Louvain School of Management

To what extent can deep out-of-the-money put options enhance downside protection for equity portfolios while preserving cost efficiency under different market regimes?

Author(s): Roba Alexandre
Supervisor(s): Vrins Frédéric
Academic year 2025.-2026.
Dissertation for the Master of Business Engineering
Focus on Financial Engineering
Daytime schedule:

Abstract

This thesis examines whether deep out-of-the-money put options can improve downside protection for equity portfolios while preserving cost efficiency across different market regimes. The analysis compares an unhedged long SPX Index portfolio with a hedged portfolio that adds European put options on the same underlying over a one-year horizon, subject to a maximum option budget of 5%. DOOM puts are defined through a target probability under the risk-neutral measure. The same target probability represents the same risk-neutral tail quantile across models, even though each model produces different strikes and premia.

The analysis is conducted in two steps. First, the model-based benchmark compares three pricing frameworks, Black-Scholes-Merton, Merton jump-diffusion and Heston stochastic volatility, at three option-surface calibration dates: 16 March 2020, 8 April 2025 and 30 June 2026. For each date, physical dynamics are estimated from historical SPX returns and used to generate one-year portfolio scenarios. The corresponding risk-neutral pricing specifications are calibrated to the dated SPX option surfaces and used to construct DOOM strikes, option premia and static hedge allocations through the $\mathrm{ES}_{95\%}$-based optimization.

Second, the resulting model-based allocations are evaluated on historical SPX one-year rolling windows classified into Crash, Stress and Normal market regimes. The allocations are applied without re-optimization to these realized scenarios. This second evaluation tests whether the model-selected hedges remain effective outside the simulated environment and therefore exposes model risk.

The results show that DOOM puts can reduce downside losses, but only when the selected strike region matches the market regime. This match depends on the volatility level and downside skew priced in the option surface, which shape the cost of protection across strikes. The best hedge does not necessarily come from the deepest puts. Low-strike puts create a far-tail floor, while higher-strike puts can be more effective in crash or stress regimes when they pay across a wider range of realized losses. Overall, effective DOOM puts rely on the alignment between the strike region, the option-surface regime and the losses targeted by the hedge, relative to the investor's objective.

Keywords: Deep out-of-the-money puts, tail-risk hedging, downside protection, Expected Shortfall, model risk, Black-Scholes-Merton, Merton jump-diffusion, Heston stochastic volatility.

AI Declaration

UCLouvain

LOUVAIN
School of Management

DECLARATION

STATEMENTS ON YOUR USE OF GENERATIVE AI

Confirm the following statements:

Statements on responsible use of generative AI	Confirmation: Yes
The content of my/our master's thesis is my/our original input, even if I/we used generative AI to help prepare it.	☒
I/we master the theories, tools and methods that I/we use for the thesis.	☒
I/we verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc.	☒
I/we acknowledge that I/we master the final output and that I am/we are able to explain it and answer questions about it.	☒

SIGNATURE

Sign your declaration:

	Author last name, first name(s)	Date	Signature
1	ROBA ALEXANDRE	26-08-05	

NOTE

I acknowledge the use of artificial intelligence tools during the preparation of this thesis. ChatGPT was used mainly as a coding and language-support tool. It helped with Python implementation, debugging, wording, grammar and the refinement of some formulations. It was also used to support the reformulation and validation of certain ideas, as well as the verification of some mathematical developments.

The research question, structure, methodological choices, mathematical reasoning, empirical design, analysis and interpretation of the results are my own. AI-generated outputs were used only as support material and were carefully reviewed, modified and validated before being included in the thesis. I take full responsibility for the final content presented in this thesis.

Acknowledgements

I would first like to sincerely thank Professor Vrins, my thesis supervisor, for his guidance, support and valuable advice throughout my academic journey and this research project. Beyond his supervision, his courses in derivative pricing and credit risk introduced me to quantitative finance and contributed to my interest in this field.

During my internship within the Quantitative Modelling team of Belfius' Trading Floor, I had the opportunity to discuss my thesis with several professionals. I would particularly like to thank Ludovic Nogues and Alexandre Garrivier for their valuable advice, constructive feedback and for encouraging me to question my assumptions throughout this work. Their perspectives helped me improve both the development of this thesis and my understanding of the topic.

I would also like to sincerely thank my family, my girlfriend and my friends for their constant support and for always believing in me. They have encouraged me to keep improving and to pursue my ambitions throughout my academic journey.

As my studies at UCLouvain come to an end, I am deeply grateful for everything this institution has given me and for the person it has helped me become. This thesis represents the completion of an important chapter of my journey, but also the beginning of new challenges and ambitions.

I would like to conclude these acknowledgements with the words of Albert Einstein:

> ”Strive not to be a success, but rather to be of value.”

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

1.1 Motivation

The multiple high-volatility periods observed in financial markets over the past two decades, including the subprime crisis, the pandemic shock, and recent geopolitical tensions, have highlighted the structural vulnerability of equity portfolios to extreme market downturns. For instance, during the COVID-19 market turmoil, the S&P 500 (SPX) fell by 33.8% between 19 February 2020 and 23 March 2020 (S&P Dow Jones Indices, 2020). These periods, driven by abrupt shifts in risk aversion and market liquidity, emphasize the need for effective hedging strategies capable of mitigating portfolio losses during volatility spikes. Tail-risk hedging aims to offset extreme downside exposure while preserving cost efficiency and long-term return. Optimizing this trade-off remains a central challenge for both academic researchers and financial practitioners.

1.2 DOOM Puts and Pricing Models

Among derivative instruments, Deep out-of-the-money (DOOM) put options are widely used to protect equity portfolios against extreme losses. Their convex payoff structure provides asymmetric exposure during market stress, but the net benefit of this protection depends on the market conditions at initiation. The trade-off between downside protection and cost efficiency is therefore fundamental to evaluating the effectiveness of DOOM puts as a hedging strategy.

In a model-based analysis, the choice of pricing model is important because the valuation and payoff probability of a DOOM put depend on the assumed left-tail dynamics. This differs from a historical SPX backtest, where option premiums are taken from market quotes rather than priced by a model. For this reason, this thesis first presents model-based results as benchmarks, comparing how continuous diffusion, jump risk and stochastic volatility affect the selected protection. Then, the historical SPX backtest sheds light on model risk by applying the resulting model-based allocations to realized index scenarios.

1.3 Position in the Literature

The theoretical starting point is the no-arbitrage framework for option pricing. Black and Scholes (1973) and Merton (1973) provide the benchmark continuous-time framework in which derivative prices are obtained in the absence of arbitrage. Hull (2018) remains a reference for the financial interpretation of option pricing, dynamic hedging and the limitations of diffusion models. In its standard Black-Scholes-Merton (BSM) diffusion specification, the model provides a simple and interpretable benchmark with closed-form option prices. However, the assumptions of constant volatility and continuous lognormal returns limit its ability to reproduce the pronounced downside skew observed in equity index option markets.

Two extensions of the BSM benchmark are therefore relevant for analysing DOOM put strategies. Merton (1976) introduces discontinuous jumps, which provide a structural way to represent sudden price drops. Heston (1993) incorporates a mean-reverting stochastic volatility process, making option prices dependent on both the index level and the current variance state. These two models are used in this thesis to study how DOOM-put protection changes when tail risk changes. More recent rough-volatility models provide an extension by modelling volatility through a fractional structure with a Hurst exponent below one half. This generates rough, non-Markovian volatility paths that are more consistent with the behaviour of realized volatility in equity markets (Gatheral et al., 2018). These models are not considered in this thesis because their non-Markovian structure would introduce a different numerical framework and additional modelling assumptions, thereby limiting direct comparability with the other models examined.

This modelling choice is also consistent with the empirical option-pricing literature. Bates (1991, 2000) shows that S&P 500 option prices contain information about crash fears, while Bakshi et al. (1997) find that stochastic-volatility and jump specifications improve on the benchmark Black-Scholes framework in empirical option pricing. This underpins the comparison in this thesis of the following three models: Black-Scholes-Merton, Merton jump-diffusion and Heston stochastic volatility as alternative frameworks. However, model evaluation cannot rely solely on pricing accuracy, as pricing fit and hedging performance are distinct criteria. A better fit to option prices does not necessarily imply a more effective hedge or a more accurate representation of the true physical dynamics of the underlying (Lassance & Vrins, 2018).

A second strand of the literature focuses on the cost and evaluation of option-based crash protection. The issue is whether the premium paid for downside insurance is justified by the protection obtained in adverse market states. Bondarenko (2014) studies the apparent overpricing of S&P 500 put options, showing that they have historically been expensive relative to their realized payoffs and to standard asset-pricing benchmarks. Broadie et al. (2009) show that option-based strategies are difficult to evaluate with standard performance measures because their payoffs are rare, non-linear and concentrated in specific market states. For this reason, this thesis evaluates

DOOM-put strategies on a long SPX exposure using average terminal returns, downside-risk metrics and mean return drag.

A third part of the literature decomposes option premia into compensation for volatility, jump and disaster risk. Bakshi and Kapadia (2003) show that option positions still reflect a volatility risk premium after delta hedging removes their direct directional exposure to the underlying index. Carr and Wu (2009) analyse variance risk premia, while Pan (2002) identifies that option prices also contain compensation for jump risk. Bollerslev and Todorov (2011) further link option prices to rare-disaster compensation. This is relevant for DOOM puts because they pay when the index reaches the far left tail of its return distribution. Option prices can also be affected by demand for protection. Bollen and Whaley (2004) show that net buying pressure affects implied volatility surfaces, while Gârleanu et al. (2009) show that demand pressure can increase option prices when intermediaries cannot absorb investor demand. Together, these results motivate the distinction between the physical measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$. Portfolio losses are realized under the physical distribution, whereas option premia are determined under the risk-neutral measure and embed market-implied compensation for volatility and crash risk.

Finally, this thesis also relates to studies comparing different forms of downside protection. Ilmanen et al. (2021) compare put protection with trend-following overlays and emphasise the central trade-off: puts provide contractual crash convexity, but this convexity is often expensive outside crisis states. The contribution of this thesis is to compare the same protection rule across alternative sources of downside risk, in order to assess whether the resulting risk-cost profile is robust to the choice of pricing and simulation models.

1.4 Research Question and Methodology

Based on these frameworks, the research question of this thesis is as follows

Research Question

To what extent can deep out-of-the-money put options enhance downside protection for equity portfolios while preserving cost efficiency under different market regimes?

The portfolio considered in this thesis is a long position in the SPX index. The unhedged benchmark is the index position alone. The hedged portfolio combines the same SPX exposure with an allocation to European put options written on the index. The comparison measures how the hedge changes the return distribution of the SPX position, using return and downside-risk metrics.

The methodology has two parts. First, a model-based analysis is conducted and the optimal hedge weights are estimated. Historical SPX prices are obtained from yfinance, listed SPX option surfaces from Bloomberg, and SOFR rates from the Federal Reserve Bank of New York.

Physical dynamics are estimated from historical S&P 500 returns, while risk-neutral inputs are calibrated to listed SPX option prices and SOFR rates for each calibration date. The selected surfaces cover three distinct market regimes. The 16 March 2020 surface corresponds to the COVID-19 crash, with extreme demand for downside protection. The 8 April 2025 surface captures a stressed environment following the Trump tariff announcements. The 30 June 2026 surface is used as a normal reference date, observed near recent SPX record highs and with more limited downside premia. For each date and model, SPX paths are simulated under $\mathbb{P}$, put values are obtained under $\mathbb{Q}$, and a one-year hedge is optimized with a maximum option budget of 5%. The benchmark objective uses 95% Expected Shortfall, because it captures the average severity of tail losses beyond the threshold, whereas VaR ignores them.

Second, the weights obtained from the model-based optimisation are tested on historical SPX returns. The weights are applied without re-optimisation to one-year rolling windows of SPX returns, with results reported for the full sample and for crash, stress and normal regimes. This backtest is used to assess return drag, downside-risk reduction and model risk.

1.5 Structure

Chapter 2 defines the portfolios, the construction of DOOM put contracts and the risk metrics. Chapter 3 presents the BSM, Merton and Heston pricing, simulation and calibration. Chapter 4 describes the static allocation and evaluation protocols. Chapter 5 reports the model-based and historical SPX results. Chapter 6 concludes by answering the research question.

CHAPTER 2

Conceptual Framework and Risk Metrics

2.1 Reference Portfolio

We consider an investor with initial wealth $B_0 = 100$ over a one-year horizon $T = 1$. The investor holds a long position in the SPX index, with the initial index level set to $S_0 = 100$. The simulated risky underlying is therefore directly the index S_t itself, rather than an explicit basket of individual securities. This modelling assumption is important because the stochastic specifications used in this thesis do not attempt to reconstruct the internal composition of the index or the correlation structure among its constituents. The strategy can therefore be represented directly through the dynamics of the relevant state variable, namely the level of the index on which the hedging put options are written.

The unhedged benchmark is the Index portfolio. Since $B_0 = S_0$, this portfolio corresponds in the simulations to holding one unit of the SPX index.

$$P_t^I = \frac{B_0}{S_0} S_t. \tag{2.1}$$

All portfolio values are reported as discounted normalised wealth,

$$W_t = \frac{P_t}{B_0 e^{rt}}, \tag{2.2}$$

so that $W_t - 1$ measures performance relative to the money-market account. The initial value $S_0 = 100$ is a scaling convention. For each calibration experiment, the rate r in Equation (2.2) is treated as a constant flat rate, using the SOFR overnight rate observed on that date. The dividend-yield proxy, SPX physical drift and model-specific inputs are supplied by the calibration layer. The SPX physical drift is estimated from the 252 trading days ending at the calibration date and is kept common across BSM, Merton and Heston, while volatility, jump and variance parameters remain model-specific.

2.2 DOOM-Put Portfolio

The strategy combines the Index portfolio with European put options based on the same index S_t. For each contract j, let K_j denote its strike, T_j its maturity, π_j its initial premium and $V_j(t)$ its value before expiry. A static allocation is represented by a vector of budget weights

$$\mathbf{w} = (w_1, \dots, w_m), \qquad w_j \geq 0, \qquad \sum_{j=1}^{m} w_j \leq b_{\max}, \tag{2.3}$$

where w_j is the fraction of initial wealth spent on the premium of contract j. In the simulations, the maximum option budget is $b_{\max} = 5\%$ over the one-year investment horizon, with a budget-sensitivity analysis reported in Section C.4. This is an upper bound rather than a target allocation. It is large enough to allow meaningful crash insurance, but small enough to keep the strategy as an equity overlay rather than an option-dominated allocation.

The number of contracts purchased is therefore

$$n_j = \frac{w_j B_0}{\pi_j}. \tag{2.4}$$

The option premia are financed by reducing the initial exposure to the index. The number of index units held after buying protection is

$$n_S(\mathbf{w}) = \frac{\left(1 - \sum_{j=1}^{m} w_j\right) B_0}{S_0}. \tag{2.5}$$

Thus, the DOOM position is not a leveraged overlay. Each dollar spent on protection is removed from the long exposure to the index.

Before its own maturity T_j, each option is valued with the corresponding risk-neutral pricing model. If $T_j < T$, the payoff $(K_j - S_{T_j})^+$ is realized at T_j and reinvested in the index until the portfolio horizon T. The value of the DOOM-put portfolio can therefore be written as

$$P_t^D(\mathbf{w}) = n_S(\mathbf{w}) S_t + \sum_{j | T_j > t} n_j V_j(t) + \sum_{j | T_j \leq t} n_j \left(K_j - S_{T_j}\right)^+ \frac{S_t}{S_{T_j}}. \tag{2.6}$$

2.3 Probabilistic DOOM-Put Definition

Deep Out-of-the-Money Put

A European put option gives its holder the right, but not the obligation, to sell the underlying asset at a pre-specified strike K at maturity T (Hull, 2018). Its payoff at maturity is

$$(K - S_T)^+ = \max(K - S_T, 0) = \begin{cases} K - S_T, & \text{if } S_T < K, \\ 0, & \text{if } S_T \geq K. \end{cases} \tag{2.7}$$

In this thesis, a *deep out-of-the-money (DOOM) put* is a European put whose strike is selected so that, for a given risk-neutral pricing specification $\mathscr{M}$,

$$\mathbb{Q}^{\mathscr{M}}(S_T < K) = p, \tag{2.8}$$

where p is the target probability under the risk-neutral measure.

European put options are standard instruments for downside protection, but they can also be used to take directional exposure to a decline in the underlying asset. The defining feature of the DOOM contracts considered here is that they are not specified by a fixed moneyness level, but by a fixed risk-neutral exercise probability p_j. For a contract j, the maturity T_j, strike K_j and probability label p_j are linked by Equation (2.8). Hence, K_j is the risk-neutral p_j-quantile of the terminal index level S_{T_j} under the pricing specification used to construct the contract. The probability label p_j is therefore a pricing convention, not the investor's physical crash probability. It is a model-implied probability under $\mathbb{Q}^{\mathscr{M}}$. The hedge is evaluated later under $\mathbb{P}$.

This definition is key to the model-based comparisons. A put associated with the same risk-neutral tail probability does not generally have the same strike under Black-Scholes-Merton, Merton jump-diffusion and Heston stochastic volatility, because their fitted risk-neutral terminal distributions differ. Defining contracts by p_j, rather than directly by K_j/S_0, therefore ensures that the comparison is based on the same risk-neutral tail probability, rather than on an arbitrary strike level. This property however only applies to model-implied contracts. A listed option has one observed strike and, when quoted, one market premium. Differences across models therefore refer to model prices, either because the fit to market prices is imperfect or because the contract is constructed inside the model-based universe.

The model-based construction can be viewed as a mapping from a tail probability to a strike, and from that strike to a premium. Once K_j has been selected, the corresponding model premium is

$$\Pi_j^{\mathscr{M}} = P^{\mathscr{M}}\left(S_0, K_j, T_j, r, q; \Theta_{\mathscr{M}}^{\mathbb{Q}}\right). \tag{2.9}$$

The same fitted risk-neutral specification is used to select the strike and to price the put.

In constructing the admissible contract set, the strike is restricted to the range

$$K_j/S_0 \in [0.30, 0.90]. \tag{2.10}$$

If the target probability is not attainable inside this range, the contract is excluded from the grid rather than clipped to a boundary. This rule preserves the economic interpretation of p_j. A contract labelled as a 1% DOOM put must effectively correspond to a 1% risk-neutral left-tail event under the pricing specification used to construct it.

2.4 Physical and Risk-Neutral Measures

The strategy relies on a strict separation between the physical measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}^{\mathscr{M}}$. The physical measure describes the distribution of economic scenarios in which portfolio outcomes are evaluated. The risk-neutral measure describes, for a given pricing specification $\mathscr{M}$, the option premia, mark-to-market values and tail probabilities used to construct DOOM strikes.

Index paths are simulated under $\mathbb{P}$, using the physical price drift $\mu^{\mathbb{P}}$. By contrast, initial option premia, model-implied option values and probabilities such as $\mathbb{Q}^{\mathscr{M}}(S_T < K)$ are computed under the corresponding risk-neutral pricing specification.

The economic point is that reducing losses in simulated scenarios is not the same as buying protection at an attractive market price. An option may reduce losses under the physical scenarios generated by $\mathbb{P}$, while still being unattractive if its premium under $\mathbb{Q}^{\mathscr{M}}$ embeds a high price for tail protection. This distinction is relevant for crash insurance, since option-implied tail probabilities and premia may contain compensation for rare-event risk that differs from the corresponding physical tail distribution (Bollerslev & Todorov, 2011). The strategy must therefore be evaluated jointly on the protection delivered under $\mathbb{P}$ and on the cost implied by risk-neutral option premia.

2.5 Risk and Cost Metrics

All risk metrics are computed from discounted normalised wealth, defined in Equation (2.2). With the same initial wealth, the hedged and unhedged strategies are directly comparable. For each simulated path i, terminal performance is measured by

$$R_T^{(i)} = W_T^{(i)} - 1. \tag{2.11}$$

Terminal metrics such as VaR and ES describe the distribution of final discounted excess returns across simulated paths.

Let $L = -R$ denote the loss associated with a return observation. In this thesis, α denotes a lower-tail probability for returns. The corresponding confidence level for losses is $1 - \alpha$. VaR and ES are defined as

$$\mathrm{VaR}_{1-\alpha} = q_{1-\alpha}(L) = -q_\alpha(R), \tag{2.12}$$

$$\begin{aligned}\mathrm{ES}_{1-\alpha} &= \mathbb{E}[L \mid L \geq q_{1-\alpha}(L)] \\ &= -\mathbb{E}[R \mid R \leq q_\alpha(R)].\end{aligned} \tag{2.13}$$

Here $q_u(\cdot)$ denotes the quantile at probability level u. For the simulated sample $\{R_i\}_{i=1}^N$, the reported ES estimate is

$$\widehat{\mathrm{ES}}_{1-\alpha}(R) = -\frac{1}{N_\alpha} \sum_{i|R_i \leq \hat{q}_\alpha(R)} R_i, \tag{2.14}$$

where N_α is the number of observations satisfying $R_i \leq \hat{q}_\alpha(R)$. Reported VaR and ES values are floored at zero in the implementation. Higher reported values indicate larger downside losses.

The baseline tail threshold is $\alpha_{\text{base}} = 5\%$, and the deep-tail level is $\alpha_{\text{deep}} = 1\%$. The analysis reports terminal VaR and ES at both the 95% and 99% confidence levels. Expected Shortfall at the 95% confidence level is the benchmark because it measures the average loss beyond the VaR threshold, whereas VaR alone does not capture the severity of losses beyond this threshold and is not coherent in full generality (Artzner et al., 1999; Rockafellar & Uryasev, 2000). This distinction is especially important for DOOM puts because their value is concentrated in rare left-tail outcomes.

Distributional diagnostics are also computed from terminal returns. The sample mean terminal return is

$$\bar{R}_T = \frac{1}{N} \sum_{i=1}^{N} R_T^{(i)}, \tag{2.15}$$

and complements VaR and ES by describing the return term sacrificed or gained when protection is carried.

Skewness and excess kurtosis are reported to describe the shape of the simulated terminal return distribution. In the implementation, their theoretical values are estimated with finite-sample corrected estimators. The theoretical skewness of terminal returns is defined as

$$\gamma_1 = \frac{\mathbb{E}\left[(R_T - \mu_R)^3\right]}{\sigma_R^3}, \tag{2.16}$$

where $\mu_R = \mathbb{E}[R_T]$ and σ_R is the standard deviation of R_T. It measures the asymmetry of terminal returns. Negative skewness indicates a longer or heavier left tail of returns, whereas positive skewness indicates a longer or heavier right tail.

Excess kurtosis measures tail thickness rather than asymmetry. The theoretical excess kurtosis is defined as

$$\gamma_2 = \frac{\mathbb{E}\left[(R_T - \mu_R)^4\right]}{\sigma_R^4} - 3. \tag{2.17}$$

This quantity is measured relative to the Gaussian benchmark, so zero corresponds to a normal distribution. Positive values indicate fatter tails or stronger concentration of extreme outcomes, while negative values can occur when the hedge imposes a partial floor on losses.

Maximum drawdown is retained as the sole path-dependent risk metric. Let

$$\bar{W}_t^{(i)} = \max_{0 \le s \le t} W_s^{(i)} \tag{2.18}$$

be the running maximum of discounted wealth. The drawdown of path i at time t is

$$D_t^{(i)} = 1 - \frac{W_t^{(i)}}{\bar{W}_t^{(i)}}, \tag{2.19}$$

and the maximum drawdown of that path is

$$\mathrm{MDD}^{(i)} = \max_{0 \le t \le T} D_t^{(i)} = 1 - \min_{0 \le t \le T} \frac{W_t^{(i)}}{\bar{W}_t^{(i)}}. \tag{2.20}$$

Drawdowns are measured as positive losses relative to the previous wealth peak. Thus $\mathrm{MDD}^{(i)} \ge 0$, and a larger value represents a larger peak-to-trough loss along path i. The reported metric is the average of these path-level maximum drawdowns:

$$L_{\mathrm{MDD}} = \frac{1}{N} \sum_{i=1}^{N} \mathrm{MDD}^{(i)}. \tag{2.21}$$

Hence, L_{MDD} should be read as the average maximum drawdown loss across simulated paths. This distinction matters because drawdown is path-dependent and captures intra-horizon losses that terminal risk measures cannot represent (Vecer, 2007).

Cost metrics compare the DOOM-put portfolio with the Index portfolio on the same evaluation paths. Let $\bar{R}_D$ and $\bar{R}_I$ denote their sample mean terminal returns. The mean return drag is

$$\Delta\bar{R} = \bar{R}_I - \bar{R}_D. \tag{2.22}$$

Mean return drag is the difference between the mean terminal returns of the two portfolios. It includes option payoffs and the lower index exposure, not only the initial premium.

CHAPTER 3

Models, Pricing and Calibration

3.1 Measure Conventions

This chapter presents the three models used to construct and price DOOM puts under Black-Scholes-Merton, Merton jump-diffusion and Heston stochastic volatility. For each model, it gives the physical dynamics, the risk-neutral pricing dynamics and the method used to recover the strike associated with a target probability p. Technical pricing details are reported in Appendix A, while calibration diagnostics are reported in Appendix B.

The notation follows Chapter 2. The index level and initial wealth are normalised to $S_0 = B_0 = 100$. Interest rates, dividends and model parameters are taken from the calibration results described in Section 3.5. For each calibration date, the physical drift $\mu^{\mathbb{P}}$ is estimated once, kept constant over the one-year horizon, and used across BSM, Merton and Heston. The remaining parameters are model-specific.

The relation between $\mathbb{P}$ and $\mathbb{Q}$ follows the standard Girsanov change-of-measure framework (Shreve, 2004). It modifies drift terms and risk premia, while preserving the diffusion coefficients shared by the two measures. In BSM, the volatility σ is common across measures and only the index drift changes. In Merton, the diffusion volatility σ_{diff} is common, while the jump arrival intensity and jump-size parameters may differ between $\mathbb{P}$ and $\mathbb{Q}$ through the jump risk premium. In Heston, v_0, ξ and ρ are common, while the variance drift parameters κ and θ are linked across measures by the variance risk premium.

3.2 Black-Scholes-Merton

3.2.1 Dynamics and Terminal Distribution

The Black-Scholes-Merton model assumes that the index follows a geometric Brownian motion with constant volatility. The pricing argument and differential equation are given by Black and Scholes (1973). The continuous-dividend version follows Merton (1973). Under a measure $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$,

$$\frac{\mathrm{d}S_t}{S_t} = \mu^{\mathbb{M}}\,\mathrm{d}t + \sigma\,\mathrm{d}W_t^{\mathbb{M}}, \tag{3.1}$$

where $\sigma > 0$ is constant. Under $\mathbb{Q}$, $\mu^{\mathbb{Q}} = r - q$, while under $\mathbb{P}$ the price drift is $\mu^{\mathbb{P}}$. Solving Equation (3.1) gives

$$S_T = S_0 \exp\left[\left(\mu^{\mathbb{M}} - \frac{1}{2}\sigma^2\right)T + \sigma W_T^{\mathbb{M}}\right]. \tag{3.2}$$

Equivalently, since

$$Z^{\mathbb{M}} := \frac{W_T^{\mathbb{M}}}{\sqrt{T}} \overset{\mathbb{M}}{\sim} \mathcal{N}(0,1), \tag{3.3}$$

the terminal index level can be written as

$$S_T = S_0 \exp\left[\left(\mu^{\mathbb{M}} - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}\,Z^{\mathbb{M}}\right]. \tag{3.4}$$

Hence, under the risk-neutral measure,

$$\log\left(\frac{S_T}{S_0}\right) \overset{\mathbb{Q}}{\sim} \mathcal{N}\left(\left(r - q - \frac{1}{2}\sigma^2\right)T,\ \sigma^2 T\right). \tag{3.5}$$

The model has one source of randomness and one constant volatility parameter. This gives a closed-form terminal distribution under $\mathbb{Q}$. Once σ is fixed, the risk-neutral law of S_T is fully determined. The same volatility parameter drives the exercise probability, the DOOM strike and the option premium.

