/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Mathematical Framework for Dynamic Downside-Minimizing Portfolio with Exotic Option Overlay

Antarik Bansal[1*]

[1*]Independent Researcher, BASIS Chandler, 4825 S Arizona Ave, Chandler, 85286, AZ, USA.

Corresponding author(s). E-mail(s): rikbansal02@gmail.com;

Abstract

This study develops and evaluates a dynamic downside-minimizing portfolio through five stages. Starting with a 4-asset universe of equity, bonds, gold, and a volatility instrument, the portfolio is expanded to 13 and then 21 assets while incorporating mean-semivariance optimization, Asian arithmetic-average puts, American puts valued with Longstaff–Schwartz Least-Squares Monte Carlo, covered calls, momentum tilts, and regime-aware rebalancing. The final stage targets a 15% expected return. Portfolio performance is evaluated using correlated Geometric Brownian Motion with 20,000 Monte Carlo paths over a one-year horizon. The four-asset semivariance-optimized portfolio has a Sortino ratio of 2.04 and a maximum simulated loss of **−23.7%**, compared with 1.20 and **−45.9%** for SPY. Adding Asian and American puts increases the Sortino ratio to 2.26 and reduces maximum simulated loss to **−8.8%**. The 13-asset buy-and-hold portfolio reaches a mean Sortino ratio of 5.81. Among the rebalancing strategies tested, calendar rebalancing produced the highest mean Sortino ratio of 9.62. The 21-asset portfolio with a 15% return target has a Sortino ratio of 4.06 and maximum simulated loss of **−28.8%**; the full option overlay reduces the latter to **−9.1%**. These results characterize the tradeoffs among diversification, downside protection, return targets, and adaptive rebalancing within the simulated framework.

Keywords: Downside risk, semivariance optimization, portfolio optimization, exotic options, Monte Carlo simulation, dynamic rebalancing

1 Introduction

The critical balance in a portfolio is between two conflicting goals: maximize expected return and minimize risk on the downside. The Markowitz mean-variance framework Markowitz (1952) tackles this dilemma by assuming that upside and downside volatility are symmetrical, which is odd in light of the observed behavior of investors and odd in light of a concave utility function and expected-utility theory Kahneman and Tversky (1979). Sortino and Price Sortino and Price (1994) suggest that downside semivariance is a better measure of risk, as it is zero when returns exceed a minimum acceptable return τ but is only penalized when returns fall below that threshold.

The full development of one such portfolio, featuring an overlay of exotic options and progressively expanded to attain a higher return level through downside semivariance minimization, is followed. To provide insight into each design decision, we document the decision as it was made, its mathematical justification, and the simulation results that validated or challenged it. This lineage-based approach is motivated by two considerations: first, portfolio construction is an iterative process in which each step constrains the next; second, this approach keeps a clean account of what is proved versus assumed in the model.

A summary of the development of the portfolio:

- **Stage 1 — 4-asset baseline:** Equity + bonds + gold + vol-instrument, buy and hold, no options. Sets up the diversification baseline and semivariance optimization framework.
- **Stage 2 — Asian and American put overlay:** Overlays an exotic option program on top of the 4-asset base. Demonstrates analytically that, under the stated payoff and exercise assumptions, the American put overlay establishes a lower bound on terminal portfolio value after accounting for the option premium, and introduces the path-averaging variance-reduction property of the Asian option Kemna and Vorst (1990).
- **Stage 3 — 13-asset optimized portfolio:** Adds 10 large-cap proxies to the equity sleeve. The larger asset set provides a greater number of return sources that are not highly correlated, resulting in a significant increase in the optimizer's Sortino ratio at the same return level.
- **Stage 4 — Enhanced EV strategies:** Four strategies are tested against the buy-and-hold base: covered calls (premium income); momentum factor tilts (21-day trigger, dynamic trailing-return weights); calendar rebalancing (21-day trigger, fixed optimizer weights — identical trigger to momentum tilt, differing only in target weights); and regime-aware rebalancing (Frobenius-norm-triggered, fixed optimizer weights). The 13-asset buy-and-hold base is the benchmark for each comparison.
- **Stage 5 — 21-asset 15% return target:** Expands to growth asset classes (SmallCap, EM, TechGrowth, Biotech, MicroCap, REIT, Commodities, HiYield) to gain access to higher expected returns, drawing on documented size and factor premia Fama and French (1993). The numerical frontier analysis indicates that a 15% expected return carries a minimum reported semiDD of 3.70% under the model constraints, illustrating the risk-return tradeoff implied by the model.

2 Mathematical Foundations

2.1 Geometric Brownian Motion

Asset prices are modeled with the classic continuous-time equity price process, Geometric Brownian Motion (GBM) Black and Scholes (1973). The SDE and its exact closed-form solution, obtained via Itô's lemma, are:

$$dS_{i,t} = \mu_i S_{i,t}\, dt + \sigma_i S_{i,t}\, dW_{i,t} \tag{1}$$

$$S_{i,t} = S_{i,0} \cdot \exp\left[\left(\mu_i - \tfrac{1}{2}\sigma_i^2\right) t + \sigma_i W_{i,t}\right] \tag{2}$$

The variance drag, $-\frac{1}{2}\sigma_i^2$, forms a wedge between the actual compound growth rate, $\mu_i - \frac{1}{2}\sigma_i^2$, and the arithmetic mean return μ_i. This means a stock with $\mu = 10\%$ and $\sigma = 20\%$ compounds at $10\% - 2\% = 8\%$, a difference that can materially affect portfolio return targets. For N correlated assets, we perform a Cholesky decomposition $L = \text{chol}(\Sigma)$ and obtain correlated Brownian increments $dW = LZ\sqrt{dt}$, where $Z \sim N(0, I)$. Because we do not use Euler–Maruyama discretization, there is no step-size bias. We review two additional risk metrics below.

2.2 Downside Semivariance and the Sortino Ratio

The downside semivariance of a portfolio with return $R(w)$ and minimum acceptable return τ is (following the downside-risk framework of Sortino and Price Sortino and Price (1994)):

$$\sigma_d^2(w) = E\left[\min(R(w) - \tau,\, 0)^2\right] \tag{3}$$

Downside semivariance quantifies shortfall below the target. The Sortino ratio Sortino and Price (1994) is used to quantify the quality of a given portfolio, trading off low downside deviation against expected return:

$$\text{Sortino} = \frac{E[R] - \tau}{\sigma_d(w)} \tag{4}$$

where $E[R]$ is the expected portfolio return, τ is the minimum acceptable return (MAR), and $\sigma_d(w)$ is the downside deviation (the square root of the downside semivariance).

