Semi-Robust Risk Minimizing Hedging Strategies

Emmanuel Sekyere Osei Mireku

A Thesis

In the Department

of

Mathematics and Statistics

Presented in Partial Fulfillment of the Requirements

For the Degree of

Doctor of Philosophy (Mathematics) at

Concordia University

Montreal, Quebec, Canada

March 2025

CONCORDIA UNIVERSITY

SCHOOL OF GRADUATE STUDIES

This is to certify that the thesis prepared

By: Emmanuel Sekyere Osei Mireku

Entitled: Semi-Robust Risk Minimizing Hedging Strategies

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Mathematics)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

______________________________ Chair

Dr. Yuan Yi

______________________________ External Examiner

Dr. Clarence Simard

______________________________ Examiner

Dr. Frédéric Godin

______________________________ Examiner

Dr. Yang Lu

______________________________ Examiner

Dr. Xiaowen Zhou

______________________________ Thesis Supervisor

Dr. Patrice Gaillardetz

Approved by ______________________________

Dr. Lea Popovic (Graduate Program Director)

__/__/2025 ______________________________

Dr. Pascale Sicotte (Dean, Faculty of Arts and Science)

Abstract

Semi-Robust Risk Minimizing Hedging Strategies

Emmanuel Sekyere Osei Mireku, Ph.D.

Concordia University, 2025

This thesis explores robust risk-minimizing hedging strategies for contingent claims in incomplete markets with transaction costs, offering a spectrum of tools to balance risk and cost-effectiveness. Robust technique applications to finance and insurance have recently gained popularity due to their ability to mitigate model risk. Model risk arises when strategies (or models) become in and out of sync with the market. A model is robust if it can adapt to a wide range of market-dependent factors. However, robust models can be costly and computationally demanding, especially for complex financial and insurance products. Using a multidimensional event tree model, we employ the asymmetric norm as a semi-robust risk measure, integrating asymmetry for customized risk profiles. Three main strategies are developed: a super-replicating approach ensuring full claim coverage at a higher cost, the norm as constraint, which introduces controlled losses to reduce costs, and the norm as objective, minimizing losses directly to enhance capital efficiency. Additionally, self-financing strategies, which require no additional capital injections, offer cost-effective hedging, while portfolio value as state variable strategies allow real-time adjustments, enhancing robustness under volatile conditions. Testing on European call options show that semi-robust strategies - especially norm-constrained and self-financing approaches - maintain low tail risk with minimized cost, demonstrating versatility in adapting to diverse market conditions, investor goals, and risk tolerances while upholding robust risk control.

Acknowledgments

I would like to thank God for seeing me through this research.

I would like to express my deepest gratitude to my supervisor, Dr. Patrice Gaillardetz, who has carefully guided me and inspired me through my research. I treasure this opportunity to study an interesting topic and bring something new to the field. Moreover, his passion, persistence and commitment have influenced me and taught me many great lessons in life.

I also want to thank Dr. Frédéric Godin for sharing with me research experience and directions. And my last thanks go to my family for their continuous support.

This research was supported by the Concordia International Award of Excellence, International Association of Black Actuaries, Jacques Goulet Graduate Scholarship in Actuarial Mathematics, Institut de Sciences Mathématiques (ISM), and the Department of Mathematics at Concordia University.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

Risks come in many forms and are caused by various circumstances, including financial, liquidity, and insurance-related risks. Uncertainty from some exposures that influence an investment or an issuer's ability to fulfill a claim gives rise to several liabilities. Due to the unpredictability of future outcome variance from expectations, investors face risk and tend to accept some uncertainty in exchange for a potential return on their investment.

The concept of managing or minimizing risk has evolved for various applications. In insurance, in case of an incidence covered by the insurance policy, the insurer sells its assurance that it would pay the policyholder or an injured party on the policyholder's behalf. The insurer tries to mitigate risk by setting aside a fund known as a reserve for future payments of incurred claims that are yet to be settled. Traditional finance involves setting up an alternative account or investment known as a hedging portfolio to offset potential losses from a polarizing position.

Various hedging strategies have been proposed in the literature. If certain payoffs cannot be replicated by trading in underlying securities, a market is said to be incomplete. As a result, perfect risk transfer is not achievable. The traditional no-arbitrage theory of valuation in a complete market, which is based on the unique price of a self-financing replicating portfolio, is reiterated by Staum (2007) as being insufficient for irreplicable payoffs in incomplete markets.

Uncertainty surrounding various pricing and valuation parameters makes establishing a hedge portfolio difficult. For instance, European options can be affected by fluctuations in market volatility, which poses a challenge when pricing such products. Jaroszkowski and Jensen (2022) develop a model to assess the impact of uncertainty of market price and volatility risk on the valuation of European options. They utilize a Hamilton-Jacobi-Bellman (HJB) approach in a Heston model to quantify the best and worst-case scenarios under uncertainty, and solve the resulting nonlinear equations using a finite element method. As a result, they highlight that the relationship between option price sensitivity and uncertainty is nonlinear, with significant variation across different parameter regimes.

In insurance, most products are not fit to be priced under market completeness because they involve mortality, financial risk, and, in some circumstances, surrender risk. Møller (1998) notes that such claims cannot be fully hedged by trading equities and risk-free assets alone, owing to the incompleteness of the market. To find self-financing hedging methods, they expand the model to a scenario where they can entirely remove the risk. As a result, they suggest risk-minimizing techniques and their related inherent risk processes. Schweizer (1995) provides a self-financing hedging strategy under ℓ_2 norm or quadratic minimization with one risk asset in discrete time. Rémillard et al. (2012) extends the concept to develop an optimal self-financing hedging strategy with multiple underlying assets by minimizing the mean square hedging error. To create a least-cost optimal replicating strategy known as ϵ-arbitrage, Bertsimas et al. (2001) uses the square root of the mean-squared replication error as an efficiency metric. In their situation, they stipulate that ϵ may be interpreted as the level of market incompleteness that gauges how expensive it is to replicate a portfolio.

Financial crises highlight the necessity of methods to manage risk and reduce market uncertainty. Sadly, practical methods to minimize risks are complicated by intricate systems of market activity. For any given predicted future return, Lisewski and Lichtarge (2010) demonstrates that investors may navigate this complexity by using global risk-minimizing strategies in portfolio models. They show that risk is reduced for markets where stocks, futures and other financial transactions are completed through margin accounts, provided the margin account requirement stays below a crucial empirically validated value. They claim that maintaining margins that are narrow enough would be an effective stabilizing technique for markets with centrally regulated risk margin requirements. They proceed to show that this innovative technique is robust to noise in empirical data and may also be generally applicable to complex networks across other disciplines.

On the concept of local risk minimization, Schweizer (2008) defines a square-integrable strategy as locally risk-minimizing for a payment stream if, for minor deviations, the discounted risk process expressed as the expectation of squared errors from the cost process is non-negative almost surely, for all increasing sequence of partitions of the time horizon. To demonstrate the continued validity of the fundamental martingale characterization for local risk-minimization, they expand their concept and methodology to a generic multi-dimensional framework.

Additionally, Lamberton et al. (1998) consider pricing and hedging contingent claims in discrete time under transaction costs in a general incomplete market. They demonstrate the existence of a local risk-minimizing strategy that includes transaction costs for each square-integrable contingent claim. They show that their approach is robust for any non-degenerate model with finite state space if the transaction cost parameter is sufficiently minimal.

Another approach to address market uncertainty is to explore robust methods. Robust optimization safeguards a decision-maker against uncertainty and ambiguous model parameters. The application of robust techniques to finance and insurance has recently gained popularity due to their ability to immunize against model risk and characterize market dynamics using tractable uncertainty sets instead of probability distributions.

An uncertainty set in robust optimization is a representation of all possible variations in the model parameters of an optimization problem. It defines the range or region within which parameters can vary while ensuring that the solution remains feasible and performs well under all scenarios within the set. Uncertainty sets that allow model parameters to vary within a defined range are referred to as box or interval uncertainty. One can also allow model parameters to vary within an ellipsoid (ellipsoidal uncertainty), convex hull (polyhedral uncertainty) or ensure the parameters remain withing a specific distance (norm-based uncertainty) from the a nominal value.

Ben-Tal and Nemirovski (1998) extensively studied the concept of robust optimization and laid the foundation for their formulations in generic convex optimization problems. By considering semi-definite programs whose data rely on certain constrained but unknown fluctuating parameters, El Ghaoui et al. (1998) explore robust solutions that minimize the worst-case objective while meeting the constraints for all feasible parameters inside the stated boundaries.

In their paper, Bandi and Bertsimas (2014) propose an alternative approach to price options by identifying a self-financing dynamic portfolio that closely resembles an option's payout utilizing the ϵ-arbitrage technique. However, they employ polyhedral uncertainty sets and norms to simulate the price dynamics. To identify the portfolio that minimizes the worst-case replication error for a given uncertainty set, they formulate the problem as a robust optimization problem and argue that this approach scales the dimension of the problem polynomially instead of exponentially.

By concentrating on robust optimization techniques, Scutella and Recchia (2013) employs robust analogues of the traditional mean-variance and minimum-variance portfolio optimization to address uncertainty in portfolio asset allocation. They also examine the connection between the robustness and convex risk measure concepts. Fonseca and Rustem (2012) assume that returns are unpredictable and aim to maximize portfolio return for the worst-case scenario using robust optimization. Both the uncertainty set and the objective function are reformulated as semi-definite problems that result in a tractable model. However, their suggested models are bi-linear and non-convex since the overall return in their framework is the sum of local and currency return.

Ceria and Stubbs (2006) introduced robust optimization to mitigate some of the negative effects of optimization introduced by estimation errors in expected return estimates after demonstrating how these errors can result in portfolios with weights significantly different from the optimal portfolio.

Sözüer and Thiele (2016) summarizes a general review of the advancements in robust optimization. They emphasize the growth of knowledge in robust optimization theory and its applications. They discuss fresh results in static and multi-stage decision-making, the relationship with stochastic optimization, distributional robustness, and robust nonlinear optimization from a theoretical perspective. Finally, they offer recommendations for researchers who want to protect their model from parameter uncertainty.

Bertsimas and Sim (2004) presents a strategy that aims to accept sub-optimal solutions and make trade-offs more desirable by looking at methods to lower conservatism of robustness. They modify the degree of conservatism of robust solutions in terms of probabilistic boundaries with constraint violations. By placing a cap on the number of parameters that can change, they established a robust method to regulate conservatism. Their robust formulation is a linear optimization problem, which is an appealing feature, and the methodology can be naturally and effectively extended to discrete optimization problems in a tractable way.

Liu et al. (2016) offer a theoretical justification for why the approach by Bertsimas and Sim (2004) exhibits excessive conservatism. They also note that when the cap is lower than the total number of nonzero components in the optimal solution, the robust approach does not achieve an extremely conservative result. They continue to state that one has to be cautious when modifying the cap's value since doing so can provide results that are more conservative than planned.

Using ellipsoidal, polyhedral, and interval uncertainty on the mean and covariance of projected returns, Lotfi and Zenios (2018) develop robust models for optimizing risk measures.

They provide an algorithm to get around the conservative nature of robust optimization models and demonstrate how the robust models reduce the well-known sensitivity of coherent risk measures to misspecifications of the first four moments of the expected return distribution.

Roos and den Hertog (2020) recommend a different robust formulation that combines all uncertainty into a single constraint and ties the worst-case predicted loss to the original constraints of the problem to reduce conservatism. They demonstrate that the suggested formulation may protect against most uncertainty at a minimal cost objective value for situations with an unfeasible robust counterpart.

1.2 Motivation

Any model that works relatively well most of the time is considered robust. Practically, strategies will move in and out of sync with the market, but this does not necessarily mean the strategy is ineffective. When there are market changes, a model's parameters may also vary, leading to models doing well in some circumstances while struggling in others. Model risk arises when a model is in and out of sync with the market. A model is considered robust if its performance remains stable even when the quality of market-dependent parameters are no longer adequate. Robust optimization is an effective method for handling such uncertainty in optimization.

A conservative model or strategy stays on the side of caution or presents the problem it addresses as the "worst-case scenario". Diverse disciplines have investigated robust optimization for market and model parameter uncertainties. However, the selection of constraints and uncertainty sets may impact robust methods. In most cases, attempts to fix these problems make the overall model complex or render the problem unsolvable. Furthermore, robust techniques and computational costs can be expensive, particularly for more complicated financial and insurance products. When robust optimization results in a solution with an objective value significantly worse than the nominal solution, it is considered conservative. Although robust techniques address model risk issues, they may result in less-than-ideal outcomes because of the model's excessive conservatism. This poses a challenge when it causes robust solutions to be overlooked in favour of the nominal answer. We seek to develop semi-robust, model-independent, and cost-effective strategies when minimizing the hedge portfolio for some contingent claims. We define a model as "semi-robust" when the quantification of uncertainty directly affects the objective of the problem rather than the model's parameters. In essence, we are interested in the region created by the overall model mismatch rather than studying each model parameter independently. This characterization distinguishes our approach from traditional robust strategies that rely on the construction of uncertainty sets for the model's parameters. Additionally, we assume no distribution on the model framework which is different from other studies and lastly, we develop a technique to obtain cheaper strategies using norm constraint optimization.

Using constrained robust optimization, we explore models that minimize risk exposure locally and globally. We consider filtration under the lattice/tree model as the framework for optimization. This grid filtration specification is different from other studies where they consider one step or range for their filtration. Our goal is to identify a semi-robust model that controls the problem of conservatism in financial and insurance products and attempts to address some of the drawbacks in robust optimization literature. We adopt convex optimization principles for our objective functions and constraints to develop feasible hedging strategies. We extend our models to include transaction costs and estimate closed-form solutions to reduce the model and computational complexity of robust hedging strategies.

1.3 Thesis overview

Chapter 2 presents the model framework, risk measures, hedging portfolio and loss function. We define a discrete state space, specifically a lattice model with recombining trees as the model framework. We further introduce the notations that are used to define our losses. Since we are interested in risk-minimizing strategies, we discuss two known risk measures in addition to norms and describe how they can be adapted to suit our model. Lastly, we define a typical loss function for a hedging portfolio. In Chapter 3, we discuss local hedging strategies and propose our semi-robust models that need to be optimized. In Chapter 4, we provide a detailed application of the semi-robust hedging strategies to European call options and compare their performance across different scenarios and model parameters. Chapter 5 introduces the framework for global hedging strategies and develops an approach to establishing a self-financing hedging strategy featuring self-financing super-replication and quadratic hedging. We further provide numerical analysis to compare the performance of such strategies to their equivalent local hedging strategies. We conclude with a summary and recommendations for our work.

Chapter 2

Model Framework and Risk Measures

In this chapter, we introduce the mathematical notations used and discuss the structure of our model framework. The risk measures and norms employed in the subsequent chapters are discussed, along with their characteristics and coherence. We define the portfolio loss function used in our proposed algorithms for a given hedging portfolio.

2.1 Model framework

We develop our model framework by assuming a multi-dimensional discrete process with finite number of realizations. Defining the structure as an event tree is convenient for both stochastic and deterministic dynamic optimization. The tree consists of nodes and branches. The nodes represent the discrete realizations of the process at each period, and the branches indicate transitions between nodes from one period to the other with positive transition probabilities. Specifying the branches in this manner makes it easy to characterize possible continuation of history, which is useful in developing path-dependent models.

The lattice tree can either be recombining or non-recombining. In a recombining tree, many branches could lead to a node, which implies many possible histories. In contrast, a non-recombining tree has only one branch leading to a node, which depicts a long dependence on history. Both trees can, however, model more general Markovian stochastic processes where the evolution of an underlying asset or security is strongly path-dependent. We can significantly reduce the computational complexity of the model by describing the discrete process as a recombining tree as opposed to a non-recombining tree.

Due to their versatility and tractability, lattice models are widely used to depict the evolution of stocks, indexes, interest rates, and other financial securities, see Pliska (1997). Furthermore, they arise when dealing with contingent lives since they can also be described using counting processes. For our framework, we assume a fixed planning horizon $T \in \mathbb{N}^+$, divided into successive dates, $\mathcal{T} = \{0, 1, 2, \ldots, T\}$. We define $\{Y_t\}_{t\in\mathcal{T}}$ as a random process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and let $\{\mathcal{F}_t\}_{t\in\mathcal{T}}$ be a filtration under which Y_t is adapted. Ω is the sample space and the process Y_t is defined as $Y : \mathcal{T} \times \Omega \to \mathbb{Y}$, where $\mathbb{Y}$ is the state space of Y_t. $\mathcal{F}_t$ denotes the information set that includes process values that have been observed up to and including time t, and $\mathbb{P}$ represents the physical/real-world probability measure.

It is possible to construct a tree with time-varying transition lengths, however, we assume that for a given date, the duration of the next transition is constant, regardless of the value of the random process. This assumption results in a tree with finite number of nodes arranged in a finite number of levels related to dates. The transitions between levels represent the branches. In general, Y_t represents the price process. For Guaranteed Investment Certificates (GICs), European, American and Barrier options, Y_t is the stock price. On the other hand, a contract can depend on multiple random processes. For example, in Equity-Indexed Annuities (EIAs), the stochastic process can include both the stock price and another dimension for the buyers' cohort.

Suppose i_t is the index number for node i at time t. Then, $y_{i_t} \in \Omega$ is the value of the random process relative to node i_t. At $t = 0$, we have a unique root node associated with the value of the stochastic process $Y_0 = y_0$. At $t = 1$, there are as many nodes as possible values of y_{i_1} which are all connected to the root node y_0 by branches. Figure 2.1 illustrates three-period recombining and non-recombining trees with three branches at each node during each transition.

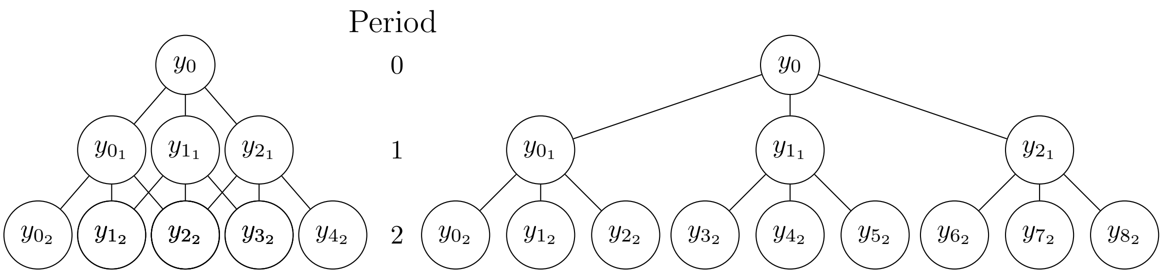

Figure 2.1: Recombining and non-recombining trinomial trees.

Let J_t denote the set of all possible indexes i_t for outcomes of Y_t at time t. For simplicity of notation, we let i_{t-1} represent the event where $Y_{t-1} = y_{i_{t-1}}$ for conditional sets. Thus, $J_t|i_{t-1}$ is the subset of indexes of Y_t, each connected to node $y_{i_{t-1}}$. For example, conditional on y_{1_1} in Figure 2.1, $J_2|1_1 = \{1_2, 2_2, 3_2\}$, corresponding to outcomes $\{y_{1_2}, y_{2_2}, y_{3_2}\}$ for the recombining tree and $J_2|1_1 = \{3_2, 4_2, 5_2\}$, corresponding to outcomes $\{y_{3_2}, y_{4_2}, y_{5_2}\}$ for the non-recombining tree. Hence, we can denote $p_{j_t|i_{t-1}} = \Pr[Y_t = y_{j_t}|Y_{t-1} = y_{i_{t-1}}] = \Pr[Y_t = y_{j_t}|i_{t-1}]$ as the conditional transition probability from node i_{t-1} to node j_t, where $j_t \in J_t|i_{t-1}$. The conditional expected value of Y_t given Y_{t-1} is,

$$\mathbb{E}[Y_t|i_{t-1}] = \sum_{j_t \in J_t|i_{t-1}} p_{j_t|i_{t-1}} y_{j_t}. \tag{2.1}$$

We also let r be the force of interest; that is, r is a nominal interest rate compounded continuously.

2.2 Risk measures

A risk measure is a function for summarizing the degree of risk associated with a random variable into a discrete value or real-valued function. Financial institutions, including insurance and investment firms, frequently utilize risk measures to assess risk associated with various business segments. The efficiency and meaningfulness of risk measures in a professional and regulatory setting are largely responsible for their extensive use. They can be used in incomplete markets to develop tractable option pricing models. Artzner et al.

(1999) defines a coherent risk measure as translation invariant, monotonic, sub-additive and positively homogeneous. Despite not being coherent, the Value-at-Risk (VaR) is one of the most popular and well-known risk measures. VaR is widely used due to regulatory requirements and ease of interpretation in risk management. On the other hand, the Conditional Value-at-Risk (CVaR) is a coherent risk measure that has recently gained recognition as more efficient in quantifying risk. For transitions from node i_{t-1}, the VaR at level $\alpha \in [0, 1]$ is defined as the α-quantile of the process Y_t. That is

$$\mathrm{VaR}_\alpha(Y_t|i_{t-1}) = \inf\{y \in \mathbb{R} : \Pr[Y_t \leq y|i_{t-1}] \leq 1 - \alpha\} = \zeta_\alpha. \tag{2.2}$$

The CVaR risk measure represents the expected value of the worst $(1 - \alpha)$ realizations of the conditional transitions. In this context, Rockafellar and Uryasev (2000) propose the definition as

$$\begin{aligned} \mathrm{CVaR}_\alpha(Y_t|i_{t-1}) &= \zeta_\alpha + \frac{1}{1-\alpha}\mathbb{E}[(Y_t - \zeta_\alpha)^+|i_{t-1}] \\ &= \zeta_\alpha + \frac{1}{1-\alpha} \sum_{j_t \in J_t|i_{t-1}} p_{j_t|i_{t-1}}(y_{j_t} - \zeta_\alpha)^+. \end{aligned} \tag{2.3}$$

They proved that this formulation of the CVaR is convex for ζ_α. Although it may not always be differentiable with respect to ζ_α, it can readily be minimized in either the objective or constraint of linear programming or line search minimization techniques.

2.3 Robust assumptions

The distribution of a model's transition probabilities can increase the model's risk when out of sync with the market. We model the outcomes of our process without assigning specific transition probabilities. This approach can be employed to construct model-independent frameworks that prioritize flexibility in achieving the modelling objective. By refraining from specifying a probability distribution, the model avoids embedding subjective beliefs about the likelihood of different scenarios, making it suitable for situations where a hedger is skeptical about market dynamics. Instead, the focus shifts to monitoring and hedging against all possible outcomes, ensuring that the strategy is robust regardless of how the process evolves. For example, in a super-replicating strategy, a hedger constructs a portfolio that dominates a future claim and is only interested in possible outcomes rather than the distribution. In this case, the hedger finds the worst-case scenario or maximum observable loss and controls it rather than optimizing for expected returns under a given distribution of losses.

We can define Y_t as a non-probabilistic space that corresponds to the image of the original probability space. As such, unless one is induced, there is no probability measure on this space. We let $Y_{t|i_{t-1}}$ be a non-probabilistic vector representation with components $\{y_{j_t}\}_{j_t \in J_t|i_{t-1}}$. By treating the stochastic process as a collection of possible paths without assigning probabilities, the hedger adopts a conservative approach that accommodates maximum uncertainty. It is important to mention that we do not employ conditional probabilities in the development of our proposed hedging strategies.

2.4 Norms

The norm is another way of measuring the riskiness of a financial position. Unlike the VaR and CVaR, which measure the riskiness based on a quantile, the norm measures specific sizes or magnitudes of mismatches from the position. Due to its convex property, the ℓ_p norm ($1 \leq p < \infty$), defined discretely as

$$\|Y_{t|i_{t-1}}\|_p = \left(\sum_{j_t \in J_t|i_{t-1}} |y_{j_t}|^p \right)^{1/p}, \tag{2.4}$$

is sometimes used in mathematical finance (specifically the ℓ_2 norm) to assign a numerical value that represents the risk associated with a financial position. We note that (2.4) is not a function of weights or probabilities. Minkowski's inequality ensures the sub-additive property of the ℓ_p norm is preserved and hence a convex risk measure, see Föllmer and Schied (2011).

The value of p represents the investor's tolerance for risk. If an investor has a high tolerance for risk, they will choose a low value of p (such as $p = 1$). This causes a risk measure to be more sensitive to the portfolio's extreme values. In contrast, a high value of p (such as $p = \infty$) is used if an investor has a low tolerance for risk, resulting in a risk measure that is less sensitive to the portfolio's extreme values.

In some situations, we modify the definition of the norm when hedging. For instance, when using a symmetric norm, equal weights are assigned to positive and negative realizations or losses. This is sometimes a drawback and can affect the feasibility of the problem. By generalizing the concept of a norm, we can introduce an asymmetric parameter to ensure the problem's feasibility while satisfying the norm's properties.

Definition 2.4.1. ***(Asymmetric norm)*** *The norm ℓ_p is asymmetric if for a given asymmetric parameter $q \in [0, 1]$ and $1 \leq p < \infty$,*

$$\|Y_{t|i_{t-1}}\|_{p,q} = \left(\sum_{j_t \in J_t|i_{t-1}} |y_{j_t}^+ - q y_{j_t}^-|^p \right)^{1/p}, \tag{2.5}$$

where $y_{j_t}^+, y_{j_t}^- \geq 0$, $y_{j_t}^+ = max(y_{j_t}, 0)$, $y_{j_t}^- = max(-y_{j_t}, 0)$ and $y_{j_t} = y_{j_t}^+ - y_{j_t}^-$.

We denote $\ell_{p,q}$ as the asymmetric norm. Here, q can be perceived as the measure of asymmetry. In practice, a hedger can set $q = 0$ to be more conservative or choose $0 < q \leq 1$ to allow a degree of flexibility during optimization. As such the choice of asymmetry can be useful for controlling conservatism. García-Raffi et al. (2002) proves that the asymmetric norm also satisfies the properties of a norm, and Conradie (2015) investigates the convex properties of asymmetric norms.

2.5 Hedging portfolio and loss function

A portfolio of assets whose value can partially or fully offset potential losses is a hedge or replicating portfolio. If there are no injections or withdrawals from a replicating portfolio after it has been created, it is considered self-financing. In other words, re-balancing the portfolio can be done by simply moving the assets around. Nevertheless, a non-self-financing portfolio permits injections or withdrawals from the hedger to guarantee that the targeted local requirement is satisfied during re-balancing. Without necessarily selling assets to raise cash, the investor can add more capital to the portfolio to purchase further assets or take money out. Non-self-financing portfolios have more flexibility since they can raise additional funding, but self-financing portfolios must carefully manage their current assets to finance new purchases.

Definition 2.5.1. ***(Hedging strategy)*** *For a hedging portfolio with $n+1$ assets, $x_{i_t,k}$ is the value of the position in asset k held by a hedger for the outcome at node i_t. The hedging strategy $X_{i_t} = (x_{i_t,k})_{k=0,1,\ldots,n}$ is the vector that comprise the hedging portfolio. Furthermore, let $f_{i_t}(X_{i_t})$ denote the value of the hedging portfolio at node i_t after re-balancing with the hedging strategy X_{i_t} such that*

$$f_{i_t}(X_{i_t}) = \sum_{k=0}^{n} x_{i_t,k}. \tag{2.6}$$

The set of assets could have many components, such as cash, options or multiple risky assets. For example for a two asset portfolio, $x_{i_t,0}$ and $x_{i_t,1}$ can be the value of the position in cash and stocks respectively held by the hedger for the outcome at node i_t. Without loss of generality, we denote $X_{t|i_{t-1}}$ as the collection of vectors of hedging strategies $X_{j_t} \in \mathbb{R}^n$ for $j_t \in J_t|i_{t-1}$.

Next, let $C_{t|i_{t-1}}(X_{t|i_{t-1}})$ denote a vector of the hedged claim $c_{j_t}(X_{j_t})$, for $j_t \in J_t|i_{t-1}$. For example, an European call option with a unit strike price will have

$$c_{i_t}(X_{i_t}) = (y_{i_t} - 1)^{+}\mathbf{1}_{\{t=T\}} + f_{i_t}(X_{i_t})\mathbf{1}_{\{t<T\}}, \tag{2.7}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. Since the option is not exercised until maturity, c_{i_t} is the value of the hedging portfolio for at node i_t for $t < T$. In products like American options, c_{i_t} can represent the maximum of the continuation value and the payoff of the option at node i_t if the option was exercised. However, we consider only European options in our analysis.

It is possible to have different properties affecting the dynamics at each node. We let $\mathbf{Z}_{t|i_{t-1}}$ denote the state vector with components $\mathbf{z}_{j_t}$ for $j_t \in J_t|i_{t-1}$ and define $H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ as a multivariate state function that describes dynamic state evolution. For example, $H_{i_{t-1}}$ could be used to introduce transaction costs for trading assets. If we let $x_{i_{t-1},1}$ be the value of the position in some stock, we can distinguish between the stock already held with the a state vector $\mathbf{z}_{i_{t-1}}$, and some variation of $x_{i_{t-1},1}$ which represents newly bought ($x^+_{i_{t-1},1}$) or sold ($x^-_{i_{t-1},1}$) stocks since transaction fees affect only the variations. Thus, the function $H_{i_{t-1}}$ can be expressed as $\mathbf{z}_{j_t} - \mathbf{z}_{i_{t-1}} - x^+_{i_{t-1},1} + x^-_{i_{t-1},1}$ to account for transaction cost in the optimization. Many state variables and equations would be required if the hedge portfolio includes many assets with transaction fees, and the value of the strategy would be adjusted to account for the variations. $H_{i_{t-1}}$ could also be defined to introduce limits or bounds to state vectors or portfolio values to achieve a desired level of robustness or ensure the feasibility of the hedging problem.

In our study, we consider different objective functions for our risk minimization problems and use dynamic programming to optimize backwards, starting at maturity. We let $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ denote the vector of backward objective (cost-to-go) function of the dynamic programming problem with components $v_{j_t}(\mathbf{z}_{j_t})$ for $j_t \in J_t|i_{t-1}$.