3.2.2 Put Pricing and DOOM Strike

For a European put with strike K and maturity T, the Black-Scholes-Merton price is

$$P_{\text{BSM}}(S_0, K, T, r, q, \sigma) = Ke^{-rT}\Phi(-d_2) - S_0 e^{-qT}\Phi(-d_1), \tag{3.6}$$

where

$$d_1 = \frac{\log(S_0/K) + (r - q + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \tag{3.7}$$

$$d_2 = d_1 - \sigma\sqrt{T} = \frac{\log(S_0/K) + (r - q - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}. \tag{3.8}$$

The term $\Phi(-d_2)$ is the risk-neutral probability that the put finishes in the money,

$$\mathbb{Q}(S_T < K) = \Phi(-d_2). \tag{3.9}$$

This identity is especially useful for DOOM puts because it gives the strike directly. Define $z_p := \Phi^{-1}(p)$. For a target probability p,

$$K_{\text{BSM}}(p, T) = S_0 \exp\left[\left(r - q - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}\,z_p\right]. \tag{3.10}$$

This expression is the p-quantile of the risk-neutral terminal index level. It has the same form as the BSM solution for S_T, with the random normal shock replaced by the deterministic quantile z_p.

A lower target probability implies a lower value of z_p, and therefore a strike further below spot. The same volatility parameter σ determines both the DOOM strike and the option premium.

3.2.3 Tail-Protection Implications

The Black-Scholes-Merton model is useful as a benchmark because it is restrictive. It produces continuous paths, no jumps, constant volatility and a flat implied-volatility surface. It provides a reference case in which downside protection is priced only through lognormal diffusion risk.

This restriction matters for DOOM puts. Equity-index option markets usually have a downside volatility skew, with deep out-of-the-money puts priced at higher implied volatilities than at-the-money options. BSM cannot reproduce this skew because the same volatility parameter is used for all strikes. As a result, it can underprice deep out-of-the-money puts when market prices include downside skew or a tail-risk premium. It can also understate the probability of large losses if the true left tail is heavier than the lognormal distribution. The later models remove some of these restrictions by introducing jump risk or stochastic volatility.

3.3 Merton Jump-Diffusion

3.3.1 Dynamics with Discontinuous Returns

The Merton jump-diffusion model extends the lognormal diffusion by adding a compound Poisson jump component (Merton, 1976). Under a measure $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$,

$$\frac{\mathrm{d}S_t}{S_{t-}} = \left(\mu^{\mathbb{M}} - \lambda^{\mathbb{M}}\bar{\kappa}^{\mathbb{M}}\right)\mathrm{d}t + \sigma_{\text{diff}}\,\mathrm{d}W_t^{\mathbb{M}} + \left(e^{J^{\mathbb{M}}} - 1\right)\mathrm{d}N_t^{\mathbb{M}}. \tag{3.11}$$

Here $N_t^{\mathbb{M}}$ is the jump-counting process, with jump arrival intensity $\lambda^{\mathbb{M}}$. The parameter σ_{diff} is the diffusion volatility. At each jump time, the log-jump size is denoted $J_\ell^{\mathbb{M}}$, with

$$J_\ell^{\mathbb{M}} \overset{\mathbb{M}}{\sim} \mathcal{N}\left(\mu_J^{\mathbb{M}}, (\sigma_J^{\mathbb{M}})^2\right), \qquad \ell = 1, 2, \ldots \tag{3.12}$$

and the jump sizes are independent across arrivals. The diffusion volatility σ_{diff} is common across measures. The jump arrival intensity and jump-size parameters may differ between $\mathbb{P}$ and $\mathbb{Q}$ through the jump risk premium.

The expected relative jump size under $\mathbb{M}$ is

$$\bar{\kappa}^{\mathbb{M}} = \mathbb{E}^{\mathbb{M}}\left[e^{J^{\mathbb{M}}} - 1\right] = \exp\left(\mu_J^{\mathbb{M}} + \frac{1}{2}(\sigma_J^{\mathbb{M}})^2\right) - 1. \tag{3.13}$$

The compensator $\lambda^{\mathbb{M}}\bar{\kappa}^{\mathbb{M}}$ uses the jump law under the same measure. Under $\mathbb{Q}$, $\mu^{\mathbb{Q}} = r - q$, while under $\mathbb{P}$, the price drift is $\mu^{\mathbb{P}}$.

The terminal log-return is

$$\log\left(\frac{S_T}{S_0}\right) = \left(\mu^{\mathbb{M}} - \lambda^{\mathbb{M}}\bar{\kappa}^{\mathbb{M}} - \frac{1}{2}\sigma_{\text{diff}}^2\right)T + \sigma_{\text{diff}}\sqrt{T}\,Z^{\mathbb{M}} + \sum_{\ell=1}^{N_T^{\mathbb{M}}} J_\ell^{\mathbb{M}}, \tag{3.14}$$

where

$$Z^{\mathbb{M}} = \frac{W_T^{\mathbb{M}}}{\sqrt{T}} \overset{\mathbb{M}}{\sim} \mathcal{N}(0,1). \tag{3.15}$$

For pricing, set $\mathbb{M} = \mathbb{Q}$. Conditional on $N_T^{\mathbb{Q}} = n$, the risk-neutral log-return is Gaussian,

$$\log\left(\frac{S_T}{S_0}\right) \mid N_T^{\mathbb{Q}} = n \overset{\mathbb{Q}}{\sim} \mathcal{N}(m_n^{\mathbb{Q}}, (v_n^{\mathbb{Q}})^2), \tag{3.16}$$

with

$$m_n^{\mathbb{Q}} = \left(r - q - \lambda^{\mathbb{Q}}\bar{\kappa}^{\mathbb{Q}} - \frac{1}{2}\sigma_{\text{diff}}^2\right)T + n\mu_J^{\mathbb{Q}}, \tag{3.17}$$

$$(v_n^{\mathbb{Q}})^2 = \sigma_{\text{diff}}^2 T + n(\sigma_J^{\mathbb{Q}})^2. \tag{3.18}$$

This conditional representation in Equation (3.16) follows the compound-Poisson structure of Merton (1976). The numerical truncation of the Poisson mixture is reported in Appendix A.

3.3.2 Tail Probability and Strike Inversion

For pricing and strike selection, the risk-neutral parameters are used. From Equation (3.16), the conditional exercise probability is

$$\mathbb{Q}(S_T < K \mid N_T^{\mathbb{Q}} = n) = \Phi\left(\frac{\log(K/S_0) - m_n^{\mathbb{Q}}}{v_n^{\mathbb{Q}}}\right). \tag{3.19}$$

Since $N_T^{\mathbb{Q}} \overset{\mathbb{Q}}{\sim} \text{Poisson}(\lambda^{\mathbb{Q}}T)$, the unconditional exercise probability is

$$\mathbb{Q}(S_T < K) = \sum_{n=0}^{\infty} e^{-\lambda^{\mathbb{Q}}T}\frac{(\lambda^{\mathbb{Q}}T)^n}{n!}\Phi\left(\frac{\log(K/S_0) - m_n^{\mathbb{Q}}}{v_n^{\mathbb{Q}}}\right). \tag{3.20}$$

For a target probability p, the DOOM strike is defined by $\mathbb{Q}(S_T < K_{\text{MJD}}(p,T)) = p$. It is therefore the solution of

$$\sum_{n=0}^{\infty} e^{-\lambda^{\mathbb{Q}}T}\frac{(\lambda^{\mathbb{Q}}T)^n}{n!}\Phi\left(\frac{\log(K_{\text{MJD}}(p,T)/S_0) - m_n^{\mathbb{Q}}}{v_n^{\mathbb{Q}}}\right) = p. \tag{3.21}$$

Unlike BSM, this quantile has no closed-form expression. In the implementation, the probability function in Equation (3.20) is evaluated under $\mathbb{Q}$, and $K_{\text{MJD}}(p,T)$ is recovered by one-dimensional root search over the admissible strike interval. The root search returns $K_{\text{MJD}}(p,T)$, the strike for which the model-implied probability $\mathbb{Q}(S_T < K_{\text{MJD}}(p,T))$ equals the target probability p.

In the Merton model, a large loss can occur through a jump rather than through a large diffusive move. A strike that appears out-of-the-money under BSM can have a higher risk-neutral probability once jump arrivals are included.

3.3.3 *Option Price as a Mixture*

The put price can be represented as a weighted sum of Black-Scholes-Merton prices with jump-count dependent parameters. A standard representation is

$$P_{\text{MJD}} = \sum_{n=0}^{\infty} e^{-\bar{\lambda}T} \frac{(\bar{\lambda}T)^n}{n!} P_{\text{BSM}}(S_0, K, T, r_n, q, \sigma_n), \tag{3.22}$$

where

$$\bar{\lambda} = \lambda^{\mathbb{Q}} \exp\left(\mu_J^{\mathbb{Q}} + \frac{1}{2}(\sigma_J^{\mathbb{Q}})^2\right), \tag{3.23}$$

$$\sigma_n^2 = \sigma_{\text{diff}}^2 + \frac{n(\sigma_J^{\mathbb{Q}})^2}{T}, \tag{3.24}$$

$$r_n = r - \lambda^{\mathbb{Q}}\bar{\kappa}^{\mathbb{Q}} + \frac{n(\mu_J^{\mathbb{Q}} + \frac{1}{2}(\sigma_J^{\mathbb{Q}})^2)}{T}. \tag{3.25}$$

The probability mixture in Equation (3.20) and the price mixture in Equation (3.22) are related but not identical. The exercise probability uses the jump-count distribution under $\mathbb{Q}$,

$$N_T^{\mathbb{Q}} \overset{\mathbb{Q}}{\sim} \text{Poisson}(\lambda^{\mathbb{Q}}T). \tag{3.26}$$

The price series uses the adjusted intensity $\bar{\lambda}$. Conditional on n jumps, the sum of lognormal jump sizes changes the mean and variance of the log-return, so each term can be written as a Black-Scholes-Merton price with adjusted parameters r_n and σ_n. This is the lognormal-jump mixture formula of Merton (1976).

3.3.4 *Tail-Protection Implications*

The Merton model introduces a crash dynamic. For the same moneyness, the risk-neutral probability of exercise can be higher than in a pure diffusion. For a given target probability p, the DOOM strike can differ from the Black-Scholes-Merton strike because the left tail has a different shape. The premium also changes because the put is exposed to diffusion volatility, jump risk and the jump risk premium embedded in $\mathbb{Q}$.

This distinction matters because the jump law used for pricing is not necessarily the same as the jump law used for physical scenarios. Under $\mathbb{Q}$, option prices and DOOM strikes depend on the risk-neutral jump arrival intensity and jump-size distribution. Under $\mathbb{P}$, portfolio losses depend on the physical jump distribution. A Merton put can therefore appear expensive relative to a diffusion benchmark because it prices crash states and jump risk premia that BSM does not include. Conversely, if risk-neutral jump risk is high relative to the physical jump risk used in the simulations, the strategy may pay too much for crash states that occur rarely under $\mathbb{P}$. The cost-benefit trade-off of DOOM protection is linked both to jump risk itself and to the gap between physical and risk-neutral jump risk.

3.4 Heston Stochastic Volatility

3.4.1 *Variance Dynamics and Risk Premium*

The Heston model keeps index paths continuous but replaces constant volatility with a stochastic variance process. The stochastic-variance dynamics, pricing PDE and characteristic-function solution are developed by Heston (1993). Under a measure $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$,

$$\frac{\mathrm{d}S_t}{S_t} = \mu^{\mathbb{M}}\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^{S,\mathbb{M}}, \tag{3.27}$$

$$\mathrm{d}v_t = \kappa^{\mathbb{M}}(\theta^{\mathbb{M}} - v_t)\mathrm{d}t + \xi\sqrt{v_t}\,\mathrm{d}W_t^{v,\mathbb{M}}, \tag{3.28}$$

$$\mathrm{d}\langle W^{S,\mathbb{M}}, W^{v,\mathbb{M}}\rangle_t = \rho\,\mathrm{d}t. \tag{3.29}$$

The parameter κ controls the speed of mean reversion, θ is the long-run variance, ξ is the volatility of variance, and ρ is the instantaneous correlation between index shocks and variance shocks. The initial variance is denoted v_0. The coefficients v_0, ξ and ρ are common under $\mathbb{P}$ and $\mathbb{Q}$. The variance drift parameters may differ across measures through a variance risk premium. In the implementation, this premium is specified through the function (Cheridito et al., 2007)

$$\lambda_v(v_t) = \frac{\lambda_0}{\sqrt{v_t}} + \lambda_1\sqrt{v_t}. \tag{3.30}$$

Combined with

$$\mathrm{d}W_t^{v,\mathbb{Q}} = \mathrm{d}W_t^{v,\mathbb{P}} + \lambda_v(v_t)\,\mathrm{d}t, \tag{3.31}$$

this gives

$$\kappa^{\mathbb{Q}} = \kappa^{\mathbb{P}} + \xi\lambda_1, \tag{3.32}$$

$$\kappa^{\mathbb{Q}}\theta^{\mathbb{Q}} = \kappa^{\mathbb{P}}\theta^{\mathbb{P}} - \xi\lambda_0. \tag{3.33}$$

In the implementation, $(\kappa^{\mathbb{P}}, \theta^{\mathbb{P}})$ are estimated from historical returns and $(\kappa^{\mathbb{Q}}, \theta^{\mathbb{Q}})$ are calibrated separately from option prices. The coefficients λ_0 and λ_1 are recovered from the difference between the two variance dynamics. The variance risk premium is not an additional calibration constraint.

For equity-index protection, ρ matters. A negative correlation means that market declines tend to coincide with rising variance. This increases put values in downside states and contributes to the downward slope of the equity-index implied-volatility curve. The left tail is affected without jumps. The index remains continuous, but negative price moves tend to occur in high-variance states.

The variance process is of square-root type. A sufficient condition for strict positivity is the Feller condition

$$2\kappa^{\mathbb{M}}\theta^{\mathbb{M}} \geq \xi^2. \tag{3.34}$$

This condition is used as an interpretative diagnostic for the variance process. It comes from the square-root diffusion studied by Cox et al. (1985).

3.4.2 Characteristic Function and Pricing

The Heston model does not provide a closed-form density for S_T. The distribution and option prices are recovered from the characteristic function of the terminal log-price under $\mathbb{Q}$,

$$\phi(u,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{iu\log S_T}\right]. \tag{3.35}$$

This characteristic function can be inverted to recover the distribution of S_T, and therefore the risk-neutral tail probabilities used to define DOOM strikes. Option prices are obtained from the same transform by applying Fourier-pricing formulas to the payoff. The transform used in the implementation is kept in the main text. It follows the affine solution in Heston (1993). In the pricing expressions below, $\kappa = \kappa^{\mathbb{Q}}$ and $\theta = \theta^{\mathbb{Q}}$:

$$\phi(u,T) = \exp\left(C(u,T) + D(u,T)v_0 + iu\log S_0\right), \tag{3.36}$$

$$d(u) = \sqrt{(\kappa - \rho\xi iu)^2 + \xi^2(u^2 + iu)}, \tag{3.37}$$

$$g(u) = \frac{\kappa - \rho\xi iu - d(u)}{\kappa - \rho\xi iu + d(u)}, \tag{3.38}$$

$$C(u,T) = iu(r-q)T + \frac{\kappa\theta}{\xi^2}\left[(\kappa - \rho\xi iu - d(u))T - 2\log\left(\frac{1 - g(u)e^{-d(u)T}}{1 - g(u)}\right)\right], \tag{3.39}$$

$$D(u,T) = \frac{\kappa - \rho\xi iu - d(u)}{\xi^2} \times \frac{1 - e^{-d(u)T}}{1 - g(u)e^{-d(u)T}}. \tag{3.40}$$

The terms $d(u)$ and $g(u)$ come from the solution of the Riccati system. The function $C(u,T)$ contains the part of the transform that does not depend on the initial variance, including the interest-rate and dividend term $iu(r-q)T$ and the long-run variance term. The function $D(u,T)$ multiplies v_0, and determines how the initial variance affects the terminal log-price distribution. The logarithm in $C(u,T)$ is complex-valued. Its branch must be chosen consistently to avoid discontinuities in the characteristic function, as discussed by Lord and Kahl (2010).

European prices are then obtained by Fourier inversion. Following Carr and Madan (1999), let $k = \log K$ and let $C^{\text{call}}(k,T)$ denote the call price as a function of log-strike. For $\alpha > 0$, define the modified call price

$$c(k,T) = e^{\alpha k}C^{\text{call}}(k,T). \tag{3.41}$$

Its Fourier transform is

$$\psi(v,T) = \frac{e^{-rT}\phi(v - i(\alpha+1), T)}{\alpha^2 + \alpha - v^2 + i(2\alpha+1)v} \tag{3.42}$$

where $\phi(\cdot, T)$ is the characteristic function defined in Equation (3.35). The call price is recovered as

$$C^{\text{call}}(k,T) = \frac{e^{-\alpha k}}{\pi}\int_0^\infty \Re\left[e^{-ivk}\psi(v,T)\right] \mathrm{d}v. \tag{3.43}$$

The coefficient α is not a Heston parameter. It is introduced only for Fourier pricing and makes the modified call price $c(k,T)$ integrable. The factor e^{-ivk} is the inverse-Fourier term associated with $k = \log K$, while e^{-rT} discounts the payoff expectation. This transform-pricing formula follows Carr and Madan (1999). At inception, with $k = \log K$, put prices follow from European put-call parity,

$$P(K,T) = C^{\text{call}}(k,T) - S_0 e^{-qT} + Ke^{-rT}. \tag{3.44}$$

This is the standard European put-call parity with continuous dividend yield (Merton, 1973).

3.4.3 Tail Probability and DOOM Strike

The DOOM strike requires the distribution function of S_T. Applying the Gil-Pelaez inversion theorem (Gil-Pelaez, 1951) to $X = \log S_T$ gives

$$F_{S_T}^{\mathbb{Q}}(K) = \mathbb{Q}(S_T < K) = \frac{1}{2} - \frac{1}{\pi}\int_0^\infty \Im\left[\frac{e^{-iu\log K}\phi(u,T)}{u}\right] \mathrm{d}u. \tag{3.45}$$

The Heston DOOM strike $K_{\mathrm{H}}(p,T)$ is then defined by

$$F_{S_T}^{\mathbb{Q}}(K_{\mathrm{H}}(p,T)) = p. \tag{3.46}$$

Equivalently,

$$\frac{1}{2} - \frac{1}{\pi}\int_0^\infty \Im\left[\frac{e^{-iu\log K_{\mathrm{H}}(p,T)}\phi(u,T)}{u}\right] \mathrm{d}u = p. \tag{3.47}$$

The integral in Equation (3.45) is evaluated numerically. The strike is then recovered by one-dimensional root search over the admissible strike interval. The numerical inversion and root-search conventions used with Equations (3.45) to (3.47) are reported in Appendix A.

For physical scenarios, the variance process is simulated with the Quadratic-Exponential discretisation proposed by L. Andersen (2008). The implementation conventions are reported in Appendix A because the scheme is a numerical approximation to the continuous-time variance process rather than a separate economic assumption.

3.4.4 Tail-Protection Implications

The Heston dynamic differs from Merton in an important way. In this model, crash protection can become costly even without discontinuous jumps, because market declines are associated with higher variance. A DOOM put is exposed not only to the terminal index level, but also to the volatility state observed in severe equity losses.

A jump model concentrates crash risk in rare discontinuous drops. A stochastic-volatility model allows the index to decline continuously while variance rises, so downside option values can increase before maturity. These mechanisms can produce similar initial premia, but they imply different mark-to-market behaviour along the path.

3.5 Calibration

This section describes the calibration procedure used to estimate model parameters from historical and market inputs. These parameters are then used for simulation, option pricing and portfolio construction. For each model and calibration date, historical SPX returns are used to estimate the physical dynamics under $\mathbb{P}$. These dynamics are then used to simulate one-year index paths. The dated SPX option surface is fitted separately under $\mathbb{Q}$. The simulated index values and the corresponding option values are then combined to compute strategy returns and risk measures. The resulting model-based allocations are later tested on historical SPX return windows.

The analysis is implemented on three dated SPX option surfaces, 16 March 2020, 8 April 2025 and 30 June 2026. These dates correspond to a crisis date, a stress date and a normal market date, respectively. Each date is treated as a separate calibration experiment. The surfaces are not pooled.

3.5.1 Common Calibration Inputs

The calibration has two tasks. The first task is historical and constructs the physical dynamics of the underlying index under $\mathbb{P}$. These parameters are used to generate scenarios, compute portfolio losses and evaluate how the hedge changes downside outcomes across simulated states. The second task is pricing and fits the dated option surface under $\mathbb{Q}$. These parameters are used to convert a DOOM probability into a strike, price the puts and compute model-implied option values along each simulated path. The two tasks are kept separate. A parameter under $\mathbb{P}$ describes the scenario law used for portfolio evaluation, whereas a parameter under $\mathbb{Q}$ is a pricing parameter that embeds risk premia and state prices.

Physical parameters are estimated from S&P 500 returns available up to the corresponding calibration date. Let $x_i = \log(S_i/S_{i-1})$, $\Delta = 1/252$. The common physical price drift uses the most recent $N = 252$ daily returns and is

$$\widehat{\mu}^{\mathbb{P},\mathrm{hist}} = \frac{1}{\Delta} \log\left(\frac{1}{N}\sum_{i=1}^{N} e^{x_i}\right). \tag{3.48}$$

This drift is estimated once for each calibration date and used in the physical simulation of BSM, Merton and Heston.

Risk-neutral calibration is implemented as a price-space least-squares fit. In schematic form, the parameter vector solves

$$\widehat{\Theta}_t^{\mathbb{Q}} = \arg\min_{\Theta^{\mathbb{Q}} \in \mathcal{A}} \frac{1}{|\mathcal{S}_t|} \sum_{j \in \mathcal{S}_t} \omega_j \left[P_j^{\text{model}}(\Theta^{\mathbb{Q}}) - P_j^{\text{mkt}} \right]^2. \tag{3.49}$$

Here $\mathcal{S}_t$ denotes the filtered option surface available at calibration date t, and $\mathcal{A}$ denotes the admissible parameter set. This price-space objective is in line with empirical option-model comparisons such as Bakshi et al. (1997). The weights ω_j account for bid-ask spreads and maturity spacing. Their construction is reported in Appendix B. During surface calibration, the dividend yield is estimated by expiry from put-call parity. The reported dividend yield is the average of the retained expiry estimates and is used for contract construction and pricing. The filtered option surface mainly retains out-of-the-money contracts, with puts selected below the forward and calls selected above the forward. The exact moneyness thresholds are reported in Appendix B.

RMSE and MAE are reported in Appendix B to measure option-surface fit. They are not hedging-performance measures. A model that fits option surfaces does not necessarily reproduce the physical dynamics of the underlying index or deliver better hedging performance (Lassance & Vrins, 2018). The calibrated parameter values, robustness checks and additional calibration diagnostics are reported in Appendix B.

3.5.2 *Model-Specific Calibration*

For BSM, the dated option surface estimates one volatility parameter σ through Equation (3.49).

Under Merton, the dated option surface estimates the risk-neutral parameter vector

$$\Theta_{\text{MJD}}^{\mathbb{Q}} = \left(\sigma_{\text{diff}}, \lambda^{\mathbb{Q}}, \mu_J^{\mathbb{Q}}, \sigma_J^{\mathbb{Q}} \right). \tag{3.50}$$

The estimate is obtained by applying the price-space objective in Equation (3.49) to the Merton pricing formula in Equation (3.22). The fitted diffusion volatility σ_{diff} is kept common across $\mathbb{P}$ and $\mathbb{Q}$.

The physical jump parameters $(\lambda^{\mathbb{P}}, \mu_J^{\mathbb{P}}, \sigma_J^{\mathbb{P}})$ are estimated from the most recent $N = 252$ daily returns available at the calibration date. In this time-series estimation, σ_{diff} is fixed at its risk-neutral estimate and $\mu^{\mathbb{P}}$ is fixed at the common drift estimate in Equation (3.48). If x_i denotes a daily log return, its physical density is the Poisson-normal mixture

$$f^{\mathbb{P}}(x_i) = \sum_{n=0}^{\infty} e^{-\lambda^{\mathbb{P}}\Delta} \frac{(\lambda^{\mathbb{P}}\Delta)^n}{n!} \varphi\left(x_i;\, m_{n,\Delta}^{\mathbb{P}},\, (s_{n,\Delta}^{\mathbb{P}})^2\right), \tag{3.51}$$

where $\varphi(\cdot; m, s^2)$ denotes the Gaussian density with mean m and variance s^2, and

$$m_{n,\Delta}^{\mathbb{P}} = \left(\mu^{\mathbb{P}} - \lambda^{\mathbb{P}}\bar{\kappa}^{\mathbb{P}} - \frac{1}{2}\sigma_{\mathrm{diff}}^2\right)\Delta + n\mu_J^{\mathbb{P}}, \tag{3.52}$$

$$(s_{n,\Delta}^{\mathbb{P}})^2 = \sigma_{\mathrm{diff}}^2\Delta + n(\sigma_J^{\mathbb{P}})^2. \tag{3.53}$$

The physical log-likelihood is

$$\ell^{\mathbb{P}}\left(\lambda^{\mathbb{P}}, \mu_J^{\mathbb{P}}, \sigma_J^{\mathbb{P}}\right) = \sum_{i=1}^{N} \log f^{\mathbb{P}}(x_i). \tag{3.54}$$

The physical jump parameters maximize Equation (3.54). No return threshold or exogenous jump classification is imposed. This likelihood uses the Poisson-normal mixture implied by the Merton jump-diffusion model (Merton, 1976) and follows the maximum-likelihood approach for mixed jump-diffusion processes in Jorion (1988). The comparison between $\lambda^{\mathbb{P}}$ and $\lambda^{\mathbb{Q}}$ is left unrestricted.

Under Heston, the dated option surface estimates the risk-neutral parameter vector

$$\Theta_{\mathrm{H}}^{\mathbb{Q}} = \left(v_0, \kappa^{\mathbb{Q}}, \theta^{\mathbb{Q}}, \xi, \rho\right). \tag{3.55}$$

The estimate is obtained by applying the price-space objective in Equation (3.49) to the Heston pricing formula in Equation (3.43). The initial variance v_0, the volatility-of-variance ξ and the correlation ρ are kept common across $\mathbb{P}$ and $\mathbb{Q}$.

The physical variance drift parameters $(\kappa^{\mathbb{P}}, \theta^{\mathbb{P}})$ are estimated from the historical return sample available from 2010 up to the calibration date. Since variance is latent, the implementation uses a 21-day realized-variance proxy,

$$\hat{v}_i = \frac{252}{21}\sum_{\ell=0}^{20} x_{i-\ell}^2. \tag{3.56}$$

The sample mean of $\hat{v}_i$ estimates $\theta^{\mathbb{P}}$. The lag-21 autocorrelation of $\hat{v}_i$, denoted $\widehat{\rho}_{21}$, estimates the persistence of the variance process. Using the CIR conditional-mean relation $\rho_h = e^{-\kappa h}$ with $h = 21/252$, the physical mean-reversion speed is

$$\widehat{\kappa}^{\mathbb{P}} = -\frac{\log \widehat{\rho}_{21}}{21/252}. \tag{3.57}$$

The realized-variance proxy follows the quadratic-variation approach of T. G. Andersen et al. (2003). The 21-day window is an implementation choice. The relation $\rho_h = e^{-\kappa h}$ comes from the mean reversion of the variance process and gives Equation (3.57) for $h = 21/252$. The physical estimates $(\kappa^{\mathbb{P}}, \theta^{\mathbb{P}})$ and their risk-neutral counterparts are linked through Equations (3.32) and (3.33).