A note on notation: τ in Eq. 3–4 (the MAR used to compute downside deviation and the Sortino ratio) and the "target return" in the optimizer's constraint (Eq. 5, Sec. 4.2 Eq. 10) are two distinct quantities, both of which this paper unfortunately denotes with language that can look interchangeable. In every reported result in this paper, $\tau = 0\%$ is the value actually used for the Sortino/semivariance calculation — this can be confirmed by back-solving $\tau = E[R] - \text{Sortino} \times \sigma_d$ from the mean return, downside deviation, and Sortino ratio reported in any row of Tables 3, 5, 6, 7, 8, or 13, all of which are consistent with $\tau \approx 0\%$ to within rounding. The "target return" (6.76% in Stages 1–4; 7–15% in the Stage 5 frontier of Table 10) is a separate return-floor constraint on the optimization itself (Eq. 5) and does not enter the Sortino ratio's denominator or numerator threshold.

2.3 The Constrained Optimization Problem

At each stage, portfolio weights are found by solving the mean-semivariance program following the formulation discussed by Estrada and related semivariance optimization work Estrada (2008); Ben Salah et al. (2018); Rigamonti and Lučivjanská (2024). Recent work by Ferrari et al. further addresses the numerical challenges of semivariance estimation and optimization, particularly for portfolios with many assets Ferrari et al. (2026):

$$\min_{w} \frac{1}{M} \sum_{m} \left[\min\Big(R^{(m)}(w) - \tau,\, 0 \Big) \right]^2 \quad \text{s.t.} \quad E[R(w)] \geq \text{target return}, \;\; \mathbf{1}^\top w = 1, \;\; w \geq 0 \tag{5}$$

The problem is solved using Sequential Least Squares Programming (SLSQP) with $N = 10$ random restarts, to avoid stopping at a local rather than global optimum. Because the benchmark is feasible, a true global minimizer would satisfy $\sigma_d^2(w^*) \leq \sigma_d^2(w_{\text{bench}})$. The numerical procedure used here seeks such a solution through SLSQP with multiple random restarts; convergence should therefore be interpreted as numerical evidence rather than a formal global-optimality proof.

3 Methods

The option pricing methodology includes Asian puts, American puts, and covered calls.

3.1 Asian Arithmetic-Average Puts

The Asian option payoff is $\max(K - \bar{A}, 0)$ where $\bar{A} = \frac{1}{n}\sum_i S_i$. The key property, established by Kemna and Vorst Kemna and Vorst (1990), is variance reduction via averaging:

$$\mathrm{Var}\left(\frac{1}{n}\sum_i X_i\right) = \frac{1}{n^2}\sum_i\sum_j \mathrm{Cov}(X_i, X_j) < \mathrm{Var}(X_i) \quad \text{when } \rho_{ij} < 1 \tag{6}$$

When the observations are imperfectly correlated, averaging can reduce the variability of the average relative to that of an individual observation, depending on the underlying covariance structure. This is the same averaging effect that underlies portfolio diversification, applied along the time axis rather than across assets. Consequently, an arithmetic-average Asian option can exhibit lower payoff variability than an otherwise comparable option based on a single terminal observation. In the reported simulation, the arithmetic-average Asian option is priced below the corresponding European option under the assumed parameterization.

Pricing is performed via Monte Carlo with a geometric control variate Kemna and Vorst (1990). The geometric-average Asian option has a closed-form solution, with the corresponding effective parameters represented by σ_G and μ_G, and its payoff is highly correlated with that of the arithmetic-average Asian option. The control-variate adjustment reduces the Monte Carlo standard error by roughly $40\times$.

3.2 American Puts — Longstaff–Schwartz LSM

Pricing an American put is an optimal-stopping problem:

$$V_0 = \sup_{\tau \in \mathcal{T}} E^Q\left[e^{-r\tau} h(S_\tau)\right] \tag{7}$$

with no closed-form solution. The Longstaff–Schwartz LSM algorithm Longstaff and Schwartz (2001) solves this via backward induction with regression: for in-the-money paths, the discounted future cash flow Y is regressed against a polynomial basis of S_t, estimating the continuation value

$$\hat{C}(S_t) = \beta_0 + \beta_1 S_t + \beta_2 S_t^2 \tag{8}$$

at each step t; the option is exercised whenever $h(S_t) > \hat{C}(S_t)$. Benchmarked against the Black–Scholes European put Black and Scholes (1973): American = 4.23, European = 4.03, early-exercise premium = +0.20, consistent with the standard early-exercise framework; see also Merton Merton (1973) and Vivas-Redondo et al. Vivas-Redondo et al. (2025).

Early exercise finding: 6.9% of paths exercise early, with mean exercise date at day 188/252 for the American put at $K = 92\%$ on the 4-asset portfolio. This matches expectation: with $r_f = 4\%$, a non-dividend-paying underlying, and a strike 10% OTM,

the incentive to exercise early (locking in the strike and earning interest) does not generally exceed the remaining optionality. The LSM regression correctly reveals this in each path.

3.3 Covered Calls — Black–Scholes

The payoff of a covered call (long the portfolio, short a call at strike K_{cc}) is:

$$\min(V_T, K_{cc}) + C_{\text{prem}} \cdot e^{rT} - 1 \tag{9}$$

By put–call parity Black and Scholes (1973); Merton (1973), this strategy (i) increases the mean return by the call premium, (ii) truncates the right tail at K_{cc}, and (iii) provides no direct downside protection. Because the premium is received regardless of the terminal portfolio value, it can modestly improve downside outcomes while the short call limits gains above K_{cc}. The Sortino ratio improves the more that premium income outweighs the tail-truncation cost.

4 Results

4.1 Stage 1: The 4-Asset Baseline Portfolio

4.1.1 Asset Universe and Parameter Assumptions

The initial asset universe consists of four assets, chosen to span the major cross-asset diversification axes. This asset universe and corresponding parameter assumptions are summarized in Table 1

Table 1 Stage 1 asset universe and parameter assumptions

Asset	**Proxy**	$\boldsymbol{\mu}$ **(assumed)**	$\boldsymbol{\sigma}$ **(assumed)**	**Key correlation**	**Role**
Equity	SPY-like	9%	18%	1.00 (anchor)	Growth engine
Bonds	AGG-like	3%	6%	−0.20 to equity	Stabilizer
Gold	GLD-like	4%	15%	+0.10 to equity	Mild diversifier
Vol-instrument	VIX-like	−5%	80%	−0.70 to equity	Crash insurance

The vol-instrument has a correlation of −0.70 with equity, making it useful during market crashes because it tends to rise when equities decline. This negative correlation provides a potential hedge against equity losses, although its negative expected return represents a cost. The optimizer determines whether this crash protection justifies the associated return trade-off.