Arising from the root node at inception, the hedge portfolio accumulates at each node. We define $W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ as a function that outputs a vector with components $w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$, the accumulation of hedge portfolio prior to claim payment for $j_t \in J_t|i_{t-1}$. Thus, w_{j_t} is the accumulated value at $j_t \in J_t|i_{t-1}$ for a hedge portfolio set at node i_{t-1}. For example, if an European option is hedged with $x_{i_{t-1},0}$ value in cash and $x_{i_{t-1},1}$ value in the underlying stock, with transaction cost as state variable such that $\mathbf{z}_{j_t} = \mathbf{z}_{i_{t-1}} + x^+_{i_{t-1},1} - x^-_{i_{t-1},1}$, we obtain $w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = x_{i_{t-1},0}e^r + \mathbf{z}_{j_t}\frac{y_{j_t}}{y_{i_{t-1}}}$ for $j_t \in J_t|i_{t-1}$ and $t < T$. A hedger has a surplus if the hedge portfolio value before the claim payment is greater or equal to the claim value. That is, $W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \geq C_{t|i_{t-1}}(X_{t|i_{t-1}})$. On the other hand, a hedger incurs temporary losses if this inequality is violated,

Definition 2.5.2. *For the transition between $t-1$ and t, a hedger who wants to hedge the claim, $C_{t|i_{t-1}}$ with wealth, $W_{t|i_{t-1}}$ incurs a loss,*

$$L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = C_{t|i_{t-1}}(X_{t|i_{t-1}}) - W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}). \tag{2.8}$$

Without loss of generality, the support or domain of the loss function can vary based on the inclusion or exclusion of the individual variables that make up the function. For instance, in the absence of state variables, Equation (2.8) is defined as function of only the hedging strategy $X_{i_{t-1}}$ without losing the structural properties of the function. Additionally, given i_{t-1} we let $l_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = c_{j_t}(X_{j_t}) - w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ be the loss value for $j_t \in J_t|i_{t-1}$.

Proposition 2.1. *If $W_{t|i_{t-1}}$ is concave (or linear) and $C_{t|i_{t-1}}$ is convex (or linear) then the loss function $L_{t|i_{t-1}}$ is convex.*

Proof. $W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ is concave $\implies$ $-W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ is convex. Thus, the convexity of $C_{t|i_{t-1}}(X_{t|i_{t-1}})$ implies $L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ is the sum of two convex functions and hence convex. □

2.5.1 European options

Next, we illustrate how our hedging portfolio and loss function can be applied to European options. We consider a European call option with strike price K and payoff at maturity $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+ \ \forall \ j_T \in J_T$. Since the option is not exercised until maturity, for $t < T$, the value of the portfolio is the continuation value. We compose $W_{t|i_{t-1}}(X_{i_{t-1}})$, the accumulated hedge portfolio prior to payment without state variables for $j_t \in J_t|i_{t-1}$. By default, we consider a portfolio made up of two assets. That is $X_{i_{t-1}} = (x_{i_{t-1},0}, x_{i_{t-1},1})$ where $x_{i_{t-1},0}$ is the amount invested in cash at the nominal interest rate and $x_{i_{t-1},1}$ is the amount invested in the underlying stock. Hence, in the absence of state variables, we obtain

$$w_{j_t}(X_{i_{t-1}}) = x_{i_{t-1},0}e^r + x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}},\ \forall\ j_t \in J_t|i_{t-1}, \tag{2.9}$$

and $f_{i_{t-1}}(X_{i_{t-1}}) = x_{i_{t-1},0}+x_{i_{t-1},1}$, the sum of the value of position in each asset. Furthermore, we can define the values of the associated loss functions as

$$l_{j_t}(X_{i_{t-1}}) = \begin{cases} c_{j_t}(X_{j_t}) - w_{j_t}(X_{i_{t-1}}), & j_t \in J_t|i_{t-1}, t < T, \\ (y_{j_T} - K)^+ - w_{j_T}(X_{i_{T-1}}), & j_T \in J_T|i_{T-1}, t = T. \end{cases} \tag{2.10}$$

We note that it is possible to generalize our model framework to different financial derivatives and insurance contracts such as European put options, American options and variable annuities. However, we focus on only European call options for this thesis.

Chapter 3

Local Hedging Strategies

This chapter proposes how products are hedged by minimizing some objective function subject to constraints on the loss function. We further discuss some numerical results from applying the proposed strategies to products.

Several hedging strategies have been introduced in literature, Föllmer and Schweizer (1988) propose a local risk-minimizing hedging strategy that sequentially minimizes the square of the prediction error process. Coleman et al. (2007) extend existing quadratic and piece-wise linear local risk minimization frameworks, which were traditionally used for European options, to American options. They highlight that piece-wise linear risk minimization strategies can result in larger probabilities of small costs but also larger extreme costs, indicating a trade-off between risk and cost efficiency. Gaillardetz and Moghtadai (2017) also proposes partial hedging strategies that allow some positive losses by controlling a risk measure with a given threshold. These risk-control strategies can be generalized using more constraints and linear programming techniques. Gaillardetz and Osei-Mireku (2022) explores the risk-control strategy as a constraint to obtain a worst-case robust optimal value.

All the optimization algorithms to obtain the hedging strategy use backward dynamic programming approach. Starting at maturity, we apply the optimization to find the hedge portfolio for all i_{T-1}. Based on these results, we apply the optimization recursively, moving backwards until we obtain the initial hedging strategy at time $t = 0$.

3.1 Super-replication

We begin our hedging strategies with one that ensures the hedge portfolio dominates the claim or payoff at each node. Since such a hedging strategy is not necessarily self-financing, we can construct infinitely many portfolios that satisfy this condition. The requirement that the hedge portfolio dominates the claim is expanded by Föllmer and Schied (2011) to include all time steps during the product's term. Creating a portfolio with a minimal investment that super-replicates the claim at any moment is conceivable to establish a super-replicating portfolio even when the maturity is not fixed. In their study, Davis and Clark (1994) explain why super-replicating strategies cannot be the foundation for a reasonable pricing theory. They postulate that when transaction costs are strictly positive, buying and holding a single share of the risky asset is the cheapest super-replicating approach under the assumption that trading is done at a continuous rate. By introducing proper reflecting barriers, they demonstrate that the only alternatives that could be candidates for super-replicating strategies are those that closely match a Black-Scholes portfolio. However, super-replication will fail if the sell and buy barriers are close due to high transaction costs. Chen et al. (2008) propose the possibility of a super-replicating portfolio to cost less when there is transaction cost. However, when there are no transaction costs, a super-replicating portfolio must cost at least as much as a replicating portfolio. Soner (2008) studies dynamic programming to establish super-replicating portfolios in continuous time. In a discrete-time linear programming setting, our optimization problem is given by

Algorithm 3.1. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1},

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{z}_{t|i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) \tag{3.1}$$

under the constraints

$$L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \leq 0, \tag{3.2}$$

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0. \tag{3.3}$$

In this formulation, (3.3) represents a state equation that can capture the presence of transaction costs and other market-dependent dynamics. Since $f_{i_{t-1}}$ and $L_{t|i_{t-1}}$ are convex and $H_{i_{t-1}}$ is linear, the optimization problem is convex. In addition, if $f_{i_{t-1}}$ and $L_{t|i_{t-1}}$ are defined as piece-wise linear, then the optimization problem becomes convex piece-wise linear. Convex piece-wise linear problems can be solved using a standard linear programming kit or software. In the particular case where there are no state variables and state equations, the optimal objective value is the replicating portfolio value, and our convex optimization problem becomes

Algorithm 3.2. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$c_{i_{t-1}}(X_{i_{t-1}}) = \min_{X_{i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) \tag{3.4}$$

under the constraints

$$L_{t|i_{t-1}}(X_{i_{t-1}}) \leq 0, \tag{3.5}$$

where $c_{j_T}(X_{j_T})$ is the payoff for $j_T \in J_T$.

Proposition 3.3. *Suppose $L_{t|i_{t-1}}$ can be ordered, then the number of constraints in* (3.5) *can be reduced to those relative to the extreme outcomes of l_{j_t} $\forall j_t \in J_t|i_{t-1}$.*

Proof. $L_{t|i_{t-1}}(X_{i_{t-1}}) \leq 0$ implies $W_{t|i_{t-1}}(X_{i_{t-1}}) \geq C_{t|i_{t-1}}(X_{t|i_{t-1}})$. If we let $l_{M_t}(X_{i_{t-1}}) = \max\{l_{j_t}(X_{i_{t-1}}), \forall j_t \in J_t|i_{t-1}\}$ and $l_{m_t}(X_{i_{t-1}}) = \min\{l_{j_t}(X_{i_{t-1}}), \forall j_t \in J_t|i_{t-1}\}$ be the extreme outcomes such that $l_{m_t}(X_{i_{t-1}}) \leq l_{j_t}(X_{i_{t-1}}) \leq l_{M_t}(X_{i_{t-1}})$. We can express each

$l_{j_t}(X_{i_{t-1}})$ as a convex combination of $l_{m_t}(X_{i_{t-1}})$ and $l_{M_t}(X_{i_{t-1}})$. This implies, $l_{j_t}(X_{i_{t-1}}) \leq \varsigma l_{m_t}(X_{i_{t-1}}) + (1-\varsigma) l_{M_t}(X_{i_{t-1}})$ for some $\varsigma \in [0,1]$. Thus, each constraint is dominated by a convex combination of the extreme constraints. Next, we consider the convex sets formed by the intersection of the constraints' half-spaces. For simplicity of set notation we let $S_1 = \{X_{i_{t-1}} | w_{m_t}(X_{i_{t-1}}) \geq c_{m_t}(X_{m_t}), w_{M_t}(X_{i_{t-1}}) \geq c_{M_t}(X_{M_t})\} = \{X_{i_{t-1}} | w_{e_t}(X_{i_{t-1}}) \geq c_{e_t}(X_{e_t})\}$ and $S_2 = \{X_{i_{t-1}} | w_{j_t}(X_{i_{t-1}}) \geq c_{j_t}(X_{j_t}), \forall j_t \in J_t | i_{t-1} \setminus e_t\}$, where e_t represents the extreme outcomes. Since we have established that constraint (3.5) is a subset of S_1, we are left to show that $S_1 \subseteq S_2$. The necessary and sufficient condition is to show that $\exists k \in \mathbb{R}^+$ such that

(*i*) $w_{e_t}(X_{i_{t-1}}) = k w_{j_t}(X_{i_{t-1}})$,

(*ii*) $c_{e_t}(X_{e_t}) \geq k c_{j_t}(X_{j_t})$.

Suppose there is no $k \in \mathbb{R}^+$ for (*i*) to hold. Then for some $A \in \mathbb{R}^n$, we have $w_{e_t}(A) = 0$ and $w_{j_t}(A) > 0$. For any $X_{i_{t-1}} \in S_1$ and $\alpha > 0$ we have $w_{e_t}(X_{i_{t-1}} + \alpha A) = w_{e_t}(X_{i_{t-1}}) \geq c_{e_t}(X_{e_t})$ and hence $w_{e_t}(X_{i_{t-1}} + \alpha A) \in S_1$. However, we have $w_{j_t}(X_{i_{t-1}} + \alpha A) \to \infty$ as $\alpha \to +\infty$ hence $w_{j_t}(X_{i_{t-1}} + \alpha A) \notin S_2$ which is a contradiction. Furthermore, if $k < 0$, then $S_1 = \{X_{i_{t-1}} | w_{e_t}(X_{i_{t-1}}) \geq c_{e_t}(X_{e_t})\} = \{X_{i_{t-1}} | k w_{j_t}(X_{i_{t-1}}) \geq c_{e_t}(X_{e_t})\} = \{X_{i_{t-1}} | w_{j_t}(X_{i_{t-1}}) \leq \frac{c_{e_t}(X_{e_t})}{k}\}$, also a contradiction hence $k > 0$. Lastly, $S_1 = \{X_{i_{t-1}} | w_{e_t}(X_{i_{t-1}}) \geq c_{e_t}(X_{e_t})\} = \{X_{i_{t-1}} | w_{j_t}(X_{i_{t-1}}) \geq \frac{c_{e_t}(X_{e_t})}{k}\}$. But for $S_1 \subseteq S_2$, clearly $\frac{c_{e_t}(X_{e_t})}{k} \geq c_{j_t}(X_{j_t})$, hence $c_{e_t}(X_{e_t}) \geq k c_{j_t}(X_{j_t})$ and (*ii*) is also satisfied. This implies S_1 is identical to constraint (3.5). Thompson et al. (1966) explain that the i-th constraint is redundant if and only if the convex set defined by the intersection of all constraints to the linear program is identical to the set without the i-th constraint. In linear programming, redundant constraints are not essential to the solution set of the problem and, hence, can be removed. □

In literature, the goal of most super-replicating strategies is to simply dominate the payoff at the time when the option is exercise. There is no consideration for the value of the replicating portfolio during re-balancing before maturity. As such, when market conditions deviate largely from expectation, it becomes less likely for the replicating portfolio to dominate the claim at payoff. In Algorithm 3.2, not only do we ensure that the hedging portfolio dominates the claim at payoff, but also each time we re-balance the portfolio locally. The hedger is only interested in possible outcomes and wants to monitor them without any regard for how they are distributed. This approach is a useful contribution as it adds an extra layer of risk protection while and hence can be robust even in volatile markets. The set-up can also be useful to regulators since it model-independent.

3.2 $\ell_{p,q}$ norm as constraint

In an ideal setting, the best hedge is to set up a portfolio that provides perfect loss protection. However, a super-replicating strategy is known to be expensive, and any surplus at claim payment or maturity is the loss of potential gain from other alternative investments. Moreover, transaction costs associated with creating and maintaining such portfolios may decrease the likelihood of surplus. We can overcome these significant drawbacks by allowing positive losses to some degree. Different scenarios can be used to control the losses.

We can control the loss function with a constant threshold parameter γ_0 as a constraint in our optimization problem. Alternatively, one can choose a dynamic threshold γ_t that changes with time or γ_{j_t} that changes with both the node and time. The choice of threshold depends on the desired level of conservatism and risk affinity of the hedger. Different norms can be used to limit the losses. For instance, the $\ell_{1,q}$ norm limits each component in a box, the $\ell_{\infty,q}$ norm limits the total losses with a linear constraint and the $\ell_{2,q}$ norm introduces a quadratic constraint to penalize large deviations. We propose the following optimization for general $\ell_{p,q}$ norms.

Algorithm 3.4. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) \tag{3.6}$$

under the constraints

$$\|L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\|_{p,q} \leq \gamma_0, \tag{3.7}$$

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0. \tag{3.8}$$

This algorithm aims to minimize the hedge portfolio value at each node subject to a threshold that controls the losses. Here, (3.8) is a constraint to capture the evolution of state variables like transaction costs and (3.7) is a constraint to limit a risk measure of losses by the constant threshold. If $\gamma_0 > 0$, it implies the hedger allows some degree of negative losses in the portfolio to offset the total cost of hedging. We note that since $f_{i_{t-1}}$ and $L_{t|i_{t-1}}$ are convex and $H_{i_{t-1}}$ is linear, the optimization problem is convex. In addition, if $f_{i_{t-1}}$ and $L_{t|i_{t-1}}$ are defined as piece-wise linear, then the optimization problem becomes convex piece-wise linear. Convex linear problems can then be solved using a standard linear programming kit. Setting $q = 1$ gives the ℓ_p norm. In the case of the ℓ_2 norm, we can use quadratic programming or any piece-wise linear approximation of the quadratic function. In the particular case where there are no state variables and state equations, our optimization problem becomes

Algorithm 3.5. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$c_{i_{t-1}}(X_{i_{t-1}}) = \min_{X_{i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) \tag{3.9}$$

under the constraints

$$\|L_{t|i_{t-1}}(X_{i_{t-1}})\|_{p,q} \leq \gamma_0, \tag{3.10}$$

where $c_{j_T}(X_{j_T})$ is the payoff for $j_T \in J_T$.

Instead of the asymmetric norm, we can also control a coherent risk measure in the constraints by imposing an upper bound γ_0 on the worst $(1-\alpha)$ losses. This generalizes the risk control strategy proposed by Gaillardetz and Osei-Mireku (2022). They explain that the robustness of the models can be explored by studying the dynamics of the underlying probability distributions. However, rather than imposing some distribution on the losses, Algorithm 3.5 assumes each outcome is equally likely and constraints the magnitude of the losses instead. This formulation is also model-independent and can easily be adapted to include information on the distribution of losses. Instead of using a constant threshold parameter, a hedger can express the threshold as a function of both the node and time. Hence, the strategy provides the hedger with parameters to optimize their preferred level of conservatism.

3.3 $\ell_{p,q}$ norm as objective

Another hedging strategy that seeks to control losses is directly minimizing the norm or risk measure in the objective function. Much like in prevalent portfolio selection problems, the aim is to obtain a hedge portfolio or strategy by minimizing the $\ell_{p,q}$ norm in the objective of our linear program. Furthermore, by introducing thresholds, we can include constraints in the form of asymmetric norms to adjust the optimization sequence to the desired level of conservatism. Again, the choice of $\ell_{p,q}$ norm depends on the investor's risk tolerance level. We propose the following optimization for general $\ell_{p,q}$ norms.

Algorithm 3.6. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r} \| L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \|_{p_1,q_1} \tag{3.11}$$

under the constraints

$$\| L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \|_{p_2,q_2} \leq \gamma_0, \tag{3.12}$$

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0, \tag{3.13}$$

where $c_{j_T}(X_{j_T})$ is the payoff for $j_T \in J_T$.

Again, (3.13) is a constraint to capture the evolution of state variables like transaction costs and (3.12) is a constraint to limit the asymmetric norm of losses by the constant threshold. The level of asymmetry can also be modified by changing the value of q. For instance, if $q = 1$, we get the symmetric norm. Also, $q = 0$ implies no weight on negative errors, pushing the optimization to consider only positive ones. The choice of parameter p_1, q_1 in the objective function is independent of the choice of p_2, q_2 in constraint (3.12). The norm in the constraint is not mandatory if the norm in the objective is two-sided. Constraint (3.12) is needed to control losses if the choice of asymmetry in the objective can result in unbounded solutions. Again, since $L_{t|i_{t-1}}$ and $\|.\|_{p,q}$ are convex, the optimization problem is convex and can be solved using a standard linear programming kit. It is important to note that, without any modification to the $\ell_{p,q}$ norm, choosing $p_1 = 2, q_1 = 1$ is equivalent to solving an ordinary least-square regression at every node of the model framework. Any modification may require quadratic programming techniques to solve the algorithm. In the absence of state variables and the threshold, the local minimization of the $\ell_{2,1}$ norm in the objective is equivalent to the local hedging strategy proposed by Schweizer (2008).

3.4 Portfolio value as state variable

In the previous sections, we introduced local hedging strategies with the goal centred around only local losses. In this section, we introduce another set of local hedging techniques developed by controlling both present and future estimates of losses. We refer to these local hedging techniques as dynamic local hedging. The term dynamic is used here to reference the inclusion of state variables as estimates of future portfolio value in the optimization problem. We note that our dynamic local hedging strategies also use the dynamic programming technique to determine the hedging strategies.

An alternative strategy to control losses is representing portfolio values as state variables and minimizing the sum of the hedge portfolio's local losses and future cost-to-go objective.

In other instances, the state variables can instead be used to control the total losses from the position. Once this optimization sequence is complete, the hedger sets the initial capital to finance the portfolio. The results of the loss optimization are used in addition to the state variables to derive the hedging strategies at each step. The choice of initial capital can be made by selecting the value that generates the least tail hedging errors after the hedging strategies have been derived. In this formulation, a function of the losses must be minimized, with the losses equal to zero at maturity. Norms or risk measures are utilized in the objective function to choose hedging strategies, much like in portfolio selection problems. We propose three variations for implementing such models. One is based on stochastic programming, the other on dynamic coherent risk and a barrier to future risk.

3.4.1 Overview of the dynamic algorithm

We start by providing the main outline of our proposed dynamic algorithm for developing a portfolio as state variable hedging strategy. The main target is to find a hedge portfolio that minimizes a function of discounted errors and proportion of the future cost-to-go function, subject to constraints on the losses and state variables. In general, we consider the following framework.

Algorithm 3.7. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r} O(L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}), V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})) \tag{3.14}$$

under the constraints

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0, \tag{3.15}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, and $c_{j_T}(X_{j_T})$ is the payoff for $j_T \in J_T$.

$O(.,.)$ is a function of the loss vector $L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ and the cost-to-go vector

$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ which is the future objective function evaluated at the vector of state variables $\mathbf{Z}_{t|i_{t-1}}$. For example, $O(.,.)$ can be the norm of the sum of $L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ and $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$.

At the onset of the algorithm, we need at least the portfolio value as a state variable. Since we do not know this intrinsic value, we modify the algorithm to include a search process. At maturity, we set $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$ in the objective function and minimize $O(L_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}), \mathbf{0})$ subject to additional state equations (3.15) that defines a polyhedral of the portfolio values. The polyhedral is bounded by a state variable grid $\nu_{i_{T-1}}$ defined as an interval that contains the payoff c_{j_T}. Thus, each state variable, $\mathbf{z}_{j_T} \in \nu_{i_{T-1}}$ for $j_T \in J_T|i_{T-1}$. We solve the optimization at maturity to obtain the cost-to-go values $v_{j_{T-1}}(\mathbf{z}_{j_{T-1}})\ \forall\ j_{T-1} \in J_{T-1}|i_{T-2}$ that make up the vector $V_{T-1|i_{T-2}}(\mathbf{Z}_{T-1|i_{T-2}})$. We also obtain the slope and intercepts for each optimal value in the state variable grid to be used to estimate the hedging strategy.

Conditional on i_{T-2}, we minimize $O(L_{T-1|i_{T-2}}(X_{i_{T-2}}, \mathbf{z}_{i_{T-2}}, \mathbf{Z}_{T-1|i_{T-2}}), V_{T-1|i_{T-2}}(\mathbf{Z}_{T-1|i_{T-2}}))$ by searching another polyhedral of the payoff at the node subject to the state equation constraint $f_{i_{T-2}}(X_{i_{T-2}}) = \mathbf{z}_{i_{T-2}}$, which represents a component of constraint (3.15). Consequently, we obtain a set of optimal slopes, corresponding intercepts and proximal estimates of the future objective function $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ by repeating the search sequentially at each node until time $t = 0$. Finally, to obtain the hedging portfolio, we start at time $t = 0$ and set the initial capital $f_0(X_0)$, to finance the strategy. Using the estimated slopes, intercepts and state variables recorded from the previous step, we re-evaluate the algorithm without the search process to obtain the optimal portfolio strategy $X^*_{i_{t-1}}$ at each node.

3.4.2 Stochastic programming

This approach aims to minimize the discounted cumulative local and future losses at each time $t < T$. Starting at maturity, we set the cost-to-go function to zero. This ensures that the claim or payoff is locally hedged. At each time $t < T$ and for every node i_{t-1}, we minimize the objective function, which comprises the discounted local losses and a fraction of estimated future objective function subject to state variables as constraints for $j_t \in J_t|i_{t-1}$. The proposed strategy minimizes the discounted average future risk such that the state variable is the continuation value and can be equal to the portfolio value at each step. We summarize the strategy using the algorithm below.

Algorithm 3.8. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r}\{\|L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\|_{p,q} + \beta\|V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})\|_{p,q}\} \tag{3.16}$$

under the constraints

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0, \tag{3.17}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, and $c_{j_T}(X_{j_T})$ is the payoff for $j_T \in J_T$.

The objective function is the sum of a local norm and the weighted average of near-future risk. Here, the parameter β sets the degree of the future portfolio value's impact on the minimization of local losses and can be used to control the level of robustness. When $\beta = 0$, we obtain Algorithm 3.6, and the problem can be solved directly without the state variables. When $\beta \neq 0$, we formulate the problem using our dynamic algorithm. Hence, we obtain the following optimization problem,

Algorithm 3.9. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \Theta_{i_{t-1}}} e^{-r}\{\|L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\|_{p,q} + \beta\|\Theta_{t|i_{t-1}}\|_{p,q}\} \tag{3.18}$$

under the constraints

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq \Theta_{t|i_{t-1}}, \tag{3.19}$$

$$f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}, \tag{3.20}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, *and* $c_{j_T}(X_{j_T})$ *is the payoff for* $j_T \in J_T$.

Here, $\Theta_{t|i_{t-1}}$ is a vector with component θ_{j_t} for $j_t \in J_t|i_{t-1}$ and represents the upper bound of $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$. In this formulation, $\Theta_{T|i_{T-1}} = \mathbf{0}$ at maturity and since $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, we remove (3.19) for $t = T$. For the remaining $t < T$, $\Theta_{t|i_{t-1}}$ is introduced in the objective function. If the supporting hyperplane of $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is piece-wise convex, (3.19) is equivalent to several linear constraints where each supporting hyperplane of $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is bounded above by $\Theta_{t|i_{t-1}}$.

3.4.3 Dynamic coherent risk approach

The information available during risk assessment is considered a conditional risk measure. Risk assessments are revised over time in light of new or additional information in a dynamic environment. A sequence of conditional risk measures adapted to the underlying filtration is known as a dynamic risk measure. Acciaio and Penner (2011) provides a detailed discussion of dynamic risk measures with robust representation and consistency properties.

We require the risk measure to be coherent to satisfy the convexity of our optimization problem. Introduced by Riedel (2004), they define a dynamic risk measure as coherent if the sequence of conditional risk measures is homogeneous and sub-additive. They show that when the expectations are taken over a group of probability measures that satisfy a consistency constraint, they can be expressed as the worst conditional expectation of discounted future risks. Gaillardetz and Hachem (2022) also propose a hedging strategy based on discounted recursion of a dynamic coherent risk measure to value Equity-Indexed Annuities and Guaranteed Investment Certificates.

Contrary to the stochastic programming approach, we aim to minimize the discounted coherent risk measure on local and successive near-future losses. However, the optimization approach remains the same. We set the cost-to-go function to be zero at maturity and use the dynamic algorithm to minimize the objective function at each time $t < T$. This approach lets us keep path-wise or dynamic risk control in the optimization problem. We summarize the strategy using the algorithm below.

Algorithm 3.10. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1},

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r} \| L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) + \beta V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \|_{p,q} \tag{3.21}$$

under the constraints

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0, \tag{3.22}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, *and* $c_{j_T}(X_{j_T})$ *is the payoff for* $j_T \in J_T$.

The objective function is the norm of the weighted average sum of local losses and the near-future cost-to-go function. Again, a risk measure like the CVaR can be used as the function on local losses and setting $\beta = 0$ produces Algorithm 3.6. For $\beta \neq 0$, we formulate the problem using our dynamic algorithm to obtain the following optimization,

Algorithm 3.11. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1},

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \Theta_{i_{t-1}}} e^{-r} \| L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) + \beta \Theta_{t|i_{t-1}} \|_{p,q} \tag{3.23}$$

under the constraints

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq \Theta_{t|i_{t-1}}, \tag{3.24}$$

$$f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}. \tag{3.25}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$*, and* $c_{j_T}(X_{j_T})$ *is the payoff for* $j_T \in J_T$*.*

The cost-to-go function $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is modelled through the definition of its epigraph. Again, (3.24) is equivalent to several linear constraints where each supporting hyperplane of $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is bounded above by $\Theta_{t|i_{t-1}}$ provided $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is a linear piece-wise convex function defined by its supporting hyperplane. The mathematical properties and solution techniques are the same for stochastic programming. The only difference lies in the objective function of the problems. Instead of having a composition of a risk measure and future cost-to-go functions, the stochastic programming approach sums the local risk measure and the norm of the cost-to-go function. This implies that at optimal value, path-wise solutions for dynamic coherent risk approach are bounded above by solutions for stochastic programming.

3.4.4 Barrier on future risk

Gaillardetz and Hachem (2022) tested a stochastic programming model that minimizes weighted average risk measures on Guaranteed Investment Certificates and noticed a weakness during back-testing. They found that models containing a risk measure in the objective function could induce some nodes to have excessively high values of this risk measure during numerical implementations. By introducing a barrier to future risk, we can obtain a model with the same properties as the stochastic programming approach. We address the drawback by putting a threshold on the future risk to make an additional constraint to the optimization problem. Essentially, we minimize a discounted local risk measure subject to a threshold on weighted future risk. We summarize the strategy using the algorithm below.

Algorithm 3.12. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1}*,*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r} \| L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \|_{p,q} \tag{3.26}$$

under the constraints

$$H_{i_{t-1}}(\mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = 0, \tag{3.27}$$

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq \Gamma_0, \tag{3.28}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$*, and* $c_{j_T}(X_{j_T})$ *is the payoff for* $j_T \in J_T$*.*

The parameter Γ_0 is a vector with components γ_0 for $j_t \in J_t|i_{t-1}$ and controls the issuer's permitted future risk. Since the cost-to-go function $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is now incorporated in the constraint rather than the objective function, the barrier on future risk algorithm differs fundamentally from our conventional stochastic programming approach. Alternative methods, such as expected value, exist to limit future risk measures. We favour the constraint (3.28) since it manages the path-wise risk. Again, the portfolio value must be one of the state variables. The state equation must be $f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}$ if we allow the state vector to be the intrinsic portfolio value. Formulating the optimization problem this way is analogous to that of stochastic programming, where $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is replaced by its supporting hyperplane. As a result, this algorithm uses the same optimization techniques as Algorithms 3.8 and 3.10 with similar features and levels of complexity.

Chapter 4

Application of local hedging strategies

We now apply the hedging strategies to European-style products. We use the term European-style loosely to represent products with a single payout only at maturity and do not permit claims to be redeemed before maturity. Examples of such products are Guaranteed Investment Certificates and European options. Guaranteed Investment Certificates provide some return on investment while guaranteeing the investor's principal. Specifically, a non-cashable Guaranteed Investment Certificate has a payoff comprising a guaranteed minimum investment rate. In some cases, the guaranteed minimum rate can be fixed or variable tied to the performance of an index. An issuer will likely cap the overall return of Guaranteed Investment Certificates with variable rates. For our application, we rely on only European options. We also study the behaviour of the strategies in the presence of transaction costs.

4.1 Financial model

To illustrate how hedging strategies can be applied to products, we must first design the model framework. For this purpose, we consider a financial setting and build a stock or index model. The dynamics of stocks are replicated using a variety of models. Over time, it has been demonstrated that lattice pricing models can successfully simulate stocks or indexes, interest rates, and other financial instruments. The two-state (binomial) lattice technique, first presented by Cox et al. (1979), has shown to be a useful model in the valuation of several financial securities. Their fundamental market process was demonstrated to converge to the independent log-normal model under various assumptions. The model is insufficient when long-term market fluctuations and dramatic movements are considered. This is because the model cannot capture the dramatic movements as a result of shifts in regimes over a long period. Boyle (1988) advances a three-state lattice model. They calculate the needed risk-neutral probabilities associated with transitions by comparing a discrete model's moments to a continuous log-normal distribution's mean and variance. Since then, numerous continuous-time approximation models and multi-period discrete multinomial lattice models have been developed.