CHAPTER 4

Strategy and Evaluation

This chapter explains how the framework in Chapter 2 and the calibrated models in Chapter 3 are combined to construct and evaluate the static DOOM strategy. The procedure first defines the contract universe from selected maturities and target probabilities, and obtains the corresponding strikes and option prices from the calibrated models. It then selects the option weights under a fixed premium budget by maximizing a return-risk objective and evaluates the resulting allocation on simulated model paths. This model-based evaluation provides a benchmark for strategy performance under each calibrated model. A separate historical SPX backtest applies the optimized model weights to rolling return windows in order to assess model risk.

4.1 Static Contract Universe

Following the DOOM put definition in Section 2.3, the contract universe is defined by a maturity grid and a set of risk-neutral target probabilities. For each model $\mathscr{M}$, each probability level is converted into a strike using the calibrated risk-neutral distribution from Chapter 3. The option price is then obtained from the same pricing specification.

The maturity and probability grids used for all models are

$$\begin{aligned} \mathscr{T} &= \left\{\frac{1}{52}, \frac{1}{12}, \frac{2}{12}, \frac{3}{12}, \frac{4}{12}, \frac{6}{12}, \frac{8}{12}, \frac{9}{12}, \frac{12}{12}\right\}, \\ \mathscr{P} &= \{0.005, 0.010, 0.020, 0.030, 0.050, 0.075\} \cup \{0.100 + 0.025k\}_{k=0}^{6} \\ &\quad \cup \{0.260 + 0.010k\}_{k=0}^{14}. \end{aligned} \tag{4.1}$$

The index a denotes the maturity level and the index j denotes the target probability level. For a calibration date d and a model $\mathscr{M}$, the contract universe provided to the optimizer is

$$\mathscr{C}^{\mathscr{M},d} = \left\{ \left(K_{a,j}^{\mathscr{M},d}, T_a, p_j\right) \;\middle|\; \begin{array}{c} T_a \in \mathscr{T}, \quad p_j \in \mathscr{P}, \\ \mathbb{Q}^{\mathscr{M},d}(S_{T_a} < K_{a,j}^{\mathscr{M},d}) = p_j, \\ 0.30 \le K_{a,j}^{\mathscr{M},d}/S_0 \le 0.90 \end{array} \right\}. \tag{4.2}$$

The strike construction follows Equations (3.10), (3.21) and (3.46). Contracts for which the implied strike falls outside the admissible range are removed rather than assigned to a boundary strike.

The resulting contract universe depends on the calibration date and the model. Although the same (T_a, p_j) grid is used, different calibrated risk-neutral distributions generate different strikes. Contracts whose strikes fall outside the admissible range are removed, leading to different numbers of retained contracts across models, as reported in Table 4.1.

Table 4.1: Number of retained contracts in the static universe by calibration date and model.

Model	**16 March 2020**	**8 April 2025**	**30 June 2026**
BSM	182	100	52
Merton	87	70	43
Heston	156	91	56

4.2 Static Allocation and Objective

For the contract universe defined in Section 4.1, the optimizer searches for the feasible budget-weight vector that maximizes the objective defined below. Contracts with a positive optimal weight are included in the strategy, while the others receive zero weight. The resulting weights determine the option quantities and the remaining index exposure through Equations (2.3) to (2.5). Selected puts are purchased at $t = 0$, and the allocation is kept fixed until the one-year horizon. The maximum option budget is 5%. Underlying paths are generated under $\mathbb{P}$, while contract construction and valuation are performed under the corresponding dated risk-neutral specification, consistently with the measure separation in Section 2.4 and the model-specific pricing rules in Chapter 3.

The optimization follows a return-risk criterion using $\mathrm{ES}_{95\%}$. With the positive-loss convention and the mean return definition in Equations (2.14) and (2.15), the objective is

$$\mathcal{J}(\mathbf{w}) = \bar{R}_T(\mathbf{w}) - \frac{\gamma}{2}\widehat{\mathrm{ES}}_{95\%}(\mathbf{w}), \qquad \hat{\mathbf{w}} \in \arg\max_{\mathbf{w}\in\mathcal{W}_b} \mathcal{J}(\mathbf{w}), \tag{4.3}$$

where $\mathcal{W}_b$ denotes the admissible budget-weight set. The criterion assigns higher scores to allocations with higher $\bar{R}_T(\mathbf{w})$ and lower Expected Shortfall. The reference specification uses $\gamma = 1$. Alternative objective specifications are reported in Section C.1. They are used as sensitivity checks. The principal specification in the main analysis remains $\mathrm{ES}_{95\%}$.

4.3 Numerical Optimization

The empirical objective is evaluated directly for each tentative weight vector rather than optimized with gradient-based updates. The ES term is computed from ordered terminal portfolio returns. When $\mathbf{w}$ changes, portfolio returns change on the fixed simulated paths, which can change the observations entering the 5% tail and makes the objective non-smooth in the weights.

First, the net contribution of each contract to portfolio wealth is precomputed. Let $W^I_{i,t}$ denote the index-only wealth on path i at time t. For contract j, $U_{j,i,t}$ is the wealth contribution associated with one unit of budget weight allocated to that contract. For any tentative weight vector $\mathbf{w}$, wealth is computed as

$$W_{i,t}(\mathbf{w}) = W^I_{i,t} + \sum_{j=1}^{M} w_j U_{j,i,t}, \tag{4.4}$$

where M is the number of retained contracts. This expression allows the portfolio wealth of each tentative allocation to be computed from the precomputed contract contributions, without rebuilding the full portfolio after each move.

Second, the algorithm constructs the initial weight vectors. The first vector is the zero-option allocation. The remaining twenty vectors are sparse starts. Each sparse start assigns positive trial weights to a subset of contracts. The number of selected contracts lies between one and $\min\{M, \max(2, \lfloor M/8 \rfloor)\}$. All other contracts receive zero weight, and the vector may leave part of the budget unused.

Third, each tentative vector $\widetilde{\mathbf{w}}$ is projected onto the feasible set before $\mathcal{J}$ is computed. With $b_{\max} = 0.05$, the feasible set is

$$\mathscr{W}_b = \left\{ \mathbf{w} \in \mathbb{R}^M_+ \mid \sum_{j=1}^{M} w_j \le b_{\max} \right\}. \tag{4.5}$$

The projection removes negative components. If the sum of positive weights exceeds $b_{\max}$, a common threshold is subtracted from the positive weights until the budget constraint is satisfied, following Duchi et al. (2008). The projected initial vectors are then evaluated, and the one with the highest value of $\mathcal{J}$ is retained as the starting point of the coordinate search.

Fourth, the selected starting vector is refined by changing one weight at a time. At each pass, the weight coordinates are considered in a pseudo-random order. For each coordinate, the algorithm tests $w_j + h$ and $w_j - h$, projects the resulting vector and computes $\mathcal{J}$. A move is accepted only if it increases $\mathcal{J}$. Otherwise, the current vector is kept and the algorithm continues with the next coordinate. The initial step is $h = 0.02$. After a full pass without any accepted move, the step size is reduced by 20%. The search stops after 60 passes or when $h < 10^{-6}$.

Finally, the resulting vector is projected again and the objective is recomputed.

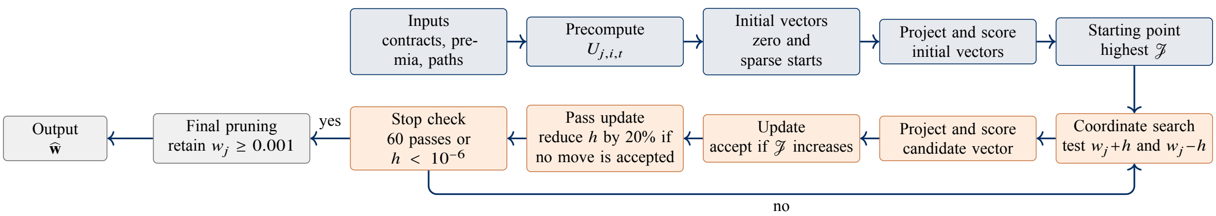

Figure 4.1: Numerical procedure used to obtain the optimized static DOOM allocation.

4.4 Model-Based Evaluation

For each model and calibration date, two independent populations of simulated physical paths are used. The in-sample population is used in Section 4.3 to select the optimized allocation $\widehat{\mathbf{w}}$. The out-of-sample population is then used to evaluate this fixed allocation. BSM and Merton use 8,000 in-sample paths and 30,000 out-of-sample paths. Heston uses 5,000 and 20,000 paths. Each path contains 252 daily intervals, and each population is generated with antithetic innovations.

On the out-of-sample paths, the optimized weights, strikes, maturities and quantities are kept unchanged. The Index and DOOM-put portfolios are evaluated on the same simulated underlying trajectories. Hedging performance is reported using the metrics defined in Section 2.5.

These experiments provide model-based benchmarks. They rely on fixed calibrated parameters over a one-year horizon, meaning that the calibrated model specification is kept unchanged throughout the simulation period. This is a strong assumption, as the simulations do not capture changes in market conditions or deviations between the model dynamics and observed option prices. They also depend on the quality of the option-surface fit at the calibration date. The benchmark is therefore useful to compare the strategy across calibrated models and tail-risk assumptions, but it is not a substitute for a real-market evaluation. This motivates the historical SPX backtest in the next section.

4.5 Historical SPX Evaluation

The historical evaluation applies each optimized static allocation $\widehat{\mathbf{w}}$ to historical SPX return windows. The model is not recalibrated and the weights are not reoptimized. The objective is to assess how allocations selected under each calibrated model perform on realized SPX return paths.

For each calibration date, the selected model-implied contracts are matched to the listed SPX option chain available on that date. The matching first selects the listed expiry closest to the model maturity and then, within that expiry, the strike closest to the model-implied strike ratio K_j/S_0. The listed mid quote is used as the transaction price. Since w_j represents the budget share allocated to position j, it is kept fixed when moving from model prices to listed prices. The option quantity is then recomputed from the listed premium. If S_0^{mkt}, K_j^{mkt} and π_j^{mkt} denote the SPX level, listed strike and listed mid premium on the calibration date, the normalized contract is

$$K_j^* = 100\frac{K_j^{\mathrm{mkt}}}{S_0^{\mathrm{mkt}}}, \qquad \pi_j^* = 100\frac{\pi_j^{\mathrm{mkt}}}{S_0^{\mathrm{mkt}}}, \qquad n_j^* = \frac{w_j B_0}{\pi_j^*}. \tag{4.6}$$

The quantity n_j^* makes the initial cost of position j equal to $w_j B_0$.

The historical SPX series uses closing prices and runs from 30 December 1927 to 22 July 2026. Each overlapping window contains 253 observations, corresponding to 252 trading-day intervals. For each window beginning at τ, the normalized SPX path is

$$S^*_{\tau,t} = 100\frac{S_{\tau+t}}{S_\tau}, \qquad t = 0, \dots, 252. \tag{4.7}$$

This normalization allows the historical paths to be evaluated using the same normalized contract characteristics rather than the same absolute SPX strike. This gives 24,503 one-year windows. The lowest 5% of terminal SPX returns define the Crash sample, the next 15% define the Stress sample and the remaining windows define the Normal sample.

Each listed expiry is mapped to the closest daily evaluation node d_j on the one-year backtest grid. Historical terminal wealth is

$$P^D_{\tau,T} = \left(1 - \sum_j w_j\right) B_0 \frac{S^*_{\tau,252}}{100} + \sum_j n^*_j \left(K^*_j - S^*_{\tau,d_j}\right)^+ \frac{S^*_{\tau,252}}{S^*_{\tau,d_j}}. \tag{4.8}$$

The corresponding discounted terminal return is

$$R^D_{\tau,T} = \frac{P^D_{\tau,T}}{B_0 e^{rT}} - 1, \tag{4.9}$$

consistent with Equations (2.2) and (2.11). The Index return is

$$R^I_{\tau,T} = \frac{S^*_{\tau,252}}{100e^{rT}} - 1. \tag{4.10}$$

These returns are used to compute the performance and risk metrics defined in Section 2.5. Performance and risk metrics are computed from $\{R^D_{\tau,T}\}$ and $\{R^I_{\tau,T}\}$ for the full sample and for each regime. Since all fixed allocations are evaluated on the same historical SPX windows, differences in results come from the model-selected allocation and the calibration date, not from different historical paths.

CHAPTER 5

Results and Empirical Analysis

This empirical analysis combines the different components developed throughout this thesis. The framework and risk metrics introduced in Chapter 2 define how portfolio outcomes and risk reductions are measured. The model specifications developed in Chapter 3 provide the dynamics used to construct and price the DOOM-put contracts. The simulation and evaluation protocol described in Chapter 4 specifies how the allocations are selected and assessed.

This chapter is organized around two empirical evaluations. The first part presents the model-based benchmarks generated by BSM, Merton jump-diffusion and Heston stochastic volatility. It analyses how model dynamics affect contract selection, portfolio allocation, option costs and downside protection across calibration dates. The second part turns to the historical SPX evaluation, where these fixed allocations are applied to realized return windows to assess their robustness and resulting model risk.

5.1 Model-Based Results

The model-based benchmark results are shown in Table 5.1. Panel A compares portfolio performance, option cost and terminal downside risk across models and calibration dates. Panel B reports path and distributional metrics. Panel C reports the selected DOOM-put positions.

Table 5.1: One-year model-based benchmark results and selected DOOM-put positions by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	**Model**	**Portfolio**	**Budget**	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	5.00	-9.04	-3.73	25.87	25.87	25.87	25.87
	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	5.00	-12.91	0.12	37.13	42.78	40.62	45.12
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	5.00	-14.33	1.68	35.13	35.13	35.13	35.13

Table 5.1 continued

2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.88	-4.05	-3.15	19.13	19.35	19.26	19.44
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	5.00	-9.02	1.90	20.05	20.63	20.41	20.91
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.34	-7.33	0.20	19.96	19.96	19.96	19.96
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	0.00	–	–	–	–	–	–
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	3.00	15.25	2.92	16.48	16.72	16.65	16.74

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are reported in percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation or not applicable.

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	29.65	2.32	7.54
	Merton	Index	33.16	0.81	1.20
		DOOM-put	28.96	1.42	2.88
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	24.37	0.22	-1.06
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	21.49	12.45	256.67
	Merton	Index	21.14	0.46	0.33
		DOOM-put	18.31	1.32	1.59
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.77	0.53	-0.58
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	–	–	–
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	11.82	-0.08	-0.39

Note. L_{MDD} is reported in percentage points. Dashes indicate zero allocation or not applicable.

Table 5.1 continued

Panel C. Selected DOOM-put positions

Date	Model	Portfolio	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	1.00	0.320	77.81	5.127	2.43	0.474
		DOOM-put	2	1.00	0.330	78.65	5.398	2.57	0.476
	Merton	DOOM-put	1	1.00	0.370	88.97	13.014	5.00	0.384
	Heston	DOOM-put	1	1.00	0.200	67.58	5.075	3.87	0.763
		DOOM-put	2	1.00	0.225	72.07	6.028	1.13	0.187
2025-04-08	BSM	DOOM-put	1	0.50	0.005	65.35	0.016	0.54	34.033
		DOOM-put	2	0.75	0.005	59.41	0.017	0.53	30.689
		DOOM-put	3	0.75	0.010	62.50	0.039	0.47	11.952
		DOOM-put	4	1.00	0.310	89.31	3.542	1.82	0.515
		DOOM-put	5	1.00	0.320	89.91	3.721	1.51	0.407
	Merton	DOOM-put	1	1.00	0.260	89.24	5.689	5.00	0.879
	Heston	DOOM-put	1	1.00	0.225	87.44	4.541	4.34	0.957
2026-06-30	BSM	DOOM-put	–	–	–	–	–	–	–
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	1.00	0.175	89.00	3.045	3.00	0.985

Note. w is reported as percentage of B_0. Panel C uses $S_0 = B_0 = 100$. Quantities follow $n_j = w_j B_0 / \pi_j$, with w_j expressed in decimal form. Dashes indicate zero allocation or not applicable.

The model-based results show that DOOM puts reduce downside risk in stressed calibrations. As expected, the reduction is most visible for $\mathrm{ES}_{95\%}$, because this is the optimized risk measure. The key issue is the cost of this protection. The selected puts mitigate losses in stressed conditions, but their premia are funded by reducing initial index exposure. The effect of DOOM puts on mean return therefore depends on the model, the market regime and the price of the selected puts.

The market regime behind each calibration is reflected not only in the level of volatility, but also in the downside skew embedded in the option surface, which affects the relative cost of protection across strikes. The 2020 calibration corresponds to a crash regime, the 2025 calibration to a stressed regime and the 2026 calibration to a normal regime. In 2020, all three models use the full option budget and produce large reductions in $\mathrm{ES}_{95\%}$. BSM lowers it from 63.00% to 25.87%, Merton from 55.31% to 40.62%, and Heston from 62.65% to 35.13%. In 2025, all models again select contracts, and the protected $\mathrm{ES}_{95\%}$ is close to 20% across specifications. In 2026, BSM and Merton select no put, whereas Heston allocates a 3.00% budget and reduces $\mathrm{ES}_{95\%}$ from 26.76% to 16.65%. Despite this $\mathrm{ES}_{95\%}$ reduction Heston's $\mathrm{VaR}_{95\%}$ increases slightly because the benchmark objective is based on the mean loss beyond the 5% tail threshold rather than the threshold itself. The selected put provides stronger loss protection in the deeper tail, but the premium paid, hence the lower index exposure, make the empirical 5% quantile slightly worse.

The selected puts also affect path losses and the shape of terminal returns. In every active model-date case, L_{MDD} falls after protection is added. Under Heston in 2020, the hedge reduces downside losses by creating a partial floor over the region where the puts finish in the money. Skewness moves from -0.55 to 0.22, and excess kurtosis falls from -0.07 to -1.06. The 2025 BSM case is different. Skewness reaches 12.45 and excess kurtosis 256.67, because very cheap DOOM puts generate rare but large payoffs. BSM produces an overly favorable protected distribution by exaggerating the convexity benefit of deep protection. This reflects a limitation of the constant-volatility benchmark, which tends to underprice the deep protection.

The 2020 comparison between Merton and Heston further shows that a crash observed on the same date corresponds to different model-implied market environment. Heston fits the crash surface much more closely, with an RMSE of 12.73 against 72.90 for Merton in Table B.4. This better fit is especially relevant for DOOM puts, because their prices are driven by the downside part of the option surface and by the skew observed at low strikes. This does not establish that Heston has the true physical dynamics or the best hedging allocation. It means that its risk-neutral pricing specification is closer to the observed pricing of downside protection at this date. By contrast, Merton represents the 2020 crash surface less closely, which affects both the strike associated with each p_j and the premium paid for that strike.

Under this calibration, the Merton model selects a single one-year put with $K_j = 88.97$, $p_j = 0.370$ and a premium of 13.014. The full budget buys only 0.384 contracts. In the Merton benchmark, the risk-neutral jump component makes downside insurance expensive. Buying deeper puts would therefore require giving up more return for protection that pays only in the most severe states. Their lower exercise probability in the simulated physical distribution does not justify the premium paid. Higher-strike puts are more expensive, but they are closer to the relevant loss region and have a higher probability of finishing in the money, allowing their payoff to be realized across a broader set of loss scenarios. Although the premium reduces initial index exposure, the higher frequency of realized payoffs limits the drag on mean return. The resulting allocation favors a small position in a higher-strike put whose payoff is more likely to contribute to the reduction in $\text{ES}_{95\%}$, rather than a larger position in lower-cost deeper puts that may expire unused in many physical scenarios.

Figure 5.1 illustrates this mechanism along one adverse path. The put has value throughout the path, but its protection becomes economically material when the index enters the strike region.

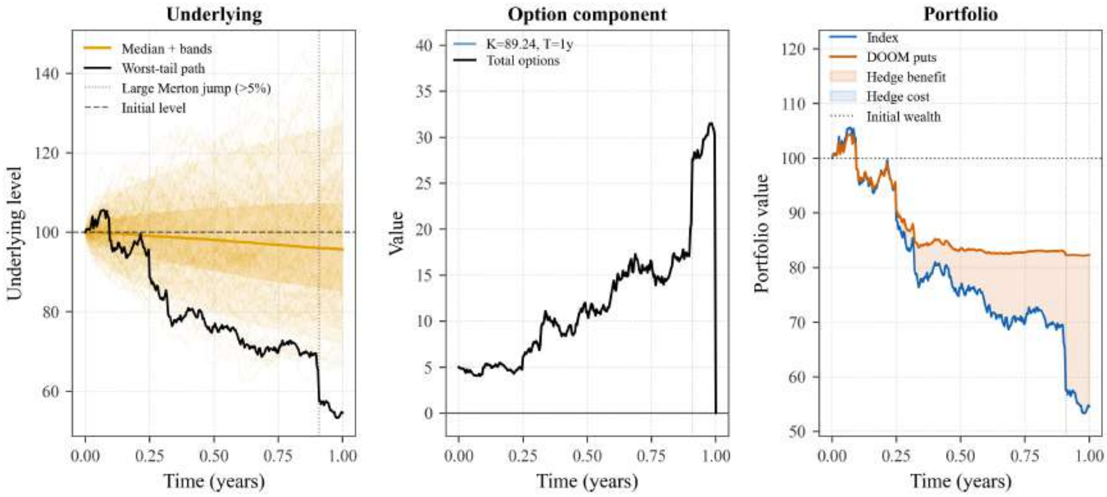

Figure 5.1: Merton static protection along one jump-diffusion path for the 8 April 2025 calibration.

Heston produces a contrasting allocation in 2020. Because its fit is closer to the crash surface, its selected strikes better reflect the severe downside conditions priced at that date. The model selects $p_j = 0.200$ and 0.225, which correspond to strikes of 67.58 and 72.07. In the crash benchmark, lower-strike DOOM puts become more attractive because their probability of finishing in the money remains sufficient in severe loss scenarios to contribute to $\mathrm{ES}_{95\%}$. They are also cheaper than higher-strike protection and provide a closer hedge for extreme losses. This differs from the Merton benchmark, where the premium required for lower-strike puts is not compensated by their contribution to the targeted tail-risk reduction.

Figure 5.2 shows the resulting shape of the protected distribution, with the partial floor visible in the DOOM-put portfolio panel.

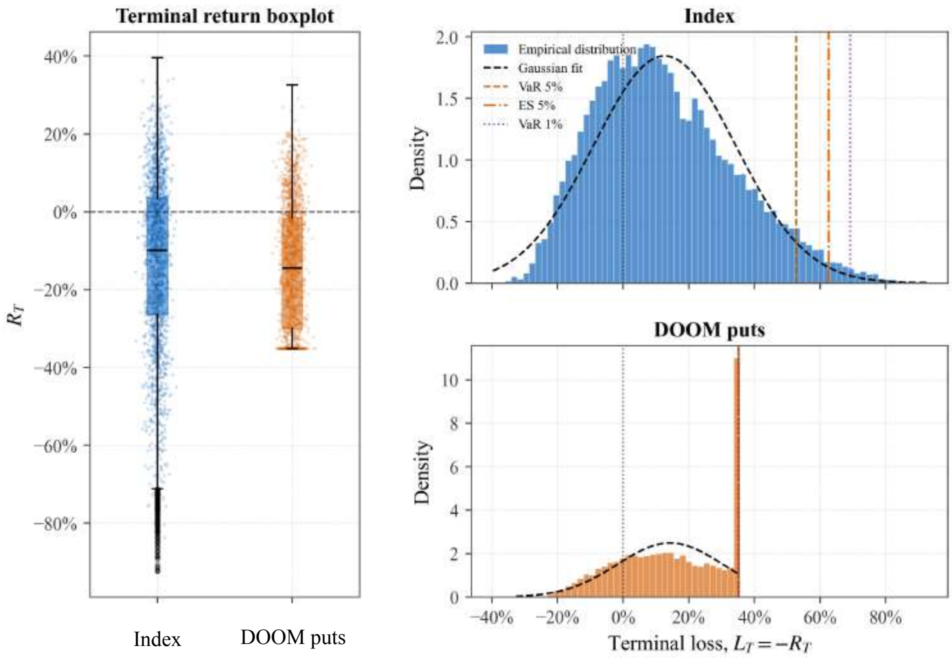

Figure 5.2: Heston terminal return distributions for the 16 March 2020 calibration.

The 2025 and 2026 results show how this model-driven divergence in strike selection changes outside the crash regime. In 2025, Heston moves closer to the Merton allocation pattern and selects a higher strike of 87.44. This is consistent with a stressed but less extreme calibration, in which very deep DOOM puts provide limited $\mathrm{ES}_{95\%}$ reduction relative to their cost. In 2026, Merton selects no protection because the premium required for downside protection is too high relative to the hedge's expected contribution in a normal regime. Heston remains active with one put at $K = 89.00$, because its stochastic-volatility specification still assigns value to downside protection relative to its cost.

Outside the crash regime, the BSM allocation reflects another limitation of the constant-volatility benchmark. In 2025, it combines very cheap deep puts with two higher one-year puts. The low premia on the deepest contracts allow large quantities, such as 34.033 and 30.689, but these puts pay only in rare states. This produces a highly convex protected distribution and explains the extreme skewness and excess kurtosis. It also explains why the BSM hedge can improve mean return despite option premia. The improvement is driven by rare model-implied payoffs from cheap deep puts, not by a robust estimate of their market cost. The interpretation of the BSM hedge must therefore remain cautious.

Across regimes, the model-based results do not identify a single optimal protection rule. They show that the allocation depends on how each model represents volatility, downside skew and tail risk in the selected strikes, premia and hedge weights. Lower strikes provide protection against more extreme losses, but their payoff is only realized when the market decline is large enough. Higher strikes are more expensive, hence bought in smaller quantities, but they protect a broader range of loss scenarios and are more likely to finish in the money. Maturity also affects this trade-off. Shorter maturities provide more targeted protection against sudden market moves. Longer maturities require a higher initial premium but provide protection over a longer horizon and give more time for the downside scenario to occur. The optimizer therefore selects the strike and maturity combination that contributes most to $\mathrm{ES}_{95\%}$ after accounting for the model-implied premium.

5.2 Option and Risk Premia

The previous section shows that Merton buys a small quantity of a relatively high-strike put in the 2020 crash because lower strikes are too expensive relative to their exercise probability. This section explains the source of that expensive pricing. The risk-neutral jump law prices crash risk very differently from the physical jump law, generating a risk premium that affects put prices and therefore influences the allocation selected by each model.

For the lognormal jump law defined in Equation (3.12), the expected jump multiplier is given by Equation (3.13) as $\mathbb{E}^{\mathbb{M}}[e^{J^{\mathbb{M}}}] = \exp(\mu_J^{\mathbb{M}} + \frac{1}{2}(\sigma_J^{\mathbb{M}})^2)$. The expected simple loss per jump

$\ell^{\mathbb{M}}$ is therefore one minus this multiplier when the average jump is negative. The annual net jump-loss intensity $\mathscr{L}^{\mathbb{M}}$ multiplies this average loss by the jump arrival intensity. Finally, $\mathscr{V}_J^{\mathbb{M}}$ is defined here as the annualized second moment of log jump sizes. The resulting diagnostics are

$$\ell^{\mathbb{M}} = \max\left\{1 - \exp\left(\mu_J^{\mathbb{M}} + \frac{1}{2}(\sigma_J^{\mathbb{M}})^2\right), 0\right\}, \tag{5.1}$$

$$\mathscr{L}^{\mathbb{M}} = \lambda^{\mathbb{M}}\ell^{\mathbb{M}}, \tag{5.2}$$

$$\mathscr{V}_J^{\mathbb{M}} = \lambda^{\mathbb{M}}\left[(\mu_J^{\mathbb{M}})^2 + (\sigma_J^{\mathbb{M}})^2\right], \qquad \mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}. \tag{5.3}$$

Panel A of Table 5.2 reports these quantities for 16 March 2020, together with $\mathscr{L}_-^{\mathbb{M}}$, the downside-only loss intensity, and $\lambda_{0.10}^{\mathbb{M}}$ and $\mathscr{L}_{0.10}^{\mathbb{M}}$, the jump intensity and loss intensity restricted to jumps below the -10% log threshold. Panel B reports the affine loadings λ_0 and λ_1 of the Heston variance-risk premium defined in Equation (3.30), together with the resulting expected one-year average variance difference $\mathbb{E}[\bar{V}^{\mathbb{P}} - \bar{V}^{\mathbb{Q}}]$, where $\bar{V}^{\mathbb{M}} = T^{-1}\int_0^T v_t^{\mathbb{M}}\,\mathrm{d}t$.