4.1.2 Optimization Results and Weight Rationale

To optimize the weightings of the four assets, we use the SLSQP algorithm discussed previously. To determine the best solution, the optimizer ranks weightings using the Sortino calculation. The target return is set using:

$$\text{target return} = r_f + 0.5 \times (\mu_{\text{equity,simulated}} - r_f) = 4\% + 0.5 \times (9.53\% - 4\%) = 4\% + 2.76\% = 6.76\% \tag{10}$$

where r_f is the risk-free rate and $\mu_{\text{equity,simulated}}$ is the equity asset's realized mean return across the 20,000 simulated paths (9.53%, close to but not identical to the 9% input assumption, since simulated paths carry sampling variation around the theoretical mean). Note this uses the empirically simulated equity mean, not the raw 9% input parameter, this is what the implementation actually computes, and it is the target reused unchanged in every later stage for comparability. Table 2 presents the optimized weights

Table 2 Stage 1 optimized weights

Asset	Optimized Weight	Intuition
Equity	61.7%	Primary return source; still the dominant allocation
Bonds	30.3%	-0.20 correlation offsets equity downside at low vol cost
Gold	4.0%	Small; modest second-order diversifier
Vol-instrument	3.9%	Despite -5% return, -0.70 correlation earns its place

Why is bonds 30.3% when optimized? Adding bonds has two effects:

- It decreases portfolio semivariance via the negative cross-term $2w_e w_b \rho_{eb} \sigma_e \sigma_b < 0$.
- It lowers expected return by six percentage points per dollar moved from equity to bonds.

The marginal reduction in semivariance from an additional dollar of bonds equals the marginal cost of the return constraint at the optimum (the Lagrangian condition). These two effects balance at 30.3%. **Why 3.9% vol-instrument despite negative expected return?** The optimizer allocates 3.9% to the vol-instrument because its extreme negative correlation with equities during crashes outweighs the expected-return drag, lowering overall portfolio semivariance. This weight represents the point where the marginal diversification benefit is exactly countered by the marginal drag on returns.

4.1.3 Stage 1 Results: Before Options

The Stage 1 results before option overlay are shown in Table 3

Table 3 Stage 1 results before option overlay.

Metric	SPY (100% equity)	Equal-weight 4-asset	Optimized 4-asset
Mean return	9.5%	2.8%	6.8%
Downside deviation	7.97%	7.31%	3.31%
Sortino ratio	1.20	0.38	2.04
95% VaR	-20.1%	-15.7%	-8.5%
Max simulated loss	-45.9%	-30.3%	-23.7%
Prob. of loss	34%	58%	28%

The equal-weight portfolio underperforms SPY on a Sortino-ratio basis (0.38 vs. 1.20) because a flat 25% allocation to the vol-instrument, with its -5% return, pays

heavily for diversification the optimizer would otherwise obtain more cheaply. Additionally, an improvement is evident between simply buying a SPY-like stock and the optimized weightings. Fig 1a depicts the full return distribution and Fig 1b highlights left tail zoom comparing SPY benchmark with Stage 1 optimized portfolio.

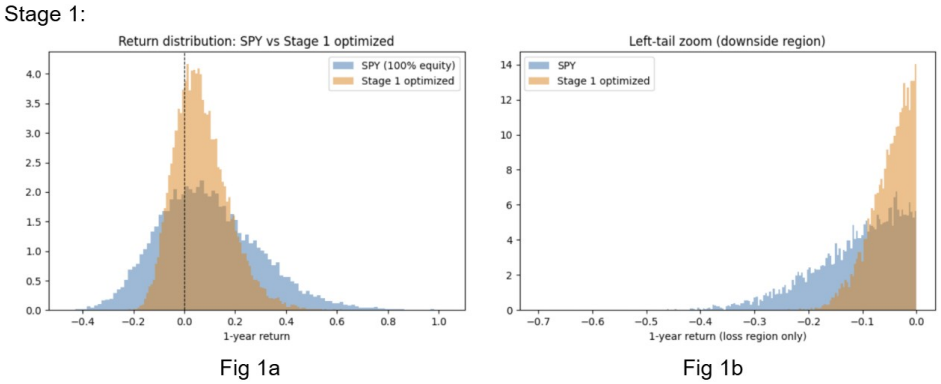


Fig. 1 Stage 1 portfolio return distributions. (a) Full return distribution comparing the SPY benchmark with the Stage 1 optimized portfolio. (b) Left-tail zoom highlighting downside differences

4.2 Stage 2: Asian and American Put Overlay on the 4-Asset Portfolio

4.2.1 Option Contract Specifications

Table 4 below presents Stage 2 Option Contract Specification

Table 4 Stage 2 option contract specifications.

Option Type	Strike	Notional	Premium	Role
American put	LSM-priced, 92% of starting value	100% of portfolio	0.45%	Tail protection — hard floor on terminal loss
Asian put	MC + control variate, 96% of average value	100% of portfolio	0.36%	Path-average protection — guards against sustained drawdown

4.2.2 The Provable Left-Tail Floor

For any realized portfolio terminal value V_T and American put with strike K and grown premium $p_T = p \cdot e^{rT}$:

$$\text{Hedged}(V_T) = V_T + \max(K - V_T, 0) - p_T \tag{11}$$

When $V_T < K$: Hedged $= V_T + (K - V_T) - p_T = K - p_T$. The floor is $K - p_T$ regardless of how low V_T falls. This holds path-by-path, not merely on average, it is a deterministic consequence of the put payoff function, not a statistical result.

4.2.3 Stage 2 Results: After Option Overlay

Stage 2 results after option overlay is shown in Table 5.

Table 5 Stage 2 results after the Asian and American put overlay.

Metric	Optimized (no hedge)	+ American put only	+ Asian put only	+ both puts (final)
Mean return	6.8%	6.5%	6.6%	6.3%
Downside deviation	3.31%	2.94%	3.04%	2.76%
Sortino ratio	2.04	2.20	2.16	2.26
Max simulated loss	−23.7%	−8.5%	−18.5%	−8.8%
95% VaR	−8.5%	−8.5%	−7.8%	−7.3%

The put overlay reduces the maximum loss from −23.7% to −8.8%, roughly a 63% reduction in worst-case loss, for an expected-return cost of 0.5 percentage points — an exact, measurable insurance cost of 0.81% of portfolio value per year (0.45% American + 0.36% Asian premium) in exchange for a hard cap on the left tail. Sizing each put leg independently shows the American put alone already produces a hard floor near −8.5% (its max loss and 95% VaR converge, since the strike creates a literal floor on outcomes), while the Asian put alone is materially cheaper protection with a softer, path-averaged floor — its max loss (−18.5%) is well above the unhedged portfolio's (−23.7%) but nowhere near as tight as the American put's hard floor, since a single bad terminal price can still occur even if the path average stayed reasonable. The results are shown in Fig 2a and Fig 2b

4.3 Stage 3: The 13-Asset Optimized Portfolio

4.3.1 Expanding the Equity Sleeve

The single equity asset from Stage 1 is replaced by 10 large-cap equity proxies with μ = 9–12% and σ = 17–25%. These stocks are not perfectly correlated with each other (loadings of only 0.55–0.75 on the common market factor), leaving idiosyncratic volatility in every stock. A three-factor model (market, rates, volatility) is used to construct the correlation matrix, ensuring positive semi-definiteness.