We assume a discrete lattice where Y_t is defined by the stock process whose values are known through time. At the root node $t = 0$, the stock process has only one value, thus $Y_0 = y_0$. At time $t > 0$, each node i_{t-1} is connected to $N + 1$ possible nodes corresponding to distinct values y_{j_t} in the vector $Y_{t|i_{t-1}}$ for $j_t \in J_t|i_{t-1}$. To model the transition from one time to the next, we assume the stock value goes up by a factor u or down by a factor d. As such, our lattice model has $N + 1$ distinct outcomes of $Y_{t|i_{t-1}}$ given by $y_{i_{t-1}} u^{N-j} d^j$ for $j = 0, 1, \ldots, N$. If we assume the stock process converges to the log-normal distribution of a geometric Brownian motion with drift μ and volatility σ, we can define $u = e^{\sigma(\Delta N)^{-0.5}}$ where Δ represents the number of splits between two-time steps. To reduce the computational intensity, we construct a recombining tree (see Figure 2.1) by setting $d = u^{-1}$, although it is not a necessary condition for our framework.

4.2 Hedging with super-replication (SR)

Here, we apply a super-replication hedging strategy to European options and review the optimization problem to find close-form solutions where possible. For the application, we consider Algorithm 3.2 with stock outcomes $y_{0_t} \geq y_{1_t} \geq \ldots \geq y_{N_t} \geq 0$ at a fixed time t. This implies $c_{0_T}(X_{0_T}) \geq c_{1_T}(X_{1_T}) \geq \ldots \geq c_{N_T}(X_{N_T}) \geq 0$ for the corresponding European option payoffs. A similar relationship can be defined for subsequent $c_{j_t}(X_{j_t})$ as we evaluate the tree dynamically. With the help of Proposition 3.3, we can formulate our optimization problem in terms of only the extreme values of the stock and payoff process as below.

Example 4.1. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$c_{i_{t-1}}(X_{i_{t-1}}) = \min_{X_{i_{t-1}}} x_{i_{t-1},0} + x_{i_{t-1},1} \tag{4.1}$$

such that

$$c_{0_t}(X_{0_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{0_t}}{y_{i_{t-1}}} \leq 0, \tag{4.2}$$

$$c_{N_t}(X_{N_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{N_t}}{y_{i_{t-1}}} \leq 0, \tag{4.3}$$

where $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+$, $\forall j_T \in J_T$.

Furthermore, since the objective function and constraints are linear with respect to the decision variables, we obtain a boundary solution for our optimization at each time step. This stems from constraint qualification analysis for convex optimization. We let $\nabla_{X_{i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}})$ be the partial derivative of $f_{i_{t-1}}(X_{i_{t-1}})$ with respect to the vector $X_{i_{t-1}}$. Using linear independence constraint qualification (LICQ), we obtain a boundary solution since an interior solution would imply $\nabla_{X_{i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{0}$ but $\nabla_{X_{i_{t-1}}} f_{i_{t-1}}(X_{i_{t-1}}) = (1, 1) \neq \mathbf{0}$. For instance, Example 4.1 has two variables, and the maximum number of linearly independent constraints is two. This confirms that the constraints can be reduced to at least the two extremes with optimal boundary values $x_{i_{t-1},0} = \frac{c_{N_t}(X_{N_t})y_{0_t} - c_{0_t}(X_{0_t})y_{N_t}}{e^r(y_{0_t} - y_{N_t})}$ and $x_{i_{t-1},1} = \frac{y_{i_{t-1}}\left(c_{0_t}(X_{0_t}) - c_{N_t}(X_{N_t})\right)}{y_{0_t} - y_{N_t}}$. This closed-form solution can be extended to having multiple assets in the portfolio provided $f_{i_{t-1}}(X_{i_{t-1}})$ remains linear. The algorithm simplifies solving the system of linear inequalities containing extreme constraints. For instance, solutions for a three-asset portfolio can be obtained by solving three inequalities simultaneously. Two are extremes, and another is centred between the extremes to preserve feasibility and convexity. Limiting the solution to only the extremes, in this case, is similar to the risk-neutral approach.

4.3 Hedging with $\ell_{p,q}$ norm as constraint (NC)

Gaillardetz and Osei-Mireku (2022) explored the use of CVaR as a controlled risk constraint in their application to EIAs. In this section, we apply the $\ell_{p,q}$ norm as controlled constraints to losses. For the application, we consider Algorithm 3.5 with the same stock dynamics as hedging with super-replication. That is, $y_{0_t} \geq \ldots \geq y_{N_t} \geq 0$ at a fixed time t and $c_{0_T}(X_{0_T}) \geq \ldots \geq c_{N_T}(X_{N_T}) \geq 0$ for the corresponding European option payoffs. Our optimization transforms to the following example for an asymmetric $\ell_{1,0}$ norm.

Example 4.2. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$c_{i_{t-1}}(X_{i_{t-1}}) = \min_{X_{i_{t-1}}} x_{i_{t-1},0} + x_{i_{t-1},1} \tag{4.4}$$

such that

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.5}$$

$$\sum_{j_t \in J_t|i_{t-1}} l^+_{j_t}(X_{i_{t-1}}) \leq \gamma_0,, \tag{4.6}$$

where $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

Based on the definition of ℓ_1 norm, constraint (4.6) should strictly be $\sum_{j_t \in J_t|i_{t-1}} |l_{j_t}(X_{i_{t-1}})| \leq \gamma_0$. However, using this characterization renders the linear programming problem infeasible at some nodes for an unbounded stock tree. As such, we consider an asymmetric norm with $q = 0$ to allow errors to be strictly positive. This makes the problem feasible at all nodes of an unbounded stock tree.

Also, for $\ell_{2,0}$ norm, we replace constraint (4.6) with the squared loss function defined as

$\sum_{j_t \in J_t|i_{t-1}} \left(l^+_{j_t}(X_{i_{t-1}})\right)^2 \leq \gamma_0^2$. Similar to the ℓ_1, this formulation can be infeasible for an unbounded stock tree at some nodes. As such, we consider the asymmetric $\ell_{2,0}$ norm for errors to be strictly positive. Thus, we obtain a quadratic constraint programming problem that can be solved with most quadratic programming kits. In our numerical analysis, we use the IBM ILOG Concert Technology for Linear Programming and Extensions (CPLEX) kit for our quadratic optimization.

For $\ell_{\infty,1}$ norm, we replace constraint (4.6) with a threshold on all losses. The goal is to constrain the maximum norm of loses with the threshold γ_0. In this case, we can deduce that if the maximum piece-wise absolute loss is bounded above by our threshold, then each absolute loss must be less than the defined threshold. Hence $|l_{j_t}(X_{i_{t-1}})| \leq \gamma_0, \ \forall\ j_t \in J_t|i_{t-1}$. Since we are only concerned with positive losses, another modification is to constrain the maximum loss instead of the maximum absolute loss. In this case, $L_{t|i_{t-1}}(X_{i_{t-1}}) \leq \Gamma_0$. Also, we can replace the constant threshold γ_0 with a time-dependent threshold γ_t or a node-dependent threshold γ_{i_t}. Setting $\gamma_{i_t} = 0$ for all time steps yields an algorithm similar to the super-replicating strategy, and the problem can be reduced to constraints based on only the extremes. However, allowing some degree of positive losses during optimization implies that any of the nodes or a combination of them contributes to active constraints, and the full algorithm has to be solved to arrive at a solution.

4.4 Hedging with $\ell_{p,q}$ norm as objective (NO)

We consider an example for hedging with $\ell_{p,q}$ norm as objective. Again, we adopt the stock dynamics from the super-replicating example and consider Algorithm 3.6 for application. We slightly modify our formulation, which allows us to introduce bounds on the sum of errors to make the minimization finite. In some situations, an asymmetric norm requires an extra constraint to be formulated to make the problem feasible. The presence of a cost-to-go function in the objective is feasible but would require more dynamic optimization techniques, which we explore under portfolio value as state variable. Our optimization transforms to the following example for an $\ell_{1,1}$ norm.

Example 4.3. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1}*,*

$$X^*_{i_{t-1}} = \operatorname*{argmin}_{X_{i_{t-1}}} e^{-r} \sum_{j_t \in J_t | i_{t-1}} |l_{j_t}(X_{i_{t-1}})| \tag{4.7}$$

such that

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.8}$$

$$|\sum_{j_t \in J_t | i_{t-1}} l_{j_t}(X_{i_{t-1}})| \le \gamma_0, \tag{4.9}$$

where $X^*_{i_{t-1}}$ *is the optimal strategy and* $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$*.*

Again, we can control the losses by setting the sum of errors to be within γ_0, although this constraint is not mandatory when the objective norm is two-sided. This is done to centre the losses.

We obtain the $\ell_{2,1}$ version of this strategy by substituting the objective function in (4.7) with the quadratic version where

Example 4.4. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1}*,*

$$X^*_{i_{t-1}} = \operatorname*{argmin}_{X_{i_{t-1}}} e^{-r} \sum_{j_t \in J_t | i_{t-1}} l_{j_t}(X_{i_{t-1}})^2 \tag{4.10}$$

such that

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.11}$$

$$|\sum_{j_t \in J_t | i_{t-1}} l_{j_t}(X_{i_{t-1}})| \le \gamma_0, \tag{4.12}$$

where $X^*_{i_{t-1}}$ *is the optimal strategy and* $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$*.*

Without the loss control feature in constraint (4.12), the optimization is reduced to a least squares problem. Again, constraint (4.12) is not mandatory as the optimization is designed to be centred by second-order condition. In this case, we seek the optimum hedging strategy that minimizes the squared errors of the hedge portfolio. The control parameter in (4.12) differs from regularization techniques. This is because, while regularized linear models tend to shrink or penalize the coefficients of regression, in this case, the hedging strategy, (4.12) tends to control the magnitude of the loss region with the parameter γ_0. On the other hand, for a $\ell_{\infty,1}$ norm, since the aim is to minimize the maximum possible absolute error, our optimization transforms to the example below.

Example 4.5. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$X^*_{i_{t-1}} = \underset{X_{i_{t-1}}}{\operatorname{argmin}}\, e^{-r} b \tag{4.13}$$

such that

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0} e^r - x_{i_{t-1},1} \frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t | i_{t-1}, \tag{4.14}$$

$$|l_{j_t}(X_{i_{t-1}})| \leq b,\ \forall j_t \in J_t | i_{t-1}, \tag{4.15}$$

$$|\sum_{j_t \in J_t | i_{t-1}} l_{j_t}(X_{i_{t-1}})| \leq \gamma_0, \tag{4.16}$$

where $X^*_{i_{t-1}}$ is the optimal strategy and $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

By setting $|l_{j_t}(X_{i_{t-1}})| \leq b$ for all $j_T \in J_T | i_{T-1}$, we ensure that $\max_{j_t \in J_t | i_{t-1}} \{|l_{j_t}(X_{i_{t-1}})|\} \leq b$ and hence we can minimize b in our objective. We introduce a symmetric modification of losses in (4.16) to control the level of conservatism and ensure the feasibility of the problem.

4.5 Hedging with stochastic programming (SP)

This section considers a typical application of the stochastic programming technique using Algorithm 3.9 and the same stock dynamics from the super-replication example. The goal is to minimize the cumulative losses over time while considering current and future losses. The algorithm minimizes discounted average future risk, with the state variable representing the continuation value, which can equal the portfolio value at each step. We obtain the following optimization for $\ell_{1,1}$ norm.

Example 4.6. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \Theta_{t|i_{t-1}}} e^{-r} \sum_{j_t \in J_t|i_{t-1}} |l_{j_t}(X_{i_{t-1}})| + \beta|\theta_{j_t}| \tag{4.17}$$

under the constraints

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.18}$$

$$v_{j_t}(\mathbf{z}_{j_t}) \leq \theta_{j_t},\ \forall j_t \in J_t|i_{t-1}, \tag{4.19}$$

$$x_{i_{t-1},0} + x_{i_{t-1},1} = \mathbf{z}_{i_{t-1}}, \tag{4.20}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, and $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

In this example, we minimize the $\ell_{1,1}$ norm of current losses and a portion of the $\ell_{1,1}$ norm of future cost-to-go function bounded above by θ_{j_t}. The desired level of future losses in the objective function is controlled by the parameter β, which can be interpreted as a measure of robustness in this setting. The $\ell_{2,1}$ norm representation for stochastic programming is similar to the $\ell_{1,1}$ norm where the objective function is replace with $\sum_{j_t \in J_t|i_{t-1}} l_{j_t}(X_{i_{t-1}})^2 + \beta\theta_{j_t}^2$. However, an for $\ell_{\infty,1}$ representation, we obtain the following optimization,

Example 4.7. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, b_1, b_2} e^{-r}\{b_1 + \beta b_2\} \tag{4.21}$$

under the constraints

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.22}$$

$$|l_{j_t}(X_{i_{t-1}})| \leq b_1,\ \forall j_t \in J_t|i_{t-1}, \tag{4.23}$$

$$|v_{j_t}(\mathbf{z}_{j_t})| \leq b_2,\ \forall j_t \in J_t|i_{t-1}, \tag{4.24}$$

$$x_{i_{t-1},0} + x_{i_{t-1},1} = \mathbf{z}_{i_{t-1}}, \tag{4.25}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, and $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

This algorithm aims to find the hedging strategy that minimizes the maximum absolute current losses and a proportion of maximum future losses. Here, we set b_1 and b_2 as the maximum of local losses and future cost-to-go functions. Thus, by minimizing b_1 and b_2, we obtain a hedging strategy that leans towards the worst-case scenario for the specified trading path. Again, setting $\beta = 0$ gives us a problem equivalent to Example 4.5.

4.6 Hedging with dynamic coherent risk (DC)

Similarly, we consider an application of the dynamic coherent risk technique using Algorithm 3.11. Unlike stochastic programming, our goal is to minimize the norm of local and near-future losses, which tends to produce a more conservative hedging strategy. We obtain the following optimization for $\ell_{1,1}$ norm.

Example 4.8. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, U_{t|i_{t-1}}, \Theta_{t|i_{t-1}}} e^{-r} \sum_{j_t \in J_t|i_{t-1}} |u_{j_t}| \tag{4.26}$$

under the constraints

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} + \beta\theta_{j_t} = u_{j_t},\ \forall j_t \in J_t|i_{t-1}, \tag{4.27}$$

$$v_{j_t}(\mathbf{z}_{j_t}) \leq \theta_{j_t},\ \forall j_t \in J_t|i_{t-1}, \tag{4.28}$$

$$x_{i_{t-1},0} + x_{i_{t-1},1} = \mathbf{z}_{i_{t-1}}, \tag{4.29}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, *and* $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

Here, $U_{t|i_{t-1}}$ is a vector with components u_{j_t} for $j_t \in J_t|i_{t-1}$. In this example, local and future losses are minimized at each optimization. The parameter β controls future losses' impact during minimization. A higher value for β leads to stricter hedging where we penalize future losses more than local losses, leading to a robust strategy. For $\ell_{2,1}$ norm representation the objective function is replace with $\sum_{j_t \in J_t|i_{t-1}} u_{j_t}^2$. On the other hand, $\ell_{\infty,1}$ optimization is formulated as follows,

Example 4.9. *For all* $t = T, T-1, \ldots, 1$ *and all* i_{t-1},

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, U_{t|i_{t-1}}, b} e^{-r}b \tag{4.30}$$

under the constraints

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} + \beta\theta_{j_t} = u_{j_t},\ \forall j_t \in J_t|i_{t-1}, \tag{4.31}$$

$$|u_{j_t}| \leq b,\ \forall j_t \in J_t|i_{t-1}, \tag{4.32}$$

$$v_{j_t}(\mathbf{z}_{j_t}) \leq \theta_{j_t},\ \forall j_t \in J_t|i_{t-1}, \tag{4.33}$$

$$x_{i_{t-1},0} + x_{i_{t-1},1} = \mathbf{z}_{i_{t-1}}, \tag{4.34}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, *and* $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

Like stochastic programming, the $\ell_{\infty,1}$ dynamic coherent risk strategy leans towards the worst-case scenario for the specified trading path. The following section compares the two methods to observe their impact on the VaR and CVaR of hedging costs.

4.7 Hedging with the barrier on future risk (BF)

Rather than having the future loss variable in the objective function, we can control the impact by limiting future risk to a threshold. The barrier to future risk hedging strategy allows the hedger to limit all future losses below a single or path-dependent threshold. For $\ell_{1,1}$ norm, Algorithm 3.12 leads to the following optimization.

Example 4.10. *For all $t = T, T-1, \ldots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} e^{-r} \sum_{j_t \in J_t|i_{t-1}} |l_{j_t}(X_{i_{t-1}})| \tag{4.35}$$

under the constraints

$$c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} = l_{j_t}(X_{i_{t-1}}),\ \forall j_t \in J_t|i_{t-1}, \tag{4.36}$$

$$v_{j_t}(\mathbf{z}_{j_t}) \leq \gamma_0,\ \forall j_t \in J_t|i_{t-1}, \tag{4.37}$$

$$x_{i_{t-1},0} + x_{i_{t-1},1} = \mathbf{z}_{i_{t-1}}, \tag{4.38}$$

where $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = \mathbf{0}$, and $c_{j_T}(X_{j_T}) = (y_{j_T} - K)^+,\ \forall j_T \in J_T$.

The $\ell_{2,1}$ and $\ell_{\infty,1}$ norm variations follow the same format except for changes in their corresponding objective functions. The static threshold parameter γ_0 can be replaced with the dynamic threshold γ_t that changes with time or γ_{j_t} that changes with both time and node depending on the hedger's preference. It is worth noting that using a static threshold often leads to infeasible solutions when the size of the lattice tree is large. As such, a more favourable threshold to ensure feasibility depends on the size of the lattice tree.

4.8 Numerical examples

In this section, we analyze the outputs of the model examples produced by taking various parameter adjustments into account. We first consider the market to be frictionless. That is, there are no transaction fees or taxes. Later in this section, we consider the addition of transaction cost and analyze its impact on the baseline models.

4.8.1 Hedging costs

To assess the effectiveness of our proposed strategies, we consider the discounted value of path-dependent mismatches incurred by the hedging strategies. The nature and scope of hedging errors that occur when a portfolio is re-balanced in discrete time was first investigated by Boyle and Emanuel (1980). Their paper analyzes the distribution of hedging errors and suggests procedures to address the mismatches during re-balancing. Since then, extensive investigations have suggested that while hedging errors cannot be eliminated, frequent re-balancing of the replicating portfolio tends to minimize these mismatches. Albeit at the expense of transaction costs, which can be optimized. This concept of minimizing mismatches through frequent re-balancing does not entirely hold under robust optimization. We examine the total cost incurred by an issuer or hedger. In effect, the initial replicating portfolio value and the errors incurred by the hedging strategy constitute the hedger's entire cost. To obtain the hedging errors, we first solve the optimization problem to obtain the hedging strategy at each node for all time $t \in \mathcal{T}$. Then, using induced probabilities of the input process, we simulate a path from the tree and use the path's corresponding strategies to compute the temporary loss values as the errors. Thus, we can define the present value of hedging errors as

$$h = \sum_{j_t \in \{i_1, \ldots, i_T\}} e^{-rt} l_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}), \tag{4.39}$$

where $l_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ is the loss from hedging the claim $c_{j_t}(X_{j_t})$ with accumulated portfolio value $w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})$ using the hedging strategy for $j_t \in J_t|i_{t-1}$. From the optimal hedging strategies, we can compute f_0, the initial value of the hedging portfolio. We can also simulate several possible paths and compute a collection of hedging errors h for each simulated path. We present a level comparison of the strategies by allowing the collection of paths to be controlled by the constant drift parameter μ to capture the evolution under physical measure $\mathbb{P}$. We introduce the induced probabilities $p_{j_t|i_{t-1}}$ to do this. At each time t, we obtain the induced probabilities using the binomial distribution since this simulates the geometric Brownian motion of the underlying stock process. Hence for some $\rho = \frac{e^{\mu/(\Delta N)}-d}{u-d}$ and given i_{t-1}, the conditional probabilities $p_{j_t|i_{t-1}}\ \forall\ j_t \in J_t|i_{t-1}$ can be expressed as,

$$p_{j|i_{t-1}} = \binom{N}{j} \rho^{N-j}(1-\rho)^j, \text{ for } j = 0, 1, \ldots, N. \tag{4.40}$$

Gaillardetz and Osei-Mireku (2022) explore using MCMC sampling techniques to develop non-homogeneous probabilities that vary at each time step in the tree.

Definition 4.8.1. *The hedger's total cost $\mathcal{H}$, is the sum of the initial value of the hedge portfolio f_0, and the simulated hedging errors h.*

Thus, for any given hedging strategy, we compute the hedger's total cost and estimate the riskiness of a strategy using our proposed risk measures.

4.8.2 Simulation setup

In our analysis, we hedge an at-the-money European option with a unit stock/index price at $t = 0$ (i.e., $Y_0 = K = 1$). We define our default hedge portfolio as comprising three assets at each time step. The first two $(x_{i_{t-1},0}, x_{i_{t-1},1})$ are the cash and stock investments respectively. The third asset, $x_{i_{t-1},2}$, is invested in a one-period at-the-money call option. We denote the node j_t Black-Scholes price of a one-period call option with strike κy_{j_t} as $\phi_{j_t}(\kappa)$, where κ is a scalar to set the moneyness of the option so that $\kappa = 1$ implies the option is at-the-money. Alternatively, we can consider the option prices implied by the tree; however, we use the one-period Black-Scholes price for simplicity. Thus our loss function can be redefined as

$l_{j_t}(X_{i_{t-1}}) = c_{j_t}(X_{j_t}) - x_{i_{t-1},0}e^r - x_{i_{t-1},1}\frac{y_{j_t}}{y_{i_{t-1}}} - x_{i_{t-1},2}(y_{j_t} - \kappa y_{i_{t-1}})^+,\ \forall j_t \in J_t|i_{t-1}$. To compare the strategies, we simulate hedging errors under market evolution governed by a constant drift μ that can be set to capture physical market expectation and the binomial transition probabilities defined in (4.40). In the absence of any indication to the contrary, we let $T = 1$ for a 1-year European call option, $\Delta = 12$ to represent monthly trading/re-balancing, $N = 20$, $r = 4\%$, $\mu = 8\%$, $\sigma = 20\%$ and $\gamma_0 = 1\%$. We compute the mean and standard deviation of $100,000$ simulated hedging costs $\mathcal{H}$ for each strategy. Recall $\mathcal{H}$ is the collection of the sum of the initial value of the hedge portfolio f_0 and each of the simulated hedging errors h. We also estimate the VaR and CVaR at $\alpha = 99.5\%$ to assess the tail performance of the strategies.

4.8.3 Analysis of hedging error distribution for SR, NC and NO strategies

We begin our analysis by comparing the hedging error distribution of the various strategies. Our goal is to assess the robustness of the models based on their tail risk and evaluate their suitability for worst-case consideration using the super-replication strategy as the baseline. By examining key metrics such as the initial portfolio value (f_0), the expected value of hedging error ($\mathbb{E}[\mathcal{H}]$), and the Conditional Value at Risk at the 99.5% confidence level ($\text{CVaR}_{99.5\%}$), we gain insights into the effectiveness and risk profiles of these strategies.

Since the super-replication strategy requires the hedging portfolio to dominate the claim or payoff throughout the investment horizon, we expect our hedging errors to be less than zero. That implies the simulated values of $\mathcal{H}$ are less than the estimate of initial portfolio value f_0. Figure 4.1 shows a histogram of hedging costs simulated with the super-replication technique in Algorithm 3.2. Here, the hedging error distribution is left-skewed with a $\text{CVaR}_{99.5\%}$ of 14.43% less than the initial portfolio value estimate of 15.97%. From the histogram, we can deduce that although the hedge portfolio dominates the payoffs, to some extent, the lower VaR and CVaR indicate a favourable tail or worst-case behaviour from an investor's

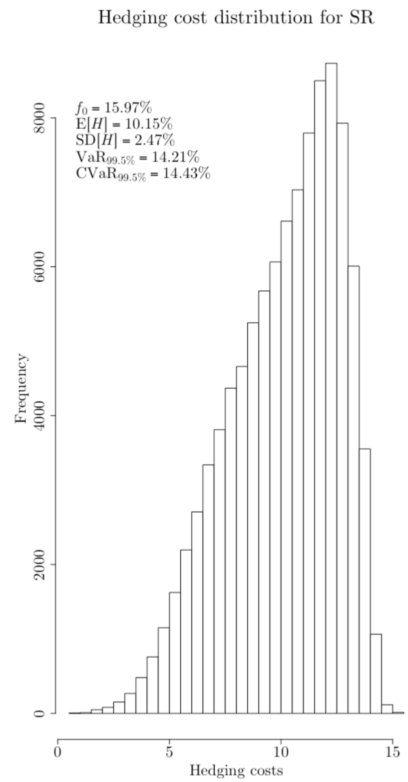

Figure 4.1: Distribution of simulated percent hedging costs for an at-the-money call option with super-replication strategy.

perspective.

In Figure 4.2, we show a histogram of hedging costs simulated with $\ell_{1,0}$, $\ell_{2,0}$ and $\ell_{\infty,1}$ norms as constraints. We obtain the hedge portfolio at each node by setting the sum of absolute errors for $\ell_{1,0}$ and $\ell_{2,0}$ norms and maximum error for $\ell_{\infty,1}$ norm to be less than $\gamma_0 = 1\%$ in Algorithm 3.5. The histograms are similar in shape and have the same expected values, variances and tail risks of hedging costs as the super-replication strategy. This illustrates the closeness of the strategies to our worst-case baseline. On the other hand, the initial portfolio values are less than that of super-replication. The strategy with $\ell_{1,0}$ norm is 0.25% less, $\ell_{2,0}$ norm is 0.29% less and $\ell_{\infty,1}$ norm is 0.41% less than the super-replication strategy. The fact that we can establish a portfolio that is cheaper than the super-replication method while having a similar hedging error distribution and tail risk by managing a small percentage of the losses indicates the robustness of the strategies. Therefore, to gain a lower cost of

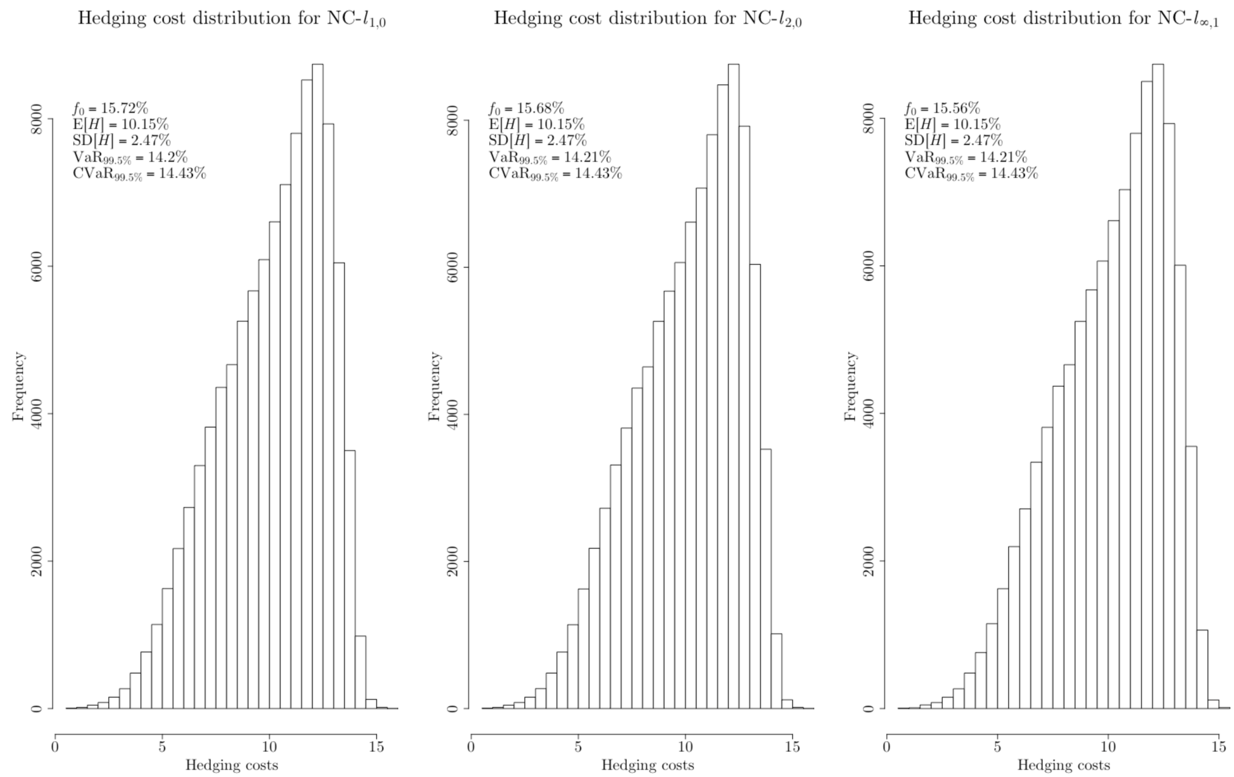

Figure 4.2: Distribution of simulated percent hedging costs for at-the-money call option with $\ell_{1,0}$ (left), $\ell_{2,0}$ (middle) and $\ell_{\infty,1}$ (right) norms as constraint.

initial portfolio value while being exposed to a similar tail risk, an investor who chooses to adopt the super-replicating portfolio may instead limit asymmetric losses by a threshold. For instance, having $\ell_{\infty,1}$ norm as a constraint yields the same value for $\text{CVaR}_{99.5\%}$ as that of super-replication but requires less initial amount to set-up than even $\ell_{1,0}$ and $\ell_{2,0}$ norms as a constraint.