Table 5.2: Jump and variance-risk-premium diagnostics by calibration date.

Panel A. Merton jump diagnostics for 16 March 2020

Date	$\mathbb{M}$	λ	ℓ	$\mathscr{L}$	$\mathscr{L}_-$	$\lambda_{0.10}$	$\mathscr{L}_{0.10}$	$\mathscr{V}_J$
2020-03-16	$\mathbb{P}$	13.0941	0.0179	0.2347	0.4353	1.1229	0.1336	0.0502
	$\mathbb{Q}$	0.4251	0.5345	0.2272	0.2272	0.4251	0.2272	0.2562

Panel B. Heston variance-risk-premium diagnostics

Date	λ_0	λ_1	$\mathbb{E}[\bar{V}^{\mathbb{P}} - \bar{V}^{\mathbb{Q}}]$
2020-03-16	0.0183	-4.6584	-0.1486
2025-04-08	0.0103	-2.4429	-0.0333
2026-06-30	0.1101	-5.8920	-0.0199

Note. $\lambda^{\mathbb{M}}$ and $\lambda_{0.10}^{\mathbb{M}}$ are annual jump intensities. The remaining Panel A quantities are annual decimal rates, and a negative value in Panel B means that the expected average variance is higher under $\mathbb{Q}$ than under $\mathbb{P}$.

Under the physical measure, jumps occur far more often than under the risk-neutral measure, but each physical jump is markedly less severe. The risk-neutral calibration changes the composition of jump risk. Under $\mathbb{P}$, jump losses are generated by frequent and smaller jumps, whereas under $\mathbb{Q}$, risk-neutral jump losses are concentrated in rarer but more severe events. The aggregate loss intensity $\mathscr{L}^{\mathbb{M}} = \lambda^{\mathbb{M}}\ell^{\mathbb{M}}$ remains of similar magnitude across measures, with $\mathscr{L}^{\mathbb{Q}}/\mathscr{L}^{\mathbb{P}} = 0.968$. The difference becomes more pronounced in the tail. Jumps beyond the -10% log threshold arrive less often under $\mathbb{Q}$, yet their associated loss intensity $\mathscr{L}_{0.10}$ is 1.70 times higher and the log-jump second-moment rate $\mathscr{V}_J$ is 5.11 times higher under $\mathbb{Q}$.

The risk-neutral parameters should not be read as a forecast of jump frequency and severity. They are pricing parameters, and their gap with the physical parameters arises from the compensation for jump risk rather than reflecting expected real-world jump frequencies. This compensation is demanded by the option seller, who must hold or delta-hedge the position. A discontinuous price move cannot be fully replicated by continuous trading in the underlying. Part of the jump risk therefore remains on the seller's book and is priced into the option premium. The larger risk-neutral weight assigned to severe jump outcomes therefore does not mean that the market expects more crashes. It means that the price of insuring against a crash reflects more than its physical frequency and average severity, consistent with the crash-o-phobia premium documented by Bates (1991). This separation between physical frequency and risk-neutral severity is also consistent with Pan (2002), who shows that option-implied jump risk premia are governed by the risk-neutral jump-size distribution rather than by jump frequency. The larger risk-neutral compensation assigned to rare and severe jump outcomes, rather than to frequent but smaller losses, is consistent with the rare-disaster risk-premium literature of Bollerslev and Todorov (2011). Figure 5.3 breaks down the one-year Merton put price into its diffusive, physical-jump and risk-neutral components under a given sequence of substitutions. The green line is the resulting total Merton price under $\mathbb{Q}$. The blue layer is the option value obtained from the diffusive component alone, without jumps. The orange layer adds the effect of the physical jump parameters, showing how jump risk changes the option price under the physical jump distribution. The red layer captures the change from physical to risk-neutral jump parameters. This red component is large across the out-of-the-money range, showing that a large part of the premium on deep, low-strike puts comes from this risk-neutral loading rather than from the diffusive or physical-jump components. Because of this additional risk-neutral premium, lower strikes in the 2020 Merton benchmark are expensive relative to their simulated contribution to $\mathrm{ES}_{95\%}$. This explains why Merton selects a small quantity of a higher strike.

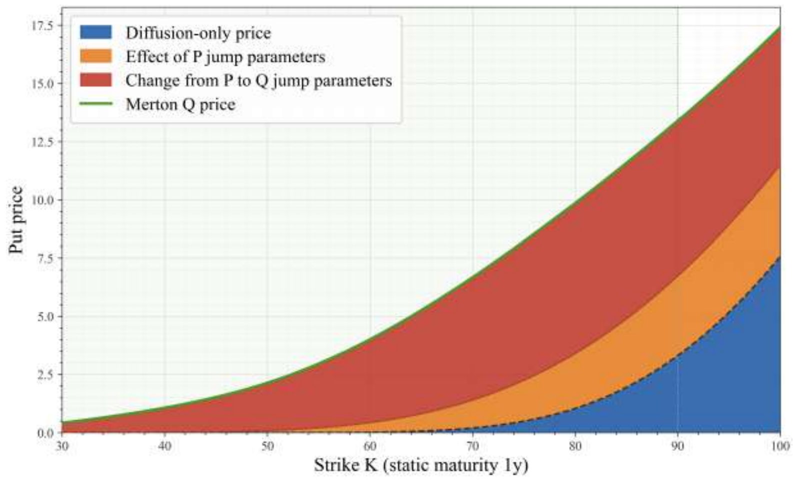

Figure 5.3: Merton put-premium decomposition (16 March 2020 calibration).

Heston prices downside protection through stochastic volatility and the associated variance-risk premium rather than through an explicit jump component, and this premium reflects compensation for variance risk that cannot be fully hedged away. On 16 March 2020, expected one-year average variance is higher under $\mathbb{Q}$ than under $\mathbb{P}$, with $\mathbb{E}[\bar{V}^{\mathbb{P}} - \bar{V}^{\mathbb{Q}}] = -0.1486$, so higher-volatility states are priced into the option surface. This contrasts with Merton, where downside protection is strongly affected by the risk-neutral pricing of jump losses. The gap in Merton's log-jump second-moment rate, $\mathscr{V}_J^{\mathbb{Q}} - \mathscr{V}_J^{\mathbb{P}} = 0.2060$, indicates that jump losses receive a large risk-neutral pricing weight. This helps explain why low-strike protection is more attractive under Heston in the 2020 benchmark, while the same region is less attractive under Merton's jump-risk pricing.

5.3 Rolling Protection

Monthly, quarterly and semi-annual rolling rules test whether resetting protection improves the static allocation. The rolling strategy keeps the total budget b selected by the static optimizer, sets the initial index exposure to $1 - b$, and allocates the reserved budget across reset dates. The reserve is invested at the risk-free rate and used only to finance future option purchases, making each reset independent of the realized SPX path. The empirical results indicate that rolling does not improve $\mathrm{ES}_{95\%}$. Splitting the same budget across shorter maturities and paying new option premia at each reset does not outperform the static allocation. The implementation details and complete results are reported in Sections C.2 and C.2.1.

5.4 Transaction Costs

A transaction-cost sensitivity is tested by increasing the purchase price of each put by 1%, so that $\pi_j^{\mathrm{tc}} = 1.01\pi_j$. The contract universe, simulated paths, objective and option budget are kept unchanged. This adjustment does not change the main results across the three calibration dates. Merton and Heston retain the same static contracts, with slightly smaller quantities because the put premium is higher. BSM sometimes reallocates weight across nearby contracts, but the selected strike region and protection profile remain similar. In these cases, transaction costs affect quantities more than the optimizer's contract choice. Detailed results are reported in Appendix C.3.

5.5 Model Risk

The dispersion of allocations across models and dates is itself a model-risk result. It shows that DOOM-put performance depends on the pricing model used to construct strikes, premia and hedge weights. The results should therefore be read as model-implied benchmarks, not as a single robust allocation rule.

This limitation affects both static and rolling strategies. In the static benchmark, the one-year paths are simulated with the same parameters calibrated at the initial date. The model therefore assumes that volatility, jump risk or variance dynamics remain governed by this initial calibration throughout the whole horizon. The allocation is evaluated in this fixed diffusion or jump-diffusion environment. In the rolling benchmark, the assumption is stronger because new puts are priced at future reset dates with the risk-neutral specification calibrated on the initial calibration date. This creates a forward-smile assumption. Under BSM, the forward smile remains flat. Under Merton, future put prices inherit the same risk-neutral jump law. Under Heston, future premia react to the simulated variance state, but the risk-neutral mean reversion, volatility of variance and spot-variance correlation remain fixed. Rolling therefore isolates the effect of each model mechanism, but it does not forecast the true future option surface.

This matters because fitting the current option surface does not determine the future smile relevant for rolling or forward-starting protection (Bergomi, 2004, 2009). Future option premia may fall if market conditions stabilize, or rise if stress, liquidity pressure or demand for protection increase. The models also price a single-index state variable. They do not capture changes in market liquidity, option supply and demand, transaction costs, cross-market effects or the correlation structure inside the index. The allocations and risk reductions reported above are therefore not unconditional properties of DOOM puts. They are conditional on the pricing and simulation framework used to construct them. These limits motivate the historical SPX backtest, which applies the model-selected allocations to realized index paths rather than to paths generated by the same model.

5.6 Historical SPX Backtest Results

The historical SPX backtest applies the model-selected allocations to realized one-year SPX return windows. The windows are grouped by terminal index return: the lowest 5% form the Crash sample, the next 15% form the Stress sample and the remaining 80% form the Normal sample. The model contracts are matched to the listed option surface at the calibration date and then evaluated on normalized SPX paths, as defined in Equations (4.6) and (4.7).

Table 5.3 reports terminal metrics only. Index results differ across calibration dates because terminal wealth is discounted with the dated risk-free rate in Equations (4.9) and (4.10).

Table 5.3: Historical one-year SPX results by calibration date, model, portfolio and market regime.

Panel A. Full regime

Date	**Model**	**Portfolio**	$\bar{R}_T$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	7.91	24.93	43.63	37.37	53.43
		DOOM puts	3.02	26.87	32.50	30.61	35.44
	Merton	Index	7.91	24.93	43.63	37.37	53.43
		DOOM puts	3.36	22.55	31.96	28.81	36.90
	Heston	Index	7.91	24.93	43.63	37.37	53.43
		DOOM puts	2.84	28.69	35.43	33.96	35.47
2025-04-08	BSM	Index	3.53	27.98	45.92	39.91	55.32
		DOOM puts	-0.05	21.58	26.09	24.34	27.49
	Merton	Index	3.53	27.98	45.92	39.91	55.32
		DOOM puts	0.29	18.49	18.74	18.65	18.77

Table 5.3 continued

	Heston	Index	3.53	27.98	45.92	39.91	55.32
		DOOM puts	0.60	19.93	20.06	20.01	20.07
2026-06-30	BSM	Index	4.33	27.42	45.50	39.45	54.98
		DOOM puts	–	–	–	–	–
	Merton	Index	4.33	27.42	45.50	39.45	54.98
		DOOM puts	–	–	–	–	–
	Heston	Index	4.33	27.42	45.50	39.45	54.98
		DOOM puts	3.06	16.57	16.70	16.65	16.72

Panel B. Crash regime

Date	**Model**	**Portfolio**	$\bar{R}_T$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	-37.37	61.08	67.36	64.68	69.63
		DOOM puts	-30.61	37.74	39.63	38.82	40.31
	Merton	Index	-37.37	61.08	67.36	64.68	69.63
		DOOM puts	-28.81	40.75	43.92	42.56	45.06
	Heston	Index	-37.37	61.08	67.36	64.68	69.63
		DOOM puts	-33.96	35.48	35.50	35.49	35.50
2025-04-08	BSM	Index	-39.91	62.66	68.69	66.11	70.86
		DOOM puts	-22.52	27.98	30.63	29.36	31.49
	Merton	Index	-39.91	62.66	68.69	66.11	70.86
		DOOM puts	-16.62	17.59	17.64	17.62	17.64
	Heston	Index	-39.91	62.66	68.69	66.11	70.86
		DOOM puts	-19.20	19.63	19.65	19.64	19.65
2026-06-30	BSM	Index	-39.45	62.37	68.44	65.85	70.64
		DOOM puts	–	–	–	–	–
	Merton	Index	-39.45	62.37	68.44	65.85	70.64
		DOOM puts	–	–	–	–	–
	Heston	Index	-39.45	62.37	68.44	65.85	70.64
		DOOM puts	-15.62	16.12	16.15	16.14	16.15

Panel C. Stress regime

Date	**Model**	**Portfolio**	$\bar{R}_T$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	-14.64	22.93	24.61	23.98	24.75
		DOOM puts	-18.84	26.27	26.78	26.59	26.82
	Merton	Index	-14.64	22.93	24.61	23.98	24.75
		DOOM puts	-17.14	21.54	22.38	22.07	22.46

Table 5.3 continued

	Heston	Index	-14.64	22.93	24.61	23.98	24.75
		DOOM puts	-18.91	26.78	28.38	27.78	28.52
2025-04-08	BSM	Index	-18.11	26.06	27.66	27.06	27.81
		DOOM puts	-19.09	21.60	22.04	21.87	22.10
	Merton	Index	-18.11	26.06	27.66	27.06	27.81
		DOOM puts	-18.06	18.75	18.80	18.78	18.80
	Heston	Index	-18.11	26.06	27.66	27.06	27.81
		DOOM puts	-19.01	20.06	20.08	20.08	20.08
2026-06-30	BSM	Index	-17.47	25.48	27.10	26.50	27.25
		DOOM puts	–	–	–	–	–
	Merton	Index	-17.47	25.48	27.10	26.50	27.25
		DOOM puts	–	–	–	–	–
	Heston	Index	-17.47	25.48	27.10	26.50	27.25
		DOOM puts	-16.09	16.71	16.73	16.72	16.74

Panel D. Normal regime

Date	**Model**	**Portfolio**	$\bar{R}_T$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	14.97	4.47	6.84	5.97	7.10
		DOOM puts	9.22	9.25	11.50	10.67	11.74
	Merton	Index	14.97	4.47	6.84	5.97	7.10
		DOOM puts	9.22	9.25	11.50	10.67	11.74
	Heston	Index	14.97	4.47	6.84	5.97	7.10
		DOOM puts	9.22	9.25	11.50	10.67	11.74
2025-04-08	BSM	Index	10.30	8.34	10.62	9.79	10.87
		DOOM puts	4.92	12.81	14.98	14.19	15.22
	Merton	Index	10.30	8.34	10.62	9.79	10.87
		DOOM puts	4.79	12.93	15.08	14.30	15.32
	Heston	Index	10.30	8.34	10.62	9.79	10.87
		DOOM puts	5.51	12.33	14.50	13.71	14.74
2026-06-30	BSM	Index	11.16	7.64	9.92	9.09	10.18
		DOOM puts	–	–	–	–	–
	Merton	Index	11.16	7.64	9.92	9.09	10.18
		DOOM puts	–	–	–	–	–
	Heston	Index	11.16	7.64	9.92	9.09	10.18
		DOOM puts	7.82	10.41	12.63	11.82	12.87

Note. Entries are percentage points. $\bar{R}_T$ *is the mean discounted terminal return. VaR and ES are reported as positive losses. Dashes indicate that no put was selected.*

The Crash sample provides the clearest evidence that active DOOM puts reduce VaR and ES when severe historical losses reach their payoff region. The strongest protection is obtained with Heston 2026, which reduces $ES_{95\%}$ to 16.14%, followed by Merton 2025 at 17.62%. Both allocations use higher strikes, which are more expensive and bought in smaller quantities, showing that the best protection-cost trade-off does not necessarily come from the deepest DOOM puts. It comes from a strike-quantity combination that pays in the realized loss region.

The Stress sample confirms this point. Lower-strike puts remain too far out of the money when losses are severe but not extreme, making their premium a direct cost. Higher-strike puts can then provide a better trade-off because their payoff is reached across a wider range of loss scenarios. Heston 2026 delivers the strongest Stress result, reducing $ES_{95\%}$ to 16.72%, while Merton 2025 also performs well at 18.78%. The optimal choice is to align the strike region with the market regime, taking into account the volatility level and downside skew priced in the option surface, and to select the put that gives the largest marginal reduction in tail losses relative to its premium.

Normal windows show the cost of carrying protection. When puts do not generate payoffs, the premium paid to acquire the hedge reduces the capital invested in the index. Since the insured downside scenario does not materialize, this cost is not compensated by protection benefits, so the protected portfolio can have lower mean return and worse downside metrics. This explains why BSM and Merton select no protection in the 2026 calibration, while Heston maintains a limited hedge because its calibration still assigns value to downside protection.

Overall, the historical backtest shows that no model provides a dominant DOOM-put allocation across all market conditions, since models cannot perfectly represent the dynamics that govern downside risk and option premia. Different specifications lead to different strikes, prices and allocations, but some model-implied allocations are better suited to specific regimes. In the historical SPX backtest, Heston 2026 provides the strongest crash and stress protection not because the 2026 surface represents a crash regime, but because its selected strike is reached in the realized loss windows while the hedge remains limited in cost. This contrasts with Heston 2020, where deeper strikes target more extreme downside states but do not achieve the same protection-cost trade-off in the historical sample. This does not imply that Heston correctly represents the true physical market dynamics or delivers the best allocation in general. Instead, it shows that the Heston-selected strike-budget combination is more effective for these historical loss windows, relative to Merton allocations influenced by a high jump-risk premium and to the BSM constant-volatility benchmark.

DOOM puts should therefore be viewed as conditional downside insurance. Deep low-strike puts are useful when the objective is to insure only the most extreme tail, but they protect only if the market falls sufficiently far. Higher strikes can be more effective in moderate-to-severe loss windows because they are more likely to pay, even though a smaller quantity is bought for the same budget. The effectiveness of the hedge therefore depends on the market-implied volatility level, the downside skew, the selected strike and the realized loss region.

CHAPTER 6

Conclusion

The research question addressed by this thesis was: to what extent can deep out-of-the-money put options enhance downside protection for equity portfolios while preserving cost efficiency under different market regimes? The results show that DOOM puts can reduce downside losses, but only when the selected strike region matches the market regime. This alignment depends on the volatility level and downside skew priced in the option surface. These two factors determine the cost of protection across strikes. The hedge is cost-efficient when the selected strike is close enough to the loss region implied by the market regime, ensuring a sufficiently high probability of finishing in the money relative to the premium paid. Since the option premium is financed by reducing the initial index position, a put that expires out of the money creates return drag relative to the unhedged portfolio, while a put that reaches the relevant loss region can offset this drag and improve the downside-risk profile. The key issue is therefore to select the put that gives the largest marginal reduction in tail losses relative to its premium. The historical SPX backtest reveals that the best protection does not necessarily come from the deepest puts. Deep low-strike puts create a far-tail floor, but they protect only if losses reach that region. In some regimes, a higher-strike put bought in smaller quantity is more robust because it is closer to the relevant loss region and pays in a wider range of realized losses.

The model-based benchmarks show that different models select different strikes and weights, so there is no single DOOM-put allocation that is optimal across market regimes. This reflects model risk, as each specification captures only part of the dynamics driving equity losses and option prices. Further research could usefully explore the impact of other regime indicators, such as the VIX, or the use of pricing models that better represent changes in market dynamics.

The use of DOOM puts also varies across investor profiles. Traders may use short-maturity puts to obtain stronger protection against sudden market moves, as their value reacts more quickly to changes in the underlying. Longer maturities provide greater exposure to changes in implied volatility and more time for the downside scenario to materialize, while reducing the need for frequent rolls. A long-only asset manager must assess whether the reduction in $ES_{95\%}$ justifies the lower initial index exposure and the risk of weaker performance when the downside scenario fails to materialize. The contribution of this thesis has been to show that the effectiveness of DOOM puts depends on the alignment between the strike region, the option-surface regime, and the loss region targeted by the hedge, with the investor's objective guiding the choice of alignment.

References

Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance*, *11*(3), 1–42. https://doi.org/10.21314/JCF.2008.189

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, *71*(2), 579–625. https://doi.org/10.1111/1468-0262.00418

Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, *9*(3), 203–228. https://doi.org/10.1111/1467-9965.00068

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, *52*(5), 2003–2049. https://doi.org/10.1111/j.1540-6261.1997.tb02749.x

Bakshi, G., & Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, *16*(2), 527–566. https://doi.org/10.1093/rfs/hhg002

Bates, D. S. (1991). The crash of '87: Was it expected? the evidence from options markets. *Journal of Finance*, *46*(3), 1009–1044. https://doi.org/10.1111/j.1540-6261.1991.tb03775.x

Bates, D. S. (2000). Post-'87 crash fears in S&P 500 futures options. *Journal of Econometrics*, *94*(1-2), 181–238. https://doi.org/10.1016/S0304-4076(99)00021-4

Bergomi, L. (2004). Smile dynamics. *Risk*, *17*(9), 117–123. https://www.risk.net/1530377

Bergomi, L. (2009). Smile dynamics IV. *Risk*, *22*(12), 94–100. https://www.risk.net/1564129

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

Bollen, N. P. B., & Whaley, R. E. (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance*, *59*(2), 711–753. https://doi.org/10.1111/j.1540-6261.2004.00647.x

Bollerslev, T., & Todorov, V. (2011). Tails, fears, and risk premia. *Journal of Finance*, *66*(6), 2165–2211. https://doi.org/10.1111/j.1540-6261.2011.01695.x

Bondarenko, O. (2014). Why are put options so expensive? *Quarterly Journal of Finance*, *4*(3), 1450015. https://doi.org/10.1142/S2010139214500153

Broadie, M., Chernov, M., & Johannes, M. (2009). Understanding index option returns. *Review of Financial Studies*, *22*(11), 4493–4529. https://doi.org/10.1093/rfs/hhp032

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, *2*(4), 61–73. https://doi.org/10.21314/JCF.1999.043

Carr, P., & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, *22*(3), 1311–1341. https://doi.org/10.1093/rfs/hhn038

Cheridito, P., Filipović, D., & Kimmel, R. L. (2007). Market price of risk specifications for affine models: Theory and evidence. *Journal of Financial Economics*, *83*(1), 123–170. https://doi.org/10.1016/j.jfineco.2005.09.008

Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, *53*(2), 385–407. https://doi.org/10.2307/1911242

Duchi, J., Shalev-Shwartz, S., Singer, Y., & Chandra, T. (2008). Efficient projections onto the ℓ_1-ball for learning in high dimensions. *Proceedings of the 25th International Conference on Machine Learning*, 272–279. https://doi.org/10.1145/1390156.1390191

Gârleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. *Review of Financial Studies*, *22*(10), 4259–4299. https://doi.org/10.1093/rfs/hhp005

Gatheral, J., Jaisson, T., & Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, *18*(6), 933–949. https://doi.org/10.1080/14697688.2017.1393551

Gil-Pelaez, J. (1951). Note on the inversion theorem. *Biometrika*, *38*(3/4), 481–482. https://doi.org/10.1093/biomet/38.3-4.481

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. https://doi.org/10.1093/rfs/6.2.327

Hull, J. C. (2018). *Options, futures, and other derivatives* (10th ed.). Pearson.

Ilmanen, A., Thapar, A., Tummala, H., & Villalon, D. (2021). Tail risk hedging: Contrasting put and trend strategies. *Journal of Systematic Investing*, *1*(1), 111–124. https://doi.org/10.52354/jsi.1.1.vi

Jorion, P. (1988). On jump processes in the foreign exchange and stock markets. *Review of Financial Studies*, *1*(4), 427–445. https://doi.org/10.1093/rfs/1.4.427

Lassance, N., & Vrins, F. (2018). A comparison of pricing and hedging performances of equity derivatives models. *Applied Economics*, *50*(10), 1122–1137. https://doi.org/10.1080/00036846.2017.1352080

Lord, R., & Kahl, C. (2010). Complex logarithms in Heston-like models. *Mathematical Finance*, *20*(4), 671–694. https://doi.org/10.1111/j.1467-9965.2010.00416.x

Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, *4*(1), 141–183. https://doi.org/10.2307/3003143

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, *3*(1-2), 125–144. https://doi.org/10.1016/0304-405X(76)90022-2

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, *63*(1), 3–50. https://doi.org/10.1016/S0304-405X(01)00088-5

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*(3), 21–41. https://doi.org/10.21314/JOR.2000.038

Shreve, S. E. (2004). *Stochastic calculus for finance II: Continuous-time models*. Springer.

S&P Dow Jones Indices. (2020, September 21). *SPIVA U.S. Mid-Year 2020*. https://www.spglobal.com/spdji/en/spiva/article/spiva-us-mid-year-2020/

Vecer, J. (2007). Preventing portfolio losses by hedging maximum drawdown. *Wilmott Magazine*, 98–104. https://www.wilmott.com/preventing-portfolio-losses-by-hedging-maximum-drawdown-wilmott-magazine-article-jan-vecer/

Appendices

APPENDIX A

Model Implementation Details

A.1 Scope of the Appendix

This appendix complements Chapter 3 by documenting the numerical specifications required when the DOOM construction cannot be implemented by a closed-form strike formula. In the Merton model, the risk-neutral tail probability and option value are evaluated through truncated Poisson mixtures. In the Heston model, option prices and tail probabilities are recovered through Fourier and distribution inversion. The appendix details the series truncation, integration grids, root-search procedures and Heston simulation conventions used to compute DOOM strikes, option values and physical scenarios.

A.2 Merton Series Evaluation

The exercise probability and the option price are evaluated as Poisson mixtures with different Poisson means. In Equation (3.20), the mixture is based on the risk-neutral jump count $N_T^{\mathbb{Q}}$, with mean $\lambda^{\mathbb{Q}}T$. In Equation (3.22), the price formula uses the adjusted intensity $\bar{\lambda}$, as explained in Section 3.3.3. Its Poisson mean is therefore $\bar{\lambda}T$.

Both mixtures are infinite sums, so they must be truncated in the numerical implementation. The truncation is chosen so that the omitted Poisson tail is negligible. If m is the Poisson mean of the series being evaluated, the last retained term is

$$n_{\max} = \max\left\{20,\ F^{-1}_{\mathrm{Pois}(m)}(1 - 10^{-10}) + 2\right\}. \tag{A.1}$$

Terms with Poisson weight below 10^{-15} are omitted. Poisson weights are computed in logarithmic form to avoid numerical overflow in the factorial terms. This gives a finite approximation while keeping the omitted mass negligible.

The strike equation in Equation (3.21) is solved with Brent's bracketing method applied to $\mathbb{Q}(S_T < K) - p$ on the admissible strike interval. This method searches for the strike K that makes the model probability equal to the target probability p, using an interval where the target is bracketed by the lower and upper strike bounds. The absolute and relative tolerances are

10^{-10}, and the iteration cap is 200. If the target probability cannot be reached in the admissible interval, the contract is excluded.

A.3 Heston Fourier and Distribution Inversion

The characteristic function in Equation (3.36) is evaluated with the stable formulation commonly used for the Heston model. The square-root branch is chosen with a non-negative real part, and the logarithm is evaluated continuously. This avoids numerical jumps in the characteristic function, as discussed by Lord and Kahl (2010).