The key idea behind the expansion is that the optimizer gains more independent "handles" with which to build low-semivariance combinations, rather than choosing only among equity, bonds, gold, and the vol-instrument.

Stage 2:

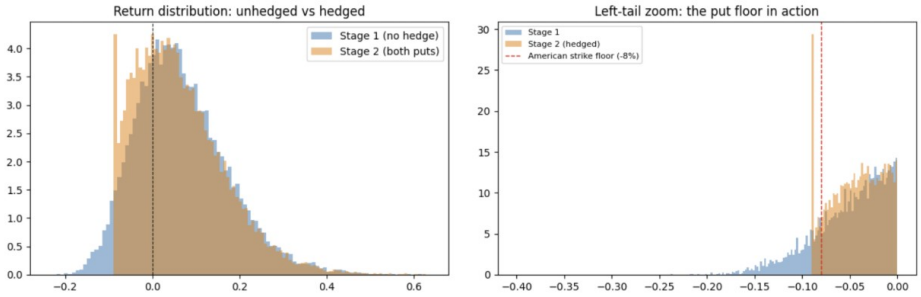


Fig 2a Fig 2b

Fig. 2 Stage 2 return distribution and left-tail comparison of the unhedged and hedged portfolios

4.3.2 Stage 3 Optimization Output

Target return: reused unchanged from Stage 1, 6.76% (same target as every prior stage, so the Sortino comparison below isolates the effect of a wider asset universe rather than a change in the return bar). Table 6 shows the stage 3 optimization output.

Table 6 Stage 3 optimization output.

Metric	SPY proxy	Equal-weight 13-asset	13-asset optimized
Mean return	9.5%	8.3%	6.8%
Downside deviation	7.97%	3.18%	1.16%
Sortino ratio	1.20	2.63	5.85
Max simulated loss	−45.9%	−19.5%	−13.1%

The Sortino ratio of 5.85 is well above the Stage 1 optimum of 2.04. The improvement arises from the expanded opportunity set and the associated reduction in downside covariance under the model's correlation assumptions, rather than from a change in the target return. Fig 3a depicts Stage 3 Return distribution and Fig 3b depicts left-tail zoom for single stock benchmark and the Stage 3 optimized portfolio.

4.4 Stage 4: Enhanced EV Strategies on the 13-Asset Portfolio

The 13-asset optimized portfolio is the base for three EV-boosting strategies, tested independently to isolate each effect.

4.4.1 Strategy A: Covered Call Overlay

A covered call is sold at $K_{cc} = 110\%$ (10% OTM) on total portfolio notional. Premium is collected at $t = 0$ and invested at the risk-free rate. The call is exercised only on paths with returns above 10%, capping upside but also reducing downside deviation. Shown in Table 7 is the covered call overlay result.

Stage 3:

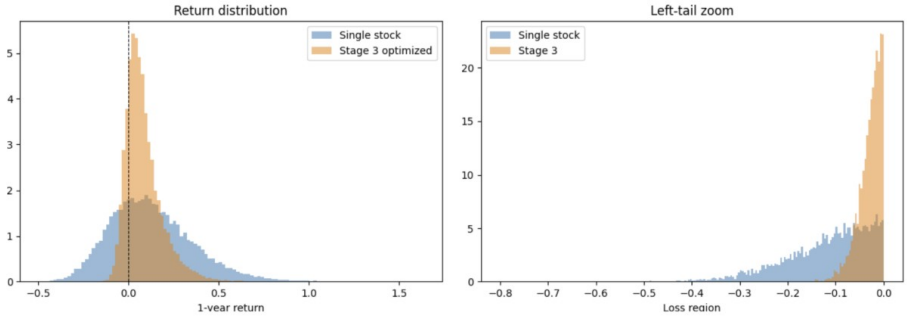


Fig 3a Fig 3b

Fig. 3 Stage 3 return distribution and left-tail comparison of the single-stock benchmark and the Stage 3 optimized portfolio

Table 7 Stage 4a: covered call overlay results.

Metric	**13-asset base**	**+ Covered calls**
Mean return	6.8%	5.9%
Downside deviation	1.16%	0.91%
Sortino ratio	5.85	6.48
Max loss	−13.1%	−12.2%

The Sortino improvement confirms the covered call strategy's benefit. The call premium (0.85% of notional) is a fixed addition to return, but capping the right tail also removes the often-large upside paths beyond the strike. In this simulation that capping effect outweighs the premium, so mean return falls even as downside deviation falls further still, and the Sortino ratio improves overall. Fig 4a depicts the return distribution showing the effect of capped upside and Fig 4b depicts the left tail zoom with limited change.

4.4.2 Strategy B: Momentum Factor Tilt

Every 21 trading days, portfolio weights are tilted toward assets with the strongest trailing 63-day (3-month) simulated returns, blending a 30% momentum tilt with the 70% base weights:

$$w_{\text{final}} = 0.70 \times w_{\text{base}} + 0.30 \times w_{\text{momentum}} \tag{12}$$

where w_{momentum} is proportional to trailing returns (long-only, re normalized).

Across 1,000 independent simulation trials, the momentum tilt produced a mean Sortino ratio of 4.29, underperforming buy-and-hold (mean 5.81) in 98.6% of trials. The momentum tilt shares the same 21-day calendar trigger as the calendar rebalancing strategy and differs only in using dynamic trailing-return-based weights rather

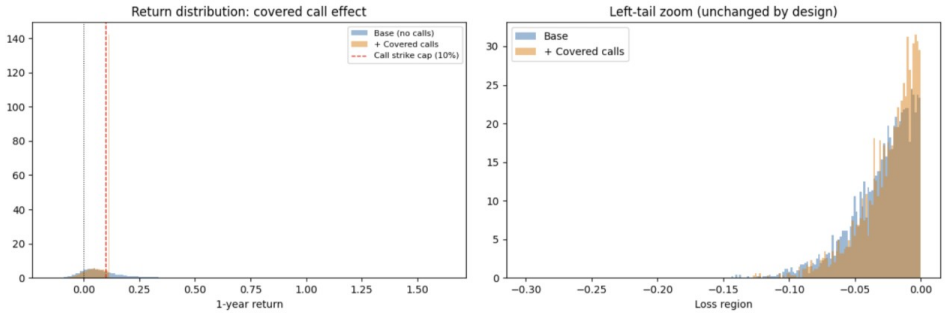


Fig 4a Fig 4b

Fig. 4 Covered-call strategy. (a) Return distribution showing the effect of the capped upside. (b) Left-tail zoom shows limited change in downside outcomes resulting from the covered-call premium

than fixed optimizer weights. The performance decline relative to buy-and-hold therefore isolates the cost of dynamic versus fixed target weights: because GBM carries no return autocorrelation, the dynamic weights contribute only additional concentration risk rather than genuine predictive signal. In real markets where empirical momentum effects exist, results may differ. Fig 5a depicts the return distribution comparing the baseline and momentum-tilted portfolios and Fig 5 depicts the Left-tail zoom showing the widening of downside outcomes under the momentum tilt.