Figure 4.3 is a histogram of hedging costs simulated with $\ell_{1,1}$, $\ell_{2,1}$ and $\ell_{\infty,1}$ norms as objective functions in Algorithm 3.6. Again, we define the strategy to set the sum of losses within the threshold $\gamma_0 = 1\%$. As a result, we obtain bell-shaped histograms with initial portfolio values f_0 less than the other strategies proposed. However, this is offset by a higher tail risk as observed by the $\text{CVaR}_{99.5\%}$. As such, by setting the asymmetric norm as an objective, an investor has to decide between a cheaper cost of establishing the portfolio with higher exposure to tail risk or having a fairly robust but expensive initial portfolio value. Also,

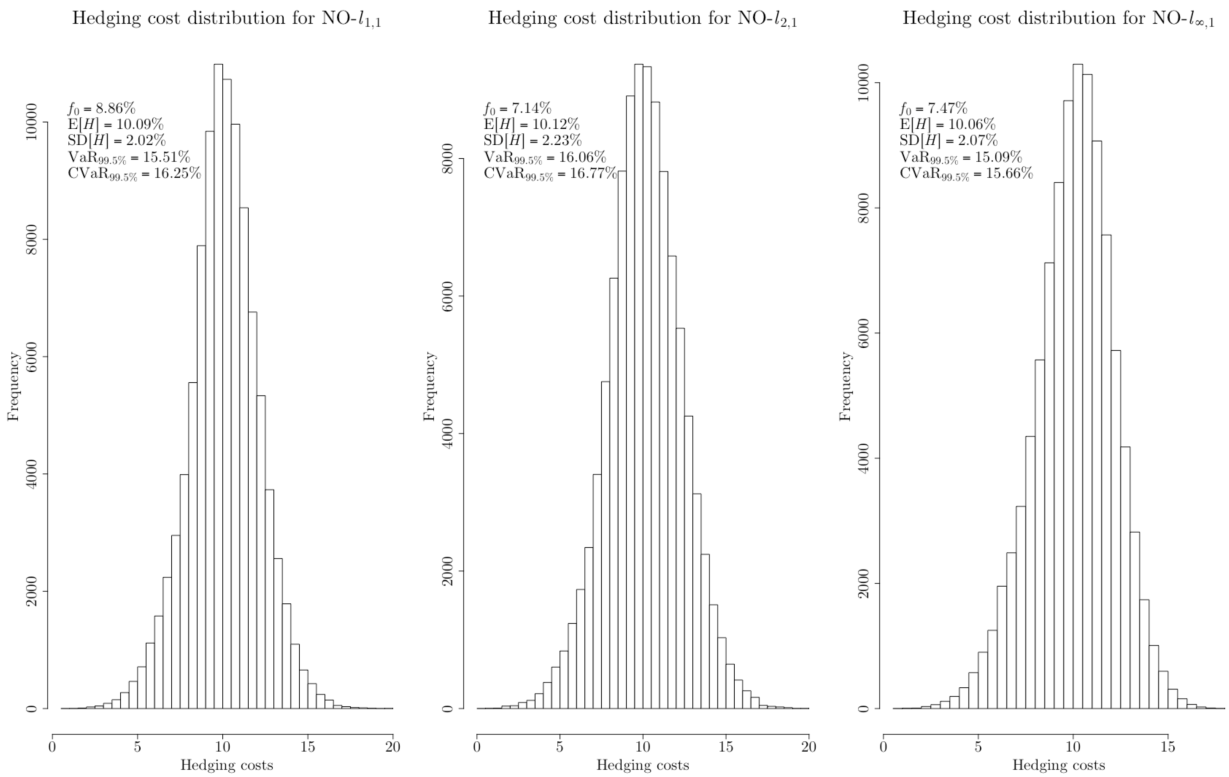

Figure 4.3: Distribution of simulated percent hedging costs for at-the-money call option with $\ell_{1,1}$ (left), $\ell_{2,1}$ (middle) and $\ell_{\infty,1}$ (right) norms as objective.

minimizing the maximum absolute loss at each node produces an initial portfolio value of 1.39% lower than minimizing the sum of absolute errors.

4.8.4 Sensitivity to number of trades and nodes for SR, NC and NO strategies

To assess the stability of the hedging strategies to dynamics in the model framework, we estimate the initial value of the portfolio f_0, $\mathbb{E}[\mathcal{H}]$ and the $\text{CVaR}_{99.5\%}$ for a varying number of trades or re-balancing and the number of nodes per branch.

Table 4.1 provides a detailed comparative analysis of the hedging strategies - super-replication (SR), norm as constraint (NC), and norm as objective (NO) - across different number of trades Δ, and the number of nodes N. The super-replication strategy serves as a baseline, offering a conservative approach to hedging by ensuring that the portfolio value always ex-

		$\Delta = 6$			$\Delta = 12$			$\Delta = 24$		
		$N = 10$	$N = 20$	$N = 30$	$N = 10$	$N = 20$	$N = 30$	$N = 10$	$N = 20$	$N = 30$
SR	f_0	13.58%	15.05%	15.84%	14.04%	15.97%	17.15%	14.24%	16.37%	17.75%
	$\mathbb{E}[\mathcal{H}]$	10.15%	10.15%	10.17%	10.15%	10.15%	10.17%	10.15%	10.15%	10.17%
	$\mathrm{CVaR}_{99.5\%}$	13.58%	15.02%	15.80%	13.23%	14.43%	15.36%	13.11%	14.37%	15.17%
NC-$\ell_{1,0}$	f_0	13.18%	14.66%	15.49%	13.79%	15.72%	16.91%	14.11%	16.23%	17.62%
	$\mathbb{E}[\mathcal{H}]$	10.16%	10.15%	10.16%	10.15%	10.15%	10.16%	10.17%	10.13%	10.17%
	$\mathrm{CVaR}_{99.5\%}$	13.77%	15.12%	15.85%	13.21%	14.43%	15.37%	13.09%	14.35%	15.16%
NC-$\ell_{2,0}$	f_0	13.09%	14.56%	15.36%	13.77%	15.68%	16.87%	14.24%	16.37%	17.75%
	$\mathbb{E}[\mathcal{H}]$	10.16%	10.15%	10.16%	10.15%	10.15%	10.17%	10.17%	10.13%	10.17%
	$\mathrm{CVaR}_{99.5\%}$	13.73%	15.11%	15.86%	13.22%	14.43%	15.37%	13.11%	14.37%	15.17%
NC-$\ell_{\infty,1}$	f_0	12.77%	14.24%	15.03%	13.63%	15.56%	16.74%	14.03%	16.16%	17.55%
	$\mathbb{E}[\mathcal{H}]$	10.16%	10.15%	10.17%	10.15%	10.15%	10.17%	10.17%	10.13%	10.17%
	$\mathrm{CVaR}_{99.5\%}$	13.58%	15.02%	15.80%	13.23%	14.43%	15.36%	13.11%	14.37%	15.17%
NO-$\ell_{1,1}$	f_0	9.97%	9.64%	9.40%	9.70%	8.86%	8.19%	9.60%	7.85%	7.00%
	$\mathbb{E}[\mathcal{H}]$	10.14%	10.08%	10.07%	10.16%	10.09%	10.05%	10.15%	10.10%	10.07%
	$\mathrm{CVaR}_{99.5\%}$	15.51%	16.71%	17.37%	14.35%	16.25%	17.32%	13.23%	15.77%	16.99%
NO-$\ell_{2,1}$	f_0	9.58%	7.76%	5.95%	9.52%	7.14%	4.49%	9.49%	6.63%	3.00%
	$\mathbb{E}[\mathcal{H}]$	10.15%	10.10%	10.08%	10.16%	10.12%	10.11%	10.17%	10.13%	10.14%
	$\mathrm{CVaR}_{99.5\%}$	15.53%	17.22%	18.02%	14.49%	16.77%	18.23%	13.54%	16.00%	18.16%
NO-$\ell_{\infty,1}$	f_0	9.62%	7.62%	5.65%	9.85%	7.47%	4.77%	10.07%	7.65%	4.17%
	$\mathbb{E}[\mathcal{H}]$	10.13%	10.06%	10.03%	10.13%	10.06%	10.03%	10.14%	10.07%	10.02%
	$\mathrm{CVaR}_{99.5\%}$	15.07%	16.40%	17.06%	13.87%	15.66%	16.76%	12.82%	14.39%	15.66%

Table 4.1: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\mathrm{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SR, NC-$(\ell_{1,0}, \ell_{2,0}, \ell_{\infty,1})$ and NO-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with $\Delta = 6, 12, 24$ number of trades and $N + 1 = 11, 21, 31$ number of nodes.

ceeds the option's payoff. Table 4.1 shows that as the number of nodes and the trading frequency increase, there is a corresponding increase in the initial portfolio value. Specifically, f_0 rises slightly with more nodes and more frequent trading, suggesting that larger tree size and finer discretization introduce additional complexity to the hedging process. The expected hedging error remains remarkably stable across all values of Δ and N. This constancy indicates that the average performance of the super-replication strategy does not improve with more frequent trading or a higher number of nodes. However, the $\text{CVaR}_{99.5\%}$, which measures the risk of extreme losses, also increases with higher N but decreases with higher Δ. This decrease suggests that while the super-replication strategy maintains a stable average error, it also reduces exposure to higher tail risks as trading frequency increase.

The norm as constraint strategies, categorized based on different norms - $\ell_{1,0}, \ell_{2,0}$ and $\ell_{\infty,1}$ - show performance similar to those of super-replication but with subtle variations. For instance, the initial portfolio value in the norm as constraint strategies is comparable to super-replication but tends slightly lower, particularly for $\ell_{\infty,1}$. This indicates that norm-based constraints might be more effective in ensuring a cheaper initial portfolio value. The expected value of hedging error for the norm as constraint strategies remains consistent and like super-replication, highlighting that these strategies do not alter the average performance of the hedge portfolio. These findings imply that while norm as constraint strategies might not drastically change the expected outcome, they could offer a slight edge in ensuring a cheaper initial portfolio while having the same exposure to tail risk, making them a potential alternative to super-replication for risk-averse hedgers.

The norm as objective strategies, which include $\ell_{1,1}, \ell_{2,1}$ and $\ell_{\infty,1}$ present a distinct approach by focusing on minimizing the risk with specific norms as objectives. One of the most notable findings is that the initial portfolio value is lower than for both super-replication and norm as constraint strategies, especially as Δ and N increase. This suggests that norms as objective strategies are particularly effective in minimizing the initial portfolio value, which could be critical in scenarios where the cost of establishing the hedge portfolio is paramount. However, this reduction in initial portfolio value comes with a trade-off. The $\text{CVaR}_{99.5\%}$ is generally higher, indicating that while these strategies reduce the set-up cost and expected costs, they also introduce higher tail risks when exposed to extreme market movements. Therefore, while norm as objective strategies offer clear benefits in reducing the cost of setting up the portfolio, they may not be suitable for all hedgers, particularly those with a low tolerance for extreme risks.

Across all strategies, we see that a hedger must balance the benefits of increased trading frequency with the potential for higher cost of capital and reduce tail risk. The super-replication strategy offers a conservative approach with expensive initial portfolio value and lower tail risks, making it a reliable choice for risk-averse hedgers. Norm as constraint strategies slightly improve tail risk management while maintaining similar expected costs, making them a viable alternative for hedgers seeking a more robust approach. Norm as objective strategies, on the other hand, excel in minimizing the initial portfolio value but at the cost of higher tail risks, making them suitable for hedgers who prioritize early cost and can tolerate increased exposure to tail risk.

4.8.5 Sensitivity to market parameters for SR, NC and NO strategies

Next, we analyze our hedging strategies' sensitivity to financial model parameters. The drift parameter μ calculates the probabilities when simulating hedging errors and not used in estimating the hedging portfolio. Hence, we only test the risk-free interest rate r and the market volatility σ.

While the conservative approach of super-replication offers a high degree of protection, Table 4.2 shows considerable sensitivity to interest rate and market volatility changes. As interest rates increase from 2% to 8%, the initial cost required to set up the super-replicating portfolio rises . For instance, at a volatility of 15%, the initial cost increases from 11.59% to 14.56%. This trend highlights the direct relationship between the cost of capital and the initial outlay required for a super-replicating portfolio. Similarly, an increase in market volatility exacerbates this effect. At a fixed interest rate of 2%, the initial cost escalates from 11.59% at $\sigma = 15\%$ to 21.94% at $\sigma = 30\%$. This indicates that in more volatile markets, a larger initial investment is necessary to maintain the conservative posture of the super-replication strategy. The expected hedging error under super-replication increases with higher interest rates and volatility. Although the super-replication aims to minimize risk, these results suggest

		$r=2\%$			$r=4\%$			$r=8\%$		
		$\sigma=15\%$	$\sigma=25\%$	$\sigma=30\%$	$\sigma=15\%$	$\sigma=25\%$	$\sigma=30\%$	$\sigma=15\%$	$\sigma=25\%$	$\sigma=30\%$
	f_0	11.59%	18.52%	21.94%	12.54%	19.37%	22.75%	14.56%	21.12%	24.40%
SR	$\mathbb{E}[\mathcal{H}]$	7.23%	11.09%	13.01%	8.25%	12.06%	13.97%	10.40%	14.06%	15.92%
	$\mathrm{CVaR}_{99.5\%}$	10.43%	16.64%	19.74%	11.37%	17.49%	20.54%	13.41%	19.24%	22.19%
	f_0	11.26%	18.32%	21.77%	12.22%	19.17%	22.58%	14.25%	20.92%	24.23%
NC-$\ell_{1,0}$	$\mathbb{E}[\mathcal{H}]$	7.23%	11.09%	13.01%	8.24%	12.06%	13.97%	10.40%	14.06%	15.92%
	$\mathrm{CVaR}_{99.5\%}$	10.43%	16.64%	19.73%	11.38%	17.49%	20.53%	13.42%	19.25%	22.19%
	f_0	11.21%	18.29%	21.75%	12.17%	19.15%	22.56%	14.20%	20.9%	24.21%
NC-$\ell_{2,0}$	$\mathbb{E}[\mathcal{H}]$	7.23%	11.09%	13.01%	8.24%	12.06%	13.97%	10.40%	14.06%	15.92%
	$\mathrm{CVaR}_{99.5\%}$	10.43%	16.64%	19.73%	11.38%	17.49%	20.54%	13.42%	19.25%	22.2%
	f_0	11.04%	18.19%	21.67%	12.00%	19.05%	22.48%	14.03%	20.80%	24.13%
NC-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	7.23%	11.09%	13.01%	8.25%	12.06%	13.97%	10.40%	14.06%	15.92%
	$\mathrm{CVaR}_{99.5\%}$	10.43%	16.64%	19.74%	11.37%	17.49%	20.54%	13.41%	19.24%	22.19%
	f_0	6.13%	9.34%	11.01%	7.23%	10.43%	11.96%	9.82%	12.56%	14.14%
NO-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	7.10%	11.08%	13.06%	8.15%	12.04%	13.99%	10.47%	14.07%	15.93%
	$\mathrm{CVaR}_{99.5\%}$	11.86%	19.10%	22.69%	12.77%	19.85%	23.43%	14.80%	21.50%	24.83%
	f_0	4.81%	7.31%	8.57%	5.94%	8.37%	9.61%	8.41%	10.61%	11.79%
NO-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	7.11%	11.11%	13.09%	8.16%	12.08%	14.03%	10.48%	14.12%	15.99%
	$\mathrm{CVaR}_{99.5\%}$	12.20%	19.61%	23.27%	13.15%	20.40%	24.00%	15.22%	22.09%	25.55%
	f_0	5.17%	7.67%	8.94%	6.25%	8.70%	9.95%	8.60%	10.85%	12.04%
NO-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	7.10%	11.07%	13.06%	8.12%	12.02%	13.99%	10.36%	14.01%	15.90%
	$\mathrm{CVaR}_{99.5\%}$	11.32%	18.24%	21.69%	12.28%	19.06%	22.45%	14.37%	20.79%	24.05%

Table 4.2: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\mathrm{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SR, NC-$(\ell_{1,0}, \ell_{2,0}, \ell_{\infty,1})$ and NO-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with interest rate $r = 2\%, 4\%, 8\%$, and volatility $\sigma = 15\%, 25\%, 30\%$.

that average costs in hedging are influenced by market conditions, reflecting the challenges of maintaining a perfect hedge in dynamic environments. Moreover, the tail risk, as measured by $\mathrm{CVaR}_{99.5\%}$, rises from 10.43% to 22.19% as both interest rate and volatility increase. This demonstrates that the portfolio may still be exposed to significant losses in extreme market scenarios, especially when both interest rates and volatility are high.

The norm as constraint strategies, which impose constraints based on different norms, show performance closely aligned with the super-replication strategy but with some notable differences. The initial portfolio value for the norm as constraint strategies is generally slightly lower than for super-replication, particularly for the $\ell_{\infty,1}$ variant. For instance, at 2% interest rate and 15% volatility, the initial cost for $\ell_{\infty,1}$ is 11.04%, compared to 11.59% for super-replication. This suggests that norms as constraint strategies might offer marginal cost savings, potentially making them more appealing in certain contexts. The expected hedging costs are also nearly identical to those of super-replication, indicating that these strategies maintain comparable average performance despite introducing norm constraints. This consistency implies that norm as constraint strategies can be employed without compromising the effectiveness of the hedge.

The norm as objective strategies present a distinct performance profile, particularly regarding the initial cost. One of the most striking features of norm as objective strategies is their consistently lower initial costs than super-replication and norm as constraint strategies. For instance, at 2% interest rate and 15% volatility, the initial portfolio value is just 4.81%, lower than the 11.59% required for super-replication. This cost efficiency makes norm as objective strategies particularly attractive in markets with lower volatility or when managing portfolios with tighter budget constraints. Despite these lower initial costs, the expected hedging costs are comparable to super-replication and norm as constraints. This indicates that norms as objective strategies can achieve similar average performance while requiring less capital to set up. However, the lower initial cost comes with a trade-off: higher tail risks. For instance, an 8% interest rate and 30% volatility have a $\text{CVaR}_{99.5\%}$ of 25.55%, higher than the 22.19% for super-replication. This suggests that while norm as constraint strategies are more cost-efficient, they may expose the portfolio to greater risk of extreme losses, particularly in volatile market conditions.

The analysis of these hedging strategies underscores the significant impact of interest rates and market volatility on their performance. As the interest rate increases, there is a consistent rise in the initial cost, expected hedging error, and $\text{CVaR}_{99.5\%}$. This suggests that higher interest rates increase the cost of hedging and the associated risks. Similarly, increasing volatility leads to higher values for all three metrics. This is expected, as greater volatility typically results in more significant price swings, making effective hedging more challenging. The effect is particularly pronounced in the $\mathrm{CVaR}_{99.5\%}$, where extreme market movements elevate tail risks.

4.8.6 Sensitivity to hedger's preference for SR, NC and NO strategies

Another way to assess the robustness of super-replication, norm as constraint and norm as objective strategies is to analyze their performance to changing model parameters based on the hedger's preference. These hyper-parameters are classified as hedger's preference since they are not dependent on market dynamics and can be tuned individually or collectively to obtain the desired level of performance and robustness. Specifically, we analyze the asymmetric parameter of the norms, the threshold and the inclusion of additional assets in the hedging portfolio.

Figure 4.4 shows nine histograms of percent hedging error distributions for different norms as objective strategies. Here, we compare hedging error distributions across different norms and asymmetric parameter values. For lower values of q, we observe an increase in the skewness of hedging error distributions similar to that of super-replication and norm as constraint strategies in Figures 4.1 and 4.2 respectively. On the other hand, for higher values of q, we observe bell-shaped histograms similar to the norm as objective strategies in Figure 4.3. Generally, strategies with higher initial costs, such as $\ell_{p,0}$, tend to exhibit lower tail risks, as evidenced by a relatively lower $\mathrm{CVaR}_{99.5\%}$. For instance, the $\ell_{\infty,0}$ strategy, which has the highest initial cost of 14.13%, demonstrates the least variability in hedging costs among $\ell_{p,0}$ strategies, with a standard deviation of 2.16%, and a moderate $\mathrm{CVaR}_{99.5\%}$ of 15.16%. This suggests that strategies with higher initial costs may offer greater protection against significant potential losses for hedgers, prioritizing the minimization of extreme losses. Conversely, strategies with lower initial costs, such as the $\ell_{2,1}$ strategy, present a different risk

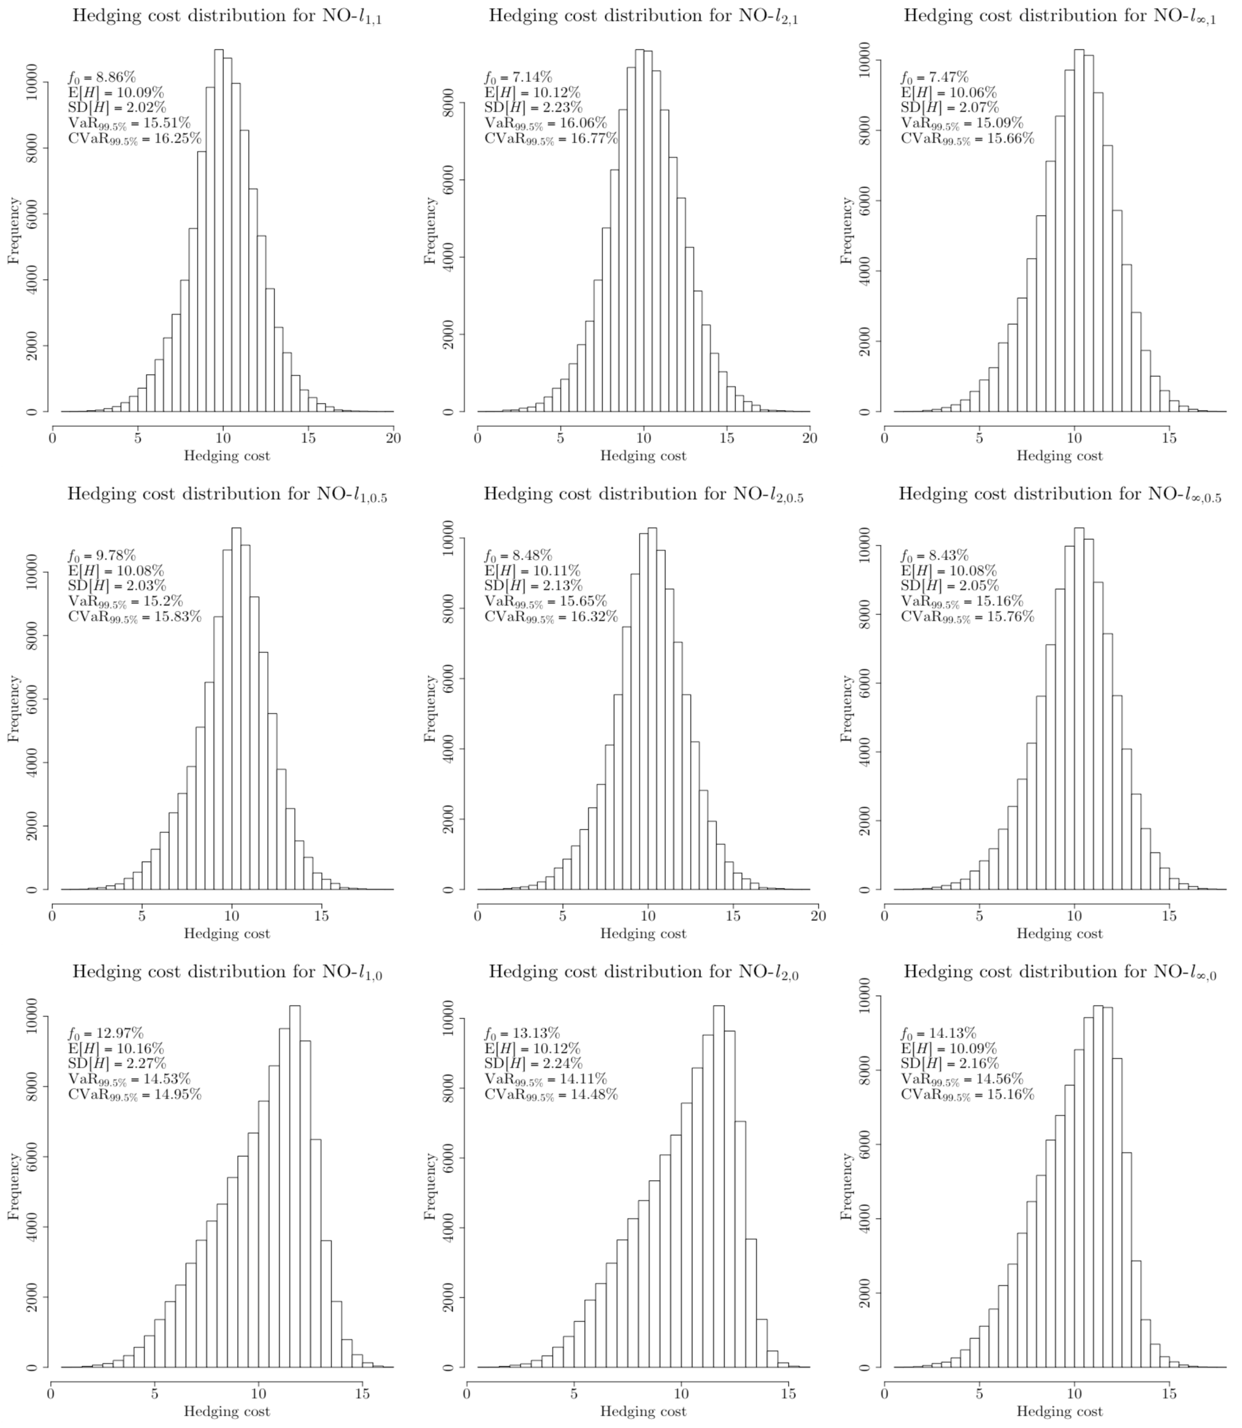

Figure 4.4: Distribution of simulated percent hedging costs for at-the-money call option with $\ell_{1,q}$ (left), $\ell_{2,q}$ (middle) and $\ell_{\infty,q}$ (right) norms as objective for $q = 1, 0.5, 0$

profile. With the lowest initial cost of 7.14%, this strategy also exhibits the highest tail risk ($\text{CVaR}_{99.5\%}$ = 16.77%) and a higher standard deviation (SD = 2.23%). While it may be more cost-effective to implement initially, the $\ell_{2,1}$ strategy exposes the portfolio to greater potential losses, making it more suitable for investors with a higher risk appetite.

Tail risks are critical in understanding the exposure to extreme outcomes. In Figure 4.4, strategies such as $\ell_{\infty,1}$ and $\ell_{\infty,0.5}$ show relatively low tail risks, with $\mathrm{CVaR}_{99.5\%}$ values of 15.66% and 15.76%, respectively. These strategies, which balance initial cost and risk, may appeal to hedgers seeking a middle ground between cost efficiency and risk control. The standard deviation of hedging costs provides further insights into the consistency of a strategy's performance. Strategies with lower asymmetric parameter values q, tend to have slightly higher standard deviations, indicating high variability in hedging costs and, consequently, less predictable outcomes. For example, the $\ell_{1,0}$ strategy, with a standard deviation of 2.27%, offers the highest variability. However, the corresponding low tail risk makes it a strong candidate for investors seeking a balance between risk and reward.

The choice of norm asymmetry ultimately depends on the hedger's objectives and risk tolerance. Strategies with higher asymmetric parameters offer higher tail risks in exchange for lower initial costs. On the other hand, strategies with lower asymmetric parameters offer higher initial costs while minimizing extreme losses and yielding consistent or robust performance. By carefully considering the trade-offs between cost, tail risk, and hedging error variability, hedgers can tune the asymmetric parameter to align with their goals and risk appetite.

In Table 4.3, the conservative super-replication strategy does not depend on the risk tolerance threshold γ_0. As such, the initial cost of the strategy remains consistent at 15.97% across all thresholds. This invariance suggests that the super-replication strategy maintains a fixed cost to achieve its robust hedging objective, irrespective of the risk tolerance level the investor sets. However, when additional assets are incorporated into the hedging portfolio, the initial cost of the strategy decreases . For instance, adding five options with strike prices defined by the moneyness parameter $\kappa = 0.7, 0.9, 1, 1.1, 1.3$, reduces the initial cost to 11.47%. This reduction highlights the impact of diversification, where expanding the asset base in

		$\gamma_0 = 1\%$	$\gamma_0 = 3\%$	$\gamma_0 = 5\%$	$\gamma_0 = \infty$	1 option	3 options	5 options
	f_0	15.97%	15.97%	15.97%	15.97%	15.97%	14.64%	11.47%
SR	$\mathbb{E}[\mathcal{H}]$	10.15%	10.15%	10.15%	10.15%	10.15%	10.05%	10.01%
	$\text{CVaR}_{99.5\%}$	14.43%	14.43%	14.43%	14.43%	14.43%	13.32%	11.28%
	f_0	15.72%	15.27%	14.86%	$-\infty\%$	15.72%	14.36%	11.28%
NC-$\ell_{1,0}$	$\mathbb{E}[\mathcal{H}]$	10.15%	10.14%	10.14%	10.78%	10.15%	10.05%	10.01%
	$\text{CVaR}_{99.5\%}$	14.43%	14.45%	14.51%	$\infty\%$	14.43%	13.29%	11.24%
	f_0	15.68%	15.15%	14.64%	$-\infty\%$	15.68%	14.34%	11.23%
NC-$\ell_{2,0}$	$\mathbb{E}[\mathcal{H}]$	10.15%	10.15%	10.14%	10.84%	10.15%	10.05%	10.01%
	$\text{CVaR}_{99.5\%}$	14.43%	14.45%	14.48%	$\infty\%$	14.43%	13.3%	11.25%
	f_0	15.56%	14.75%	13.95%	$-\infty\%$	15.56%	14.23%	11.06%
NC-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.15%	10.15%	10.15%	10.15%	10.15%	10.05%	10.01%
	$\text{CVaR}_{99.5\%}$	14.43%	14.43%	14.43%	14.43%	14.43%	13.32%	11.28%
	f_0	8.86%	8.54%	8.49%	8.50%	8.86%	9.27%	9.73%
NO-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.09%	10.10%	10.10%	10.10%	10.09%	10.03%	9.99%
	$\text{CVaR}_{99.5\%}$	16.25%	16.38%	16.38%	16.44%	16.25%	14.9%	11.68%
	f_0	7.14%	7.14%	7.14%	7.15%	7.14%	9.04%	9.75%
NO-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	10.12%	10.12%	10.12%	10.12%	10.12%	9.97%	10.00%
	$\text{CVaR}_{99.5\%}$	16.77%	16.77%	16.77%	16.76%	16.77%	14.80%	11.41%
	f_0	7.47%	8.03%	8.13%	7.81%	7.47%	9.69%	10.62%
NO-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.06%	10.07%	10.08%	10.07%	10.06%	9.92%	10.06%
	$\text{CVaR}_{99.5\%}$	15.66%	15.43%	15.36%	15.39%	15.66%	14.17%	12.31%

Table 4.3: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SR, NC-$(\ell_{1,0}, \ell_{2,0}, \ell_{\infty,1})$ and NO-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with constant threshold $\gamma_0 = 1\%, 3\%, 5\%$, and 1 call option with $\kappa = 1$, 3 call options with $\kappa = 0.8, 1, 1.2$, and 5 call options with $\kappa = 0.7, 0.9, 1, 1.1, 1.3$ as assets to the hedging portfolio.

the portfolio lowers the required capital for effective hedging. Regarding expected hedging error, the super-replication strategy shows stability, with a consistent value of 10.15% across all thresholds, indicating that the average performance is not influenced by γ_0. Moreover, including additional options slightly improves the expected error, reducing it to 10.01% with five options, suggesting that diversification can enhance the accuracy of the hedge. The $\text{CVaR}_{99.5\%}$ decreases with the addition of more assets in the portfolio. For example, the $\text{CVaR}_{99.5\%}$ reduces from 14.43% to 11.28% with five additional options, demonstrating the effectiveness of diversification in mitigating extreme risks within the super-replication strategy.