Option prices are computed from the Carr-Madan representation in Equation (3.43) (Carr & Madan, 1999). The Fourier integral is not evaluated over the infinite interval directly. It is computed on a finite grid of u-values, from zero to u_{max}. Composite Simpson quadrature is then used to approximate the integral on this grid. In practice, the integral is replaced by a weighted sum over points that are equally spaced between zero and u_{max}. Surface calibration uses 512 intervals. Final inception prices use 2,048 intervals. The interpolation surfaces used for intermediate option values use 384 intervals. These are numerical settings.

The distribution function in Equation (3.45) is recovered with the Gil-Pelaez inversion formula (Gil-Pelaez, 1951). The integral is computed on $u \in [0, 220]$ with 2,048 Simpson intervals. The first grid point is set to 10^{-12} instead of zero to avoid evaluating the term $1/u$ at $u = 0$. The strike $K_{\text{H}}(p, T)$ is found by repeatedly narrowing the admissible strike interval until the model probability in Equation (3.46) matches the target probability p. The tolerance is 10^{-8}, with an iteration cap of 80. If the target probability cannot be reached in the admissible interval, the contract is excluded.

A.4 Heston Physical Simulation

The physical variance process is simulated with the Quadratic-Exponential scheme of L. Andersen (2008). This scheme is used because it preserves the non-negativity of variance while matching the first two conditional moments of the process. The one-year horizon is divided into 252 daily steps. The threshold $\psi_c = 1.5$ selects the variance approximation used at each step. The values $\gamma_1 = \gamma_2 = 0.5$ give equal weight to the variance at the beginning and at the end of the step in the log-price update. This allows the dated Heston estimates to be simulated even when the sufficient Feller condition in Equation (3.34) is not satisfied.

APPENDIX B

Calibration and Robustness Details

B.1 Scope of the Appendix

This appendix reports the calibration details used in the empirical analysis. Section B.2 describes the dated option chains and filtering rules. Section B.3 defines the price-space weights used in the calibration objective. Section B.4 reports the parameter domains and calibration procedure. Sections B.5 and B.6 report the dated parameter estimates and the option-surface fit. Section B.7 reports the local robustness tests and identification diagnostics. The model equations are given in Chapter 3.

B.2 Calibration Inputs and Surface Filtering

The calibrations use the option chains dated 16 March 2020, 8 April 2025 and 30 June 2026. Each date is treated as a separate experiment.

Calibration is performed on observed SPX option prices. A contract is retained only if the underlying level, strike, maturity and calibration price are positive. The calibration price is the reported mid price when available, because bid and ask prices include the execution side of the quote. Using the mid price avoids calibrating the surface to either the selling or buying price. If the mid price is missing, the bid-ask midpoint is used when both quotes are valid. Contracts with no positive usable price are removed.

The same surface domain is applied to all models:

$$0.20 \leq \frac{K}{S_0} \leq 1.40, \qquad 0 < T \leq \tau_t^+, \qquad \tau_t^+ = \min\{\tau \text{ listed at } t \mid \tau \geq 1\}. \tag{B.1}$$

The maturity cutoff τ_t^+ is the first listed expiry that reaches or exceeds one year. If no contract expires exactly at one year, the first expiry beyond one year is retained. Later expiries are excluded.

After this moneyness and maturity filter, an additional option-type filter is applied relative to the forward $F_T = S_0 e^{(r-q_T)T}$:

$$K_j \leq 1.03F_T \quad \text{for puts,} \qquad K_j \geq 0.97F_T \quad \text{for calls.} \tag{B.2}$$

This filter keeps the calibration focused on out-of-the-money options. Deep in-the-money option prices are mostly driven by intrinsic value and have lower vega, meaning lower sensitivity to changes in volatility. They are less informative for the calibration of volatility and tail parameters. Including them could add noise and distort the fitted parameters. The calibration surface includes puts up to three percent above the forward and calls down to three percent below the forward, keeping a small buffer around the at-the-money region. All contracts remaining after these filters are used.

Table B.1: Calibration-surface filtering summary by date.

Date	**Raw Quotes**	**Valid Quotes**	**Calibrated Surface**	**Expiries**
2020-03-16	5,526	5,525	2,027	12
2025-04-08	8,774	8,774	3,761	13
2026-06-30	10,981	9,981	5,041	14

For each expiry, the dividend yield is estimated from put-call parity using the matched call and put closest to the money, as in Equation (3.44). Estimates below -2%, above 10%, or based on maturities shorter than 30 calendar days are replaced by the median valid estimate for the same calibration date. If no valid parity pair is available for an expiry, q_T is set to zero. The scalar dividend yield q used for contract construction and pricing is then the simple average of the q_T estimates across retained expiries.

These filters are used to build a calibration sample that is comparable across dates and models. They remove contracts with unreliable prices, exclude maturities outside the one-year calibration horizon and keep the surface focused on options that are informative about volatility and tail risk. This reduces the influence of quotes that are mainly driven by intrinsic value or by data quality issues.

B.3 Surface Weights

The price-space calibration objective is defined in Equation (3.49). The weights give more importance to reliable quotes while preventing an expiry or contract type with many listed contracts from dominating the calibration. Without this adjustment, the model could fit mainly the most populated part of the surface instead of balancing information across maturities and strikes.

For contract j, let P_j^{mkt} denote the observed market price used in the calibration. The uncertainty assigned to this price is

$$d_j = \max\left\{a_j - b_j,\ 0.02P_j^{\mathrm{mkt}},\ 10^{-6}\right\}, \qquad \widetilde{\omega}_j = d_j^{-2}, \tag{B.3}$$

where a_j and b_j are the ask and bid prices. The first term captures the bid-ask spread. The second term imposes a minimum uncertainty equal to two percent of the market price, so that quotes with very small spreads do not receive excessive weight. The last term is a numerical floor. The weight is then set to d_j^{-2}, so quotes with lower pricing uncertainty receive more weight in the calibration.

The contract weights are first rescaled separately within each expiry. If $\mathscr{E}_m$ denotes the set of contracts with maturity τ_m, the rescaled weight of contract j is

$$\frac{\widetilde{\omega}_j}{\sum_{\ell\in\mathscr{E}_m}\widetilde{\omega}_\ell}, \qquad j \in \mathscr{E}_m. \tag{B.4}$$

The weights of all contracts with the same maturity therefore sum to one. This prevents an expiry with many contracts from receiving more weight only because more quotes are listed at that maturity.

Each expiry is then assigned a maturity weight $\Delta\tau_m$. Let $\bar{\tau}_m = \min(\tau_m, 1)$ denote the retained maturity capped at one year, with maturities ordered increasingly. The maturity weights are defined by the midpoints between neighbouring retained maturities:

$$\Delta\tau_m = \begin{cases} \dfrac{\bar{\tau}_1 + \bar{\tau}_2}{2}, & m = 1, \\ \dfrac{\bar{\tau}_{m+1} - \bar{\tau}_{m-1}}{2}, & 1 < m < M, \\ 1 - \dfrac{\bar{\tau}_{M-1} + \bar{\tau}_M}{2}, & m = M. \end{cases} \tag{B.5}$$

An expiry therefore receives more weight when it represents a wider gap in the retained maturity grid.

For a contract in expiry set $\mathscr{E}_m$, the resulting calibration weight is

$$\omega_j^{(0)} = \Delta\tau_m \frac{\widetilde{\omega}_j}{\sum_{\ell\in\mathscr{E}_m}\widetilde{\omega}_\ell}, \qquad j \in \mathscr{E}_m. \tag{B.6}$$

This weight combines the contract weight within its expiry with the maturity weight assigned to that expiry. Finally, all weights are multiplied by the same constant so that their sum equals the number of retained contracts. Since the same multiplier is applied to every contract, this does not change which contracts receive more or less weight in the calibration.

This weighting scheme reflects the idea that quotes with wider effective pricing uncertainty should have less influence on the calibration, while each expiry should contribute to the surface fit without being overweighted only because more contracts are listed at that maturity.

B.4 Calibration Domains and Numerical Procedure

The physical estimators and risk-neutral parameter vectors are described in Sections 3.5.1 and 3.5.2. This section only reports the numerical choices used for the calibrations. BSM is calibrated with a bounded scalar optimization. Merton and Heston are calibrated from several starting values. The retained solution is the converged solution with the lowest finite objective value.

Table B.2 reports the finite numerical domains used to avoid invalid variances, undefined correlations and unstable transforms.

Table B.2: Numerical domains used in the dated calibrations.

Model block	**Parameter**	**Domain**
BSM	σ	$[0.01, 3]$
Merton common	σ_{diff}	$[0.005, 1.50]$
Merton $\mathbb{Q}$	$\lambda^{\mathbb{Q}}$	$[0, 40]$
	$\mu_J^{\mathbb{Q}}$	$[-5, 1]$
	$\sigma_J^{\mathbb{Q}}$	$[10^{-6}, 1]$
Merton $\mathbb{P}$	$\lambda^{\mathbb{P}}$	$[0, \infty)$
	$\mu_J^{\mathbb{P}}$	$[-5, 1]$
	$\sigma_J^{\mathbb{P}}$	$[10^{-6}, 2]$
Heston	v_0	$[10^{-5}, 3]$
	$\kappa^{\mathbb{Q}}$	$[0.02, 40]$
	$\theta^{\mathbb{Q}}$	$[10^{-5}, 3]$
	ξ	$[0.005, 6]$
	ρ	$[-0.999, 0.999]$

Note. Volatility, intensity and mean-reversion parameters are annualized.

B.5 Dated Calibration Estimates

Table B.3 reports the calibrated parameters used for 16 March 2020, 8 April 2025 and 30 June 2026. The estimates are shown by model and by implementation block.

Table B.3: Calibration parameter estimates by model and date.

Model	Block	Parameter	2020-03-16	2025-04-08	2026-06-30
BSM	Common	σ	0.3847	0.2345	0.1500
	$\mathbb{P}$	$\mu^{\mathbb{P}}$	-0.1339	-0.0307	0.2025
	$\mathbb{Q}$	r	0.0026	0.0440	0.0363
	$\mathbb{Q}$	q	-0.0005	0.0133	0.0013
Merton	Common	σ_{diff}	0.1942	0.1483	0.1182
	$\mathbb{P}$	$\mu^{\mathbb{P}}$	-0.1339	-0.0307	0.2025
	$\mathbb{P}$	$\lambda^{\mathbb{P}}$	13.0941	5.3789	21.7784
	$\mathbb{P}$	$\mu_J^{\mathbb{P}}$	-0.0198	-0.0351	-0.0041
	$\mathbb{P}$	$\sigma_J^{\mathbb{P}}$	0.0586	0.0158	0.0117
	$\mathbb{Q}$	r	0.0026	0.0440	0.0363
	$\mathbb{Q}$	q	-0.0005	0.0133	0.0013
	$\mathbb{Q}$	$\lambda^{\mathbb{Q}}$	0.4251	0.3478	0.1123
	$\mathbb{Q}$	$\mu_J^{\mathbb{Q}}$	-0.7698	-0.4172	-0.4548
	$\mathbb{Q}$	$\sigma_J^{\mathbb{Q}}$	0.1002	0.2340	0.3276
Heston	Common	v_0	0.6778	0.2818	0.0190
	Common	ξ	1.8071	2.4625	1.6799
	Common	ρ	-0.9990	-0.7816	-0.6258
	$\mathbb{P}$	$\mu^{\mathbb{P}}$	-0.1339	-0.0307	0.2025
	$\mathbb{P}$	$\kappa^{\mathbb{P}}$	12.1578	13.4173	13.6669
	$\mathbb{P}$	$\theta^{\mathbb{P}}$	0.0229	0.0300	0.0300
	$\mathbb{Q}$	r	0.0026	0.0440	0.0363
	$\mathbb{Q}$	q	-0.0005	0.0133	0.0013
	$\mathbb{Q}$	$\kappa^{\mathbb{Q}}$	3.7394	7.4017	3.7690
	$\mathbb{Q}$	$\theta^{\mathbb{Q}}$	0.0655	0.0509	0.0596
	Risk premium	λ_0	0.0183	0.0103	0.1101
	Risk premium	λ_1	-4.6584	-2.4429	-5.8920
	Risk premium	$\mathbb{E}[\bar{V}^{\mathbb{P}} - \bar{V}^{\mathbb{Q}}]$	-0.1486	-0.0333	-0.0199

Note. Parameters are annualized. Variance levels are in decimal units and λ is an annual jump intensity.

B.6 Option-Surface Fit

Table B.4 reports RMSE and MAE on the retained contracts in observed SPX option-price units. They measure how well each model fits the dated option surface.

Table B.4: Option-surface price fit by model and calibration date.

Date	Model	RMSE	MAE
2020-03-16	BSM	70.88	57.35
	Merton	72.90	53.09
	Heston	12.73	7.63
2025-04-08	BSM	60.92	48.17
	Merton	57.33	42.22
	Heston	8.46	5.94
2026-06-30	BSM	47.90	33.19
	Merton	13.27	9.81
	Heston	6.53	5.13

Note. RMSE and MAE are expressed in SPX option-price units.

B.7 Parameter Robustness and Identification

B.7.1 Calibration Robustness Test

The robustness test is a Monte Carlo recovery exercise around each dated calibration. It uses a representative subset of the retained surface rather than all contracts, so that recalibrations remain feasible while keeping the main maturity and moneyness regions. In each run, contracts are randomly selected within maturity and moneyness groups. They are priced with the calibrated reference parameters, small random noise is added to mimic quote uncertainty, and the model is recalibrated to this perturbed sample.

The recovered parameters are then compared with the reference parameters. Parameter RMSE and MAE report how much the estimates move across the perturbed samples. Price errors report whether these parameter changes lead to different fitted option prices, using the reference model prices on the capped subset as the benchmark. Physical parameters, risk-free rates and dividend yields are kept fixed. The robustness figures report the three calibration dates, with the dashed line showing the reference estimate.

B.7.2 Recovery Results by Calibration Date

Table B.5 reports parameter recovery errors. Each recovered parameter vector is compared with the parameter vector obtained from the full option surface. The RMSE and MAE are reported in each parameter's own units.

Table B.6 reports price recovery errors. For each recovered parameter vector, the model prices are compared with the prices generated by the reference calibrated vector on the capped subset. These errors are expressed in SPX option-price units.

Table B.5: Parameter-recovery errors on stratified contract slices by model and calibration date.

Date	Model	Parameter	Recovery RMSE	Recovery MAE
2020-03-16	BSM	σ	0.001063	0.000872
2020-03-16	Merton	σ_{diff}	0.005625	0.004281
		$\lambda^{\mathbb{Q}}$	0.033825	0.023962
		$\mu_J^{\mathbb{Q}}$	0.036025	0.027042
		$\sigma_J^{\mathbb{Q}}$	0.060210	0.048614
2020-03-16	Heston	v_0	0.031691	0.025635
		$\kappa^{\mathbb{Q}}$	0.326695	0.263912
		$\theta^{\mathbb{Q}}$	0.005999	0.004914
		ξ	0.096033	0.076466
		ρ	0.017719	0.009751
2025-04-08	BSM	σ	0.000277	0.000222
2025-04-08	Merton	σ_{diff}	0.002323	0.001958
		$\lambda^{\mathbb{Q}}$	0.032167	0.027044
		$\mu_J^{\mathbb{Q}}$	0.031967	0.026496
		$\sigma_J^{\mathbb{Q}}$	0.017133	0.013440
2025-04-08	Heston	v_0	0.007493	0.006124
		$\kappa^{\mathbb{Q}}$	0.352324	0.290370
		$\theta^{\mathbb{Q}}$	0.000827	0.000684
		ξ	0.095230	0.080793
		ρ	0.007500	0.006240
2026-06-30	BSM	σ	0.000074	0.000059
2026-06-30	Merton	σ_{diff}	0.000254	0.000204
		$\lambda^{\mathbb{Q}}$	0.001775	0.001432
		$\mu_J^{\mathbb{Q}}$	0.007395	0.005971
		$\sigma_J^{\mathbb{Q}}$	0.004175	0.003344
2026-06-30	Heston	v_0	0.000207	0.000166
		$\kappa^{\mathbb{Q}}$	0.081402	0.065925
		$\theta^{\mathbb{Q}}$	0.000306	0.000255
		ξ	0.021901	0.017729
		ρ	0.002172	0.001732

Note. Errors are in each parameter's own units.

Table B.6: Option-price stability of the recovered parameter vectors by model and calibration date.

Date	Model	Price RMSE	Price MAE
2020-03-16	BSM	0.5408	0.3726
	Merton	1.5316	1.1173
	Heston	2.0571	1.5210
2025-04-08	BSM	0.2837	0.1835
	Merton	0.6281	0.4518
	Heston	0.9521	0.6741
2026-06-30	BSM	0.0985	0.0604
	Merton	0.1579	0.1089
	Heston	0.2226	0.1509

Note. Price errors are expressed in SPX option-price units.

The two tables provide complementary diagnostics. Table B.5 highlights how much the calibrated parameters move when the option sample and quotes are perturbed. Table B.6 shows whether these parameter changes lead to different option prices.

The recovery errors are lower in 2026, which is consistent with a calmer option surface. Errors are higher in 2020 and 2025, where the crisis and stressed surfaces make parameter recovery more difficult. Part of the difference can also come from bid-ask spreads and from the number of retained contracts.

The price errors in Table B.6 remain much smaller than the original surface fit errors in Table B.4. This shows that the recovered parameter vectors preserve the fitted price surface reasonably well. The robustness exercise supports local price stability, but the separate interpretation of Merton and Heston parameters remains cautious.

B.7.3 BSM

The BSM results show limited dispersion around the reference volatility at each calibration date. The boxplots in Figure B.1 are narrow and centred close to the reference estimates. The recovered volatility is higher in 2020, lower in 2025 and lowest in 2026, which is consistent with the crisis, stressed and calmer calibration dates.

This stability is expected because BSM has only one parameter to recover from the option surface. Small recovery errors indicate that the calibration procedure recovers the reference volatility when the input quotes are perturbed. They do not imply that the constant-volatility model provides a good fit to the full cross-section of option prices.

Table B.5 measures the local stability of the recovered parameter, while Table B.4 measures the fit of the model to the observed option surface. The BSM volatility can be stable in the recovery test even when the constant-volatility assumption fails to reproduce the option smile.

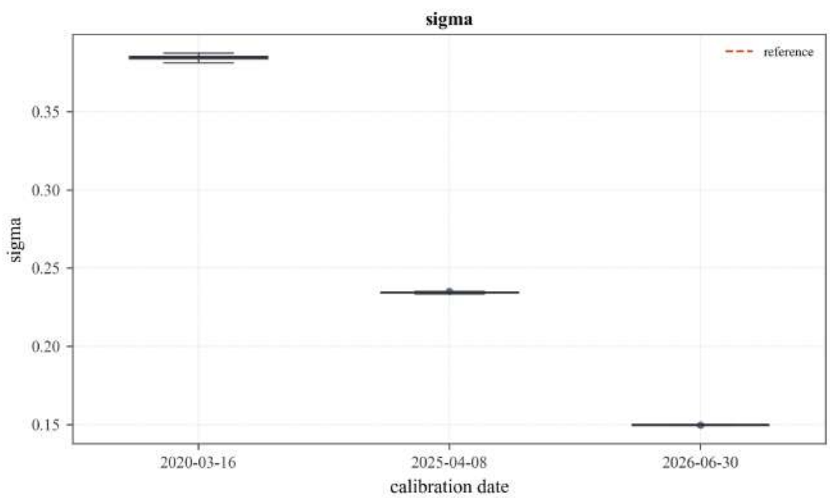

Figure B.1: BSM volatility recovery by calibration date.

B.7.4 Merton

The Merton recovery results depict a clear difference across calibration dates. The 2026 boxes are narrow for σ_{diff}, $\lambda^{\mathbb{Q}}$, $\mu_J^{\mathbb{Q}}$ and $\sigma_J^{\mathbb{Q}}$. The recovered values remain close to the reference estimates, which indicates local stability under the perturbation design. The 2020 and 2025 boxes are wider, especially for the jump parameters, which is consistent with the crisis and stressed option surfaces.

The diffusion volatility σ_{diff} is recovered with limited dispersion at all dates. It is highest in 2020, lower in 2025 and lowest in 2026, in line with the three market environments. The jump intensity $\lambda^{\mathbb{Q}}$ and the mean jump size $\mu_J^{\mathbb{Q}}$ also remain close to their reference estimates, although they move more in 2020 and 2025. This reflects the fact that the option surface can fit similar jump effects with different combinations of jump frequency and jump severity.

The jump-dispersion parameter $\sigma_J^{\mathbb{Q}}$ is the least stable jump parameter in 2020. Its recovery distribution is wide. It indicates that $\sigma_J^{\mathbb{Q}}$ is weakly identified on the 2020 crash surface. In particular, different combinations of $\lambda^{\mathbb{Q}}$, $\mu_J^{\mathbb{Q}}$ and $\sigma_J^{\mathbb{Q}}$ can generate similar jump contributions to option prices. The wide dispersion suggests that $\sigma_J^{\mathbb{Q}}$ is difficult to identify separately from the other jump parameters.

The 2026 recovery results are more stable, but they do not prove that the four Merton parameters are separately identified. The recovered jump parameters co-move strongly in 2026. The correlations are 0.995 between $\lambda^{\mathbb{Q}}$ and $\mu_J^{\mathbb{Q}}$, 0.918 between $\lambda^{\mathbb{Q}}$ and $\sigma_J^{\mathbb{Q}}$, and 0.946 between $\mu_J^{\mathbb{Q}}$ and $\sigma_J^{\mathbb{Q}}$. The low recovery RMSE values of 0.000105 for the jump-loss intensity and 0.000115 for the jump second moment indicate that the main jump effects are stable, even if the individual jump parameters move together.

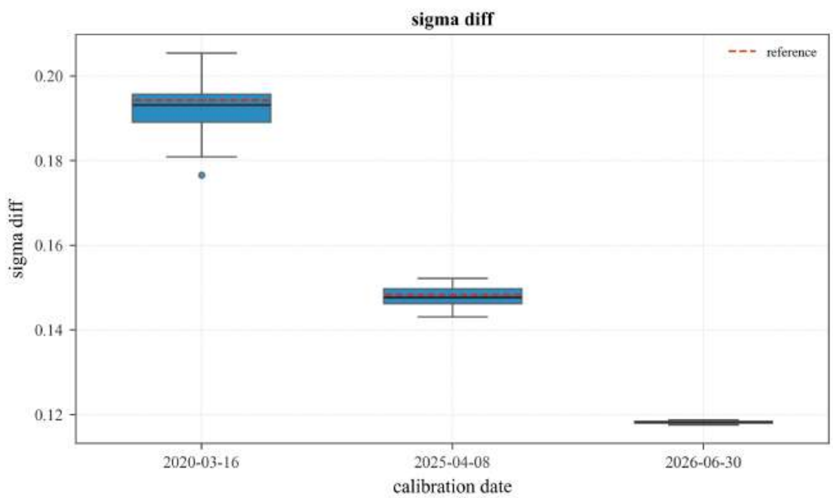

Figure B.2: Merton diffusion-volatility recovery by calibration date.

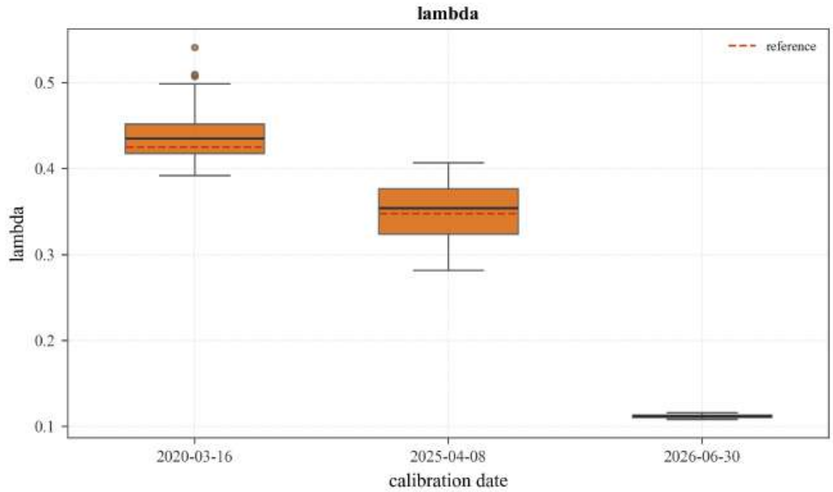

Figure B.3: Merton risk-neutral jump-intensity recovery by calibration date.

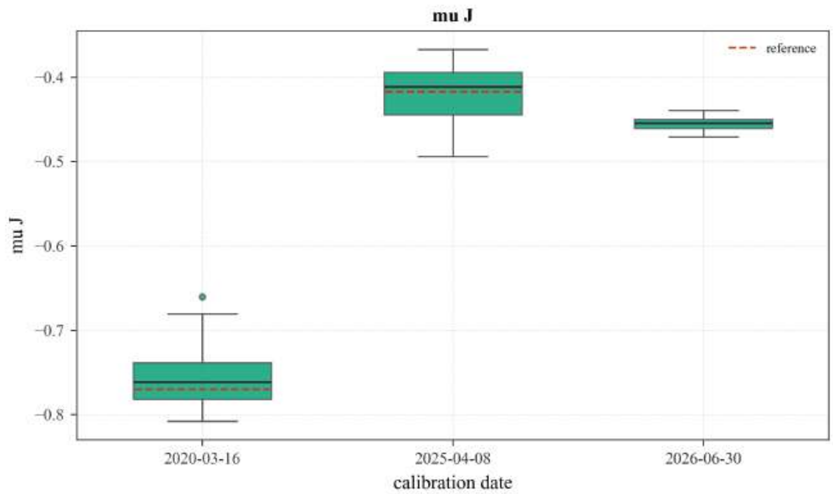

Figure B.4: Merton risk-neutral mean jump-size recovery by calibration date.

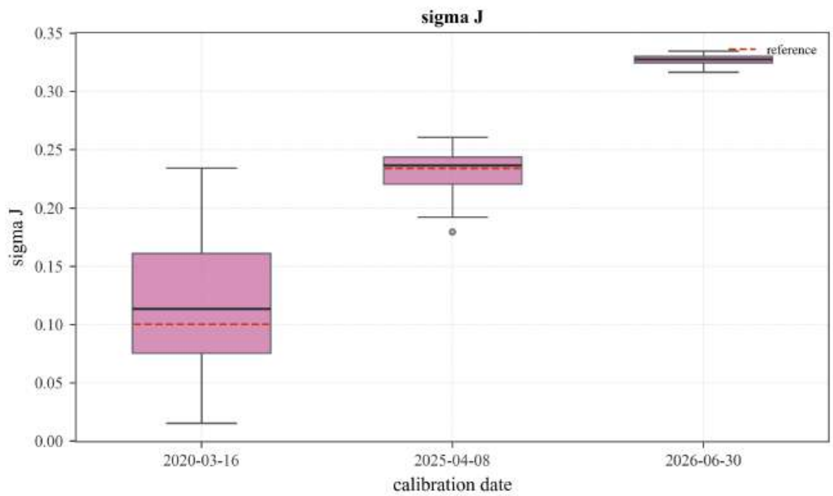

Figure B.5: Merton risk-neutral jump-dispersion recovery by calibration date.

B.7.5 Heston

The Heston recovery results differ across calibration dates. In 2026, the recovered Heston parameters stay close to the reference estimates for v_0, $\kappa^{\mathbb{Q}}$, $\theta^{\mathbb{Q}}$, ξ and ρ. The 2020 and 2025 results are more dispersed, especially for $\kappa^{\mathbb{Q}}$ and ξ. This shows that variance parameters are harder to recover on crisis and stressed surfaces.

The initial variance v_0 is highest in 2020, lower in 2025 and lowest in 2026, which is consistent with the volatility level of the three dates. The long-run variance $\theta^{\mathbb{Q}}$ remains more stable, while $\kappa^{\mathbb{Q}}$ and ξ move more. This is expected because different combinations of mean reversion and volatility of variance can produce similar option prices.