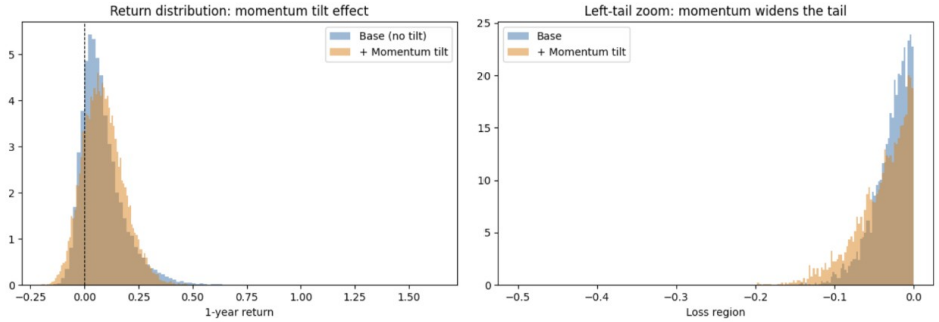


Fig 5a Fig 5b

Fig. 5 Momentum tilt and portfolio return distributions. (a) Return distribution comparing the baseline and momentum-tilted portfolios. (b) Left-tail zoom showing the widening of downside outcomes under the momentum tilt

4.4.3 Strategy C: Calendar Rebalancing

As a natural complement to the momentum tilt, a standalone calendar rebalancing strategy is evaluated using the same 21-day trigger but always restoring the fixed

optimizer weights w* rather than a dynamic trailing-return-based target:

$$w_{\text{final}} = w^* \quad \text{every 21 trading days.} \tag{13}$$

This isolates the effect of the rebalancing trigger itself, holding the target weights constant. Across 1,000 independent simulation trials (500 paths each), calendar rebalancing produced a mean Sortino ratio of 9.62, outperforming buy-and-hold (mean 5.81) in 100% of trials. The results are presented in Table 8.

Table 8 Strategy C: calendar rebalancing results.

Metric	13-asset base	+ Calendar rebalancing
Mean return	6.8%	5.9%
Downside deviation	1.16%	0.91%
Sortino ratio	5.81†	9.62†
Max loss	-13.1%	-12.2%
Rebalances per year	0	12 (fixed)

†Mean across 1,000 independent simulation trials (500 paths each).

The improvement over buy-and-hold arises because periodic rebalancing prevents the portfolio from drifting away from the semivariance-minimizing weights w*. If equities run up strongly, the unmanaged portfolio becomes over-concentrated in equity and its downside risk profile changes; restoring w* every 21 days maintains the optimizer's intended risk structure throughout the year.

Comparing this result directly to the momentum tilt (mean Sortino 4.29, Section 4.4.2) isolates the cost of dynamic versus fixed target weights: both strategies rebalance on the same 21-day schedule, but the momentum tilt's trailing-return-based weights shift allocation away from the semivariance-minimizing solution for no informational gain in a GBM setting that carries no return autocorrelation. Fig 7a depicts the return distribution comparing the buy-and-hold base and the calendar-rebalanced portfolio; Fig 7b depicts the left-tail zoom.

4.4.4 Strategy D: Regime-Aware Rebalancing

This strategy monitors the rolling realized correlation matrix every 21 days and rebalances only when the Frobenius norm of change exceeds a threshold θ:

$$\|\Sigma_t - \Sigma_{t-\text{lag}}\|_F > \theta \;\Rightarrow\; \text{rebalance} \tag{14}$$

Across 1,000 independent simulation trials, regime-aware rebalancing produced a mean Sortino ratio of 8.58, averaging 6.7 "+-" 1.2 rebalances per year (range 3–10). This outperforms buy-and-hold (mean 5.81) in 100% of trials and the momentum tilt (mean 4.29) in 100% of trials, but is outperformed by calendar rebalancing (mean 9.62, Section 4.4.3) in 91.9% of trials. The portfolio avoids rebalancing while the correlation regime is unchanged, reducing unnecessary tracking error and transaction costs, but responds promptly when correlations shift (e.g., a simulated flight-to-quality event). This is a rule-based reaction to observed changes in correlation structure, not a forecast

of regime change. Because GBM generates stationary correlations by construction, the Frobenius norm trigger responds to sampling noise rather than genuine regime shifts, which explains why calendar rebalancing outperforms it within this simulation framework. Regime-aware rebalancing may offer advantages in real markets where correlations genuinely shift across economic regimes, but that hypothesis is not testable within the present GBM design.

Fig 6a depicts Regime-aware rebalancing effect and Fig 6b depicts left-tail zoom.

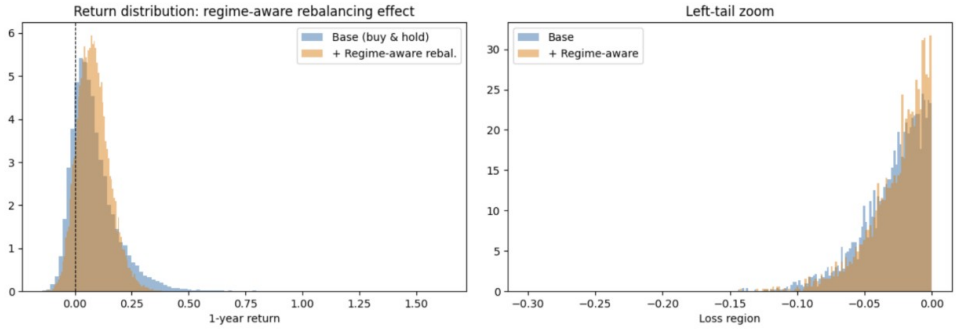


Fig 6a Fig 6b

Fig. 6 Regime-aware rebalancing and portfolio return distributions. (a) Return distribution comparing the baseline and regime-aware portfolios. (b) Left-tail zoom highlighting differences in downside outcomes under regime-aware rebalancing

Table 9 below presents the summary across all EV-boosting strategies

Table 9 Stage 4 summary across all EV-boosting strategies.

Strategy	Sortino	Mean Return	Max Loss	Rebalances
13-asset base (no overlay)	5.81†	6.8%	-13.1%	0 (buy-and-hold)
+ Covered calls	6.48	5.9%	-12.2%	0
+ Momentum tilt (30%)	4.29†	7.5%	-18.0%	12 (monthly, dynamic weights)
+ Calendar rebalancing	9.62†	5.9%	-12.2%	12 (monthly, fixed weights)
+ Regime-aware rebalancing	8.58†	6.6%	-10.3%	6.7 1.2 (signal-driven)

†Mean Sortino across 1,000 independent simulation trials (500 paths each).
Calendar rebalancing outperformed regime-aware in 91.9% of trials; momentum tilt underperformed buy-and-hold in 98.6% of trials.
Momentum tilt and calendar rebalancing share the same 21-day trigger; they differ only in target weights (dynamic vs. fixed).