Norm as constraint strategies introduce a risk management framework that sets thresholds to constrain risk exposure. These strategies show a more dynamic relationship between the initial cost and the threshold γ_0. As the threshold increases, the initial portfolio value for the norm as constraint strategies generally decrease. For example, in the $\ell_{\infty,1}$ strategy, the initial cost drops from 15.56% at $\gamma_0 = 1\%$ to 13.95% at $\gamma_0 = 5\%$. This trend reflects the trade-off between risk tolerance and the cost of setting up the portfolio - higher risk tolerance allows for a lower initial cost. However, when the threshold is set to infinity, the initial cost for the norm as constraint strategies falls to negative infinity, indicating that the strategies become infeasible or excessively risky without a proper constraint. This outcome underscores the importance of carefully selecting an appropriate threshold to maintain the viability of norms as constraint strategies. Like the super-replication strategy, adding more options to the norm as constraint portfolios reduces the initial cost. For instance, in the $\ell_{1,0}$ strategy, the initial cost decreases from 15.72% to 11.28% when five options are added. This reduction demonstrates the positive impact of diversification, which not only lowers the cost but also enhances the overall performance of the hedge. Regarding expected hedging error, the norm as constraint strategies exhibit high stability, with values hovering around 10.15% across different thresholds. This consistency suggests that norm as constraint strategies deliver reliable hedging performance, regardless of the specific conditions or risk constraints imposed. The $\text{CVaR}_{99.5\%}$ shows slight variations with changes in the threshold, generally increasing as the threshold rises. In some cases, such as with $\gamma_0 = \infty$, the $\text{CVaR}_{99.5\%}$ can escalate to infinity, indicating the potential for extreme losses if no constraint is applied. However, like the super-replication strategy, adding more assets to the norm as constraint portfolios reduces the $\text{CVaR}_{99.5\%}$, demonstrating that diversification is crucial in controlling tail risk.

Table 4.3 also shows the norm as objective strategies that prioritize minimizing the portfolio setup's cost while achieving the desired hedging outcome. Among the strategies analyzed, norm-as-objective strategies exhibit much lower initial costs than super-replication and norm-as-constraint strategies, highlighting the cost-efficiency of this approach. Interestingly, the initial cost for norm as objective strategies slightly increases when additional options are added to the portfolio, which contrasts the super-replication and norm as constraint strategies. For instance, in the $\ell_{\infty,1}$ strategy, the initial cost rises from 7.47% to 10.62% when five options are added. This increase might reflect the increased complexity in the hedging strategy as more assets are included, suggesting that diversification can reduce risk and may also introduce additional costs in certain scenarios. The expected hedging error remains relatively stable, around 10.10%, with minor variations across thresholds and additional assets. This performance is comparable to the super-replication and norm as constraint strategies, indicating that norm as objective strategies can achieve similar average hedging outcomes despite their lower initial costs. However, regarding tail risk, the norm as objective strategies tend to exhibit higher values, which suggests that they expose the portfolio to greater extreme risks. Despite this, adding more assets to the norm as objective strategies help to reduce the $\text{CVaR}_{99.5\%}$, with the $\ell_{\infty,1}$ strategy experiencing a decrease from 15.66% to 12.31% when five options are added. This reduction reinforces the benefit of diversification, even within cost-sensitive strategies like norm as objective, in managing tail risk.

The comparative analysis of the three hedging strategies reveals distinct characteristics and trade-offs. While conservative and consistent, the super-replication strategy requires a relatively high initial cost, which can be reduced through diversification. Norm as constraint strategies offers a balanced approach, requiring careful management of thresholds to avoid excessive risks and benefit from diversification in terms of cost and tail risk reduction. Norm as objective strategies stand out for their cost efficiency, making them suitable for scenarios where minimizing initial costs is a priority, though they come with higher exposure to extreme risks. Ultimately, the choice of hedging strategy should be guided by the specific goals and robustness of the portfolio. For conservative risk management, super-replication and norm-as-constraint strategies may be preferable, while norm-as-objective strategies are better suited for capital-efficient scenarios. Across all strategies, diversification consistently emerges as a critical factor in enhancing hedging performance, particularly in reducing tail risks and managing extreme market conditions, albeit in the absence of transaction costs.

4.8.7 Analysis of hedging error distribution for portfolio as state variable strategies

We continue our analysis by comparing the hedging error distribution of portfolios as state variable strategies. To assess the models' adaptability for worst-case scenarios, we proceed with our analysis by looking at their expected hedging costs and tail risk. We compare the portfolio as a state variable strategy to the other hedging techniques for various model parameters and norms for a fixed initial portfolio value, $f_0 = 0.1$. The hedging portfolio is established here for a specific initial cost or fund, from which the hedger can cover the future payoff through the accumulated wealth process. Also, in order the reduce computational complexity, we set the number of trades, $\Delta = 6$ for this set of simulations.

Figure 4.5 analyzes the performance of three stochastic programming strategies for hedging an at-the-money call option, each employing different norms: $\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1}$. All three stochastic programming strategies analyzed have a predefined initial cost of 0.1, ensuring a level playing field in comparing their risk-return profiles. The expected hedging costs are relatively close across the strategies, ranging from 10.00% to 10.19%, suggesting that each strategy maintains a similar level of performance in terms of minimizing the average error. However, the slight differences in expected hedging costs indicate that the norm's choice can

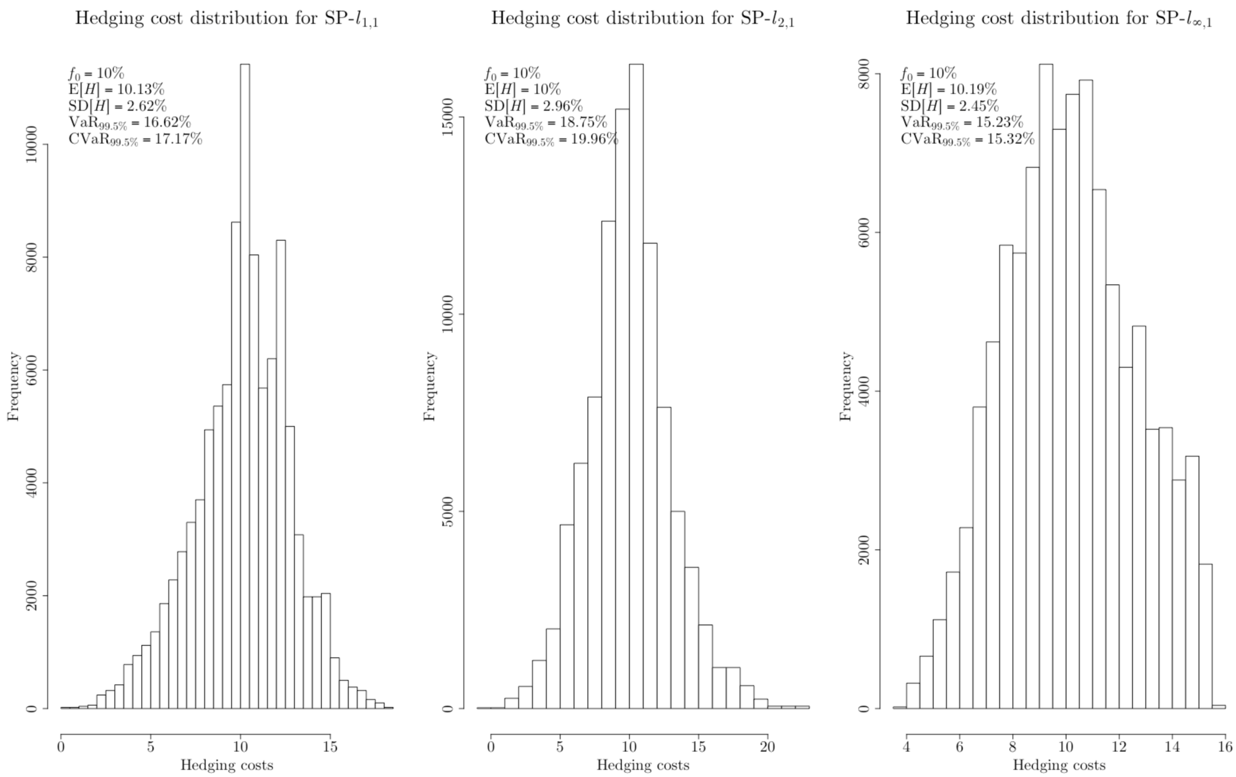

Figure 4.5: Distribution of simulated percent hedging costs for at-the-money call option with stochastic programming strategy.

influence the hedging strategy's efficiency. The $\ell_{2,1}$ norm strategy has the lowest expected hedging error of 10.00% but the highest standard deviation of 2.96%. This implies that the strategy may yield slightly lower average costs but with a low degree of consistency, making it a good choice for hedgers seeking cheaper initial hedging costs. The $\ell_{\infty,1}$ norm strategy stands out with the lowest tail risk ($\text{VaR}_{99.5\%} = 15.23\%$, $\text{CVaR}_{99.5\%} = 15.32\%$). This illustrates the strategy's appeal for risk-averse investors, offering tighter control over extreme losses. In contrast, the $\ell_{1,1}$ and $\ell_{2,1}$ norm strategies exhibit higher tail risks, with $\text{CVaR}_{99.5\%}$ values of 17.17% and 19.96%, respectively. While offering similar expected hedging costs, these strategies expose hedgers to greater potential losses in extreme market conditions. This increased risk is accompanied by a higher standard deviation for the $\ell_{2,1}$ norm strategy, indicating more variability in hedging performance.

Similar to stochastic programming, Figure 4.6 shows the histogram of simulated hedging

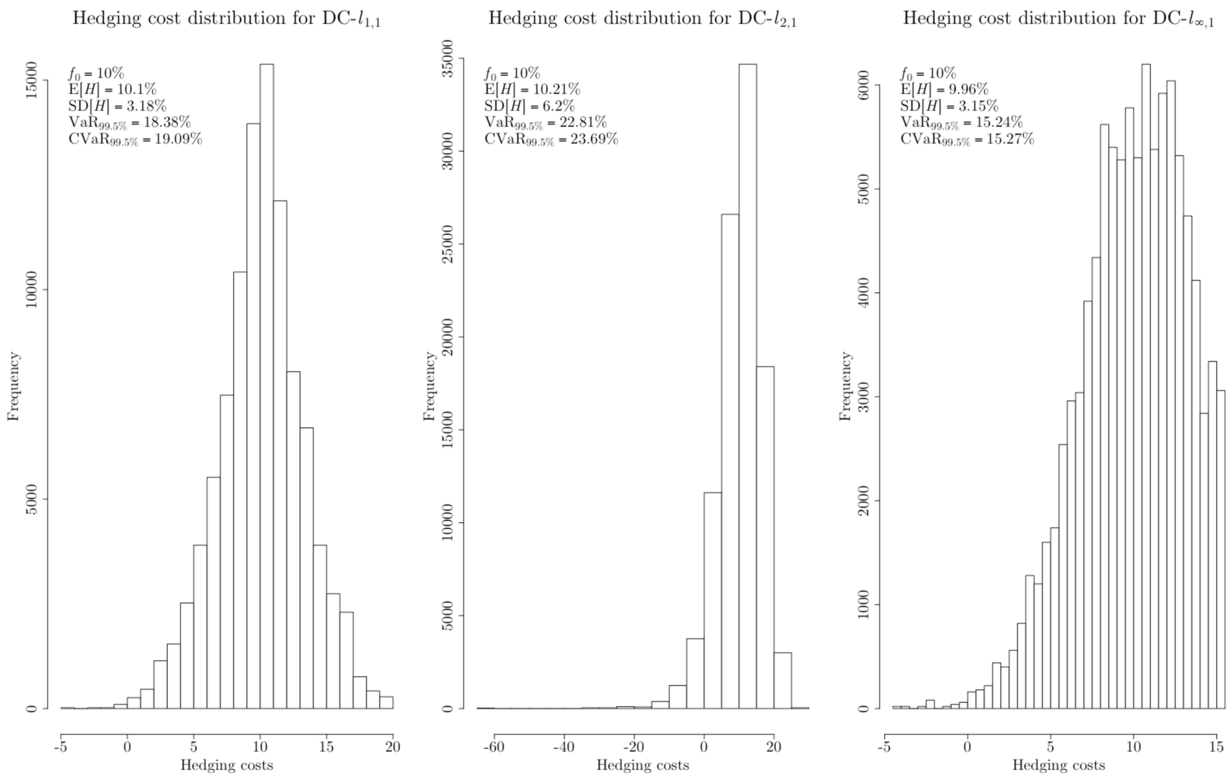

Figure 4.6: Distribution of simulated percent hedging costs for an at-the-money call option with dynamic coherent risk strategy.

costs of three dynamic coherent risk strategies for $\ell_{1,1}, \ell_{2,1}$ and $\ell_{\infty,1}$ norms. Again, all three strategies share a predefined initial cost of 0.1, ensuring comparability in their performance. The expected hedging costs across the strategies vary more and with a higher standard deviation than the stochastic programming strategies. This indicates that the choice of norm is crucial in influencing hedging efficiency for dynamic coherent risk strategies. The $\ell_{\infty,1}$ norm strategy has the lowest standard deviation among the three. Coupled with a relatively low expected error, this strategy offers a balanced performance, providing consistency and efficiency in minimizing average hedging costs. On the other hand, $\ell_{2,1}$ norm strategy shows a higher standard deviation. The $\ell_{\infty,1}$ norm strategy also maintains the lowest tail risk with $\mathrm{VaR}_{99.5\%}$ of 15.24% and $\mathrm{CVaR}_{99.5\%}$ of 15.27%. This aligns with its relatively lower standard deviation, reinforcing its position as a conservative choice for risk-averse investors. It offers a robust solution for minimizing risk without exposing the portfolio to extreme losses. The

$\ell_{1,1}$ norm strategy, while showing a lower expected hedging error than $\ell_{2,1}$ at 10.10%, still carries considerable tail risks. A standard deviation of 3.18% for the $\ell_{1,1}$ norm strategy reflects its middle ground performance in terms of variability, suggesting that it may serve as a compromise between the low-risk $\ell_{\infty,1}$ norm strategy and high-risk $\ell_{2,1}$ norm strategy.

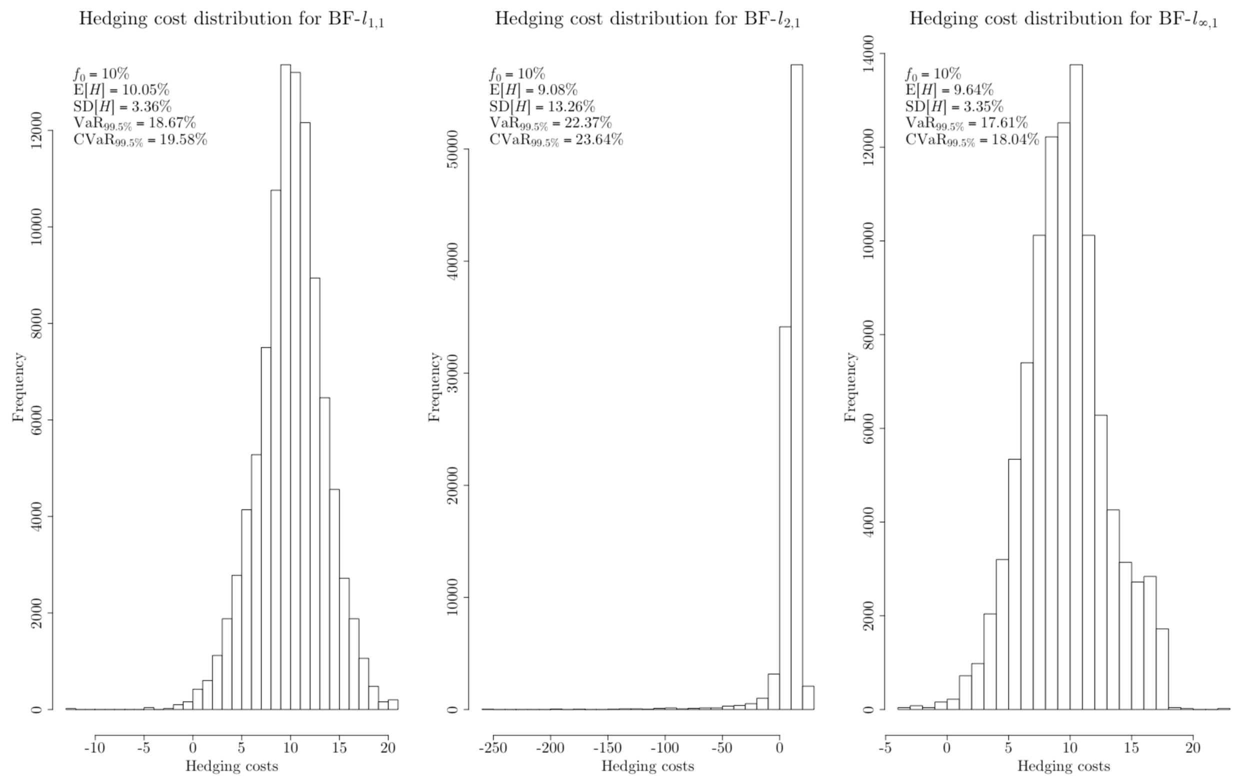

Figure 4.7: Distribution of simulated percent hedging costs for an at-the-money call option with barrier on future risk strategy.

The barrier on future risk strategies limits risk exposure by setting predefined barriers or thresholds. In Figure 4.7, we show the histogram of simulated hedging costs for $\ell_{1,1}, \ell_{2,1}$ and $\ell_{\infty,1}$ norms and analyze their performance. The choice of threshold is crucial to the performance of the barrier of future risk strategies. Without loss of generality, we set a dynamic threshold $\gamma_{i_t} = (1 + \gamma_0/i_t)^{i_t} - 1$ for a predefined initial cost of 0.1 to compare the strategies. The expected hedging costs across the strategies show some variation, ranging from 9.08% to 10.05%, reflecting the differing impact of each norm on hedging efficiency. In this case, the $\ell_{2,1}$ norm strategy has the lowest expected hedging error but a higher standard deviation, suggesting that while it may reduce the average error, it does so at the cost of very high variability.

On the other hand, the $\ell_{1,1}$ norm strategy shows a slightly higher expected hedging error but a much lower standard deviation compared to the $\ell_{2,1}$ norm. This implies that while it is slightly less efficient in minimizing average costs, it offers more consistent performance with reduced variability in hedging outcomes. Also, despite its lower average hedging error, the $\ell_{2,1}$ norm strategy exhibits higher tail risks. This shows that the strategy exposes the portfolio to higher potential losses in the tail. The $\ell_{\infty,1}$ norm strategy emerges as the most suitable barrier on future risk strategy for risk-averse hedgers, offering consistent performance with the lowest tail risks. The strategy $\ell_{1,1}$ norm offers a middle ground, balancing tail risk management at the expense of higher variability. We highlight the complexity of barriers to future risk strategies as it requires tuning both initial cost and selecting suitable thresholds, preferably at each re-balancing node to obtain lower tail risk comparable to super-replication.

4.8.8 Sensitivity to initial portfolio values for SP, DC and BF strategies

As pointed out in earlier sections, the initial portfolio value for the portfolio as a state variable strategy requires tuning to obtain a lower exposure to tail risk. In this subsection, we continue our analysis by comparing the performance of different initial costs of setting up the portfolio.

Table 4.4 presents a comparative analysis of stochastic programming, dynamic coherent risk, and barriers on future risk strategies across various initial investment costs ranging from 5% to 20%. Each strategy is evaluated using different norms: $\ell_{1,1}, \ell_{2,1}$ and $\ell_{\infty,1}$. The stochastic programming strategy displays a stable expected hedging cost across all norms and initial portfolio values. Specifically, $\ell_{1,1}$ and $\ell_{\infty,1}$ norms maintain expected costs around 10%, while $\ell_{2,1}$ norm exhibits slight variability, ranging from approximately 9.95% to 10.30% as

		$f_0 = 5\%$	$f_0 = 7\%$	$f_0 = 10\%$	$f_0 = 13\%$	$f_0 = 15\%$	$f_0 = 18\%$	$f_0 = 20\%$
SP-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.12%	10.13%	10.13%	10.13%	10.14%	10.14%	10.16%
	$\text{CVaR}_{99.5\%}$	18.51%	17.97%	17.17%	16.56%	16.33%	16.65%	16.86%
SP-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	9.95%	10.09%	10.00%	10.07%	9.98%	10.01%	10.30%
	$\text{CVaR}_{99.5\%}$	22.81%	21.69%	19.96%	16.55%	16.79%	17.31%	22.97%
SP-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.16%	10.16%	10.19%	10.15%	10.17%	10.16%	10.08%
	$\text{CVaR}_{99.5\%}$	15.05%	15.05%	15.32%	16.00%	16.11%	18.04%	19.71%
DC-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.14%	10.13%	10.10%	9.35%	9.62%	9.79%	9.18%
	$\text{CVaR}_{99.5\%}$	21.51%	20.06%	19.09%	18.91%	19.28%	19.32%	20.19%
DC-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	10.33%	10.29%	10.21%	10.05%	9.98%	9.90%	9.83%
	$\text{CVaR}_{99.5\%}$	19.54%	20.99%	23.69%	26.46%	28.33%	31.13%	32.98%
DC-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	9.96%	9.96%	9.96%	9.96%	9.96%	9.48%	9.33%
	$\text{CVaR}_{99.5\%}$	15.27%	15.27%	15.27%	15.27%	15.27%	18.14%	20.06%
BF-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.08%	10.08%	10.05%	9.83%	9.66%	9.39%	9.22%
	$\text{CVaR}_{99.5\%}$	22.64%	21.16%	19.58%	19.48%	19.44%	18.75%	20.09%
BF-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	10.05%	10.04%	9.08%	7.91%	7.38%	6.76%	6.43%
	$\text{CVaR}_{99.5\%}$	22.03%	21.56%	23.64%	26.40%	28.30%	31.08%	32.79%
BF-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	9.64%	9.64%	9.64%	9.64%	9.64%	9.62%	9.38%
	$\text{CVaR}_{99.5\%}$	18.04%	18.04%	18.04%	18.04%	18.04%	18.15%	20.21%

Table 4.4: Expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SP-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$, DC-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ and BF-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with initial investment value $f_0 = 5\%, 7\%, 10\%, 13\%, 15\%, 18\%, 20\%$.

f_0 increases. However, the CVaR for stochastic programming strategies shows more noticeable changes: $\ell_{\infty,1}$ norm has the lowest CVaR, increasing modestly from 15.05% to 19.71%, suggesting a controlled risk profile. The $\ell_{1,1}$ and $\ell_{2,1}$ norms exhibit higher CVaR values, especially for $\ell_{2,1}$ norm, which peaks at 22.97%. Thus, $\ell_{\infty,1}$ norm seems to offer a balance between stable expected costs and moderate tail risk.

The dynamic coherent risk strategy generally shows lower expected values than the stochastic programming strategy, particularly as f_0 increases, with a notable decline in expected hedging costs. Expected costs for $\ell_{\infty,1}$ norm ranges narrowly from 9.96% to 9.33%, highlighting a relatively steady hedging profile. In terms of risk, CVaR for dynamic coherent risk strategies diverges depending on the chosen norm. The $\ell_{\infty,1}$ norm displays moderate CVaR (15.27% to 20.06%), whereas $\ell_{2,1}$ norm experiences high and increasing CVaR, peaking at nearly 32.98%, indicating substantial tail risk exposure. Thus, the $\ell_{\infty,1}$ norm may appeal for its balanced risk and cost and the $\ell_{2,1}$ norm could be considered riskier despite low expected costs.

The barrier on future risk strategy appears distinct, reflecting notably lower expected costs but markedly higher tail risks (CVaR). The $\ell_{\infty,1}$ norm has consistent expected values around 9.6% with CVaR starting at 18.04% for $f_0 = 5\%$ and gradually increasing to 20.21% as f_0 rises, which may indicate risk stabilization at lower initial investments. The $\ell_{2,1}$ norm, while starting at lower expected costs than $\ell_{1,1}$ norm strategy, has the highest CVaR, with values from 22.03% up to 32.79%, making it the riskiest among all the configurations analyzed. The $\ell_{1,1}$ norm falls in between, with CVaR ranging from 18.75% to 22.64% as expected costs decrease. This pattern suggests that BF strategies may be beneficial for low-cost risk management but can present substantial tail risks. Again, the complexity of barriers to future risk strategies requires tuning both initial cost and selecting suitable thresholds.

There is a U-shape trend, particularly in the CVaR for stochastic programming and dynamic coherent risk strategies as f_0 increases. Specifically, for SP-$\ell 1, 1$, SP-$\ell_{2,1}$ and DC-$\ell 1, 1$ norms, tail risk initially decreases as f_0 increases from lower levels but eventually begins to rise again for higher values of f_0. This trend suggests that there's an optimal range of initial investment where tail risk is minimized (around $f_0 = 13\%$ and $f_0 = 15\%$). Below this optimal range, insufficient capital limits risk-mitigation capacity, while above it, additional funds may lead to riskier allocations or exposure, raising the tail risk again. As such, balancing initial investment levels is critical for optimizing risk containment in stochastic programming and dynamic coherent risk strategies. Investors might consider maintaining capital close to this optimal range to achieve effective tail risk management without excessive exposure to extreme losses at either end of the spectrum.

4.8.9 Sensitivity to hedger's preference for SP, DC and BF strategies

Next, we analyze the performance of portfolios as state variable strategies to changing model parameters based on the hedger's preference. Again, these parameters can be tuned individually or collectively to obtain the desired level of performance and robustness. Specifically, we analyze the proportion of future risk included in the objective function, the asymmetric parameter of the norms, the threshold for the barrier on future risk and the inclusion of additional assets in the hedging portfolio. We fix the initial cost to set up the portfolio at 0.1 for all the strategies.

Table 4.5 presents the expected hedging costs and $\mathrm{CVaR}_{99.5\%}$ as a performance measure for the portfolio as state variable strategies. The parameter β represents the proportion of future costs considered in the objective function of the hedging strategy. The expected error of the stochastic programming strategy remains relatively stable across different values of β with no changes. The $\mathrm{CVaR}_{99.5\%}$ also shows consistency, suggesting that these strategies are relatively insensitive to changes in the proportion of future costs. There is also no change in the expected hedging error for dynamic coherent risk strategies as β increases, except for the $\ell_{2,1}$ norm. The $\mathrm{CVaR}_{99.5\%}$ also remain the same for $\ell_{1,1}$ and $\ell_{\infty,1}$ norms, indicating that accounting for higher future costs have no significant impact on reducing tail risk under dynamic coherent risk strategies. Similar to the stochastic programming strategies, the expected hedging costs for the barrier on future risk are stable, but the $\mathrm{CVaR}_{99.5\%}$ decreases as the threshold γ_0 increases, particularly for $\ell_{1,1}$ and $\ell_{2,1}$ norms. This trend indicates that higher thresholds or barriers may reduce both the expected cost of hedging tail risk in these strategies, except for the $\ell_{\infty,1}$ norm strategy.

		β			q			Number of options	
		$\beta = 0.5$	$\beta = 0.7$	$\beta = 1$	$q = 0.7$	$q = 0.85$	$q = 1$	1	3
SP-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.13%	10.13%	10.13%	10.12%	10.12%	10.13%	10.13%	10.00%
	$\mathrm{CVaR}_{99.5\%}$	17.21%	17.18%	17.17%	16.59%	16.78%	17.17%	17.17%	14.60%
SP-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	10.37%	10.09%	10.00%	10.76%	10.03%	10.00%	10.00%	10.85%
	$\mathrm{CVaR}_{99.5\%}$	19.01%	19.81%	19.96%	18.66%	19.29%	19.96%	19.96%	20.47%
SP-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.21%	10.21%	10.19%	10.15%	10.09%	10.19%	10.19%	9.88%
	$\mathrm{CVaR}_{99.5\%}$	15.42%	15.42%	15.32%	15.04%	15.18%	15.32%	15.32%	13.03%
DC-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.10%	10.10%	10.10%	10.11%	10.10%	10.10%	10.10%	9.88%
	$\mathrm{CVaR}_{99.5\%}$	19.09%	19.09%	19.09%	18.89%	18.99%	19.09%	19.09%	15.37%
DC-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	9.97%	10.09%	10.21%	10.21%	10.21%	10.21%	10.21%	11.20%
	$\mathrm{CVaR}_{99.5\%}$	29.43%	24.16%	23.69%	23.61%	23.63%	23.69%	23.69%	22.35%
DC-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	9.96%	9.96%	9.96%	9.96%	9.96%	9.96%	9.96%	9.66%
	$\mathrm{CVaR}_{99.5\%}$	15.27%	15.27%	15.27%	15.23%	15.26%	15.27%	15.27%	13.30%
		γ_0							
		$\gamma_0 = 1\%$	$\gamma_0 = 3\%$	$\gamma_0 = 5\%$					
BF-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.05%	10.02%	10.00%	10.08%	10.06%	10.05%	10.05%	9.93%
	$\mathrm{CVaR}_{99.5\%}$	19.58%	19.42%	19.35%	19.22%	19.32%	19.58%	19.58%	15.46%
BF-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	9.08%	9.07%	9.04%	9.08%	9.07%	9.08%	9.08%	10.56%
	$\mathrm{CVaR}_{99.5\%}$	23.64%	23.58%	23.52%	23.40%	23.49%	23.64%	23.64%	20.90%
BF-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	9.64%	9.32%	9.10%	9.65%	9.64%	9.64%	9.64%	9.32%
	$\mathrm{CVaR}_{99.5\%}$	18.04%	23.32%	26.47%	17.88%	17.93%	18.04%	18.04%	18.72%

Table 4.5: Expected value $\mathbb{E}[\mathcal{H}]$ and $\mathrm{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SP-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$, DC-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ and BF-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with proportion of future cost $\beta = 0.5, 0.7, 1$, asymmetric parameter $q = 0.7, 0.85, 1$, threshold $\gamma_0 = 1\%, 3\%, 5\%$, and 1 call option with $\kappa = 1$, 3 call options with $\kappa = 0.8, 1, 1.2$.