In 2020, the recovered values of ρ are close to the lower numerical boundary. This should not be interpreted as evidence that the true spot-variance correlation is exactly -1. It shows calibration tension on the crash surface.

The robustness test confirms that the fitted option prices remain stable under quote perturbations. However, this stability does not imply that the individual Heston parameters are uniquely recovered, since different parameter combinations may generate similar option prices. This is why the recovery table and the price-stability table must be interpreted together.

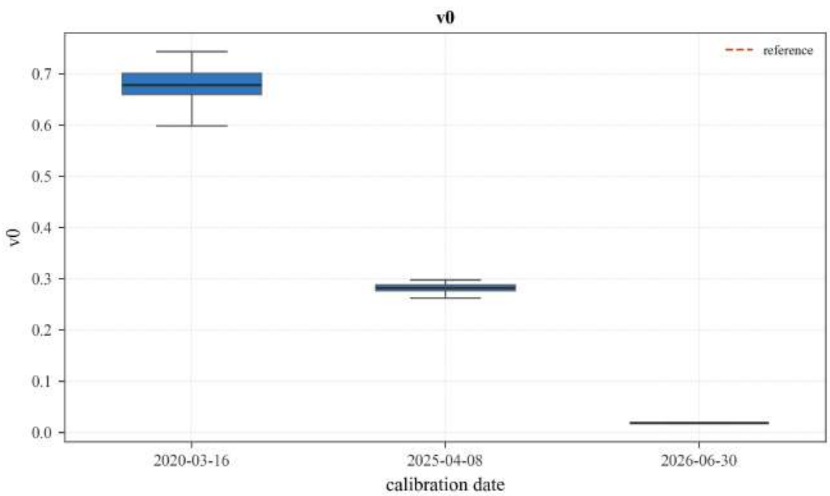

Figure B.6: Heston initial-variance recovery by calibration date.

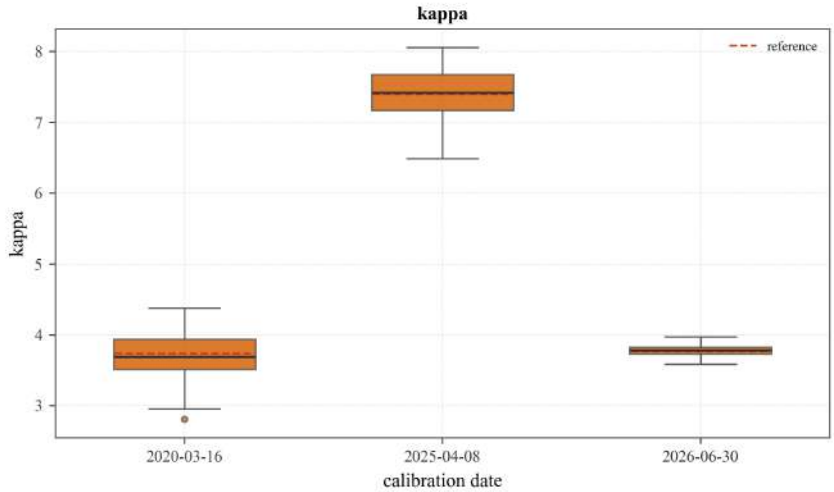

Figure B.7: Heston risk-neutral mean-reversion recovery by calibration date.

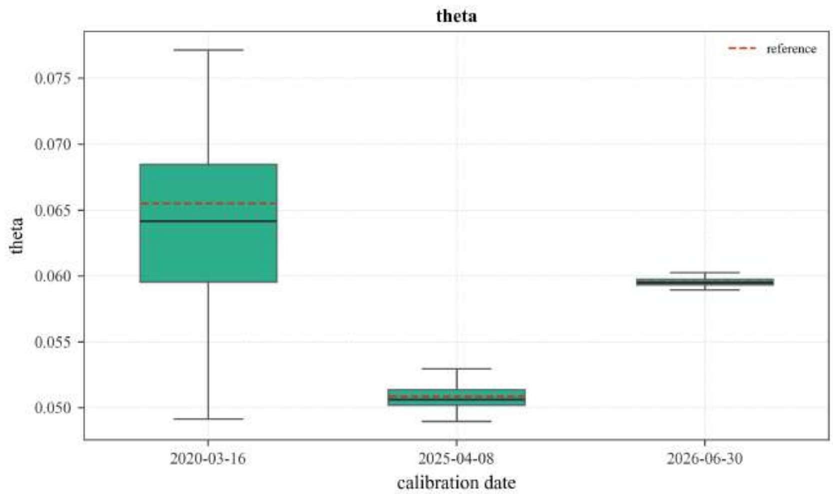

Figure B.8: Heston risk-neutral long-run variance recovery by calibration date.

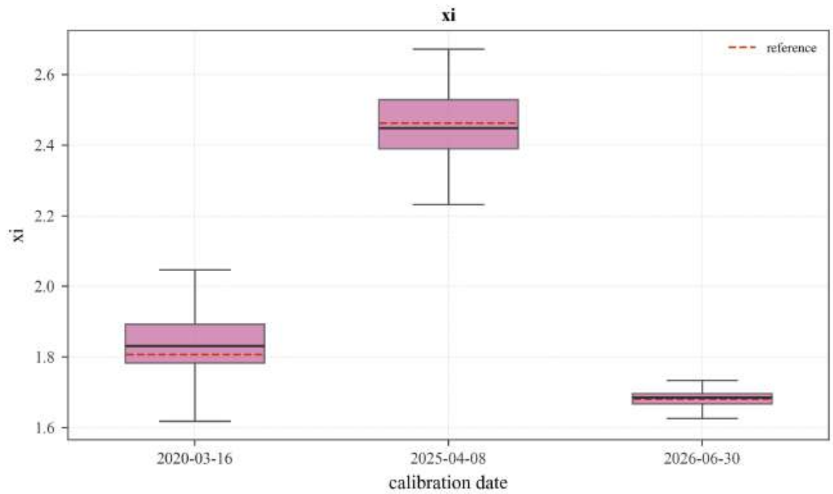

Figure B.9: Heston volatility-of-variance recovery by calibration date.

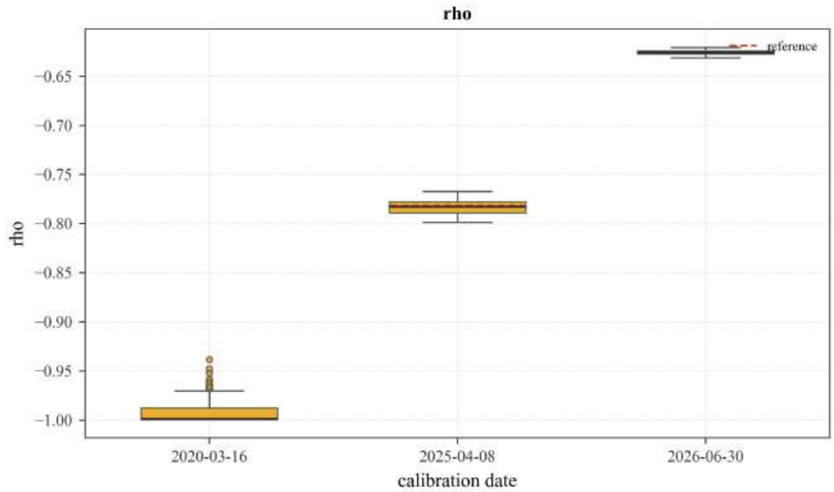

Figure B.10: Heston spot-variance correlation recovery by calibration date.

APPENDIX C

Complementary Results

This appendix provides complementary results and robustness checks that support the main analysis in Chapter 5. It first examines the sensitivity of the optimizer to alternative risk objectives, including $\mathrm{ES}_{99\%}$, $\mathrm{VaR}_{95\%}$, $\mathrm{VaR}_{99\%}$ and L_{MDD}. It then reports rolling-hedge results, a one-percent transaction-cost test and a protection-budget sensitivity analysis.

C.1 Risk-Metric Sensitivity

The benchmark objective is re-estimated after replacing only its risk term, while keeping the simulated paths, contract universe, budget limit and return term unchanged. The alternatives are $\mathrm{ES}_{99\%}$, $\mathrm{VaR}_{95\%}$, $\mathrm{VaR}_{99\%}$ and L_{MDD}. The resulting allocations are compared with the $\mathrm{ES}_{95\%}$ benchmark in Table 5.1, using the same panels, columns and reporting conventions.

The allocation sensitivity depends on the selected risk metric and the model. The largest changes relative to the $\mathrm{ES}_{95\%}$ benchmark are observed for $\mathrm{ES}_{99\%}$ and $\mathrm{VaR}_{95\%}$, which emphasize different parts of the loss distribution. $\mathrm{ES}_{99\%}$ focuses on more extreme losses and can shift the allocation towards stronger tail protection, whereas $\mathrm{VaR}_{95\%}$ targets the loss threshold itself and can lead to different strike and weight choices. $\mathrm{VaR}_{99\%}$ generally remains closer to the benchmark allocation. The L_{MDD} objective leads to a different type of allocation because it targets pathwise drawdowns rather than terminal tail losses.

C.1.1 99% Expected Shortfall

Table C.1: One-year model-based results optimized with 99% Expected Shortfall, by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	**Model**	**Portfolio**	**Budget**	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	5.00	-9.04	-3.73	25.87	25.87	25.87	25.87
	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	4.98	-15.17	2.39	38.73	38.76	38.75	38.77
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	5.00	-14.33	1.68	35.13	35.13	35.13	35.13

Table C.1 continued

2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.95	-4.12	-3.08	18.87	18.87	18.87	18.87
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	5.00	-9.02	1.90	20.05	20.63	20.41	20.91
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.34	-7.33	0.20	19.96	19.96	19.96	19.96
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	1.12	16.84	1.24	9.82	15.13	13.37	15.13
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	2.19	15.95	2.21	17.82	21.36	20.85	21.36

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation.

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	29.65	2.32	7.54
	Merton	Index	33.16	0.81	1.20
		DOOM-put	30.81	1.26	2.00
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	24.37	0.22	-1.06
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	21.32	8.93	143.42
	Merton	Index	21.14	0.46	0.33
		DOOM-put	18.31	1.32	1.59
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.77	0.53	-0.58
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	11.03	0.52	0.31
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	12.63	-0.20	-0.24

Note. L_{MDD} is reported in percentage points. γ_1 is skewness and γ_2 is excess kurtosis. Dashes indicate zero allocation.

Table C.1 continued

Panel C. Selected DOOM-put positions

Date	Model	Portfolio	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	1.00	0.320	77.81	5.127	2.43	0.474
		DOOM-put	2	1.00	0.330	78.65	5.398	2.57	0.476
	Merton	DOOM-put	1	1.00	0.225	59.25	3.837	0.83	0.217
		DOOM-put	2	1.00	0.260	63.72	4.920	1.19	0.242
		DOOM-put	3	1.00	0.270	65.17	5.305	2.20	0.414
		DOOM-put	4	1.00	0.310	72.61	7.466	0.21	0.029
		DOOM-put	5	1.00	0.360	86.54	12.127	0.55	0.045
	Heston	DOOM-put	1	1.00	0.200	67.58	5.075	3.87	0.763
		DOOM-put	2	1.00	0.225	72.07	6.028	1.13	0.187
2025-04-08	BSM	DOOM-put	1	0.50	0.005	65.35	0.016	0.33	20.530
		DOOM-put	2	0.67	0.010	64.20	0.038	0.92	24.034
		DOOM-put	3	0.75	0.010	62.50	0.039	0.17	4.398
		DOOM-put	4	1.00	0.020	61.98	0.095	0.20	2.063
		DOOM-put	5	1.00	0.290	88.12	3.199	0.59	0.185
		DOOM-put	6	1.00	0.300	88.72	3.368	0.39	0.115
		DOOM-put	7	1.00	0.310	89.31	3.542	1.06	0.300
		DOOM-put	8	1.00	0.320	89.91	3.721	1.30	0.349
	Merton	DOOM-put	1	1.00	0.260	89.24	5.689	5.00	0.879
	Heston	DOOM-put	1	1.00	0.225	87.44	4.541	4.34	0.957
2026-06-30	BSM	DOOM-put	1	1.00	0.175	89.01	1.134	1.12	0.989
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	1.00	0.125	82.79	2.157	1.84	0.855
		DOOM-put	2	1.00	0.150	86.18	2.605	0.19	0.072
		DOOM-put	3	1.00	0.175	89.00	3.045	0.16	0.052

Note. w_j is reported as percentage of B_0. All contracts use normalized $S_0 = B_0 = 100$, and $n_j = w_j B_0 / \pi_j$ with w_j expressed in decimal form. Dashes indicate zero allocation.

C.1.2 95% Value at Risk

Table C.2: One-year model-based results optimized with 95% Value at Risk, by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	Model	Portfolio	Budget	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	5.00	-9.04	-3.73	25.87	25.87	25.87	25.87

Table C.2 continued

	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	5.00	-12.91	0.12	37.13	42.78	40.62	45.12
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	4.91	-12.66	0.01	36.16	44.39	41.10	47.93
2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.94	-4.03	-3.17	18.82	18.93	18.89	18.97
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	4.99	-9.02	1.90	20.08	20.67	20.44	20.96
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.34	-7.33	0.20	19.96	19.96	19.96	19.96
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	0.00	–	–	–	–	–	–
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	1.64	16.79	1.37	13.52	27.06	21.71	34.43

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation.

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	29.65	2.32	7.54
	Merton	Index	33.16	0.81	1.20
		DOOM-put	28.96	1.42	2.88
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	23.43	0.11	-0.53
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	21.39	12.28	250.54
	Merton	Index	21.14	0.46	0.33
		DOOM-put	18.32	1.32	1.58
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.77	0.53	-0.58
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	–	–	–
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	13.29	-0.27	0.18

Note. L_{MDD} is reported in percentage points. γ_1 is skewness and γ_2 is excess kurtosis. Dashes indicate zero allocation.

Table C.2 continued

Panel C. Selected DOOM-put positions

Date	Model	Portfolio	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	1.00	0.320	77.81	5.127	2.43	0.474
		DOOM-put	2	1.00	0.330	78.65	5.398	2.57	0.476
	Merton	DOOM-put	1	1.00	0.370	88.97	13.014	5.00	0.384
	Heston	DOOM-put	1	0.17	0.175	73.20	2.741	0.11	0.039
		DOOM-put	2	1.00	0.320	86.66	9.967	0.11	0.011
		DOOM-put	3	1.00	0.340	89.35	10.849	4.69	0.432
2025-04-08	BSM	DOOM-put	1	0.50	0.005	65.35	0.016	0.55	34.454
		DOOM-put	2	0.75	0.005	59.41	0.017	0.49	28.193
		DOOM-put	3	0.75	0.010	62.50	0.039	0.50	12.814
		DOOM-put	4	1.00	0.310	89.31	3.542	1.78	0.504
		DOOM-put	5	1.00	0.320	89.91	3.721	1.61	0.433
	Merton	DOOM-put	1	1.00	0.260	89.24	5.689	4.99	0.877
	Heston	DOOM-put	1	1.00	0.225	87.44	4.541	4.34	0.957
2026-06-30	BSM	DOOM-put	–	–	–	–	–	–	–
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	0.25	0.075	89.13	0.758	0.10	0.134
		DOOM-put	2	0.50	0.125	89.69	1.693	0.44	0.258
		DOOM-put	3	0.67	0.150	89.93	2.262	1.11	0.489

Note. w_j *is reported as percentage of* B_0*. All contracts use normalized* $S_0 = B_0 = 100$*, and* $n_j = w_j B_0 / \pi_j$ *with* w_j *expressed in decimal form. Dashes indicate zero allocation.*

C.1.3 99% Value at Risk

Table C.3: One-year model-based results optimized with 99% Value at Risk, by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	Model	Portfolio	Budget	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	5.00	-9.04	-3.73	25.87	25.87	25.87	25.87
	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	5.00	-12.91	0.12	37.13	42.78	40.62	45.12
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	5.00	-14.33	1.68	35.13	35.13	35.13	35.13
2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.89	-4.05	-3.16	19.15	19.38	19.29	19.47

Table C.3 continued

	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	5.00	-9.02	1.90	20.05	20.63	20.41	20.91
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.34	-7.33	0.20	19.96	19.96	19.96	19.96
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	1.12	16.84	1.24	9.82	15.13	13.37	15.13
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	2.95	15.29	2.87	16.71	16.71	16.71	16.71

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation.

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	29.65	2.32	7.54
	Merton	Index	33.16	0.81	1.20
		DOOM-put	28.96	1.42	2.88
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	24.37	0.22	-1.06
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	21.50	12.49	257.86
	Merton	Index	21.14	0.46	0.33
		DOOM-put	18.31	1.32	1.59
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.77	0.53	-0.58
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	11.03	0.52	0.31
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	11.85	-0.08	-0.38

Note. L_{MDD} is reported in percentage points. γ_1 is skewness and γ_2 is excess kurtosis. Dashes indicate zero allocation.

Table C.3 continued

Panel C. Selected DOOM-put positions

Date	Model	Portfolio	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	1.00	0.320	77.81	5.127	2.43	0.474
		DOOM-put	2	1.00	0.330	78.65	5.398	2.57	0.476
	Merton	DOOM-put	1	1.00	0.370	88.97	13.014	5.00	0.384
	Heston	DOOM-put	1	1.00	0.200	67.58	5.075	3.87	0.763
		DOOM-put	2	1.00	0.225	72.07	6.028	1.13	0.187
2025-04-08	BSM	DOOM-put	1	0.50	0.005	65.35	0.016	0.55	34.222
		DOOM-put	2	0.75	0.005	59.41	0.017	0.54	30.814
		DOOM-put	3	0.75	0.010	62.50	0.039	0.47	12.087
		DOOM-put	4	1.00	0.310	89.31	3.542	1.83	0.516
		DOOM-put	5	1.00	0.320	89.91	3.721	1.51	0.405
	Merton	DOOM-put	1	1.00	0.260	89.24	5.689	5.00	0.879
	Heston	DOOM-put	1	1.00	0.225	87.44	4.541	4.34	0.957
2026-06-30	BSM	DOOM-put	1	1.00	0.175	89.01	1.134	1.12	0.989
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	1.00	0.175	89.00	3.045	2.95	0.970

Note. w_j is reported as percentage of B_0. All contracts use normalized $S_0 = B_0 = 100$, and $n_j = w_j B_0 / \pi_j$ with w_j expressed in decimal form. Dashes indicate zero allocation.

C.1.4 Maximum Drawdown

Table C.4: One-year model-based results optimized with average maximum drawdown, by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	Model	Portfolio	Budget	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	4.83	-8.03	-4.73	31.95	32.21	32.11	32.24
	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	5.00	-12.91	0.12	37.13	42.78	40.62	45.12
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	5.00	-12.62	-0.03	35.76	43.83	40.66	47.30
2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.66	-4.64	-2.56	17.75	18.05	17.94	18.08
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	0.00	–	–	–	–	–	–

Table C.4 continued

	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.83	-7.35	0.22	20.01	20.30	20.19	20.34
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	0.00	–	–	–	–	–	–
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	0.00	–	–	–	–	–	–

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation.

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	32.00	1.90	5.10
	Merton	Index	33.16	0.81	1.20
		DOOM-put	28.96	1.42	2.88
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	23.28	0.13	-0.52
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	18.50	1.79	4.09
	Merton	Index	21.14	0.46	0.33
		DOOM-put	–	–	–
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.44	0.61	-0.44
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	–	–	–
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	–	–	–

Note. L_{MDD} is reported in percentage points. γ_1 is skewness and γ_2 is excess kurtosis. Dashes indicate zero allocation.

Table C.4 continued

Panel C. Selected DOOM-put positions

Date	**Model**	**Portfolio**	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	0.50	0.005	47.89	0.019	0.10	5.420
		DOOM-put	2	0.75	0.075	58.69	0.578	0.25	0.430
		DOOM-put	3	1.00	0.030	45.18	0.177	0.88	4.989
		DOOM-put	4	1.00	0.050	49.47	0.346	0.14	0.391
		DOOM-put	5	1.00	0.225	69.66	2.915	1.83	0.626
		DOOM-put	6	1.00	0.250	71.86	3.437	0.34	0.100
		DOOM-put	7	1.00	0.260	72.73	3.657	0.33	0.090
		DOOM-put	8	1.00	0.280	74.44	4.119	0.96	0.233
	Merton	DOOM-put	1	1.00	0.370	88.97	13.014	5.00	0.384
	Heston	DOOM-put	1	1.00	0.340	89.35	10.849	5.00	0.461
2025-04-08	BSM	DOOM-put	1	0.75	0.225	85.99	1.995	0.13	0.065
		DOOM-put	2	0.75	0.250	87.41	2.323	0.18	0.076
		DOOM-put	3	1.00	0.050	68.22	0.294	0.37	1.273
		DOOM-put	4	1.00	0.300	88.72	3.368	0.14	0.041
		DOOM-put	5	1.00	0.310	89.31	3.542	1.12	0.316
		DOOM-put	6	1.00	0.320	89.91	3.721	2.72	0.732
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	1.00	0.225	87.44	4.541	4.83	1.065
2026-06-30	BSM	DOOM-put	–	–	–	–	–	–	–
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	–	–	–	–	–	–	–

Note. w_j *is reported as percentage of* B_0. *All contracts use normalized* $S_0 = B_0 = 100$, *and* $n_j = w_j B_0 / \pi_j$ *with* w_j *expressed in decimal form. Dashes indicate zero allocation.*

C.2 Rolling Hedge

Rolling uses the total premium budget selected by the corresponding static $\mathrm{ES}_{95\%}$ strategy. If this fraction is b_{stat} and the replacement interval is Δ, the premium amount allocated to renewal date τ_ℓ is

$$B_\ell = \frac{b_{\text{stat}}}{T} \min(\Delta, T - \tau_\ell)\, B_0. \tag{C.1}$$

These premium amounts sum to $b_{\text{stat}} B_0$ over the one-year horizon. At $t = 0$, the index position is funded with $(1 - b_{\text{stat}})B_0$. The remaining wealth is kept in a reserve account and is used only to finance future option purchases. The reserve earns the constant risk-free rate until each renewal date. This avoids treating unused premium cash as idle cash and keeps the comparison with the static strategy consistent. Without this reserve account, the rolling strategy would be penalized for not spending the full option budget at inception. Monthly, quarterly and semi-annual replacement therefore produce twelve, four and two purchases over the one-year horizon.

For each replacement frequency, the optimizer selects a fixed allocation rule across the available exercise-probability contracts. At each renewal date, the current SPX level is used as the new normalization point and the same allocation rule is applied to the premium amount available at that date. The quantity of puts purchased can therefore differ across renewals because the SPX level and option premia change over time. BSM and Merton price the new puts with the risk-neutral specification calibrated at the initial date. Heston uses the same initial risk-neutral parameters, but updates the variance state along the simulated path when repricing the puts. If a put pays off before the terminal horizon, the payoff is reinvested in the SPX until the final date.

The rolling allocation rule is determined on the optimization paths and then applied unchanged to the evaluation paths. At each renewal date, the same rule is used with the new SPX level and available premium budget. The strategy is therefore fixed in advance and does not adjust the allocation after observing the realized outcomes.

C.2.1 Rolling Results

The tables report the rolling results obtained for the three calibration dates considered in the thesis. The rolling allocations are obtained by optimizing the same $\mathrm{ES}_{95\%}$-based objective as the static benchmark.

Table C.5: One-year rolling results optimized with 95% Expected Shortfall, without transaction costs.

Panel A. Portfolio performance and terminal downside risk

Date	**Model**	**Portfolio**	**Budget**	$\bar{R}_T$	$\Delta\bar{R}$	$\mathrm{VaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{ES}_{95\%}$	$\mathrm{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put^{1M}	5.00	-11.37	-1.40	50.28	58.53	55.27	61.63
		DOOM-put^{3M}	5.00	-10.66	-2.11	45.34	51.32	48.94	53.01
		DOOM-put^{6M}	5.00	-10.03	-2.74	38.71	39.58	39.23	39.81

Table C.5 continued

	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put^{1M}	5.00	-17.14	4.35	51.68	61.17	57.53	65.10
		DOOM-put^{3M}	5.00	-14.74	1.96	45.81	52.82	50.16	55.68
		DOOM-put^{6M}	5.00	-13.95	1.16	42.24	48.16	45.94	50.69
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put^{1M}	5.00	-12.65	-0.00	48.28	63.44	57.24	69.68
		DOOM-put^{3M}	5.00	-13.10	0.45	45.52	59.08	53.69	65.11
		DOOM-put^{6M}	5.00	-12.85	0.20	42.65	54.57	49.88	59.64
2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put^{1M}	4.88	-6.19	-1.01	33.97	41.99	38.83	45.19
		DOOM-put^{3M}	4.88	-5.79	-1.41	28.92	33.65	31.84	35.38
		DOOM-put^{6M}	4.88	-5.32	-1.88	24.18	25.81	25.16	26.16
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put^{1M}	5.00	-11.67	4.54	34.61	42.32	39.36	46.09
		DOOM-put^{3M}	5.00	-10.32	3.20	30.74	36.31	34.11	38.62
		DOOM-put^{6M}	5.00	-9.65	2.53	27.75	29.74	28.99	30.54
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put^{1M}	4.34	-7.37	0.24	36.54	52.99	46.45	60.63
		DOOM-put^{3M}	4.34	-7.63	0.50	32.01	44.95	39.74	51.48
		DOOM-put^{6M}	4.34	-7.63	0.50	28.95	36.76	33.86	41.33
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put^{1M}	0.00	–	–	–	–	–	–
		DOOM-put^{3M}	0.00	–	–	–	–	–	–
		DOOM-put^{6M}	0.00	–	–	–	–	–	–
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put^{1M}	0.00	–	–	–	–	–	–
		DOOM-put^{3M}	0.00	–	–	–	–	–	–
		DOOM-put^{6M}	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put^{1M}	3.00	16.92	1.25	13.28	28.35	22.66	36.67
		DOOM-put^{3M}	3.00	16.18	1.98	11.70	24.61	19.56	31.45
		DOOM-put^{6M}	3.00	15.71	2.45	11.83	22.52	18.36	27.39

Note. Budget is the total one-year premium envelope selected by the corresponding static strategy. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. Dashes indicate that the static strategy selected no option budget.

Table C.5 continued

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put^{1M}	37.22	1.35	3.34
		DOOM-put^{3M}	35.06	1.54	4.12
		DOOM-put^{6M}	32.94	1.80	5.22
	Merton	Index	33.16	0.81	1.20
		DOOM-put^{1M}	34.36	0.81	1.20
		DOOM-put^{3M}	31.95	0.98	1.58
		DOOM-put^{6M}	30.75	1.12	1.95
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put^{1M}	26.38	-0.02	1.39
		DOOM-put^{3M}	24.67	-0.44	-0.12
		DOOM-put^{6M}	23.60	-0.33	-0.28
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put^{1M}	24.27	0.77	1.07
		DOOM-put^{3M}	21.83	1.07	1.79
		DOOM-put^{6M}	20.28	1.35	2.55
	Merton	Index	21.14	0.46	0.33
		DOOM-put^{1M}	22.81	0.46	0.33
		DOOM-put^{3M}	21.17	0.64	0.50
		DOOM-put^{6M}	19.95	0.86	0.79
	Heston	Index	22.40	-0.81	0.79
		DOOM-put^{1M}	20.04	-0.11	2.69
		DOOM-put^{3M}	17.92	-0.44	0.48
		DOOM-put^{6M}	17.03	-0.10	-0.19
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put^{1M}	–	–	–
		DOOM-put^{3M}	–	–	–
		DOOM-put^{6M}	–	–	–
	Merton	Index	9.98	0.38	0.19
		DOOM-put^{1M}	–	–	–
		DOOM-put^{3M}	–	–	–
		DOOM-put^{6M}	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put^{1M}	13.33	-0.05	1.78
		DOOM-put^{3M}	12.02	-0.20	0.29
		DOOM-put^{6M}	11.80	-0.09	-0.08

Note. L_{MDD} is reported in percentage points. γ_1 is skewness and γ_2 is excess kurtosis.