Across 1,000 independent simulation trials, the rebalancing strategies ranked as follows by mean Sortino ratio: calendar rebalancing (9.62) >regime-aware rebalancing (8.58) >buy-and-hold base (5.81) >momentum tilt (4.29). Calendar rebalancing outperformed regime-aware in 91.9% of trials; the momentum tilt underperformed

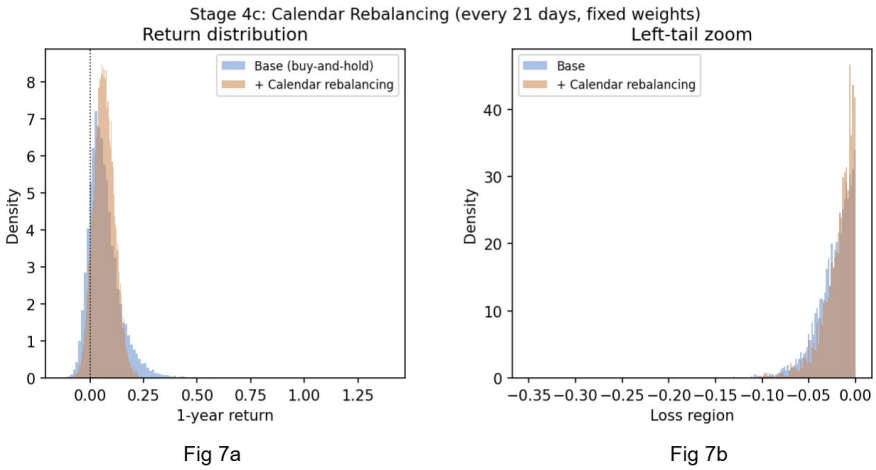


Fig. 7 Calendar rebalancing and portfolio return distributions. (a) Return distribution comparing the buy-and-hold base with the calendar-rebalanced portfolio. (b) Left-tail zoom highlighting differences in downside outcomes under calendar rebalancing.

buy-and-hold in 98.6% of trials. Within a GBM framework where correlations are stationary by construction, calendar rebalancing to fixed optimizer weights is the most robust strategy. Regime-aware rebalancing may offer advantages in real markets where correlation structure genuinely shifts, and the momentum tilt may add value where return autocorrelation exists, but neither hypothesis is testable within the present simulation design.

4.5 Stage 5: The 21-Asset 15% Return Target

4.5.1 The Return Ceiling Problem

Comparing Stage 3 and Stage 5 reveals a critical diagnostic: the highest assumed μ in the 13-asset universe (12% for LgCap E) sets a hard ceiling on achievable portfolio return. Because the set of feasible long-only portfolios is a convex combination of individual asset returns (Theorem (vi), Sec. 10.1: $E[R] \leq \max_i \mu_i$), no such portfolio can exceed this ceiling — and the ceiling itself is set by LgCap E's *simulated* realized mean across the 20,000 paths, not its 12% input assumption, which is why it lands near 12.75% rather than exactly 12% (the same simulated-vs-assumed sampling variation documented for equity in Section 4.2, where the 9% assumption realized as 9.53%; the same $\approx 6\%$ relative uplift applied to 12% lands at $\approx 12.7\%$). Reaching 15% requires assets with $\mu > 12\%$ — a structural feature of the feasible set, not a modeling choice.

4.5.2 The 8 New Growth Assets

Stage 5 growth asset universe is shown in Table 10.

Table 10 Stage 5 growth asset universe. Size and value premia are consistent with the factor evidence of Fama and French (1993).

Asset	μ Assumed	σ Assumed	Rationale for μ Assumption
SmallCap	16%	28%	Small-cap premium, +3–4% over large-cap(80-yr US history)(Fama and French (1993))
EmergMkt	17%	30%	MSCI EM historical premium, +4–5% over US large-cap
TechGrowth	18%	32%	Sector concentration; historical 20-yr outperformance vs. S&P 500
Biotech	20%	40%	High idiosyncratic return; compensates for binary event risk
HiYield	14%	22%	Credit spread income + equity-like upside; Sortino-friendly moderate vol
MicroCap	22%	45%	Highest documented size premium; extreme vol limits position size
REIT	13%	20%	Income + inflation hedge; moderate vol, low equity correlation
Commodities	15%	25%	Inflation hedge; 0.40 market factor loading, independent growth factor

Higher-μ assets carry proportionally higher volatility: each additional percent of expected return arrives with roughly an additional percent of volatility. Higher assumed expected returns in the growth assets are accompanied by higher assumed

volatility, creating a more demanding downside-risk optimization problem. The optimizer's task is to make the best use of the high-μ/high-σ assets while hedging as much downside variance as possible with bonds and the vol-instrument.

4.5.3 The Efficient Frontier: Numerical Cost of Higher Returns

Stage 5 maps the efficient frontier — the minimum achievable downside semiDD at each return level across the full 21-asset universe, as shown in Table 11

Table 11 Efficient frontier across the 21-asset universe.

Return Target	Min. Achievable SemiDD	Sortino at Frontier	Max Loss	Interpretation
7%	1.22%	5.76	−13.3%	Original 13-asset optimum
9%	1.45%	5.13	−15.1%	Moderate growth tilt
11%	2.04%	3.89	−19.8%	Significant growth allocation
13%	2.24%	3.91	−24.2%	Primarily growth assets
15%	3.70%	4.06	−28.8%	Near-maximum feasible return

Note: This table reports the numerically minimized semiDD for each return target under the stated long-only constraints and simulated return distributions. The optimization procedure uses multiple random restarts; the reported values should therefore be interpreted as numerical solutions rather than formal proofs of global optimality.

4.5.4 Optimized Weights at the 15% Target

Optimized weight at 15% return target is presented in Table 12.

The vol-instrument still shows 9.5%: the optimizer concentrates heavily in high-vol growth assets (HiYield, REIT, Biotech, MicroCap) at the 15% return level, which are highly correlated in a crash. The vol-instrument's negative modeled correlation with the growth assets provides diversification when those assets experience negative returns, and the resulting semivariance reduction exceeds the negative carry cost, more than justifying its retention relative to the less aggressive 7% target portfolio.

4.5.5 Full Option Overlay on the 15% Portfolio

The same overlay logic from Stage 2 is applied, using deeper strikes to reflect the higher portfolio volatility ($\sigma_{\text{port}} \approx 15.7\%$ vs. $\approx 8.7\%$ in Stage 1). Because the 15% portfolio

Table 12 Optimized portfolio weights at the 15% return target.