The asymmetric parameter q modifies the distribution of losses in the risk measure. Again, the expected hedging costs for stochastic programming strategies remain fairly constant across different levels of asymmetry. The $\mathrm{CVaR}_{99.5\%}$ however, shows significant fluctuations with values increasing as we increase the impact of positive losses. The expected hedging costs are stable for dynamic coherent risk strategies, but there is a noticeable increase in $\mathrm{CVaR}_{99.5\%}$ as q increases. This trend indicates that the potential for extreme losses under dynamic coherent risk strategies increases as the distribution becomes more symmetric. Similar to the other strategies, barriers on future risk strategies show little sensitivity in expected error to changes in the asymmetric parameter. The $\mathrm{CVaR}_{99.5\%}$ also follows a similar trend as stochastic programming and dynamic coherent risk, suggesting that these strategies are more affected by asymmetry.

The introduction of extra assets in the form of additional options has a notable impact on both the expected hedging error and tail risk. Generally, including three options (with strike prices $\kappa = 0.8, 1, 1.2$) instead of only one option reduces the expected hedging error and reduces $\mathrm{CVaR}_{99.5\%}$. This outcome suggests diversifying the hedging portfolio with more assets can reduce risk. The decrease in $\mathrm{CVaR}_{99.5\%}$ is particularly pronounced in the dynamic coherent strategies, indicating a substantial reduction in potential extreme losses. The analysis of Table 4.5 reveals that the stochastic programming strategies are generally robust across different parameters (β, q), showing stable expected hedging costs and controlled $\mathrm{CVaR}_{99.5\%}$. Dynamic coherent risk strategies tend to exhibit higher sensitivity to changes in these parameters, especially regarding tail risk, indicating a higher potential for extreme losses under certain conditions. While generally exhibiting higher CVaR values, barriers on future risk strategies benefit from including extra assets, which can mitigate risk and improve hedging outcomes. Overall, the choice of strategy and parameters should be carefully considered based on the hedger's specific risk tolerance and investment objectives.

4.8.10 Sensitivity to option moneyness

We perform sensitivity analysis of the models to different strike prices of the call option. Specifically, we compare the performance of the models when the option is in-the-money (ITM) with $K = 0.95$, at-the-money (ATM) with $K = 1$, and out-of-the-money (OTM) with $K = 1.05$. Again, we set the initial portfolio value to be 0.1 for the portfolio as state variable strategies.

	ITM			ATM			OTM		
	f_0	$\mathbb{E}[\mathcal{H}]$	$\text{CVaR}_{99.5\%}$	f_0	$\mathbb{E}[\mathcal{H}]$	$\text{CVaR}_{99.5\%}$	f_0	$\mathbb{E}[\mathcal{H}]$	$\text{CVaR}_{99.5\%}$
SR	15.97%	10.15%	14.43%	15.97%	10.15%	14.43%	15.97%	10.15%	14.43%
NC-$\ell_{1,0}$	18.61%	12.70%	16.89%	15.72%	10.15%	14.43%	14.10%	8.07%	12.50%
NC-$\ell_{2,0}$	18.59%	12.70%	16.90%	15.68%	10.15%	14.43%	14.08%	8.07%	12.51%
NC-$\ell_{\infty,1}$	18.49%	12.70%	16.89%	15.56%	10.15%	14.43%	13.99%	8.07%	12.53%
NO-$\ell_{1,1}$	11.96%	12.84%	18.36%	8.86%	10.09%	16.25%	6.43%	7.81%	14.18%
NO-$\ell_{2,1}$	10.28%	12.90%	19.01%	7.14%	10.12%	16.77%	4.58%	7.80%	14.79%
NO-$\ell_{\infty,1}$	10.40%	12.78%	18.09%	7.47%	10.06%	15.66%	5.01%	7.82%	13.62%
SP-$\ell_{1,1}$	10.00%	12.78%	19.67%	10.00%	10.13%	17.17%	10.00%	7.97%	14.48%
SP-$\ell_{2,1}$	10.00%	12.43%	26.33%	10.00%	10.00%	19.96%	10.00%	8.59%	35.42%
SP-$\ell_{\infty,1}$	10.00%	12.80%	18.12%	10.00%	10.19%	15.32%	10.00%	7.86%	13.81%
DC-$\ell_{1,1}$	10.00%	12.86%	22.35%	10.00%	10.10%	19.09%	10.00%	7.73%	17.12%
DC-$\ell_{2,1}$	10.00%	12.92%	21.44%	10.00%	10.21%	23.69%	10.00%	7.99%	22.27%
DC-$\ell_{\infty,1}$	10.00%	12.35%	18.49%	10.00%	9.96%	15.27%	10.00%	7.78%	13.96%
BF-$\ell_{1,1}$	10.00%	12.78%	23.08%	10.00%	10.05%	19.58%	10.00%	7.84%	17.01%
BF-$\ell_{2,1}$	10.00%	12.78%	24.06%	10.00%	9.08%	23.64%	10.00%	6.05%	22.91%
BF-$\ell_{\infty,1}$	10.00%	12.10%	21.08%	10.00%	9.64%	18.04%	10.00%	7.55%	16.34%

Table 4.6: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for ITM ($K = 0.95$), ATM ($K = 1$) and OTM ($K = 1.05$) call options for SR, NC-($\ell_{1,0}, \ell_{2,0}, \ell_{\infty,1}$), NO-($\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1}$), SP-($\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1}$), DC-($\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1}$) and BF-($\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1}$) strategies.

In Table 4.6, we provide a detailed comparison of the proposed hedging strategies across three option scenarios: in-the-money (ITM), at-the-money (ATM), and out-of-the-money (OTM). Across all strategies, there is a clear trend where initial costs decrease as the option moves from ITM to OTM. For example, NO-$\ell_{2,1}$ has an initial cost of 10.28% ITM but drops to 4.58% OTM. This indicates that as options become less risky (OTM), the cost to hedge them decreases across all strategies. Super replicating strategy stands out as an exception with its initial cost fixed across all moneyness scenarios. The expected hedging costs follow a similar trend, with costs typically lower for OTM options compared to ITM. This reflects the lower risk of hedging OTM options since they are less likely to be exercised.

The norm as objective strategies show the sharpest decline in expected costs from ITM to OTM (e.g., NO-$\ell_{1,1}$ drops from 12.84% to 7.81%), while super-replication and norm as constraint strategies show the same values within each moneyness level. The super-replicating strategy is more stable across moneyness levels. The portfolio as state variable strategies have more stable expected hedging costs across moneyness scenarios, with small variations across the different norms. The expected cost for these strategies is consistently around 10% for ATM and slightly lower for OTM. There is a general trend of decreasing CVaR as options move from ITM to OTM, highlighting that OTM options involve less extreme risk. For instance, NO-$\ell_{2,1}$ has a CVaR of 19.01% for ITM but only 14.79% for OTM. However, stochastic programming, dynamic coherent risk and barriers on future risk strategies display more variation in CVaR values across different moneyness scenarios, highlighting the impact of the choice of norm in controlling tail risk and robustness.

To measure robustness, we seek a model that minimizes the downside risks (CVaR) while maintaining cost efficiency. Super-replication is the most robust strategy across all moneyness scenarios, with consistent initial costs ($f_0 = 15.97\%$) and stable expected hedging costs ($\mathbb{E}[\mathcal{H}] = 10.15\%$). Its CVaR is moderate at 14.43%, which does not fluctuate across ITM, ATM, or OTM scenarios. This suggests that super-replication is highly resilient to market changes, delivering predictable performance and protection across different market conditions. This strategy is also conservative, reflecting its design to hedge against the worst-case scenario. However, its high cost might limit its practicality for less risk-averse investors. The norm as constraint strategy is relatively robust but demonstrates slightly higher variability in CVaR compared to super-replication. For instance, NC-$\ell_{1,0}$ has a CVaR that decreases from 16.89% (ITM) to 12.50% (OTM). The initial costs for the norm as constraint strategies are somewhat high, especially in ITM scenarios, but these costs decrease as the option moves out of the money. The robustness of norm as constraint lies in its ability to control risk across different market conditions, though it is not as stable as super-replication. The use of different norms provides slight adjustments to risk management but does not drastically change the robustness. Norm as objective strategies display the lowest initial costs, especially as options move OTM, with NO-$\ell_{2,1}$ costing only 4.58% OTM. However, the CVaR for the norm as objective strategies is higher, especially in ITM scenarios. This makes having norms as objective strategies less robust in terms of managing extreme risk. While they minimize upfront costs, norm as objective strategies expose the portfolio to high risk under extreme market conditions. However, this sensitivity to market uncertainty, particularly in ITM scenarios, shows a trade-off between cost and robustness, making them suitable for investors seeking low upfront costs but willing to take on higher tail risk.

Stochastic programming strategies have relatively high CVaR with values between 13.81% and 35.42% across all moneyness scenarios and norms. This suggests that stochastic programming is highly sensitive to extreme risks given the same initial cost in all cases. As such, stochastic programming strategies may appeal to market speculators interested in hedging strategies that are versatile across different market conditions. Dynamic coherent risk strategies show similar trends to stochastic programming, with initial costs also fixed at 10%, they produce slightly higher CVaR than stochastic programming. This suggests that dynamic coherent risk strategies have a moderate level of robustness, particularly for more extreme market conditions when we consider the $\ell_{2,1}$ norm. They are not as cost-efficient as the norm as objective strategies but comparable to tail-risk protection. Barriers on future risk strategies exhibit similar robustness to stochastic programming and dynamic coherent risk strategies. The initial cost is also set to 10%, but the CVaR fluctuates between 16.34% (OTM) and 24.06% (ITM). These strategies appear slightly more volatile in risk management compared to stochastic programming, but they offer good performance in terms of hedging costs. Barrier on future risk is therefore a robust strategy for situations where investors are willing to tolerate slightly higher extreme risk in return for lower initial costs.

4.9 Transaction costs

In this section, we will explore the impact of transaction costs on hedging strategies and portfolios. Since transaction fees are only paid for variations in the number of stocks, we distinguish between stocks that are already held and those that vary in quantity. Additionally, stock transactions do not occur at their theoretical or book value. As such, we include a proportional bid-ask spread due to the existence of proportional transaction costs.

4.9.1 Risky asset as state variable

Using state variables to describe how the portfolio composition changes over time, we can distinguish stocks already held and the amount of newly bought or sold stocks. If we let $\mathbf{z}_{i_{t-1}}$, $x^{+}_{i_{t-1},1}(\geq 0)$ and $x^{-}_{i_{t-1},1}(\geq 0)$ denote the stock already held and the amount of stocks newly bought and sold respectively, we can define the amount of stocks that are accessible for every transition with a state equation $H_{i_{t-1}}$ in our optimization for Algorithms 3.2 and 3.5. The corresponding state equation can be defined as,

$$\mathbf{z}_{j_t} = \mathbf{z}_{i_{t-1}} + x^{+}_{i_{t-1},1} - x^{-}_{i_{t-1},1},\ \forall j_t \in J_t|i_{t-1}. \tag{4.41}$$

The value of $\mathbf{z}_{j_t}$ is independent of the future amount and is dependent only on the stocks that are held at i_{t-1} and those recently traded on the market. Furthermore, there is no dependence of this state equation on the conditional transitions from node i_{t-1}.

The distinction between purchasing and selling of stocks is required since transaction costs differ for each position. More precisely, for a fix proportional transaction cost parameter $\epsilon \in [0, 1)$, we let $(1-\epsilon)y_{j_t}$ and $(1+\epsilon)y_{j_t}$ denote the bid and ask prices respectively for $j_t \in J_t|i_{t-1}$. By including the third investment in the call option with strike price κy_{j_t} and option price $\phi_{j_t}(\kappa)$ for $j_t \in J_t|i_{t-1}$, the corresponding accumulated hedge with transaction cost for our three asset portfolio can be formulated as $\forall j_t \in J_t | i_{t-1}$,

$$w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \begin{cases} x_{i_{t-1},0}e^r + \mathbf{z}_{j_t}\frac{y_{j_t}}{y_{i_{t-1}}} + \frac{x_{i_{t-1},2}}{\phi_{i_{t-1}}(\kappa)}(y_{j_t} - \kappa y_{i_{t-1}})^+, & t < T \\ x_{i_{T-1},0}e^r + \mathbf{z}_{j_T}(1-\epsilon)\frac{y_{j_T}}{y_{i_{T-1}}} + \frac{x_{i_{T-1},2}}{\phi_{i_{T-1}}(\kappa)}(y_{j_T} - \kappa y_{i_{T-1}})^+, & \mathbf{z}_{j_T} \geq 0 \\ x_{i_{T-1},0}e^r + \mathbf{z}_{j_T}(1+\epsilon)\frac{y_{j_T}}{y_{i_{T-1}}} + \frac{x_{i_{T-1},2}}{\phi_{i_{T-1}}(\kappa)}(y_{j_T} - \kappa y_{i_{T-1}})^+, & \mathbf{z}_{j_T} < 0. \end{cases} \tag{4.42}$$

4.9.2 Portfolio value and risky asset as state variables

In the case where the hedging strategy already employs state variables to represent the portfolio values - stochastic programming, dynamic coherent risk and barrier on future risk strategies - we extend our formulation to include both transaction costs and portfolio value as distinct state variables. That is, we let $\mathbf{z}^{(1)}_{i_{t-1}}$ denote the state variable for transaction costs as before in Equation (4.41), and introduce $\mathbf{z}^{(2)}_{i_{t-1}}$ to represent the portfolio value as a state variable. Thus the two-dimensional state variable equations $H_{i_{t-1}}$ for Algorithms 3.9 and 3.11 can be defined as,

$$\mathbf{z}^{(1)}_{j_t} = \mathbf{z}^{(1)}_{i_{t-1}} + x^+_{i_{t-1},1} - x^-_{i_{t-1},1}, \ \forall j_t \in J_t | i_{t-1}. \tag{4.43}$$

$$\mathbf{z}^{(2)}_{i_{t-1}} = \mathbf{z}^{(1)}_{i_{t-1}} + x_{i_{t-1},0} + (1-\epsilon)x^+_{i_{t-1},1} - (1+\epsilon)x^-_{i_{t-1},1} + (1+\epsilon)x_{i_{t-1},2} \tag{4.44}$$

The corresponding accumulated hedge with transaction cost can be formulated as $\forall j_t \in J_t | i_{t-1}$,

$$w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \begin{cases} x_{i_{t-1},0}e^r + \mathbf{z}^{(1)}_{j_t}\frac{y_{j_t}}{y_{i_{t-1}}} + \frac{x_{i_{t-1},2}}{\phi_{i_{t-1}}(\kappa)}(y_{j_t} - \kappa y_{i_{t-1}})^+, & t < T \\ x_{i_{T-1},0}e^r + \mathbf{z}^{(1)}_{j_T}(1-\epsilon)\frac{y_{j_T}}{y_{i_{T-1}}} + \frac{x_{i_{T-1},2}}{\phi_{i_{T-1}}(\kappa)}(y_{j_T} - \kappa y_{i_{T-1}})^+, & \mathbf{z}^{(1)}_{j_T} \geq 0 \\ x_{i_{T-1},0}e^r + \mathbf{z}^{(1)}_{j_T}(1+\epsilon)\frac{y_{j_T}}{y_{i_{T-1}}} + \frac{x_{i_{T-1},2}}{\phi_{i_{T-1}}(\kappa)}(y_{j_T} - \kappa y_{i_{T-1}})^+, & \mathbf{z}^{(1)}_{j_T} < 0. \end{cases} \tag{4.45}$$

We set $\epsilon = 0$ if there are no transaction costs. More generally, for some predictable processes, we could describe the bid and ask prices differently as time-dependent variables, ϵ^{bid}_t and ϵ^{ask}_t, respectively. However, we employ a symmetric bid-ask spread parameter for ease of use. We lose the feasibility of the hedging optimizations by introducing transaction costs. As such, we rely on Algorithm 3.7, where we search a polyhedral to initialize the state variables. The resulting problem has the same mathematical properties and solution techniques as that for Algorithms 3.9 and 3.11.

4.9.3 Analysis of transaction costs for SR, NC, SP and DC strategies

Due to the continuous evolution of the financial landscape, sophisticated strategies are required to hedge against potential risks. We illustrate the impact of different levels of transaction cost parameters on our proposed models and assess their sensitivities. We also increase the price of the option asset by $(1+\epsilon)$ to simplify the impact of transaction cost when trading the option component in the portfolio. In Table 4.7, we test the super-replication, norm as constraint, stochastic programming, and dynamic coherent risk strategies under different levels of transaction costs, denoted by ϵ, ranging from 0% to 1.5%. The super replication strategy, characterized by its robust approach to ensuring no shortfall in meeting liabilities, predictably has the highest initial portfolio values across all transaction cost levels. This indicates the conservative nature of super replication, which requires substantial initial capital to guarantee no losses. However, as transaction costs increase, the initial value only rises from 15.97% to 16.09%. This slight increase suggests that the super replication strategy can reduce the impact of transaction costs on initial capital. The expected hedging cost for super-replication increases across different transaction cost levels, starting at 10.15% and rising slightly to 10.53%. This trend shows it becomes more expensive to maintain expected outcomes despite the higher initial costs. Similarly, the $\mathrm{CVaR}_{99.5\%}$ for super replication gradually increases from 14.43% to 14.69%, reflecting the strategy's effectiveness in containing extreme losses, albeit at a higher initial cost.

Norm as constraint strategies, including $\ell_{1,0}, \ell_{2,0}$ and $\ell_{\infty,1}$ norms, offer a balance between

ϵ		SR	NC-$\ell_{1,0}$	NC-$\ell_{2,0}$	NC-$\ell_{\infty,1}$	SP-$\ell_{1,1}$	SP-$\ell_{\infty,1}$	DC-$\ell_{1,1}$	DC-$\ell_{\infty,1}$
	f_0	15.97%	15.72%	15.68%	15.56%	10.00%	10.00%	10.00%	10.00%
0%	$\mathbb{E}[\mathcal{H}]$	10.15%	10.15%	10.15%	10.15%	10.13%	10.19%	10.10%	9.96%
	$\text{CVaR}_{99.5\%}$	14.43%	14.43%	14.43%	14.43%	17.17%	15.32%	19.09%	15.27%
	f_0	15.99%	15.74%	15.71%	15.58%	10.00%	10.00%	10.00%	10.00%
0.3%	$\mathbb{E}[\mathcal{H}]$	10.41%	10.27%	10.27%	10.26%	15.16%	15.44%	11.84%	13.76%
	$\text{CVaR}_{99.5\%}$	14.59%	14.51%	14.51%	14.50%	19.69%	15.52%	19.07%	16.17%
	f_0	16.01%	15.76%	15.73%	15.6%	10.00%	10.00%	10.00%	10.00%
0.5%	$\mathbb{E}[\mathcal{H}]$	10.43%	10.29%	10.29%	10.28%	15.64%	15.72%	11.95%	13.94%
	$\text{CVaR}_{99.5\%}$	14.60%	14.53%	14.53%	14.52%	20.90%	15.85%	19.38%	16.91%
	f_0	16.09%	15.85%	15.81%	15.69%	10.00%	10.00%	10.00%	10.00%
1.5%	$\mathbb{E}[\mathcal{H}]$	10.53%	10.39%	10.39%	10.39%	16.72%	17.08%	12.68%	14.85%
	$\text{CVaR}_{99.5\%}$	14.69%	14.61%	14.61%	14.60%	22.85%	17.35%	27.24%	20.57%

Table 4.7: Initial value f_0, Expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SR, NC-($\ell_{1,0}, \ell_{2,0}, \ell_{\infty,1}$), SP-($\ell_{1,1}, \ell_{\infty,1}$) and DC-($\ell_{1,1}, \ell_{\infty,1}$) strategies with transaction cost parameter $\epsilon = 0\%, 0.3\%, 0.5\%$, and 1.5%.

risk and initial cost. The initial portfolio values for these strategies are slightly lower than those for super replication, indicating a less conservative approach. For instance, $\ell_{1,0}$ norm starts with an initial value of 15.72% with no transaction cost, slightly increasing to 15.85% at 1.5% transaction cost. This marginal increase across all norms as constraint strategies suggests that they are relatively more efficient in capital utilization compared to super replication. The expected hedging costs for the norm as constraint strategies are consistent with those of super replication, remaining close to 10.15% across different values of ϵ. However, the norm as constraint strategies exhibits lower $\text{CVaR}_{99.5\%}$ in the presence of transaction cost, indicating slightly better performance than super-replication under transaction costs. As transaction costs increase, the $\text{CVaR}_{99.5\%}$ for norm as constraint strategies remains relatively stable, indicating that these strategies effectively manage tail risk without incurring significant additional costs.

Stochastic programming strategies, specifically $\ell_{1,1}$ and $\ell_{\infty,1}$ norms, are designed to optimize the trade-off between risk and reward by incorporating stochastic elements into the decision-making process. The initial values for stochastic programming strategies are fixed at 10%, regardless of the transaction cost. The expected hedging costs for both norms increase with transaction costs included, starting around 10% and rising to 17%. The $\mathrm{CVaR}_{99.5\%}$ for stochastic programming strategy varies for each norm, with $\ell_{1,1}$ norm showing a higher increase to 22.85% at higher transaction costs, indicating its tendency to produce high tail risk in extreme risk scenarios when transaction costs are accounted for.

Dynamic coherent risk strategies integrate dynamic risk measures to adjust hedging strategies in real-time, particularly $\ell_{1,1}$ and $\ell_{\infty,1}$ norms. These strategies start with the same fixed initial portfolio value at 10%. The expected hedging error for $\ell_{\infty,1}$ norm increases from 9.96% to 14.85% as transaction costs rise, suggesting a slight deterioration in hedging performance with higher transaction costs. On the other hand, $\ell_{1,1}$ norm slightly increases in expected hedging cost from 10.10% to 12.68% as ϵ increases.

The analysis of the hedging strategies across varying transaction costs reveals that each strategy has its strengths and weaknesses depending on the hedger's specific financial goals and risk tolerance. The super replication strategy offers the highest security but at a significant initial cost, making it suitable for highly risk-averse investors. Norm as constraint strategies provides a balanced approach, offering efficient capital utilization while maintaining reasonable risk management. Stochastic programming and dynamic coherent risk strategies, particularly $\ell_{\infty,1}$ norms, demonstrate dynamic adaptability, showing substantial improvements in risk management as transaction costs are introduced.

Chapter 5

Self-financing Hedging Strategies

Hedging is a popular approach to control risks and lower possible losses in insurance and finance. The cost is a major deterrent for investors when it comes to hedging. Self-financing strategies are one of the best ways to hedge because they allow hedgers to obtain their objective without injecting or withdrawing cash from the portfolio. This is because the possible profits from the financial position are used to cover the cost of hedging. Hence, the portfolio's assets are used to finance the re-balancing of the portfolio, which involves the purchasing and selling of assets. In other words, there is no need for outside cash flows when adjusting the portfolio's composition. Therefore, a self-financing strategy might be preferred because of its theoretical consistency, sustainability (as it guarantees that the hedging strategy can be maintained over time without reliance on external funding), and ease of use of mathematical models for pricing and risk management. However, self-financing techniques can be challenging to implement in practice since they need ongoing portfolio monitoring and adjustment, which can be difficult to execute operationally. Furthermore, the portfolio can be exposed to model risk due to the high sensitivity of the performance of such strategies to the accuracy of model parameters. The addition of extra constraints can also introduce computational cost and the feasibility of the hedge portfolio selection problem.

Mahayni (2003) examines the effectiveness of self-financing strategies under model misspecifications and trading restrictions. They use the robustness of Gaussian hedging strategies to analyze the discrete-time errors. They observe that simply discretizing time introduces duplication due to asset price trends, which can be avoided by discretizing the whole model instead. Rudloff (2009) also uses a coherent risk measure to find a self-financing strategy using static optimization. They show that the optimal self-financing strategy that minimizes the coherent risk of the shortfall consists of super-replicating a modified claim that is the product of the original payoff and the solution of their static optimization problem.

We discuss the formulation of self-financing strategies for the proposed hedging portfolios. So far, our goal for constructing the hedging strategies has been to either ensure a negative loss in the case of super-replication or allow for some loss using thresholds and state variables. In these cases, our local hedging strategies allow possible injections to meet the desired criteria. However, to make the strategies self-financing, we need to impose an additional constraint to ensure that the compositions of assets in the hedge portfolio are adjusted to equal the state value without any external cash flows. While constructing a self-financing replicating portfolio may be easier in a complete market, the same cannot be said in an incomplete market. We extend our development to self-financing strategies under market incompleteness in this chapter.

5.1 Global super-replication (SF-SR)

In section 3.1, we introduced the super-replicating strategy where we ensure the value of the hedge portfolio supersedes the contingent claim at each node almost surely. Algorithm 3.2 minimizes the value of this hedge portfolio subject to constraints that ensure negative losses at each node throughout the optimization. This formulation renders the portfolio non-self-financing; hence, injections and withdrawals can be made to the assets to satisfy constraints at each step. A self-financing strategy with the least cost that dominates the claim or payoff at maturity is known as a self-financing super-replicating strategy. Such strategies do not allow any positive loss. Hence, the highest admissible loss is 0 for all possible outcomes. When the portfolio compositions are fixed from time $t = 0$, the strategy is known as a static super-replicating strategy. Chen et al. (2015) investigate the problems with optimizing a self-financing super-replicating strategy. They prove that one can have some flexibility when composing the optimal strategy in real market situations since the optimal solution is not unique. Dolinsky and Neufeld (2018) also developed a probabilistic approach to self-financing super-replication under market incompleteness. They demonstrate that the super-replication and model-independent super-replication prices are the same for incomplete markets. As a result, understanding the model does not lower the cost of super-replication.

Definition 5.1.1. *For an European option and* $t \in \mathcal{T}$*,*

$$\psi = \{X_{i_{t-1}} | X_{i_{t-1}} \textit{ is self-financing, } L_{T|i_{T-1}}(X_{i_{T-1}}) < 0\}.$$

Thus, ψ is the set of all self-financing strategies with no positive errors or losses at maturity such that $\{X_{i_0}, X_{i_1}, \ldots, X_{i_{t-1}}\}$ is one sequence of such strategies. In our framework, a self-financing super-replicating strategy belongs to ψ and provides the least-cost initial portfolio value. Hence, we seek the solution set to the following problem.

Algorithm 5.1.

$$\min_{\{X_{i_0}, X_{i_1}, \ldots, X_{i_{T-1}}\} \in \psi} f_0(X_0) \tag{5.1}$$

The idea is to obtain the least-cost hedge portfolio at time $t = 0$ such that, without any outside cash flow during re-balancing, the portfolio value supersedes the claim at maturity T. Alternatively, Algorithm 5.1 can be presented as an unconstrained minimization problem. That is, for a large nulling constant ϱ, we seek the solution set to the following optimization problem.

Algorithm 5.2.

$$\min_{X_{i_0}, X_{i_1}, \ldots, X_{i_{T-1}}} f_0(X_0) + \varrho \sum_{i_{T-1} \in J_{T-1}} \sum_{j_T \in J_T | i_{T-1}} \max\left(0, l_{j_T}(X_{i_{T-1}})\right), \tag{5.2}$$

where $l_{j_T}(X_{i_{T-1}}) = c_{j_T}(X_{j_T}) - \left(f_0(X_0) + \sum_{t=1}^{T} w_{j_t}(X_{i_{t-1}}) - f_{i_{t-1}}(X_{i_{t-1}})\right)$ for all paths and the component $f_0(X_0)+\sum_{t=1}^{T} w_{j_t}(X_{i_{t-1}})-f_{i_{t-1}}(X_{i_{t-1}})$ is the value of the self-financing hedging strategy at maturity. Several methodologies have been proposed in the literature to find a self-financing super-replicating strategy. Recent approaches to solving Algorithm 5.1 involve the use of deep hedging techniques. Carbonneau and Godin (2021) argue that in incomplete markets, super-replication is often not feasible due to high costs, and hedgers must accept a certain level of residual risk. They introduce the equal risk pricing framework by balancing the risk exposure of both long and short positions in a derivative. By training neural networks to minimize the risk exposure for both positions, they use deep reinforcement learning to generate dynamic hedging strategies under a chosen convex risk measure, like Conditional Value-at-Risk.

We rely on a two-step optimization to solve Algorithm 5.1 under our model framework. First, we find the set of strategies with no positive losses at maturity. Secondly, we impose the self-financing condition for each of these strategies and solve for the optimal strategies at time $t = 0$. This two-step optimization gives us the set ψ. An advantage of using linear programming to solve the problem is that we can also include another constraint to minimize the portfolio value at each node and obtain the least-cost self-financing super-replication strategy under the minimization problem. Our two-step optimization requires state equations for a path-dependent process to solve the problem. We introduce state variables and equations under our proposed dynamic algorithm, which allows us to optimize the state variables when solving the problem. Let θ_T be the maximum error at maturity. Then the self-financing super-replicating problem is equivalent to solving the following optimization problem.

Algorithm 5.3. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t} f_{i_{t-1}}(X_{i_{t-1}})\mathbf{1}_{\{t=0\}} + e^{-r}\{\theta_t \mathbf{1}_{\{t<T\}} + \varrho\theta_T \mathbf{1}_{\{t=T\}}\} \tag{5.3}$$

under the constraints

$$l_{j_T}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \leq \theta_T, \forall j_T \in J_T|i_{T-1}, t = T \tag{5.4}$$

$$v_{j_t}(\mathbf{z}_{j_t}) \leq \theta_t, \forall j_t \in J_t|i_{t-1}, t \leq T \tag{5.5}$$

$$f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}, \tag{5.6}$$

$$w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \mathbf{z}_{j_t}, \forall j_t \in J_t|i_{t-1}, t \leq T. \tag{5.7}$$

The state Equations (5.6) and (5.7) ensure the portfolio is self-financing while we seek the strategy such that $\theta_T = 0$ for super-replication. Starting at maturity, we establish a hedge portfolio that supersedes the contingent claim. Then, for all time $t < T$, we ensure the hedge portfolio remains self-financed using the state variables. This approach, coupled with the dynamic algorithm, provides a strategy that offers the least-cost portfolio value that dominates the claim at maturity.