Table C.5 continued

Panel C. Selected DOOM-put positions

Date	**Model**	**Portfolio**	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put^{1M}	1	0.08	0.050	82.81	0.184	6.7	0.152
		DOOM-put^{1M}	2	0.08	0.075	84.72	0.303	5.3	0.073
		DOOM-put^{1M}	3	0.08	0.100	86.22	0.433	59.5	0.572
		DOOM-put^{1M}	4	0.08	0.125	87.49	0.575	28.5	0.206
		DOOM-put^{3M}	1	0.25	0.100	76.78	0.647	4.2	0.082
		DOOM-put^{3M}	2	0.25	0.175	82.08	1.365	95.8	0.877
		DOOM-put^{6M}	1	0.50	0.050	61.70	0.318	2.5	0.197
		DOOM-put^{6M}	2	0.50	0.250	80.34	2.860	78.5	0.687
		DOOM-put^{6M}	3	0.50	0.260	81.02	3.034	19.0	0.156
	Merton	DOOM-put^{1M}	1	0.08	0.030	53.35	0.237	100.0	1.760
		DOOM-put^{3M}	1	0.25	0.125	87.42	4.057	100.0	0.308
		DOOM-put^{6M}	1	0.50	0.225	87.62	7.492	100.0	0.334
	Heston	DOOM-put^{1M}	1	0.08	0.310	89.80	4.560	100.0	0.091
		DOOM-put^{3M}	1	0.25	0.340	89.61	8.203	100.0	0.152
		DOOM-put^{6M}	1	0.50	0.330	89.69	9.833	100.0	0.254
2025-04-08	BSM	DOOM-put^{1M}	1	0.08	0.050	89.49	0.123	100.0	3.308
		DOOM-put^{3M}	1	0.25	0.175	89.69	0.929	100.0	1.313
		DOOM-put^{6M}	1	0.50	0.005	65.35	0.016	3.3	5.061
		DOOM-put^{6M}	2	0.50	0.250	89.57	1.999	96.7	1.181
	Merton	DOOM-put^{1M}	1	0.08	0.020	75.29	0.313	100.0	1.333
		DOOM-put^{3M}	1	0.25	0.100	89.86	1.896	100.0	0.659
		DOOM-put^{6M}	1	0.50	0.175	89.77	3.478	100.0	0.719
	Heston	DOOM-put^{1M}	1	0.08	0.175	88.92	1.814	100.0	0.200
		DOOM-put^{3M}	1	0.25	0.200	89.58	3.320	100.0	0.327
		DOOM-put^{6M}	1	0.50	0.200	87.99	3.695	100.0	0.588
2026-06-30	BSM	DOOM-put^{1M}	–	–	–	–	–	–	–
		DOOM-put^{3M}	–	–	–	–	–	–	–
		DOOM-put^{6M}	–	–	–	–	–	–	–
	Merton	DOOM-put^{1M}	–	–	–	–	–	–	–
		DOOM-put^{3M}	–	–	–	–	–	–	–
		DOOM-put^{6M}	–	–	–	–	–	–	–
	Heston	DOOM-put^{1M}	1	0.08	0.030	89.31	0.156	100.0	1.598
		DOOM-put^{3M}	1	0.25	0.075	89.13	0.758	100.0	0.989
		DOOM-put^{6M}	1	0.50	0.125	89.69	1.693	100.0	0.886

Note. w_j is the percentage of the available premium budget allocated to contract j at each renewal date. It is therefore not a percentage of B_0. n_j is the number of contracts purchased at a renewal date, based on the available premium budget and the option price at that date. Dashes indicate that no contract is selected.

C.3 Transaction Costs

The transaction-cost test increases the purchase price of each put by 1%. The adjusted premium used in the optimization and portfolio accounting is

$$\pi_j^{\text{tc}} = \pi_j(1 + c_{\text{tc}}) = 1.01\pi_j. \tag{C.2}$$

All other inputs remain unchanged, including the calibration, simulated paths, strike selection, objective function and budget limit. The adjusted premium affects the selected contracts, the number of puts purchased $n_j = w_j B_0/\pi_j^{\text{tc}}$, and the terminal portfolio value. The results below use the $\text{ES}_{95\%}$ objective and report the adjusted premiums in Panel C.

The one-percent transaction cost does not change the main conclusions. Merton and Heston keep the same selected contracts at all active dates, while the higher premium reduces the number of puts purchased and slightly affects the final metrics. BSM reallocates between nearby contracts in 2020 and 2025, but remains in the same strike regions. The zero allocations of BSM and Merton in 2026 are unchanged. Transaction costs affect the amount of protection purchased more than the selected allocation.

Table C.6: One-year model-based results optimized with 95% Expected Shortfall and 1% transaction costs, by calibration date and model.

Panel A. Portfolio performance and terminal downside risk

Date	Model	Portfolio	Budget	$\bar{R}_T$	$\Delta\bar{R}$	$\text{VaR}_{95\%}$	$\text{VaR}_{99\%}$	$\text{ES}_{95\%}$	$\text{ES}_{99\%}$
2020-03-16	BSM	Index	–	-12.77	–	57.01	66.81	63.00	70.52
		DOOM-put	5.00	-9.11	-3.66	26.03	26.03	26.03	26.03
	Merton	Index	–	-12.78	–	49.15	59.13	55.31	63.27
		DOOM-put	5.00	-12.95	0.17	37.28	42.96	40.79	45.32
	Heston	Index	–	-12.65	–	52.63	69.12	62.65	76.23
		DOOM-put	5.00	-14.38	1.73	35.36	35.36	35.36	35.36
2025-04-08	BSM	Index	–	-7.20	–	38.64	47.59	44.12	51.32
		DOOM-put	4.93	-3.91	-3.29	18.62	18.62	18.62	18.62
	Merton	Index	–	-7.12	–	31.27	39.39	36.27	43.36
		DOOM-put	5.00	-9.05	1.93	20.20	20.85	20.60	21.16
	Heston	Index	–	-7.13	–	42.09	60.83	53.23	68.57
		DOOM-put	4.39	-7.37	0.24	20.00	20.00	20.00	20.00
2026-06-30	BSM	Index	–	18.07	–	8.80	17.55	14.17	21.40
		DOOM-put	0.00	–	–	–	–	–	–
	Merton	Index	–	18.08	–	5.98	13.95	10.80	17.33
		DOOM-put	0.00	–	–	–	–	–	–
	Heston	Index	–	18.16	–	16.09	33.34	26.76	42.16
		DOOM-put	3.01	15.23	2.93	16.61	16.74	16.70	16.75

Note. Budget is reported as percentage of B_0. Returns, $\Delta\bar{R}$, VaR and ES are percentage points. $\Delta\bar{R} = \bar{R}_I - \bar{R}_D$. VaR and ES use the positive-loss convention. Dashes indicate zero allocation or an unavailable dated output.

Table C.6 continued

Panel B. Path and distributional metrics

Date	**Model**	**Portfolio**	L_{MDD}	γ_1	γ_2
2020-03-16	BSM	Index	40.98	1.25	2.91
		DOOM-put	29.72	2.31	7.49
	Merton	Index	33.16	0.81	1.20
		DOOM-put	28.99	1.41	2.86
	Heston	Index	28.91	-0.55	-0.07
		DOOM-put	24.42	0.21	-1.06
2025-04-08	BSM	Index	26.98	0.72	0.93
		DOOM-put	21.24	10.92	189.11
	Merton	Index	21.14	0.46	0.33
		DOOM-put	18.33	1.31	1.57
	Heston	Index	22.40	-0.81	0.79
		DOOM-put	15.77	0.53	-0.58
2026-06-30	BSM	Index	11.76	0.45	0.36
		DOOM-put	–	–	–
	Merton	Index	9.98	0.38	0.19
		DOOM-put	–	–	–
	Heston	Index	14.85	-0.50	0.58
		DOOM-put	11.83	-0.08	-0.39

Note. L_{MDD} *is reported in percentage points.* γ_1 *is skewness and* γ_2 *is excess kurtosis. Dashes indicate zero allocation.*

Table C.6 continued

Panel C. Selected DOOM-put positions

Date	**Model**	**Portfolio**	j	T_j	p_j	K_j	π_j	w_j	n_j
2020-03-16	BSM	DOOM-put	1	1.00	0.320	77.81	5.179	3.40	0.657
		DOOM-put	2	1.00	0.330	78.65	5.452	1.60	0.293
	Merton	DOOM-put	1	1.00	0.370	88.97	13.145	5.00	0.380
	Heston	DOOM-put	1	1.00	0.200	67.58	5.126	4.18	0.815
		DOOM-put	2	1.00	0.225	72.07	6.088	0.82	0.135
2025-04-08	BSM	DOOM-put	1	0.75	0.005	59.41	0.018	0.52	29.368
		DOOM-put	2	0.75	0.010	62.50	0.040	0.53	13.328
		DOOM-put	3	1.00	0.010	58.14	0.042	0.44	10.572
		DOOM-put	4	1.00	0.290	88.12	3.231	0.31	0.096
		DOOM-put	5	1.00	0.310	89.31	3.578	1.50	0.418
		DOOM-put	6	1.00	0.320	89.91	3.758	1.64	0.436
	Merton	DOOM-put	1	1.00	0.260	89.24	5.746	5.00	0.870
	Heston	DOOM-put	1	1.00	0.225	87.44	4.586	4.39	0.956
2026-06-30	BSM	DOOM-put	–	–	–	–	–	–	–
	Merton	DOOM-put	–	–	–	–	–	–	–
	Heston	DOOM-put	1	1.00	0.175	89.00	3.076	3.01	0.978

Note. w_j is reported as percentage of B_0. All contracts use normalized $S_0 = B_0 = 100$, and $n_j = w_j B_0 / \pi_j$ with w_j expressed in decimal form. π_j includes the one-percent transaction-cost adjustment. Dashes indicate zero allocation.

C.4 Protection-Budget Sensitivity

This sensitivity analysis studies how the maximum option budget affects the static $\mathrm{ES}_{95\%}$ allocation. For each model and calibration date, the strategy is reoptimized for different budget limits while keeping the calibration, contract universe and simulation framework unchanged. The objective is to determine whether the reference 5% budget constraint limits the protection obtained or whether the optimizer naturally selects a lower budget. The 5% level provides a compromise between meaningful downside protection and the cost of reducing initial index exposure, while remaining consistent with a realistic portfolio allocation.

C.4.1 16 March 2020

The crash calibration is the case where additional budget provides the largest benefit. As shown in Figures C.1 to C.3, the 5% reference budget still limits the protection obtained for all models. BSM continues to improve over the tested range, although gains become smaller at higher budgets. Merton and Heston reach their highest objective values only at higher allocations, around 11.45% and 9.79%, respectively.

C.4.2 8 April 2025

The benefit of increasing the budget is more limited in the stressed calibration. As reported in Figures C.4 to C.6, Merton and Heston reach their plateau close to or below the reference budget, at approximately 5.38% and 4.34%. The BSM objective continues to improve at higher budgets, but the additional gain becomes limited beyond 12.5%. Increasing the budget does not necessarily reduce $\mathrm{ES}_{95\%}$, since the objective also rewards mean return.

C.4.3 30 June 2026

The calmer calibration leads to a different outcome. As shown in Figures C.7 to C.9, BSM and Merton do not allocate any budget over the tested range, while Heston uses approximately 2.99%. The 5% budget limit therefore does not restrict the allocation selected by these models. Across the three dates, the budget acts as an upper limit rather than a target allocation, with additional protection mainly valuable in stressed market conditions.

The following figures illustrate the sensitivity of the optimized objective to the maximum protection budget across models and calibration dates.

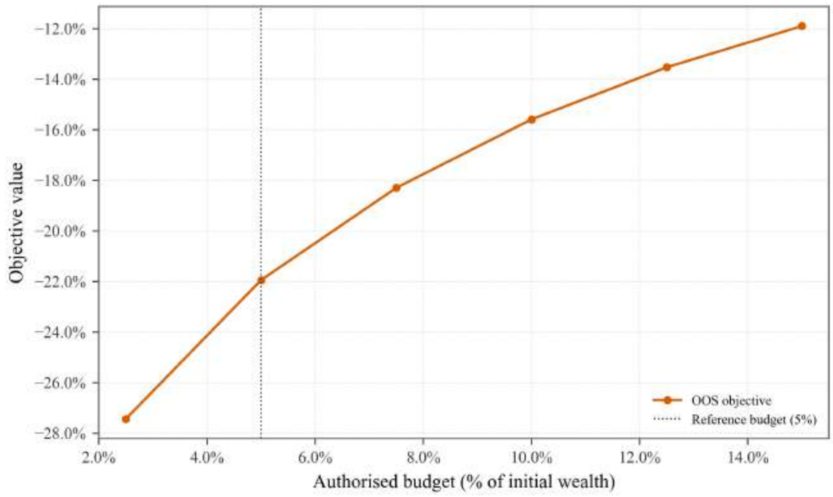

Figure C.1: BSM objective-value sensitivity to the maximum protection budget, 16 March 2020.

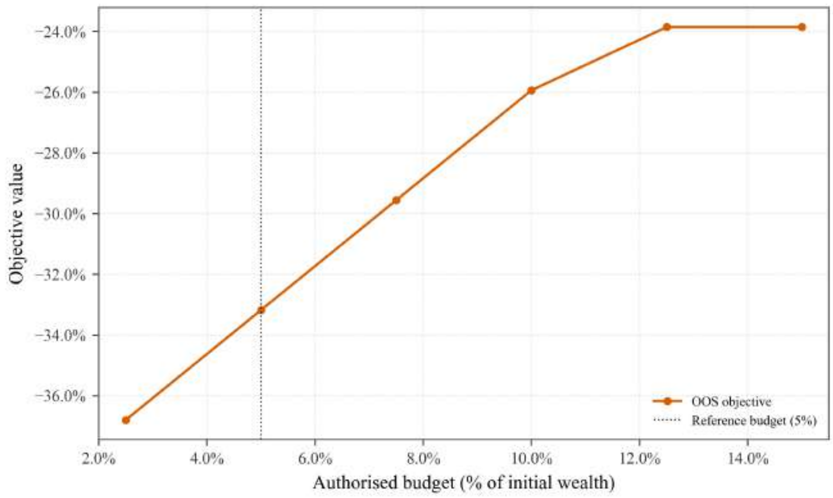

Figure C.2: Merton objective-value sensitivity to the maximum protection budget, 16 March 2020.

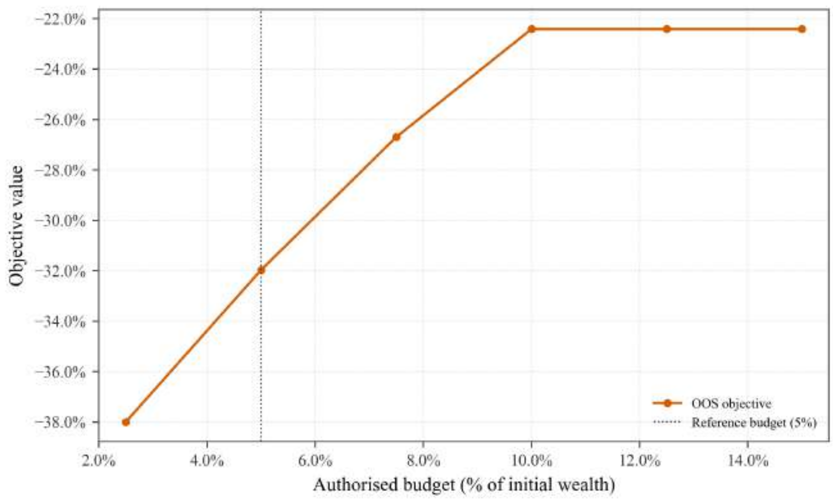

Figure C.3: Heston objective-value sensitivity to the maximum protection budget, 16 March 2020.

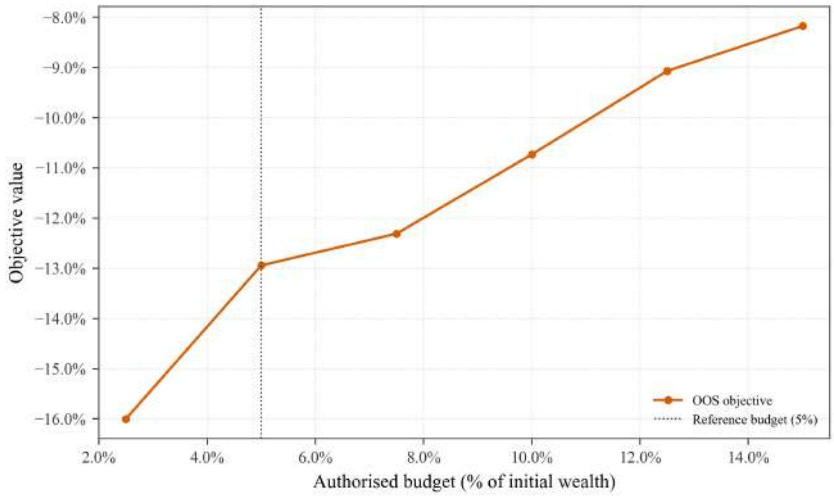

Figure C.4: BSM objective-value sensitivity to the maximum protection budget, 8 April 2025.

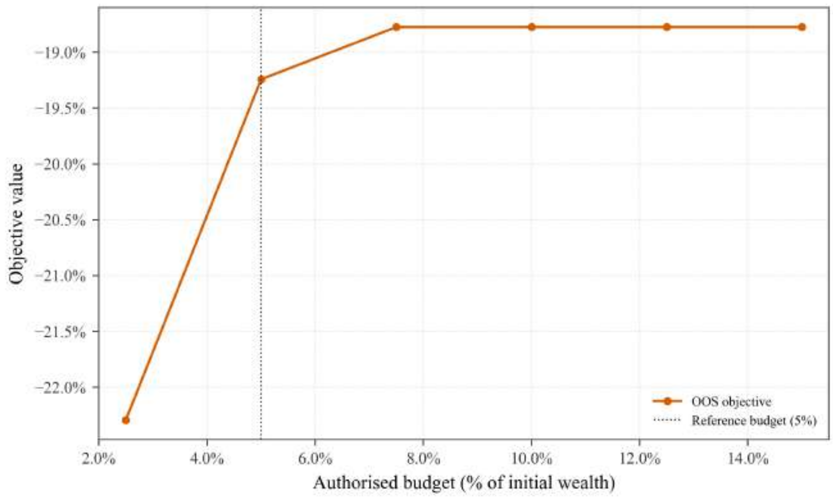

Figure C.5: Merton objective-value sensitivity to the maximum protection budget, 8 April 2025.

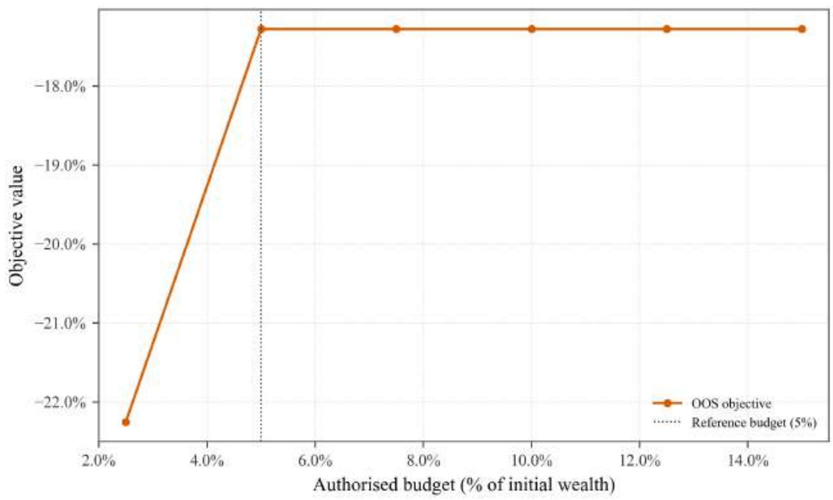

Figure C.6: Heston objective-value sensitivity to the maximum protection budget, 8 April 2025.

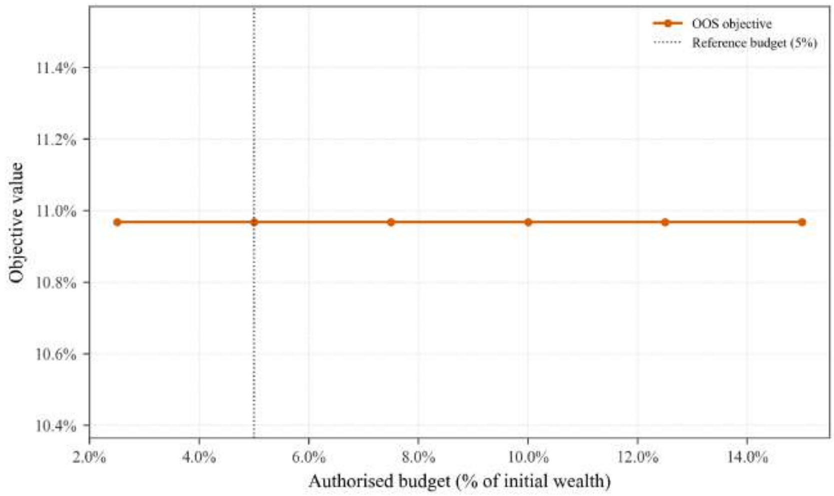

Figure C.7: BSM objective-value sensitivity to the maximum protection budget, 30 June 2026.

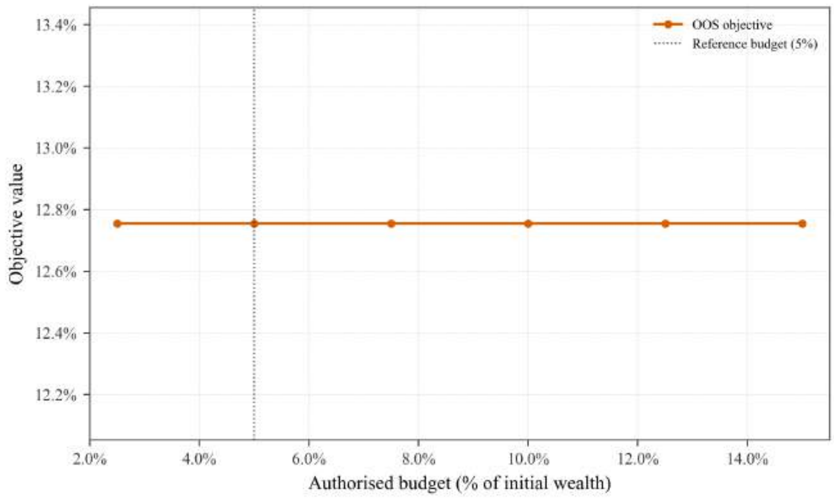

Figure C.8: Merton objective-value sensitivity to the maximum protection budget, 30 June 2026.

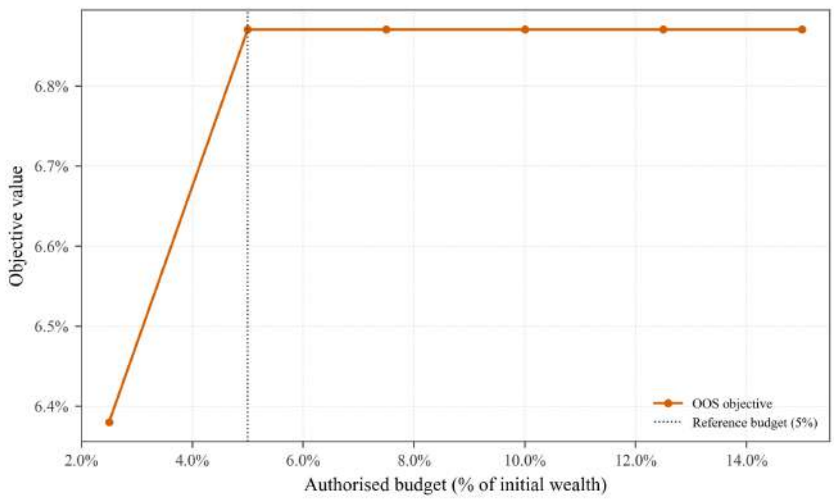

Figure C.9: Heston objective-value sensitivity to the maximum protection budget, 30 June 2026.

APPENDIX D

Additional Figures

The figures report the static one-year strategies optimized with $ES_{95\%}$ without transaction costs.

D.1 16 March 2020

D.1.1 BSM

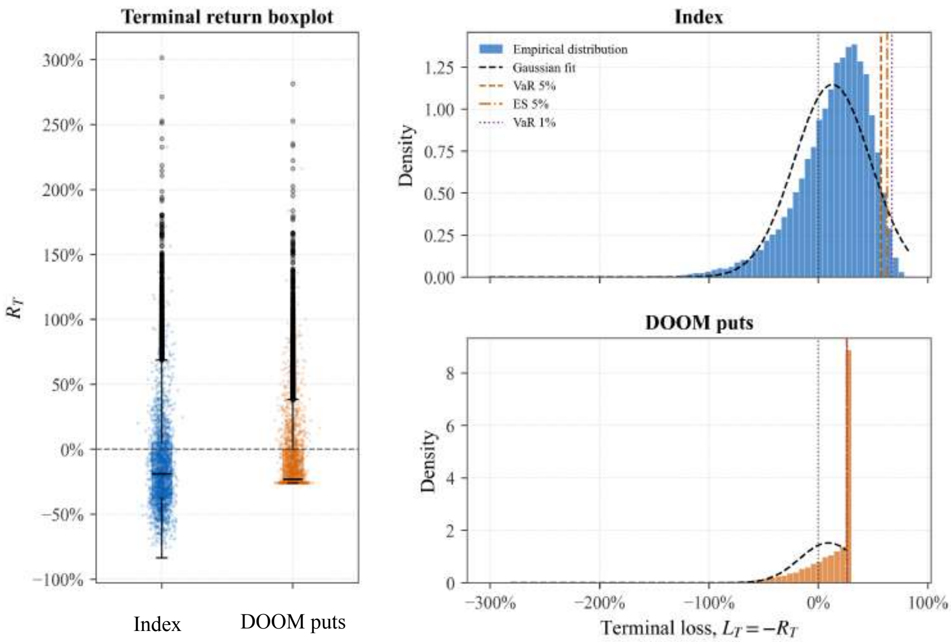

Figure D.1: BSM terminal distributions, 16 March 2020.

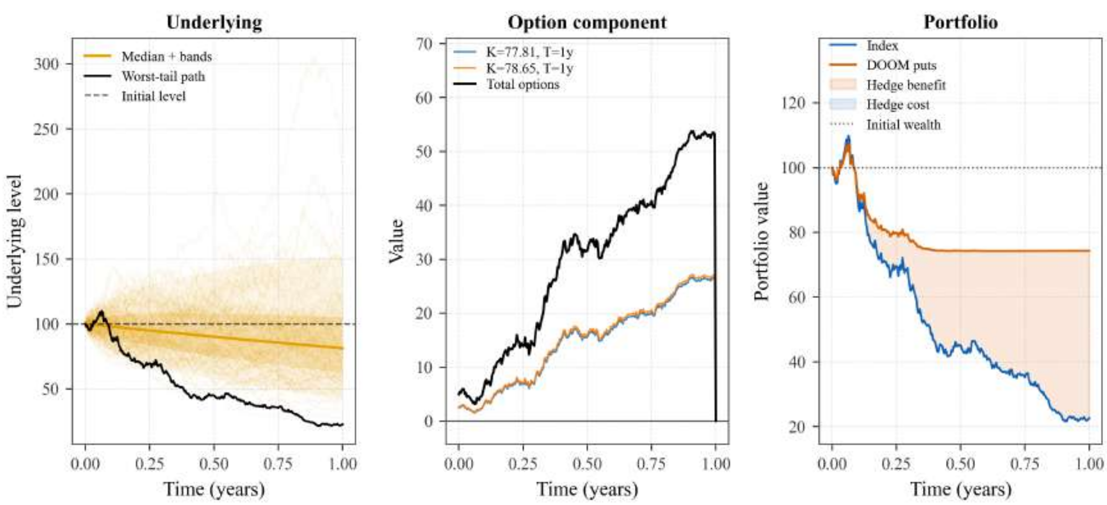

Figure D.2: BSM selected path, 16 March 2020.