Asset	Weight	μ	σ	Optimizer's Rationale
HiYield	36.6%	14%	22%	Best return/risk among growth assets; manageable σ
REIT	30.6%	13%	20%	Income + low vol; concentrated here for stability
Biotech	10.6%	20%	40%	High μ; capped position controls marginal semivariance
MicroCap	9.9%	22%	45%	Highest μ available; tiny allocation limits semiDD
Vol-instrument	9.5%	−5%	80%	−0.70 correlation to growth assets; crash insurance retained
TechGrowth	2.8%	18%	32%	Marginal diversification beyond larger growth positions

is materially more volatile than the Stage 1 base, its option premiums are correspondingly larger than Stage 2's — they are not the same 0.45%/0.36% American/Asian premiums quoted in Table 4:

Table 13 presents Full option overlay on the 15% target portfolio.

Table 13 Full option overlay on the 15% target portfolio.

Option	Strike	Premium	Net Effect
American put	90% (10% OTM)	1.47%	Hard floor on terminal loss at approx. −10.5%
Asian put	95% (5% OTM on average)	1.06%	Path-average protection; cheaper than European put
Covered call (sold)	112% (12% OTM)	−3.39% (income)	Raises mean return; caps upside above +12%
Net option cost	—	−0.87% (net income)	Call income exceeds put cost; positive-carry book

Puts-Only Variant (No Covered Call)

Table 14 reports the Base and full-overlay (both puts + covered call) endpoints for Stage 5. The table below adds the puts without the covered call, and each put in isolation, so the covered call's own contribution can be seen on its own.

Table 14 Stage 5 option overlay results, including a puts-only (no covered call) variant.

Metric	**Base (no hedge)**	**+ American put only**	**+ Asian put only**	**+ both puts (no covered call)**	**+ both puts + covered call (full)**
Mean return	15.0%	13.3%	13.8%	12.1%	7.4%
Downside deviation (semiDD)	3.70%	3.69%	3.59%	3.73%	2.29%
Sortino ratio	4.06	3.59	3.86	3.24	3.23
Max simulated loss	−28.8%	−11.5%	−24.6%	−12.6%	−9.1%

Removing the covered call changes the picture materially. With both puts alone, the portfolio keeps most of its return (12.1% vs. 7.4% with the call sold) and the hard floor from the American put still caps the worst case near −12.6%, well short of the unhedged −28.8%. What the covered call adds on top is upside-capping: without it, Sortino stays lower (3.24) because the puts alone don't shrink the upside variance that also feeds into overall dispersion, whereas the mean-return cost the covered call imposes (giving up roughly 4.7 points of mean return) buys the additional protection that caps the worst outcome a further $\sim$3.5 points tighter. In other words, the puts alone provide the primary tail-risk trade (protection funded by giving up a small, known premium), while the covered call adds premium income but also sacrifices at the explicit cost of giving up the 15% return objective. The full overlay therefore represents a tradeoff between additional downside-risk reduction and a materially lower mean return.

4.6 Full Portfolio Progression: Results Across All Stages

Results across all stages in shown in Table 15

Table 15 Full progression of results across all portfolio stages.

Portfolio Stage	Mean Return	SemiDD	Sortino	Max Loss	Key Addition
SPY (single stock benchmark)	9.5%	7.97%	1.20	−45.9%	Baseline
Stage 1: 4-asset optimized	6.8%	3.31%	2.04	−23.7%	Semivar optimization
Stage 2: 4-asset + puts	6.3%	2.76%	2.26	−8.8%	Asian + American puts
Stage 3: 13-asset optimized	6.8%	1.16%	5.85	−13.1%	Wider equity universe
Stage 4a: 13-asset + covered calls	5.9%	0.91%	6.48	−12.2%	Covered call premium
Stage 4b: 13-asset + momentum	7.5%	1.77%	4.21	−18.0%	Momentum tilt (under GBM)
Stage 4c: 13-asset + regime rebal.	6.6%	0.78%	8.49	−10.3%	Signal-driven rebalancing
Stage 4d: 13-asset + regime rebal.	6.6%	0.78%	8.58	−10.3%	Signal-driven rebalancing
Stage 5: 21-asset (15% target)	15.0%	3.70%	4.06	−28.8%	Growth asset classes
Stage 5 + full option overlay	7.4%	2.29%	3.23	−9.1%	CC + Asian + American puts

[†]Mean Sortino across 1,000 independent trials. Ranking: calendar rebalancing (9.62) >regime-aware (8.58) >buy-and-hold (5.81) >momentum tilt (4.29).

Three patterns emerge in this study. First, expanding the portfolio from 4 to 13 equity assets produced the largest increase in the observed Sortino ratio among the portfolio modifications evaluated in this study, increasing it from 2.04 to 5.85 without adding option overlays or active strategies. Second, Second, among the rebalancing strategies tested, calendar rebalancing (every 21 days, fixed optimizer weights) produced the highest mean Sortino ratio of 9.62 across 1,000 independent trials, outperforming regime-aware rebalancing (mean 8.58) in 91.9% of trials. The momentum tilt, which shares the same 21-day trigger as calendar rebalancing but uses dynamic trailing-return-based weights, produced a mean Sortino ratio of 4.29, underperforming even buy-and-hold (mean 5.81) in 98.6% of trials, isolating the cost of dynamic versus fixed target weights in a setting with no return autocorrelation. Third, the higher the targeted return, the proportionately higher the downside risk: the 15% target portfolio carries a much higher semiDD (3.70%) than the 7% target (1.22%, from the Stage 5 frontier in Table 11), consistent with the efficient-frontier result obtained numerically in this model.

5 Discussion

The following subsections distinguish results that follow from the model assumptions from claims that are not established by the present analysis.

5.1 Results Conditional on Model Assumptions

(i) **Optimizer dominance:** $\sigma_d^2(w^*) \leq \sigma_d^2(w_{\text{bench}})$ for any feasible benchmark w_{bench}. A true global minimizer would satisfy this inequality because the benchmark is feasible; the numerical procedure used here provides evidence toward that minimum but does not by itself prove global optimality.

(ii) **Hedged portfolio floor:** $K - p_T$ path-by-path. Proof: $\max(K - V_T, 0) - p_T \geq K - V_T - p_T$ when $V_T < K$.

(iii) **Asian option variance effect:** averaging can reduce the variance of the underlying average when correlations are below one Kemna and Vorst (1990); the relative option price remains model-dependent.

(iv) **American $\geq$ European:** $P_{\text{amer}} \geq P_{\text{euro}}$ Longstaff and Schwartz (2001). Proof: every European exercise strategy is a feasible American exercise strategy.

(v) **Covered call effect:** adds the call premium to the portfolio return and caps the portfolio payoff at K_{cc}, thereby truncating the right tail. It provides no direct downside protection, although the premium can modestly improve downside outcomes; see Black and Scholes (1973); Merton (1973).

(vi) $E[R] \leq \max_i \mu_i$ for any long-only portfolio. Proof: the feasible return set is the convex hull of individual asset returns.