Proposition 5.4. *Algorithm 5.3 is equivalent to Algorithm 5.1.*

Proof. We show that the constraints in Algorithm 5.3 are formulated to satisfy the requirements of Algorithm 5.1. Since Algorithm 5.1 is the least-cost self-financing super-replicating strategy, we minimize $f_0(X_0)$ such that $L_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \leq 0$ and pick the minimum strategy from the set ψ of self-financing strategies.

We can construct infinitely many self-financing super-replicating portfolios when we do not consider the overall cost of setting up the portfolio. We use θ_T to constraint the losses at maturity and obtain the least cost when $\theta_T = 0$. Constraints (5.6) and (5.7) are used to ensure the strategy is self-financing. Given the portfolio's value at node i_{t-1} is the sum of the positions in each asset, $f_{i_{t-1}}(X_{i_{t-1}}) = \sum_{k=0}^{n} x_{i_{t-1},k}$.

The trading strategy $X_{i_{t-1}} = (x_{i_{t-1},k})_{k=0,1,\cdots,n}$ is $\mathcal{F}_{i_{t-2}}$-measurable where $x_{i_{t-1},k}$ is the value of the position in asset k at node i_{t-1} to be determined based on information available before node i_{t-1}. The initial value of the portfolio $f_0(X_0) = \sum_{k=0}^{n} x_{i_0,k}$. Without loss of generality, we let $y_{i_{t-1},k}$ be the price of asset k at node i_{t-1}. For the portfolio to be self-financing, the change in the value at node i_{t-1} to $j_t, \forall j_t \in J_t|i_{t-1}$ should only depend on changes in asset prices without any additional withdrawal or investment. Thus, $f_{j_t}(X_{j_t}) - f_{i_{t-1}}(X_{i_{t-1}}) = \sum_{k=0}^{n} x_{i_{t-1},k}\left(\frac{y_{j_t,k}}{y_{i_{t-1},k}} - 1\right), \forall j_t \in J_t|i_{t-1}$ and $w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \sum_{k=0}^{n} x_{i_{t-1},k}\frac{y_{j_t,k}}{y_{i_{t-1},k}}, \forall j_t \in J_t|i_{t-1}$. At time t, $f_{j_t}(X_{j_t}) = \sum_{k=0}^{n} x_{j_t,k}$ and the change in portfolio value $f_{j_t}(X_{j_t}) - f_{i_{t-1}}(X_{i_{t-1}}) = \sum_{k=0}^{n} x_{j_t,k} - x_{i_{t-1},k}, \forall j_t \in J_t|i_{t-1}$. Nevertheless, under self-financing conditions, the change in the value of the position in the asset $x_{i_{t-1},k}\frac{y_{j_t,k}}{y_{i_{t-1},k}}$ and $x_{j_t,k}$ should balance to 0. That is, $x_{j_t,k} - x_{i_{t-1},k}\frac{y_{j_t,k}}{y_{i_{t-1},k}} = 0 \Rightarrow w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = f_{j_t}(X_{j_t}), \forall j_t \in J_t|i_{t-1}$. Hence $f_{j_t}(X_{j_t}) - f_{i_{t-1}}(X_{i_{t-1}}) = \sum_{k=0}^{n} x_{j_t,k} - x_{i_{t-1},k}\frac{y_{j_t,k}}{y_{i_{t-1},k}} + x_{i_{t-1},k}\frac{y_{j_t,k}}{y_{i_{t-1},k}} - x_{i_{t-1},k} = \sum_{k=0}^{n} 0 + x_{i_{t-1},k}\left(\frac{y_{j_t,k}}{y_{i_{t-1},k}} - 1\right), \forall j_t \in J_t|i_{t-1}$ as expected.

We replace $f_{j_t}(X_{j_t})$ with state vector $\mathbf{Z}_{t|i_{t-1}}$ to obtain the self-financing constraints in Equations (5.6) and (5.7). Finally, we use the duality theorem in convex optimization and Proposition A.2 to show super-replication. Let $m = \#(J_t|i_{t-1})$, the cardinality of the conditional set $J_t|i_{t-1}$ and let I_m be a vector of ones of size m. Then for $X_{i_{t-1}} \in \mathbb{R}^n$, $C_{t|i_{t-1}}(X_{t|i_{t-1}}) \in \mathbb{R}^m$, $W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \in \mathbb{R}^m$, $\mathbf{Z}_{t|i_{t-1}} \in \mathbb{R}^m$, $V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \in \mathbb{R}^m$ and $\mathbf{z}_{i_{t-1}} \in \mathbb{R}$ we have $L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) \in \mathbb{R}^m$. This implies at maturity, for $\theta_T \in \mathbb{R}$ and $f_{i_{T-1}}(X_{i_{T-1}}) = I_n^\top X_{i_{T-1}}$, we obtain

$$v_{i_{T-1}}(\mathbf{z}_{i_{T-1}}) = \min_{X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}, \theta_T} e^{-r} \varrho \theta_T \tag{5.8}$$

such that

$$C_{T|i_{T-1}}(X_{T|i_{T-1}}) - W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \leq I_m \theta_T, \tag{5.9}$$

$$I_n^\top X_{i_{T-1}} = \mathbf{z}_{i_{T-1}}. \tag{5.10}$$

Since the discounting factor e^{-r} and the nulling constant ϱ in the objective function are constants, they do not affect the optimal values and the feasibility of the problem. The scaled objective changes only the magnitude of the optimal values. We define the Lagrangian $\mathcal{L}$ of the optimization problem as a combination of both the objective function and the constraints. Thus, for non-negative $\Lambda = \{\lambda_{j_t}\}_{j_t \in J_t | i_{t-1}} \in \mathbb{R}^{m+}$ and for $\vartheta \in \mathbb{R}$,

$$\begin{aligned}\mathcal{L}(X_{i_{t-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}, \theta_T) = \theta_T &+ \Lambda^\top \left(C_{T|i_{T-1}}(X_{T|i_{T-1}}) - W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) - I_m \theta_T \right) \\ &+ \vartheta \left(I_n^\top X_{i_{T-1}} - \mathbf{z}_{i_{T-1}} \right)\end{aligned} \tag{5.11}$$

such that

$$\begin{aligned}\inf \mathcal{L}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}, \theta_T) &= \Lambda^\top C_{T|i_{T-1}}(X_{T|i_{T-1}}) + \inf_{\theta_T} \left(1 - \Lambda^\top I_m \right) \theta_T \\ &+ \inf_{X_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}} \left(\vartheta I_n^\top X_{i_{T-1}} - \Lambda^\top W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \right) \\ &- \inf_{\mathbf{z}_{i_{T-1}}} \vartheta \mathbf{z}_{i_{T-1}}.\end{aligned} \tag{5.12}$$

Next we let $\nabla_{X_{i_{T-1}}} W_{T|i_{T-1}}$ be the first partial derivative of $W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}})$ with respect to $X_{i_{T-1}}$ and $\mathbf{Z}_{T|i_{T-1}}$ provided $W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}})$ is differentiable. The respective infimums are the values of Λ and ϑ such that

$$\inf_{\theta_T} \left(1 - \Lambda^\top I_m \right) \theta_T = 1 - \Lambda^\top I_m = 1 - \sum_{j_T \in J_T | i_{T-1}} \lambda_{j_T} = 0, \tag{5.13}$$

$$\begin{aligned}\inf_{X_{i_{T-1}}} \left(\vartheta I_n^\top X_{i_{T-1}} - \Lambda^\top W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \right) &= \vartheta I_n^\top - \Lambda^\top \nabla_{X_{i_{t-1}}} W_{T|i_{T-1}} \\ &= n\vartheta - \sum_{k=0}^{n} \sum_{j_T \in J_T | i_{T-1}} \lambda_{j_T} \frac{\partial w_{j_T}(X_{i_{T-1}})}{\partial x_{i_{T-1},k}} \\ &= 0,\end{aligned} \tag{5.14}$$

and

$$\inf_{\mathbf{z}_{i_{T-1}}} \vartheta \mathbf{z}_{i_{T-1}} = \vartheta = 0. \tag{5.15}$$

By putting (5.13), (5.14) and (5.15) together, the dual counterpart to the optimization (5.8), (5.9) and (5.10) at maturity can be formulated as

$$\max_{\Lambda} \sum_{j_T \in J_T | i_{T-1}} \lambda_{j_T} c_{j_T}(X_{j_T}) \tag{5.16}$$

such that

$$\Lambda^\top I_m = 1 \tag{5.17}$$

$$\sum_{k=0}^{n} \sum_{j_T \in J_T | i_{T-1}} \lambda_{j_T} \frac{\partial w_{j_T}(X_{i_{T-1}})}{\partial x_{i_{T-1},k}} = 0 \tag{5.18}$$

$$\Lambda \geq 0. \tag{5.19}$$

If $W_{T|i_{T-1}}$ is linear, then $\nabla_{X_{i_{T-1}}} W_{T|i_{T-1}}$ is a $m \times n$ matrix with the coefficients of $X_{i_{T-1}}$ as the entries. By the variant of Farka's Lemma in Proposition A.2, if every non-negative $\Lambda \in \mathbb{R}^{m+}$ with $\Lambda^\top \nabla_{X_{i_{T-1}}} W_{T|i_{T-1}} = 0$ also satisfies $\Lambda^\top C_{T|i_{T-1}}(X_{T|i_{T-1}}) \leq 0$, then the system $W_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \geq C_{T|i_{T-1}}(X_{T|i_{T-1}})$ has a solution for every $X_{i_{T-1}} \in \mathbb{R}^n$. This implies $L_{T|i_{T-1}}(X_{i_{T-1}}, \mathbf{z}_{i_{T-1}}, \mathbf{Z}_{T|i_{T-1}}) \leq 0$. Hence, the minimum feasible upper bound of the loss variable is when $\theta_T = 0$ and thus super-replication. Also, by duality theorem $\max \Lambda^\top C_{T|i_{T-1}}(X_{T|i_{T-1}}) = \min \theta_T$, so the dual objective $\max \Lambda^\top C_{T|i_{T-1}}(X_{T|i_{T-1}})$ attains the maximum bound under super-replication. The dual problem can be interpreted as seeking the maximum expected value of payoff such that there is no expected change in the value of the replicating portfolio and, hence, self-financing.

Next, we extend our derivations to the optimization problem for when $0 \leq t < T$. After the optimization at maturity, we obtain the cost-to-go function $V_{i_{T-2}}(\mathbf{Z}_{i_{T-2}}) \in \mathbb{R}^m$ and for

$0 \leq t < T$ our optimization becomes

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t} I_n^\top X_0 \mathbf{1}_{\{t=0\}} + e^{-r}\theta_t \tag{5.20}$$

such that

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq I_m \theta_t, \tag{5.21}$$

$$I_n^\top X_{i_{t-1}} = \mathbf{z}_{i_{t-1}}, \tag{5.22}$$

$$W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \mathbf{Z}_{t|i_{t-1}}. \tag{5.23}$$

The corresponding Lagrangian $\mathcal{L}$, of the optimization problem for $\Lambda = (\lambda_{j_t})_{j_t \in J_t|i_{t-1}} \in \mathbb{R}^{m+}$, $\vartheta_1 \in \mathbb{R}$ and $\vartheta_2 \in \mathbb{R}^m$ is defined as

$$\begin{aligned} \mathcal{L}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t) = I_n^\top X_0 + \theta_t + \Lambda^\top \left(V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) - I_m\theta_t\right) + \vartheta_1 \left(I_n^\top X_{i_{t-1}} - \mathbf{z}_{i_{t-1}}\right) \\ + \vartheta_2 \left(W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) - \mathbf{Z}_{t|i_{t-1}}\right), \end{aligned} \tag{5.24}$$

such that

$$\begin{aligned} \inf \mathcal{L}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t) = \Lambda^\top V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) + \inf_{\theta_t} \left(1 - \Lambda^\top I_m\right)\theta_t \\ + \inf_{X_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}} \left(I_n^\top X_0 + \vartheta_1 I_n^\top X_{i_{t-1}} + \vartheta_2 W_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\right) \\ - \inf_{\mathbf{z}_{i_{t-1}}} \vartheta_1 \mathbf{z}_{i_{t-1}} - \inf_{\mathbf{Z}_{t|i_{t-1}}} \vartheta_2 \mathbf{Z}_{t|i_{t-1}}. \end{aligned} \tag{5.25}$$

The minimization of the Lagrangian with respect to $\mathbf{z}_{i_{t-1}}$ and $\mathbf{Z}_{t|i_{t-1}}$ implies the free variable ϑ_1 and ϑ_2 are both zero at infimums. This also means the infimum with respect to $X_{i_{t-1}}$ is also zero. As such, we obtain a simple but intuitive dual counterpart to the optimization as

$$\max_{\Lambda} \sum_{j_t \in J_t|i_{t-1}} \lambda_{j_t} v_{j_t}(\mathbf{z}_{j_t}) \tag{5.26}$$

such that

$$\sum_{j_T \in J_T | i_{T-1}} \lambda_{j_T} = 1 \tag{5.27}$$

$$\lambda_{j_T} \geq 0 \; \forall j_T \in J_T | i_{T-1}. \tag{5.28}$$

The dual objective function, $\max \Lambda^\top V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}})$ is equal to primal objective function, $\min \theta_t$ by the duality theorem. Hence, for every $0 \leq t < T$ of Algorithm 5.3, the optimization is equivalent to solving for the maximum weighted average of the future cost-to-go functions beginning with the value obtained at maturity through to time zero. This shows that Algorithm 5.3 is equivalent to Algorithm 5.1. □

5.2 Norm minimization strategies

In this section, we introduce the self-financing variant of having $\ell_{p,q}$ norm as the objective function in the optimization. Under global hedging, the goal is to obtain the least-cost self-financing strategy such that the $\ell_{p,q}$ norm of the loss at maturity is minimized. We also compare the strategies to the optimal hedging strategy in discrete time proposed by Rémillard and Rubenthaler (2013).

5.2.1 Self-financing stochastic programming (SF-SP)

We develop a global hedging strategy that minimizes the $\ell_{p,q}$ norm of losses at maturity such that any changes in the portfolio's value over time are entirely funded by the portfolio's existing assets, with no external capital injections or withdrawals. We model the self-financing condition using state variables as constraints similar to self-financing super-replication in Algorithm 5.3. We propose the following optimization for general $\ell_{p,q}$ norms.

Algorithm 5.5. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t} e^{-r}\{\|L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\|_{p,q}\mathbf{1}_{\{t=T\}} + \theta_t\mathbf{1}_{\{t<T\}}\} \quad (5.29)$$

under the constraints

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq I_m\theta_t, \quad (5.30)$$

$$f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}, \quad (5.31)$$

$$w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \mathbf{z}_{j_t}, \forall j_t \in J_t|i_{t-1}. \quad (5.32)$$

Again, the state Equations (5.31) and (5.32) ensure the portfolio is self-financing while we seek the strategy such that the $\ell_{p,q}$ norm of losses at maturity is minimized. Alternatively, the dynamic coherent risk formulation of the self-financing stochastic programming strategy is given as,

Algorithm 5.6. *For all $t = T, T-1, \dots, 1$ and all i_{t-1},*

$$v_{i_{t-1}}(\mathbf{z}_{i_{t-1}}) = \min_{X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}, \theta_t} e^{-r}\|L_{t|i_{t-1}}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}})\mathbf{1}_{\{t=T\}} + \theta_t\mathbf{1}_{\{t<T\}}\|_{p,q} \quad (5.33)$$

under the constraints

$$V_{t|i_{t-1}}(\mathbf{Z}_{t|i_{t-1}}) \leq I_m\theta_t, \quad (5.34)$$

$$f_{i_{t-1}}(X_{i_{t-1}}) = \mathbf{z}_{i_{t-1}}, \quad (5.35)$$

$$w_{j_t}(X_{i_{t-1}}, \mathbf{z}_{i_{t-1}}, \mathbf{Z}_{t|i_{t-1}}) = \mathbf{z}_{j_t}, \forall j_t \in J_t|i_{t-1}. \quad (5.36)$$

However, since θ_t is a scalar, equation (5.29) is the same as equation (5.33). As such both stochastic programming and dynamic coherent risk strategies are the same under self-financing conditions. Starting at maturity, we establish a hedge portfolio that minimizes the $\ell_{p,q}$ norm. Then, for all time $t < T$, we ensure the hedge portfolio remains self-financed using the state variables. Similar to global super-replication, the optimal strategy is obtained with the help of the dynamic algorithm in Section 3.4.1.

Typically, hedging strategies like Rémillard and Rubenthaler (2013) and Bertsimas et al. (2001) minimize the mean-square error of losses using a quadratic cost function. However, using an asymmetric $\ell_{p,q}$ norm is more tractable and robust as it provides a generalized penalty function while characterizing the risk profile of the hedger in a model-independent framework. Carbonneau and Godin (2021) also introduce an asymmetric ϵ-completeness measure to quantify the magnitude of unhedgeable risk associated with a position in a contingent claim. Different norms can be used to quantify the losses at maturity and the values of both p and q can be set to align with the hedger's risk tolerance. The $\ell_{2,q}$ norm penalizes large deviations while the $\ell_{1,q}$ norm is more sensitive to extreme values and the $\ell_{\infty,q}$ norm limits the total losses. Since at maturity $V_{T|i_{T-1}}(\mathbf{Z}_{T|i_{T-1}}) = 0$, the self-financing stochastic programming and self-financing dynamic coherent risk strategies are the same. Additionally, by minimizing the maximum losses at maturity, the self-financing $\ell_{\infty,q}$ norm as objective coincides with the super-replication strategy. Table 4.1 highlights the similarities between the $\ell_{\infty,q}$ norm and super-replication can be even under local hedging conditions.

5.2.2 Global $\ell_{2,1}$ hedging (SF-QH)

Next, we elucidate the mean-squared error minimization under self-financing conditions proposed by Rémillard and Rubenthaler (2013) to provide a basis for comparing our self-financing strategies. They begin with the univariate case proven by Schweizer (1995) and extend to the multivariate case. If options are included as assets in our framework, then without loss of generality, we define $\bar{Y}_{t|i_{t-1}}$ and $\bar{y}_{i_{t-1}}$ as matrices with stock prices and option payoffs (or prices, $\phi_{i_{t-1}}$) as columns and let $\Upsilon_{i_{t-1}}$ be the diagonal matrix of $\bar{y}_{i_{t-1}}$. Thus by assuming the price process, $\bar{Y}_t$ to be square integrable and adapted under a discrete filtration $\mathcal{F}_t$ for $t = 0, 1, \ldots, T$, they define some $\Delta_{t|i_{t-1}} = e^{-rt}\bar{Y}_{t|i_{t-1}} - e^{-r(t-1)}\bar{y}_{i_{t-1}}$. The goal is to find the initial portfolio value f_0 and hedging strategy $X_{i_{t-1}}$ such that $\Delta_{t|i_{t-1}}\Upsilon_{i_{t-1}}^{-1}X_{i_{t-1}}$ is square integrable and minimizes the $\mathbb{E}[L_{T|i_{T-1}}(X_{i_{T-1}})^2]$. Rémillard and Rubenthaler (2013) provide a closed-form solution to the following minimization problem.

Algorithm 5.7.

$$\min_{X_{i_0}, X_{i_1}, \ldots, X_{i_{T-1}}} \mathbb{E}\left[\left\{e^{-rT} C_{T|i_{T-1}}(X_{T|i_{T-1}}) - \left(f_0(X_0) + \sum_{t=1}^{T} \Delta_{t|i_{t-1}} \Upsilon_{i_{t-1}}^{-1} X_{i_{t-1}}\right)\right\}^2\right] \tag{5.37}$$

where $e^{rT}\left(f_0(X_0) + \sum_{t=1}^{T} \Delta_{t|i_{t-1}} \Upsilon_{i_{t-1}}^{-1} X_{i_{t-1}}\right)$ is the value of the self-financing strategy at maturity. Thus, they seek the least cost self-financing portfolio that minimizes the expected squared losses at maturity. In this setup, we ensure that the value of the accumulated hedge portfolio is equal to the portfolio value relative to node i_{t-1}. In our case, we set the transition probabilities to be equal and illustrate a robust variant of the closed-form solution to Algorithm 5.7 when the probabilities for all observations are equal. In essence, we set $p_{j_t|i_{t-1}} = 1/(N+1)$, $\forall j_t \in J_t|i_{t-1}$ where $N+1$ represents the number of outcomes from i_{t-1}. Hence the robust counterpart of the closed-form solution proposed by Rémillard and Rubenthaler (2013) is outlined in the following algorithm.

Algorithm 5.8. *Set* $\Pi_{T+1|i_T} = I_{N+1}$.
For $t = T, T-1, \ldots, 1,$ *and* i_{t-1},
Define:

$$\Delta_{t|i_{t-1}} = e^{-rt}\bar{Y}_{t|i_{t-1}} - e^{-r(t-1)}\bar{y}_{i_{t-1}}$$

$$A_{t|i_{t-1}} = \mathbb{E}[\Delta_{t|i_{t-1}}^{\top} \Pi_{t+1|i_t} \Delta_{t|i_{t-1}} | \mathcal{F}_{t-1}] = \tfrac{1}{N+1}(\Delta_{t|i_{t-1}}^{\top} \Pi_{t+1|i_t} I_{N+1}^{\top} \Delta_{i_t|t-1})$$

$$M_{t|i_{t-1}} = \mathbb{E}[\Delta_{t|i_{t-1}}^{\top} \Pi_{t+1|i_t} | \mathcal{F}_{t-1}] = \tfrac{1}{N+1}(\Delta_{t|i_{t-1}}^{\top} \Pi_{t+1|i_t})$$

$$B_{t|i_{t-1}} = A_{t|i_{t-1}}^{-1} M_{t|i_{t-1}}, \textit{ and } P_{t|i_{t-1}} = 1 - \Delta_{t|i_{t-1}} B_{t|i_{t-1}}$$

$$\pi_{i_{t-1}} = \mathbb{E}[\Pi_{t|i_{t-1}} | \mathcal{F}_{t-1}] = \tfrac{1}{N+1}(P_{t|i_{t-1}}^{\top} \Pi_{t+1|i_t})$$

$$\omega_{i_{t-1}} = \mathbb{E}[C_{T|i_{T-1}}(X_{T|i_{T-1}}) \Pi_{t+1|i_t} | \mathcal{F}_{t-1}] = \tfrac{1}{N+1}\left(P_{t|i_{t-1}}^{\top} C_{T|i_{T-1}}(X_{T|i_{T-1}})\right)$$

Compute $f_{i_{t-1}}(X_{i_{t-1}}) = e^{-r\frac{\omega_{i_{t-1}}}{\pi_{i_{t-1}}}}$

Retrieve:

$$\begin{aligned}\Upsilon^{-1}_{i_{t-1}} X_{i_{t-1}} &= A^{-1}_{t|i_{t-1}} \mathbb{E}[e^{-rT} C_{T|i_{T-1}}(X_{T|i_{T-1}}) P^{\top}_{t|i_{t-1}} \mathit{\Delta}_{t|i_{t-1}} | \mathcal{F}_{t-1}] - f_{i_{t-2}}(X_{i_{t-2}}) B_{t|i_{t-1}} \\ &= e^{-rT} \left[A^{-1}_{t|i_{t-1}} \frac{1}{N+1} (\omega_{i_{t-1}} I^{\top}_{N+1} \mathit{\Delta}_{t|i_{t-1}}) \right] - f_{i_{t-2}}(X_{i_{t-2}}) B_{t|i_{t-1}}.\end{aligned}$$

In their analysis, they consider regime-switching geometric random walk models governed by a Markov chain with i.i.d. Gaussian transition matrix for hedging European call options. In our case, we set the transition probabilities to be equal for a typical European call option.

5.3 Numerical examples

In this section, we apply our proposed global hedging strategies to find the initial portfolio value for the at-the-money European call option. Similar to local hedging, we analyze the outputs and sensitivities of the strategies by taking various model parameter adjustments into account. In the absence of any indication to the contrary, we let $T = 1$, $\Delta = 12$, $N = 20$, $r = 4\%$, $\mu = 8\%$, $\sigma = 20\%$, and compute the mean and standard deviation of $100,000$ simulated hedging cost $\mathcal{H}$.

5.3.1 Analysis of self-financing strategies

In this section, we consider numerical examples to compare distributions of the self-financing strategies to their non-self-financing counterparts in Chapters 3 and 4.

The histograms presented in Figure 5.1 depict the distribution of simulated percentage hedging costs for an at-the-money call option under two distinct strategies: local super-replication (SR) and global/self-financing super-replication (SF-SR). Both strategies are designed to mitigate the risk of financial exposure, but their effectiveness can vary based on their structural differences. The super-replicating strategy is a conservative approach that ensures the port-

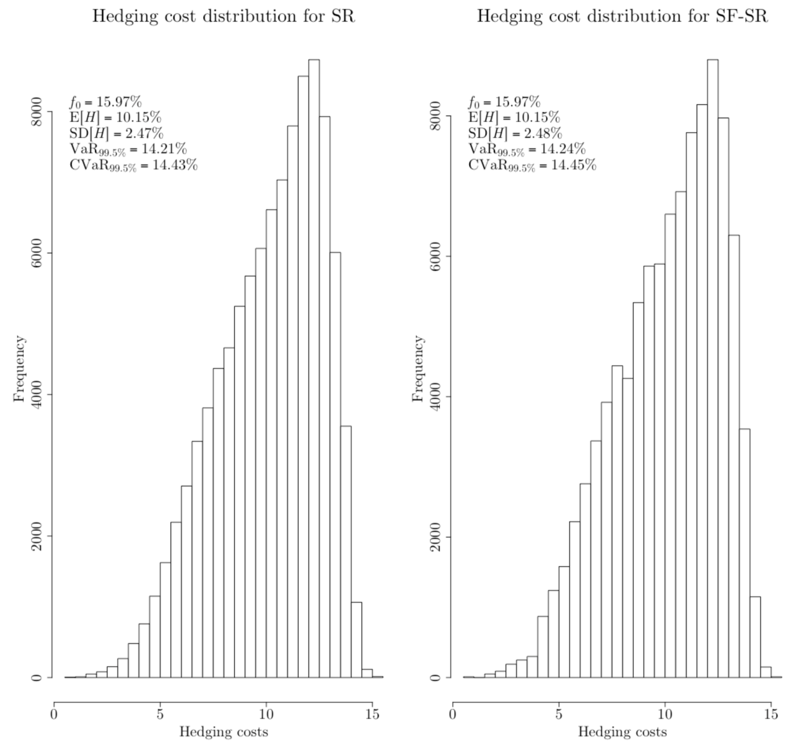

Figure 5.1: Distribution of simulated percent hedging costs for an at-the-money call option with local (left) and global (right) super-replication strategy.

folio's liabilities are fully covered under all possible scenarios. The distribution of hedging costs for the super-replication strategy is fairly skewed to the left, indicating that the hedging costs are less than the initial cost of setting up the portfolio. The global super-replicating strategy, depicted by the histogram on the right, is a variant of the local super-replicating strategy that incorporates the principle of self-financing, where the portfolio is adjusted over time without additional capital injections. The initial cost remarkably remains identical to that of super-replication at 15.97%, indicating that both strategies require the same initial investment to set up. The expected hedging error for the strategies also remains the same at 10.15%, suggesting similar average hedging performance. The standard deviations are relatively similar, indicating similar variability of hedging costs. The CVaR$_{99.5\%}$ for self-financing is 14.45%, closely matching that of local super-replication. This consistency suggests that the self-financing strategy is equally efficient in reducing average hedging costs and does not alter the tail risk profile when compared to the non-self-financing counterpart.

In summary, both super-replicating strategies exhibit similar performance characteristics in terms of initial cost, expected hedging costs, and extreme risk measures, indicating the robustness of each strategy. The differences or improvements between the two strategies are minimal, indicating that the self-financing component does not drastically enhance the overall risk management capabilities of the super-replication strategy. Ultimately, the choice between self-financing and non-self-financing may come down to considerations such as operational ease or specific market conditions rather than a substantial difference in risk management performance.

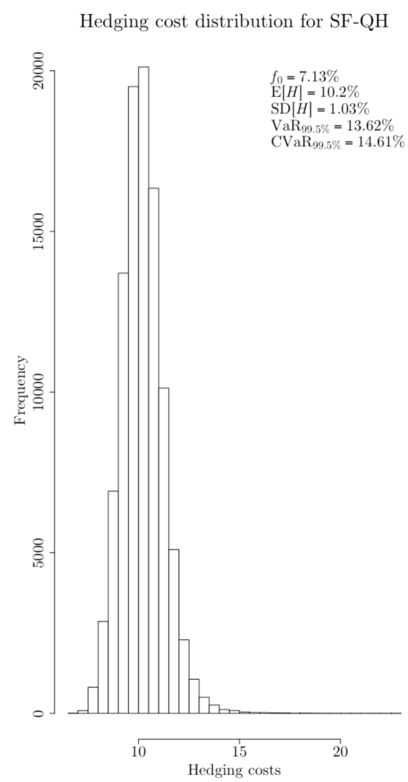

Figure 5.2: Distribution of simulated percent hedging costs for an at-the-money call option with self-financing $\ell_{2,1}$ strategy.

Figures 5.2 and 5.3 present histograms showing the distribution of hedging costs for four different self-financing hedging strategies - quadratic hedging, stochastic programming with $\ell_{1,1}$, $\ell_{2,1}$ and $\ell_{\infty,1}$ norms. These histograms provide insights into how different strategies manage hedging costs. The self-financing quadratic hedging histogram shows a narrow and peaked distribution, with most of the hedging costs clustered tightly around the expected cost of 10.20%. The low standard deviation (1.04%) indicates that the hedging costs are fairly stable, with only slight deviations from the mean. This suggests that SF-QH provides a highly predictable and relatively low-cost hedging solution. However, when looking at extreme risks, the VaR and CVaR values (13.62% and 14.61%, respectively) show that while most of the outcomes are concentrated around the mean, there is still a noticeable right tail in the distribution, signifying the presence of some higher-cost outcomes.