D.1.2 Merton

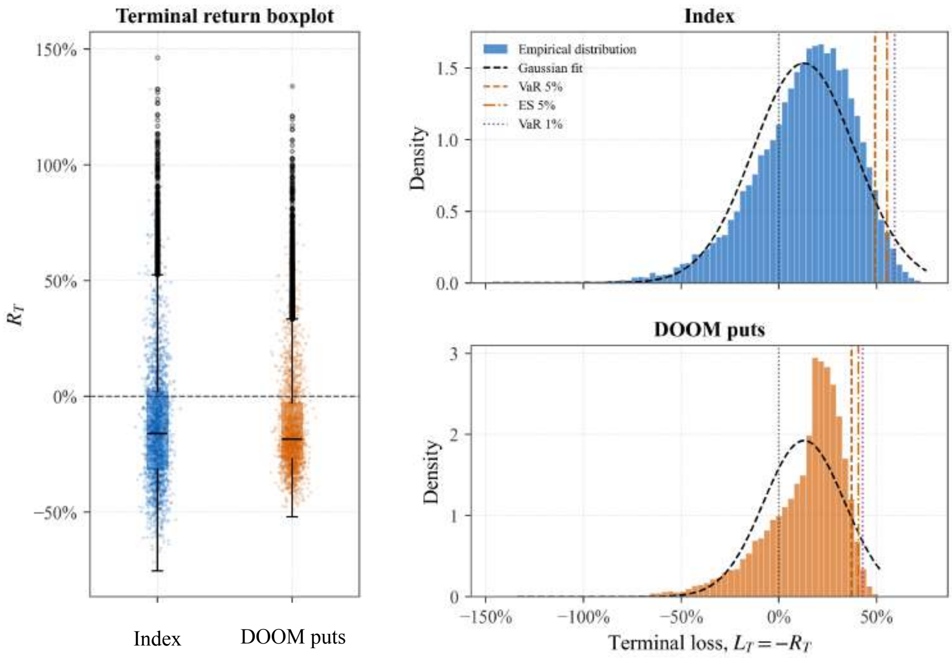

Figure D.3: Merton terminal distributions, 16 March 2020.

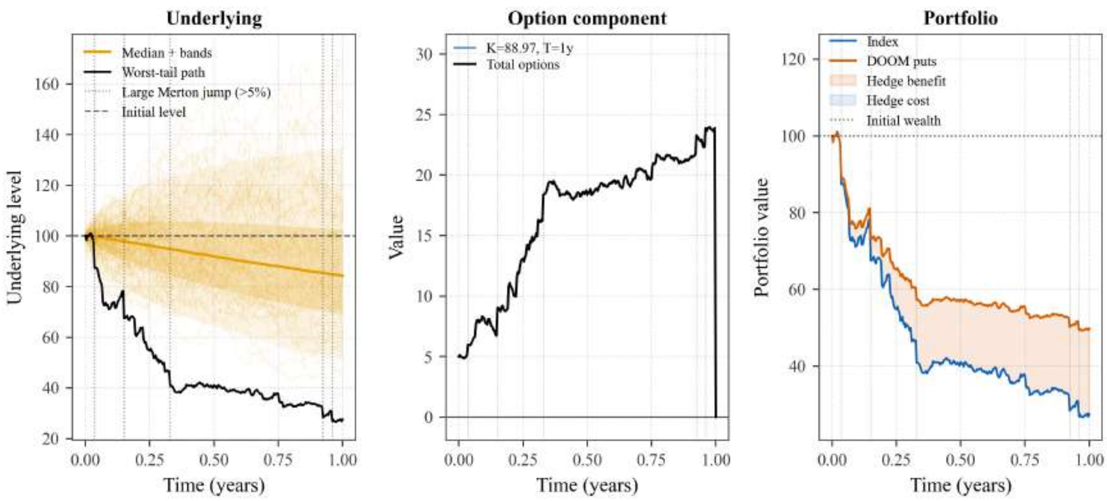

Figure D.4: Merton selected path, 16 March 2020.

D.1.3 Heston

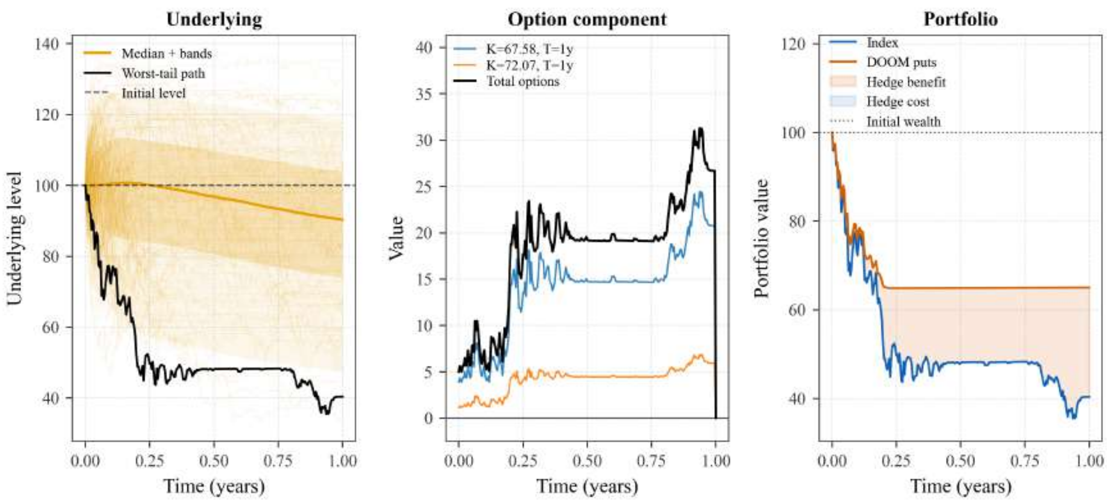

Figure D.5: Heston selected path, 16 March 2020.

D.2 8 April 2025

D.2.1 BSM

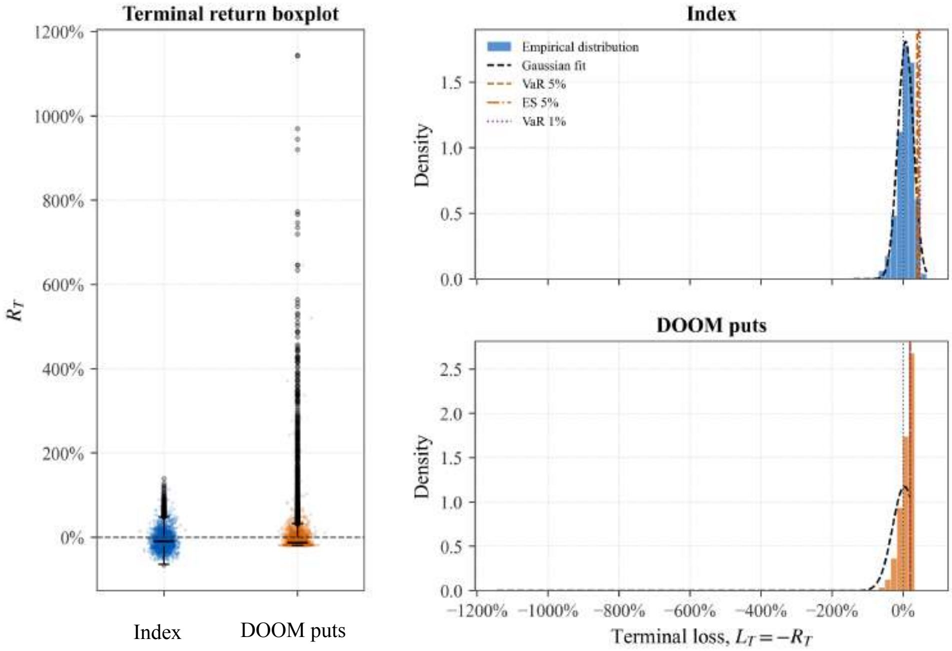

Figure D.6: BSM terminal distributions, 8 April 2025.

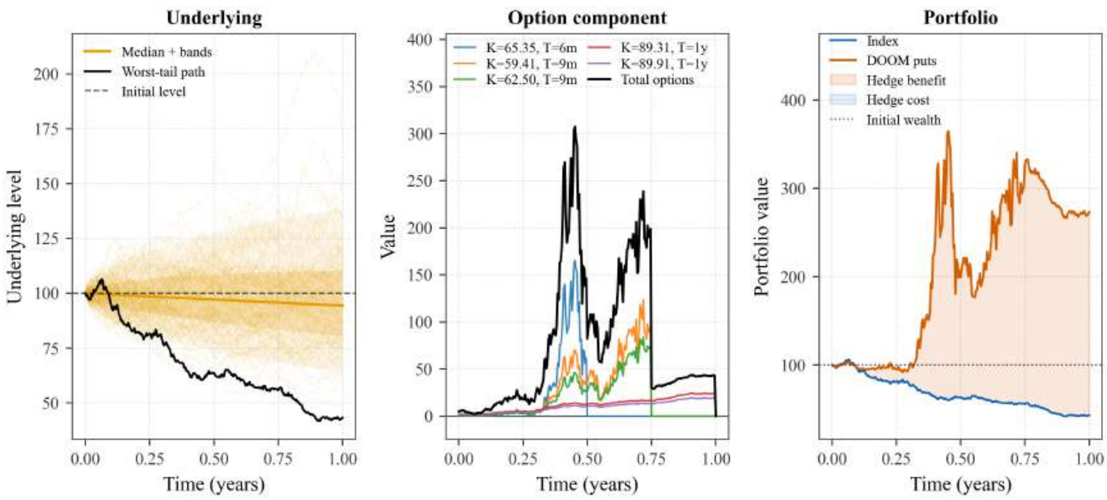

Figure D.7: BSM selected path, 8 April 2025.

D.2.2 Merton

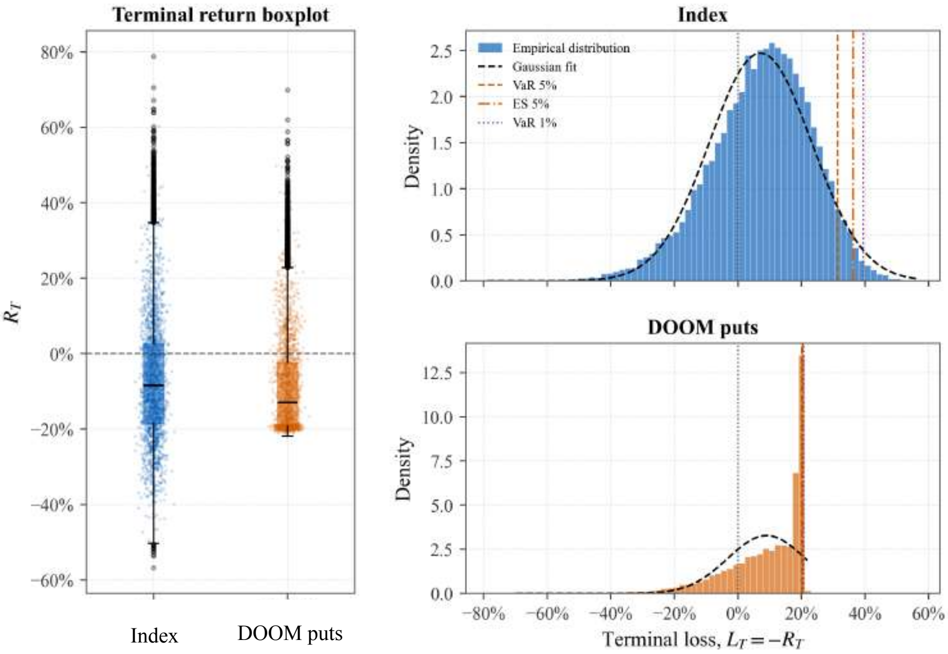

Figure D.8: Merton terminal distributions, 8 April 2025.

D.2.3 Heston

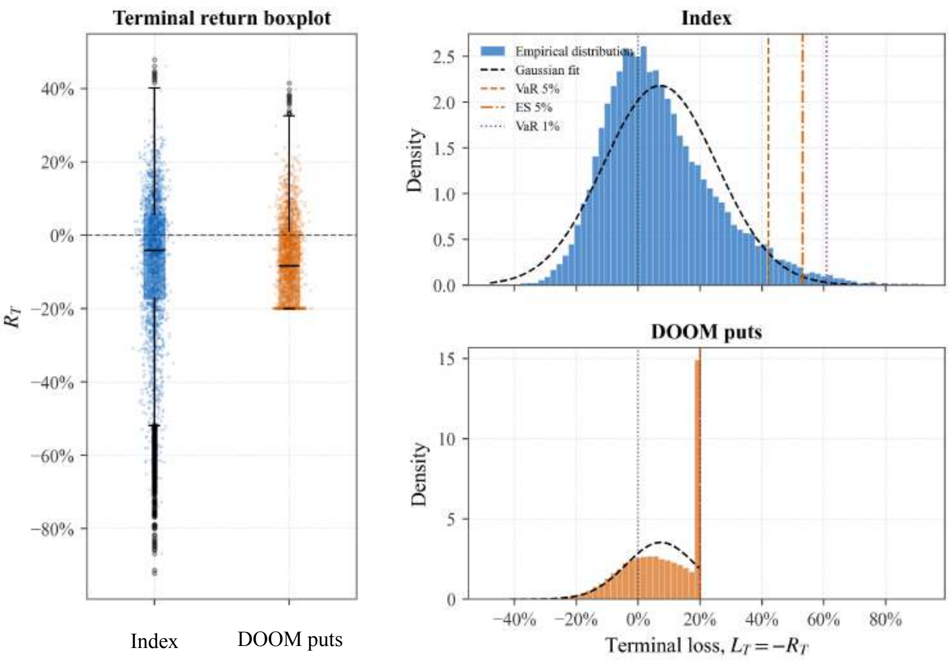

Figure D.9: Heston terminal distributions, 8 April 2025.

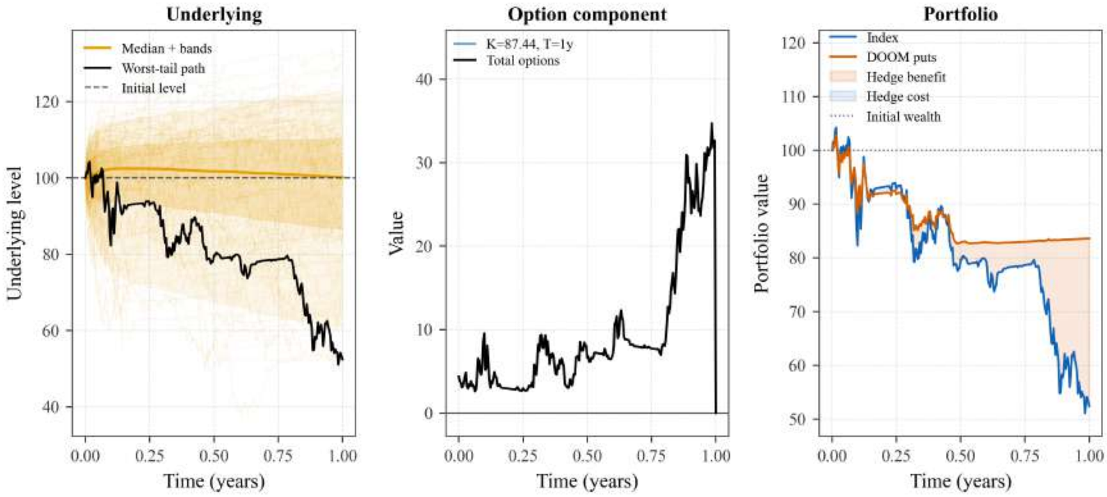

Figure D.10: Heston selected path, 8 April 2025.

D.3 30 June 2026

D.3.1 BSM

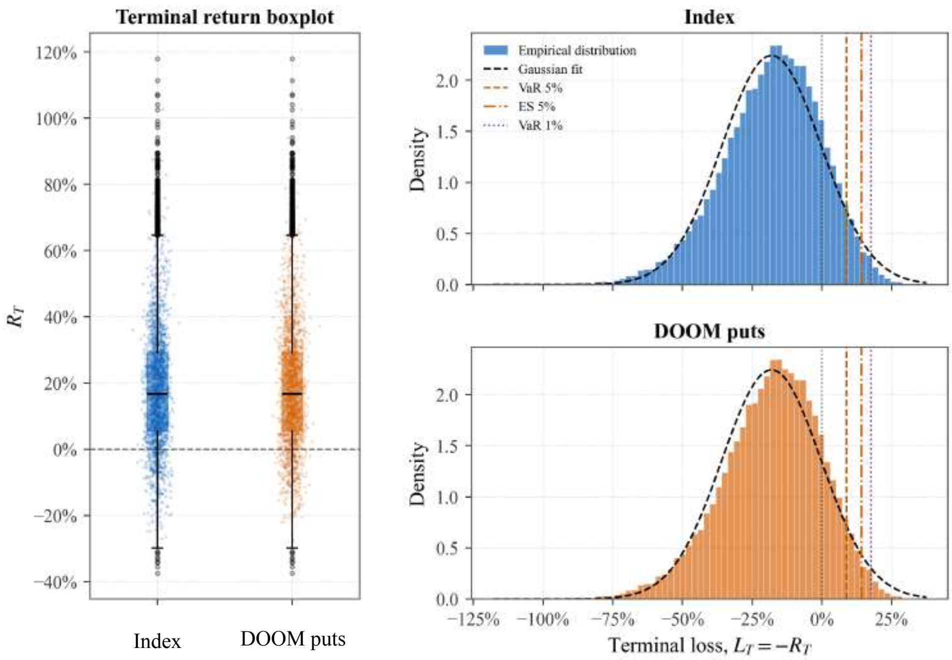

Figure D.11: BSM terminal distributions, 30 June 2026.

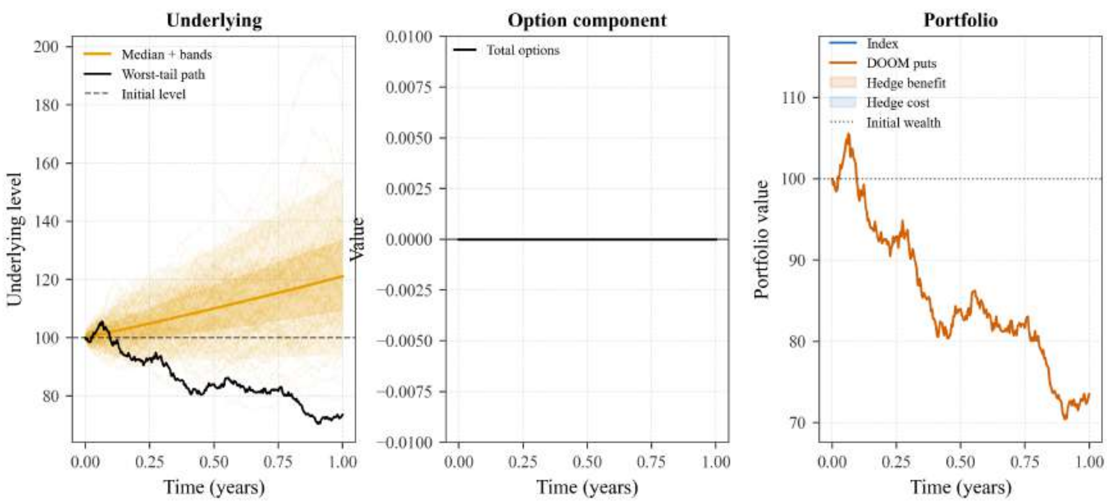

Figure D.12: BSM selected path, 30 June 2026.

Note. No put is selected for BSM at this date.

D.3.2 Merton

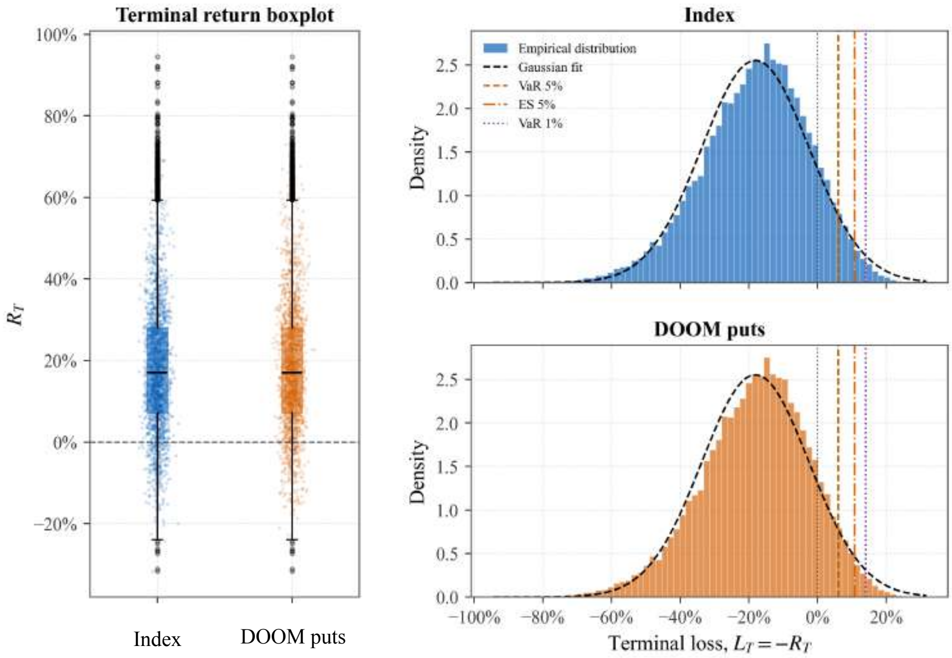

Figure D.13: Merton terminal distributions, 30 June 2026.

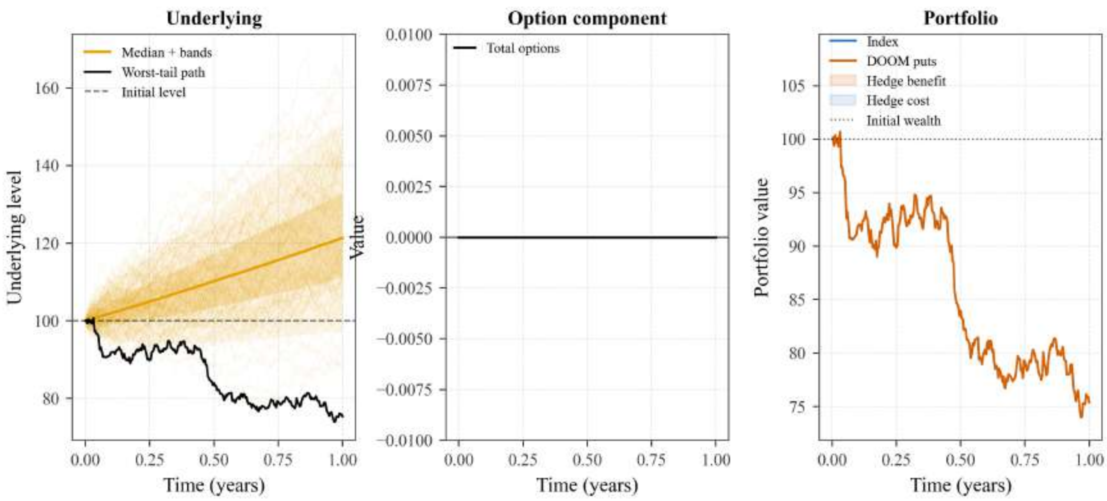

Figure D.14: Merton selected path, 30 June 2026.

Note. No put is selected for Merton at this date.

D.3.3 Heston

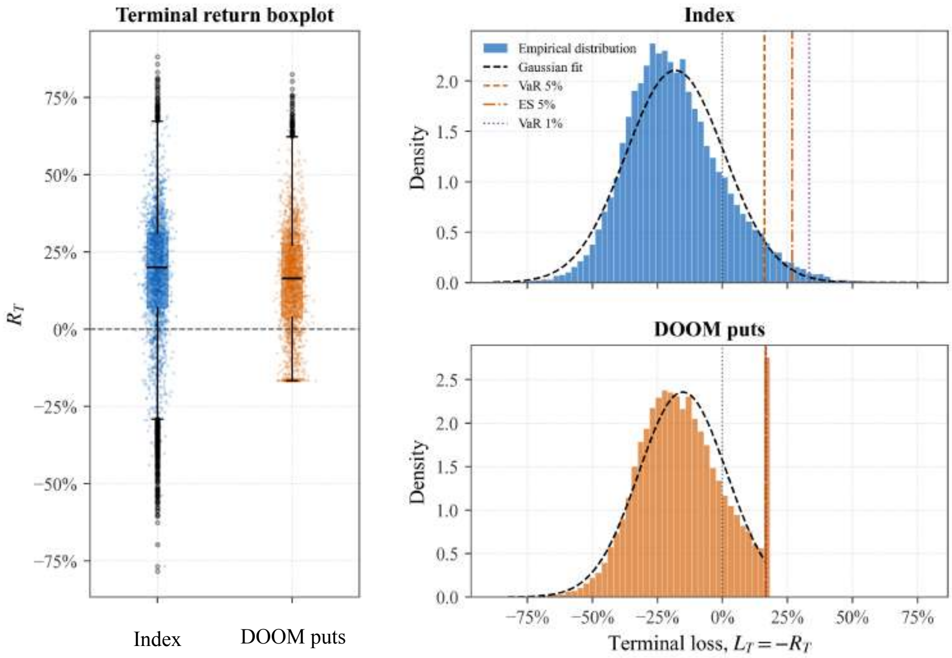

Figure D.15: Heston terminal distributions, 30 June 2026.

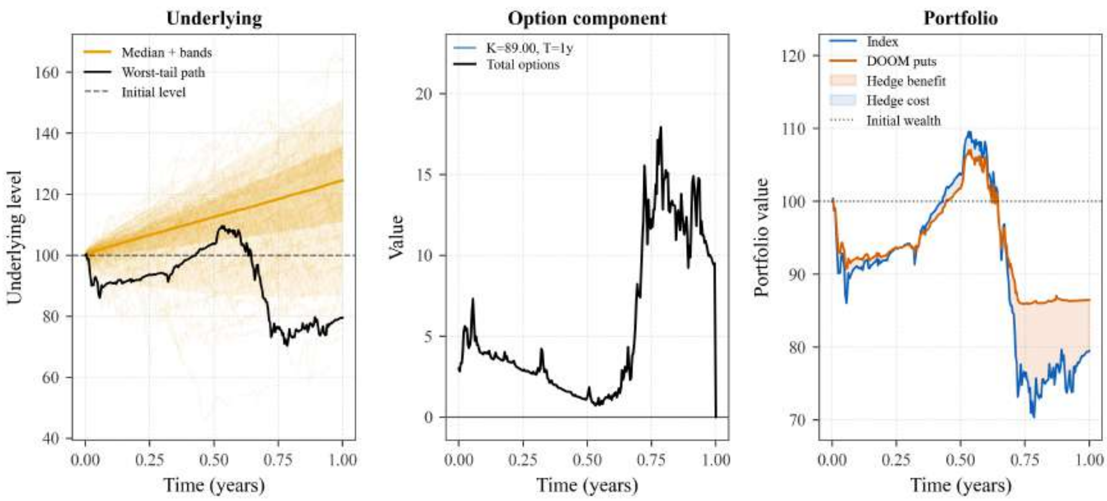

Figure D.16: Heston selected path, 30 June 2026.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management

Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique

www.uclouvain.be/lsm