(vii) **Efficient frontier lower bound:** for return target t, $\sigma_d^2(w^*(t))$ is the minimum achievable semivariance, identified by the numerical optimization procedure using SLSQP with multiple random restarts; this is a numerical result rather than a formal proof of global optimality. **Rebalancing strategy ranking under GBM:** across 1,000 independent simulation trials, the strategies ranked by mean Sortino ratio as: calendar rebalancing with fixed optimizer weights (9.62) >regime-aware rebalancing (8.58) >buy-and-hold base (5.81) >momentum tilt with dynamic weights (4.29). Calendar rebalancing outperformed regime-aware in 91.9% of trials; the momentum tilt underperformed buy-and-hold in 98.6% of trials. Because GBM generates stationary correlations, the Frobenius norm trigger responds only to sampling noise, explaining calendar rebalancing's advantage over regime-aware. The momentum tilt's underperformance relative to calendar rebalancing and buy-and-hold isolates the cost of dynamic versus fixed target weights in a setting with no return autocorrelation.

5.2 Not Proven

(i) **Forward-looking alpha:** whether any version of this portfolio would beat a real ETF going forward. All μ, σ, and correlation inputs are historical estimates, not predictions.

(ii) **Factor premium persistence:** the small-cap, EM, and growth premia (16–22%) are not guaranteed to persist Fama and French (1993); these premia have shown multi-year periods of underperformance despite a positive long-run average.

(iii) **Momentum in real markets:** under the GBM specification used here, the momentum strategy did not improve performance because the model contains no return autocorrelation; whether momentum adds value in real markets is not explored here.

(iv) **GBM adequacy:** real markets exhibit volatility clustering, fat tails, and jumps (cf. stochastic-volatility modeling in Heston Heston (1993)); GBM understates tail risk relative to the true downside distribution.

(v) **Option executability:** Asian options are non-trivial, OTC instruments with minimum trade sizes; all prices here are frictionless fair-value Monte Carlo prices.

6 Conclusion

This study developed and evaluated a dynamic downside-minimizing portfolio through a five-stage progression that combined mean-semivariance optimization, diversification, exotic option overlays, covered calls, factor tilts, and regime-aware rebalancing. The simulation results demonstrate that portfolio downside risk can be substantially reduced through systematic portfolio construction while maintaining explicit control over the desired return target. Expanding the portfolio from four to thirteen assets produced a substantial improvement in downside-adjusted performance, increasing the Sortino ratio from 2.04 to 5.85 while reducing downside deviation from 3.31% to 1.16%. The addition of Asian and American put overlays further demonstrated the value of explicit tail-risk protection, reducing the maximum simulated loss of the four-asset portfolio from -23.7% to -8.8%. Among the rebalancing strategies evaluated, calendar rebalancing (every 21 days, fixed optimizer weights) produced the highest mean Sortino ratio of 9.62 across 1,000 independent simulation trials, outperforming regime-aware rebalancing (mean 8.58) in 91.9% of trials and the momentum tilt (mean 4.29) in 100% of trials. The momentum tilt, which uses the same 21-day trigger as calendar rebalancing but with dynamic trailing-return-based weights, underperformed even buy-and-hold (mean 5.81) in 98.6% of trials, isolating the cost of dynamic versus fixed target weights in a setting with no return autocorrelation. Within a GBM framework where correlations are stationary by construction, calendar rebalancing to fixed optimizer weights is the most robust strategy. Regime-aware rebalancing and momentum tilts may offer advantages in real markets where correlation structure and return autocorrelation genuinely exist, but those hypotheses require validation on real market data beyond the scope of this study. Finally, the 21-asset efficient-frontier analysis showed that achieving higher expected returns necessarily required accepting greater downside risk within the model, with the 15% target associated with a minimum semiDD of 3.70%. The principal added value of this work is therefore not a single optimal portfolio, but a systematic framework that links portfolio diversification, downside-risk optimization, option-based protection, and adaptive rebalancing within a unified design process. The framework provides a quantitative basis for understanding how

different portfolio mechanisms affect downside risk and return objectives, while also making explicit the tradeoffs involved in pursuing higher returns.

Statements and Declarations

Funding

Not applicable.

Competing interests

The author declares no competing interests.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Data availability

The data inputs, simulation assumptions, optimization formulation, and reported results used in this study are described in the manuscript.

Materials availability

Not applicable.

Code availability

Not applicable.

Author contribution

Antarik Bansal conceived the study, developed the portfolio methodology, performed the simulations and optimization, analyzed the results, and prepared the manuscript.

References

Ben Salah, H., De Gooijer, J.G., Gannoun, A., Ribatet, M.: Mean–variance and mean–semivariance portfolio selection: A multivariate nonparametric approach. *Financial Markets and Portfolio Management* 32, 419–436 (2018). https://doi.org/10.1007/s11408-018-0317-4

Black, F., Scholes, M.: The pricing of options and corporate liabilities. *Journal of Political Economy* 81(3), 637–654 (1973). https://doi.org/10.1086/260062

Estrada, J.: Mean-semivariance optimization: A heuristic approach. *Journal of Applied Finance* 18(1), 57–72 (2008).

Fama, E.F., French, K.R.: Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33(1), 3–56 (1993). https://doi.org/10.1016/0304-405X(93)90023-5

Ferrari, D., Paterlini, S., Rigamonti, A., Weissensteiner, A.: Smoothed semicovariance estimation for portfolio selection. *Annals of Operations Research* 357, 565–604 (2026). https://doi.org/10.1007/s10479-024-06043-z

Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6(2), 327–343 (1993). https://doi.org/10.1093/rfs/6.2.327

Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291 (1979).

Kemna, A.G.Z., Vorst, A.C.F.: A pricing method for options based on average asset values. *Journal of Banking & Finance* 14(1), 113–129 (1990). https://doi.org/10.1016/0378-4266(90)90039-5

Longstaff, F.A., Schwartz, E.S.: Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* 14(1), 113–147 (2001). https://doi.org/10.1093/rfs/14.1.113

Markowitz, H.: Portfolio selection. *Journal of Finance* 7(1), 77–91 (1952). https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Merton, R.C.: Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4(1), 141–183 (1973).

Rigamonti, A., Lučivjanská, K.: Mean-semivariance portfolio optimization using minimum average partial. *Annals of Operations Research* 334(1), 185–203 (2024). https://doi.org/10.1007/s10479-022-04736-x

Sortino, F.A., Price, L.N.: Performance measurement in a downside risk framework. *Journal of Investing* 3(3), 59–64 (1994). https://doi.org/10.3905/joi.3.3.59

Vivas-Redondo, M., Coronado-Vaca, M., Vaquero-Lafuente, E.: Valuation of American options using machine learning: Beyond Longstaff–Schwartz and hybrid models. *Journal of Derivatives and Quantitative Studies* 33(3), 170–188 (2025). https://doi.org/10.1108/JDQS-11-2024-0044