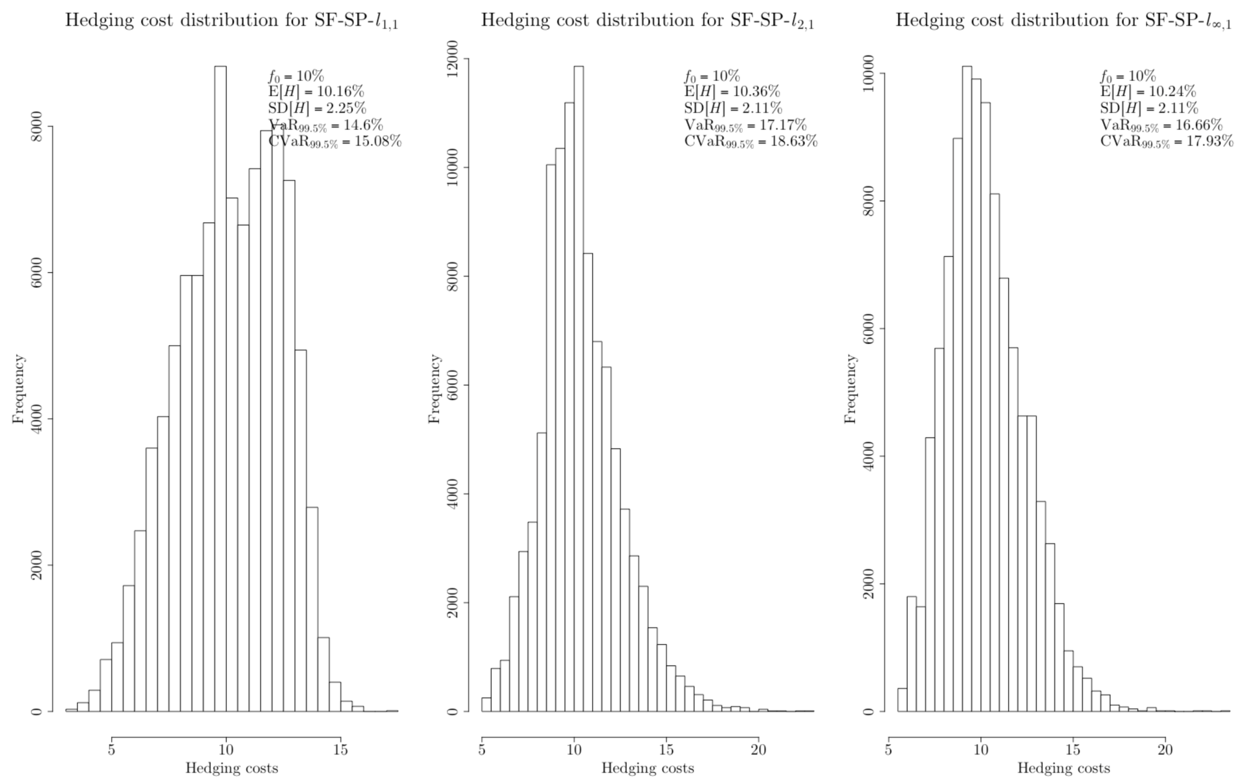

Figure 5.3: Distribution of simulated percent hedging costs for an at-the-money call option with self-financing stochastic programming strategy.

The self-financing stochastic programming with $\ell_{1,1}$ norm strategy results in a wider and less skewed distribution compared to quadratic hedging. The expected hedging cost (10.16%) is slightly lower than that of SF-QH, but the standard deviation is larger at 2.25%, indicating much greater variability in the hedging costs. This wider spread suggests that self-financing stochastic programming with $\ell_{1,1}$ norm may be less predictable and more volatile than quadratic hedging, with a notable chance of experiencing higher hedging costs. The VaR and CVaR values (14.60% and 15.08%) are relatively higher, reflecting the heavy tail risk in this strategy. However, when compared to its local hedging counterparts (NO-$\ell_{1,1}$, SP-$\ell_{1,1}$ and DC-$\ell_{1,1}$), the self-financing stochastic programming with $\ell_{1,1}$ norm remarkably has the lowest CVaR with comparable expected hedging cost. With the initial setup cost set to 0.1, the potential for lower tail outcomes on the distribution suggests that the self-financing approach may be better suited for investors willing to tolerate more volatility in exchange for potentially lower tail risk.

The distribution for self-financing stochastic programming with $\ell_{2,1}$ and $\ell_{\infty,1}$ norms are also shifted to the left like the quadratic hedging strategy, for a fixed initial setup cost of 0.1. Despite this higher initial cost, the expected hedging costs (10.36% and 10.24% respectively) are only slightly higher than the quadratic hedging strategy. This indicates that while the initial investment is higher, the resulting expected hedging costs are closer to the initial cost. The VaR and CVaR values (17.17% and 18.63%) are much lower than those of SP-$\ell_{2,1}$ and DC-$\ell_{2,1}$ norm strategies, suggesting that self-financing stochastic programming with $\ell_{2,1}$ norm also provides better tail-risk protection compared to the local hedging $\ell_{2,1}$ norm. The self-financing stochastic programming with $\ell_{2,1}$ and $\ell_{\infty,1}$ norm distributions are wider, and the tails are more pronounced, meaning that the strategies may be strong alternatives to their robust local hedging counterparts for risk management.

The choice of hedging strategy depends largely on the investor's risk tolerance and cost preferences. The SF-QH strategy strikes a balance between cost efficiency and robustness, offering low expected hedging costs with moderate risk protection. The SF-SP strategies can be cost-effective upfront but introduce considerable volatility and risk, making them suitable for risk-tolerant investors. Also, the SF-SP-$\ell_{1,1}$ norm strategy, with its relatively higher expected cost than the SF-QH strategy, delivers strong risk management and stability, appealing to hedgers who prioritize long-term robustness over short-term cost savings.

5.3.2 Sensitivity to number of trades and nodes for SF strategies

Next, we estimate the initial value of the portfolio f_0, $\mathbb{E}[\mathcal{H}]$ and the $\text{CVaR}_{99.5\%}$ for a varying number of trading periods and the number of nodes per branch.

		$\Delta = 6$			$\Delta = 12$			$\Delta = 24$		
		$N = 10$	$N = 20$	$N = 30$	$N = 10$	$N = 20$	$N = 30$	$N = 10$	$N = 20$	$N = 30$
SR	f_0	13.58%	15.05%	15.84%	14.04%	15.97%	17.15%	14.24%	16.37%	17.75%
	$\mathbb{E}[\mathcal{H}]$	10.16%	10.15%	10.17%	10.15%	10.15%	10.17%	10.17%	10.13%	10.17%
	$\text{CVaR}_{99.5\%}$	13.58%	15.02%	15.80%	13.23%	14.43%	15.36%	13.11%	14.37%	15.17%
SF-SR	f_0	13.58%	15.05%	15.84%	14.04%	15.97%	17.15%	14.24%	16.37%	17.75%
	$\mathbb{E}[\mathcal{H}]$	10.14%	10.19%	10.12%	10.14%	10.15%	10.19%	10.16%	10.15%	10.16%
	$\text{CVaR}_{99.5\%}$	13.58%	15.02%	15.80%	13.15%	14.45%	15.45%	13.08%	14.40%	15.17%
SF-QH	f_0	9.58%	7.74%	5.89%	9.52%	7.13%	4.35%	9.49%	6.62%	2.72%
	$\mathbb{E}[\mathcal{H}]$	10.12%	10.20%	10.35%	10.11%	10.20%	10.39%	10.10%	10.21%	10.48%
	$\text{CVaR}_{99.5\%}$	16.03%	16.20%	16.32%	14.04%	14.61%	15.21%	12.79%	13.94%	15.44%
SF-SP-$\ell_{1,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.13%	10.19%	10.20%	10.14%	10.16%	10.21%	10.14%	10.13%	10.14%
	$\text{CVaR}_{99.5\%}$	15.77%	16.65%	17.50%	13.95%	15.08%	15.94%	13.80%	14.76%	15.26%
SF-SP-$\ell_{2,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.32%	10.46%	10.52%	10.24%	10.36%	10.53%	10.18%	10.22%	10.33%
	$\text{CVaR}_{99.5\%}$	23.93%	22.23%	20.73%	23.20%	18.63%	18.52%	21.34%	18.51%	17.38%
SF-SP-$\ell_{\infty,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.08%	10.27%	10.17%	10.34%	10.24%	10.13%	10.18%	10.17%	10.17%
	$\text{CVaR}_{99.5\%}$	27.88%	18.92%	18.77%	23.96%	17.93%	15.81%	18.69%	16.47%	14.38%

Table 5.1: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for local and self-financing strategies with $\Delta = 6, 12, 24$ number of trades and $N + 1 = 11, 21, 31$ number of nodes.

Table 5.1 provides a detailed comparison of different hedging strategies—super-replication (SR), self-financing super-replication (SF-SR), self-financing quadratic hedging (SF-QH), and various forms of self-financing stochastic programming (SF-SP) under different norms ($\ell_{1,1}$, $\ell_{2,1}$, $\ell_{\infty,1}$). The comparison is across initial costs (f_0), expected hedging costs ($\mathbb{E}[\mathcal{H}]$), and Conditional Value at Risk at the 99.5% confidence level ($\text{CVaR}_{99.5\%}$), with varying re-balancing frequencies (Δ values of 6, 12, and 24 trades) and tree granularity (node counts $N = 10$, $N = 20$, and $N = 30$).

The self-financing super-replicating strategy follows the same pattern as the local super-replicating strategy, with the initial cost rising with the number of trades and nodes. This reflects the strategy's more conservative nature, allocating sufficient capital to ensure replication of payoff. Both local and global super-replicating strategies yield expected hedging costs around 10.12% to 10.17%, relatively stable across different numbers of trade and node configurations, indicating these strategies' robustness and consistency. Their CVaR values increase modestly with a higher number of trades and node counts, showing a risk management benefit. Both maintain a lower tail risk than self-financing stochastic programming under specific norms, reflecting a reliable approach to minimizing extreme losses. These findings imply that while the super-replicating strategies in general offer robust performance with consistent results across different investment outlooks, the differences between self-financing and non-self-financing are minimal, and the choice between these strategies may depend on other external factors.

The self-financing quadratic hedging strategy has the lowest initial costs among the strategies considered. At $\Delta = 24$ and $N = 30$, the initial cost drops to 2.72%, compared to 5.89% at $\Delta = 6$. This suggests that SF-QH is a cost-efficient strategy, as it requires minimal upfront capital while adapting flexibly to dynamic markets. Despite lower initial costs, the self-financing quadratic hedging strategy maintains expected costs comparable to the super replicating, ranging from 10.10% to 10.48%, signifying effective allocation of resources to achieve reliable performance within reasonable bounds. While generally higher than super replication, the tail risk for self-financing quadratic hedging strategy declines with higher node counts and trading frequency, with CVaR values falling from 16.03% to around 12.79%. This highlights the strategy's capacity to mitigate risk with higher adaptability in trading intervals and structure granularity.

The self-financing stochastic programming strategies ($\ell_{1,1}$, $\ell_{2,1}$, $\ell_{\infty,1}$) has initial cost fixed at 10%. The expected hedging costs generally hover around 10.08% to 10.53%, suggesting these strategies are cost-effective in achieving the desired outcomes without significant cost fluctuation across different lattice configurations. The tail risk values for $\ell_{1,1}$ norm show a slight decline with lower Δ and higher node counts, indicating better risk containment with frequent number of trades and increased lattice granularity. The $\ell_{2,1}$ norm variant experiences a significant drop in CVaR with higher nodes, reducing from 23.93% (for $N = 10$) to 17.38% (for $N = 30$), implying that a larger tree structure enhances risk containment. The $\ell_{\infty,1}$ norm shows a marked tail risk improvement from 27.88% to 14.38% as Δ increases and node count grows, suggesting that it effectively mitigates extreme losses with a highly granular trading structure.

The self-financing quadratic hedging strategy stands out as a cost-efficient strategy with comparatively low f_0, making it suitable for scenarios where capital preservation is prioritized, though it entails relatively higher risk. Super replicating strategies, with their stability in tail risk, suit investors focusing on robust risk containment. However, self-financing stochastic programming strategies, particularly under the $\ell_{\infty,1}$ norm, demonstrate more favourable risk management with higher nodes and trading frequency, supporting a more flexible adaptation to volatile conditions. For conservative investors prioritizing lower tail risk, super replication provides consistent results. SF-SP-$\ell_{2,1}$ and SF-SP-$\ell_{\infty,1}$ norms, meanwhile, offer adaptive alternatives with effective cost containment and declining CVaR as the number of trade and structural granularity increase, balancing efficiency with resilience against extreme losses.

5.3.3 Sensitivity to market parameters for SF strategies

We also consider the effectiveness of the self-financing strategies to changes in interest rates and market volatility.

In financial markets, hedging strategies are crucial for managing risk, particularly when dealing with derivatives such as options. Various market conditions, including interest rates and market volatility, can influence the effectiveness of these strategies. Table 5.2 offers a

		$r = 2\%$			$r = 4\%$			$r = 8\%$		
		$\sigma = 15\%$	$\sigma = 25\%$	$\sigma = 30\%$	$\sigma = 15\%$	$\sigma = 25\%$	$\sigma = 30\%$	$\sigma = 15\%$	$\sigma = 25\%$	$\sigma = 30\%$
SR	f_0	11.59%	18.52%	21.94%	12.54%	19.37%	22.75%	14.56%	21.12%	24.40%
	$\mathbb{E}[\mathcal{H}]$	7.23%	11.09%	13.01%	8.25%	12.06%	13.97%	10.40%	14.06%	15.92%
	$\text{CVaR}_{99.5\%}$	10.43%	16.64%	19.74%	11.37%	17.49%	20.54%	13.41%	19.24%	22.19%
SF-SR	f_0	11.59%	18.52%	21.94%	12.54%	19.37%	22.75%	14.56%	21.12%	24.40%
	$\mathbb{E}[\mathcal{H}]$	7.25%	11.09%	13.04%	8.26%	12.07%	14.00%	10.42%	14.06%	15.95%
	$\text{CVaR}_{99.5\%}$	10.44%	16.66%	19.75%	11.38%	17.51%	20.55%	13.42%	19.28%	22.22%
SF-QH	f_0	4.80%	7.29%	8.55%	5.92%	8.35%	9.59%	8.40%	10.60%	11.77%
	$\mathbb{E}[\mathcal{H}]$	7.34%	11.09%	12.96%	8.36%	12.07%	13.94%	10.49%	14.08%	15.91%
	$\text{CVaR}_{99.5\%}$	11.50%	16.76%	19.62%	11.86%	17.45%	20.35%	13.56%	19.18%	22.04%
SF-SP-$\ell_{1,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	7.10%	11.09%	13.06%	8.19%	12.06%	14.01%	10.38%	14.05%	15.95%
	$\text{CVaR}_{99.5\%}$	10.31%	17.76%	21.47%	11.75%	18.81%	22.44%	13.94%	21.00%	24.57%
SF-SP-$\ell_{2,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	7.04%	11.18%	13.05%	8.15%	12.40%	14.15%	11.39%	15.35%	17.03%
	$\text{CVaR}_{99.5\%}$	10.06%	23.56%	34.01%	11.83%	29.57%	40.96%	22.28%	46.70%	56.60%
SF-SP-$\ell_{\infty,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	7.05%	11.02%	12.94%	8.20%	12.14%	13.84%	10.99%	15.44%	16.35%
	$\text{CVaR}_{99.5\%}$	10.81%	18.74%	23.64%	12.12%	25.20%	29.66%	24.42%	56.30%	54.42%

Table 5.2: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for local and self-financing strategies with interest rate $r = 2\%, 4\%, 8\%$, and volatility $\sigma = 15\%, 25\%, 30\%$.

comparative analysis of the self-financing strategies for different interest rates and market volatility levels. The table indicates a clear trend in the initial portfolio value as both interest rate and volatility increase; the initial cost also rises across the super-replication and quadratic hedging strategies.

The super-replicating strategies both show relatively high initial costs across different interest rates and volatilities, which increase with higher volatilities and interest rates. The expected hedging cost and $\text{CVaR}_{99.5\%}$ values also rise in response to higher market risk, particularly with increasing volatility. This indicates that both super-replicating strategies are relatively conservative approaches, providing substantial coverage for risk at a correspondingly high initial cost. The similarity between local and global super-replicating values suggests that self-financing constraints do not alter the risk and return profile of the super-replicating strategy in this context.

The self-financing quadratic hedging strategy demonstrates lower initial costs compared to the super-replicating strategies, with initial values ranging from 4.8% to 11.77%, depending on the volatility and interest rate combination. This reduction in f_0 suggests SF-QH is a less conservative strategy, potentially leaving the portfolio more exposed to adverse market movements. While the expected hedging cost $\mathbb{E}[\mathcal{H}]$ is similar to super-replication, SF-QH shows slightly lower $\text{CVaR}_{99.5\%}$ values, indicating a narrower range of potential extreme losses compared to the super-replication strategies.

The self-financing stochastic programming strategies explore different norm configurations ($\ell_{1,1}$, $\ell_{2,1}$, and $\ell_{\infty,1}$) and are set with a default initial cost of 10% to demonstrate a controlled cost structure. However, doing this causes these strategies to exhibit varying levels of risk tolerance. The $\text{CVaR}_{99.5\%}$ values for SF-SP-$\ell_{1,1}$ strategy are moderate, increasing with higher volatility and interest rates, but remaining more controlled compared to the SF-SP-$\ell_{2,1}$ and SF-SP-$\ell_{\infty,1}$ strategies. This strategy shows a balanced approach, with an expected cost close to that of the super-replicating strategies but with slightly lower extreme risk. On the other hand, the $\ell_{2,1}$ norm exhibits higher $\text{CVaR}_{99.5\%}$ values, particularly at elevated interest rates and volatilities, reflecting higher potential losses in extreme market conditions. It appears that this strategy is more exposed to tail risk, indicating a less conservative hedging profile than SF-SP-$\ell_{1,1}$. The SF-SP-$\ell_{\infty,1}$ strategy presents the highest $\text{CVaR}_{99.5\%}$ values in high-volatility and high-interest scenarios. The increase in risk exposure suggests that SF-SP-$\ell_{\infty,1}$ may be suited for portfolios with a greater tolerance for extreme losses, making it a less conservative approach in stable market conditions. We highlight that using the stochastic programming strategies requires the hedger to calibrate for the minimum initial cost of the portfolio every time a parameter is adjusted.

Table 5.2 reveals that super-replicating strategies are the most conservative, with consistently higher initial costs and lower exposure to extreme tail risks. In contrast, SF-QH provides a lower initial cost and maintains a slightly lower CVaR, making it a viable choice for scenarios requiring moderate hedging. The SF-SP strategies, with their fixed initial cost values, offer cost-effective alternatives with varying levels of risk exposure. SF-SP-$\ell_{1,1}$ stands out as a balanced approach, while SF-SP-$\ell_{2,1}$ and SF-SP-$\ell_{\infty,1}$ exhibit higher tail risks, particularly in volatile markets.

5.3.4 Sensitivity to hedger's preference and moneyness for SF strategies

We also assess the robustness of the self-financing strategies to changes in the hedger's preference such as the number of assets that constitute the portfolio, the moneyness of the call option and the level of asymmetry in the $\ell_{p,q}$ norm.

Table 5.3 presents a comparative analysis of the self-financing strategies under different scenarios: adding multiple options to the portfolio, changing the strike price of the option to reflect moneyness and changing the asymmetric parameter to reflect various risk levels. The initial cost is a critical metric that reflects the upfront capital required to establish a hedging portfolio. The table compares results for hedging portfolios containing either one or three options. The number of options in the portfolio affects both the initial setup cost (f_0) and the associated risk ($\text{CVaR}_{99.5\%}$) for each strategy. With local super-replication, the initial cost slightly decreases from 15.97% to 14.64% when moving from one to three options. This decrease also applies to global super-replication and the self-financing constraint shows a similar variation in expected cost and CVaR values. This suggests that diversification within the hedging portfolio (i.e., using multiple options) reduces the initial cost of setting up the portfolio. However, SF-QH shows an increase in initial cost from 7.13% to 8.95% when moving to three options, reflecting an increase in complexity. Although still relatively low in cost compared to super-replication, the expected cost remains similar across option counts, maintaining an economical profile with a moderate increase in CVaR. The stochastic programming strategies using $\ell_{1,1}$, $\ell_{2,1}$, and $\ell_{\infty,1}$ norms have a set initial cost of 10%, which

		Number of options		Moneyness			Asymmetry		
		1	3	ITM	ATM	OTM	$q = 0.7$	$q = 0.85$	$q = 1$
SR	f_0	15.97%	14.64%	15.97%	15.97%	15.97%	15.97%	15.97%	15.97%
	$\mathbb{E}[\mathcal{H}]$	10.15%	10.05%	10.15%	10.15%	10.15%	10.15%	10.15%	10.15%
	$\text{CVaR}_{99.5\%}$	14.43%	13.32%	14.43%	14.43%	14.43%	14.43%	14.43%	14.43%
SF-SR	f_0	15.97%	14.64%	18.90%	15.97%	14.39%	15.97%	15.97%	15.97%
	$\mathbb{E}[\mathcal{H}]$	10.15%	10.05%	12.71%	10.15%	8.08%	10.15%	10.15%	10.15%
	$\text{CVaR}_{99.5\%}$	14.45%	13.30%	16.89%	14.45%	12.52%	14.45%	14.45%	14.45%
SF-QH	f_0	7.13%	8.95%	10.27%	7.13%	4.56%	7.13%	7.13%	7.13%
	$\mathbb{E}[\mathcal{H}]$	10.20%	10.01%	13.08%	10.20%	7.78%	10.20%	10.20%	10.20%
	$\text{CVaR}_{99.5\%}$	14.61%	14.23%	17.41%	14.61%	12.18%	14.61%	14.61%	14.61%
SF-SP-$\ell_{1,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.16%	10.06%	12.72%	10.16%	7.86%	10.13%	10.13%	10.16%
	$\text{CVaR}_{99.5\%}$	15.08%	13.53%	18.16%	15.08%	13.10%	15.16%	15.10%	15.08%
SF-SP-$\ell_{2,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.36%	10.03%	13.83%	10.36%	7.69%	10.42%	10.37%	10.36%
	$\text{CVaR}_{99.5\%}$	18.63%	18.77%	38.80%	18.63%	12.99%	19.61%	18.99%	18.63%
SF-SP-$\ell_{\infty,1}$	f_0	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%	10.00%
	$\mathbb{E}[\mathcal{H}]$	10.24%	10.12%	12.92%	10.24%	7.71%	10.49%	10.26%	10.24%
	$\text{CVaR}_{99.5\%}$	17.93%	22.08%	36.94%	17.93%	13.17%	27.78%	19.96%	17.93%

Table 5.3: Initial value f_0, expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for local and self-financing strategies with 1 call option with $\kappa = 1$, and 3 call options with $\kappa = 0.8, 1, 1.2$ as assets to the hedging portfolio, ITM ($K = 0.95$), ATM ($K = 1$) and OTM ($K = 1.05$) call options, and asymmetric parameter $q = 0.7, 0.85, 1$.

is independent of the number of options. However, risk profiles diverge, with SF-SP-$\ell_{\infty,1}$ showing a higher CVaR increase (from 17.93% to 22.08%) when moving to three options, reflecting heightened risk exposure under more complex portfolios. SF-SP-$\ell_{1,1}$ and $\ell_{2,1}$ see smaller CVaR increases, offering more controlled risk expansion relative to SF-SP-$\ell_{\infty,1}$. Moneyness plays a critical role in the performance and cost of each strategy. The super-replicating strategies exhibit stability across moneyness conditions, with local super-replication maintaining a consistent f_0 and CVaR at 15.97% and 14.43% respectively. Self-financing super-replication demonstrates slightly higher variability, with a higher f_0 of 18.9% for ITM options and a reduced cost of 14.39% for OTM options. The CVaR values also vary with moneyness in global super-replication, indicating that in-the-money positions increase exposure to extreme losses. The self-financing quadratic hedging shows increased sensitivity to moneyness, especially with ITM options, where the initial cost rises to 10.27% (compared to 7.13% for ATM) and the CVaR reaches 17.41%. This suggests that while SF-QH remains cost-effective overall, it becomes more expensive with ITM options due to the need for higher capital reserves. Moneyness affects CVaR for the self-financing stochastic programming strategies. SF-SP-$\ell_{2,1}$ is especially impacted, with ITM options resulting in a high CVaR of 38.80%, indicating a notable risk increase when options are ITM. SF-SP-$\ell_{1,1}$ and $\ell_{\infty,1}$ demonstrate similar patterns, though the increase in CVaR is more moderate, peaking at 18.16% and 38.8% for ITM options, respectively. Thus, self-financing stochastic programming strategies tend to carry greater risk with ITM options, particularly with the $\ell_{2,1}$ norm.

Both super-replicating strategies and quadratic hedging strategies are unaffected by changes in norm asymmetry. The choice of the asymmetric parameter has a marked effect on self-financing stochastic programming strategies, especially SF-SP-$\ell_{\infty,1}$. For instance, when $q = 0.7$, the CVaR increases to 27.78%, indicating heightened tail risk under asymmetry. This trend underscores that while self-financing stochastic programming strategies flexible, they are more vulnerable and sensitive to asymmetry changes.

Examining these three conditions highlights key trade-offs in choosing a strategy. The number of options primarily impacts initial costs and risk in SF-QH and SF-SP strategies, especially with increased CVaR for SF-SP-$\ell_{\infty,1}$. Moneyness particularly affects SF-QH and SF-SP, where ITM options substantially increase costs and CVaR. Finally, asymmetry exerts its strongest influence on SF-SP strategies, most noticeably SF-SP-$\ell_{\infty,1}$, where increased asymmetry corresponds with heightened risk. In summary, super-replicating strategies demonstrate high resilience across conditions, with both local and global configurations being the most stable but also the most costly. SF-QH offers cost-efficiency with moderate risk sensitivity, while SF-SP strategies present the most flexibility but require careful management due to sensitivity to option characteristics and asymmetry.

5.3.5 Sensitivity to initial portfolio values for SF-SP strategies

Similar to the local stochastic programming and dynamic coherent risk strategies, the initial portfolio value for the global stochastic programming strategies requires tuning to obtain a lower exposure to tail risk. In this subsection, we continue our analysis by comparing the performance of different initial costs of setting up the portfolio and their impact on expected hedging cost and tail risk for $\Delta = 6$ number of trades.

		$f_0 = 5\%$	$f_0 = 7\%$	$f_0 = 10\%$	$f_0 = 13\%$	$f_0 = 15\%$	$f_0 = 18\%$	$f_0 = 20\%$
SP-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.12%	10.13%	10.13%	10.13%	10.14%	10.14%	10.16%
	$\text{CVaR}_{99.5\%}$	18.51%	17.97%	17.17%	16.56%	16.33%	16.65%	16.86%
SP-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	9.95%	10.09%	10.00%	10.07%	9.98%	10.01%	10.30%
	$\text{CVaR}_{99.5\%}$	22.81%	21.69%	19.96%	16.55%	16.79%	17.31%	22.97%
SP-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.16%	10.16%	10.19%	10.15%	10.17%	10.16%	10.08%
	$\text{CVaR}_{99.5\%}$	15.05%	15.05%	15.32%	16.00%	16.11%	18.04%	19.71%
SF-SP-$\ell_{1,1}$	$\mathbb{E}[\mathcal{H}]$	10.16%	10.15%	10.19%	10.07%	10.05%	10.08%	10.09%
	$\text{CVaR}_{99.5\%}$	18.04%	17.63%	16.65%	15.15%	15.69%	16.38%	16.73%
SF-SP-$\ell_{2,1}$	$\mathbb{E}[\mathcal{H}]$	12.04%	10.99%	10.46%	10.16%	10.06%	10.05%	10.07%
	$\text{CVaR}_{99.5\%}$	43.23%	33.90%	22.23%	16.53%	16.38%	16.54%	16.62%
SF-SP-$\ell_{\infty,1}$	$\mathbb{E}[\mathcal{H}]$	10.72%	10.65%	10.27%	10.19%	10.20%	10.20%	10.21%
	$\text{CVaR}_{99.5\%}$	31.83%	28.40%	18.92%	16.55%	17.04%	18.59%	20.04%

Table 5.4: Expected value $\mathbb{E}[\mathcal{H}]$ and $\text{CVaR}_{99.5\%}$ of simulated hedging costs for at-the-money call option for SP and SF-SP-$(\ell_{1,1}, \ell_{2,1}, \ell_{\infty,1})$ strategies with initial investment value $f_0 = 5\%, 7\%, 10\%, 13\%, 15\%, 18\%, 20\%$.

Table 5.4 explores the self-financing stochastic programming (SF-SP) compared to the local stochastic programming (SP) strategies, aiming to calibrate the initial portfolio setup cost

(f_0) that minimizes Conditional Value at Risk across three norms: $\ell_{1,1}$, $\ell_{2,1}$, and $\ell_{\infty,1}$. We examine the expected value of hedging costs and CVaR at the 99.5% confidence level to evaluate the trade-offs in risk and expected cost across both self-financing and non-self-financing strategies. As illustrated earlier in Table 4.4, the non-self-financing stochastic programming strategies show moderate adjustments in both expected costs and CVaR with changes in f_0, which reflects the capital allocated upfront for hedging. For SP strategies, increasing f_0 generally reduces CVaR, suggesting a direct benefit in risk reduction as initial investment increases. For example, in the SP-$\ell_{1,1}$ strategy, CVaR reduces steadily from 18.51% to 16.33% as f_0 increases from 5% to 15%, from which it increases at higher values of f_0. Similarly, SP-$\ell_{2,1}$ experiences a significant CVaR reduction as f_0 increases from 5% to 13%, achieving a minimum CVaR of 16.55% and increasing thereafter. The $\ell_{\infty,1}$ norm, however, displays a constrained sensitivity to initial cost increases, limiting its risk-reduction capability in high-cost scenarios.

The self-financing stochastic programming strategies introduce more variability in outcomes, with both expected costs and CVaR showing different behaviours across norms as f_0 changes. For instance, SF-SP-$\ell_{1,1}$ exhibits a consistent decrease in CVaR as initial cost increases from 5% to 13%, achieving a minimum CVaR of 15.15% at $f_0 = 13\%$ and then increasing slightly thereafter. Interestingly, the expected cost also decreases with a higher initial portfolio value, indicating improved cost efficiency alongside tail risk reduction for this norm. The SF-SP-$\ell_{2,1}$ and SF-SP-$\ell_{\infty,1}$ norms behave differently. For SF-SP-$\ell_{2,1}$, a substantial CVaR reduction occurs between 5% and 15%, plummeting from 43.23% to 16.38%. This shows a significant improvement in risk management with moderate initial costs but less sensitivity to further increases, as CVaR stabilizes around 16.5% for higher f_0. The SF-SP-$\ell_{\infty,1}$ norm follows a similar pattern, with CVaR declining dramatically and then increasing at higher costs, highlighting a limited capacity for further risk control.

Across norms, self-financing strategies tend to exhibit higher initial CVaR levels compared to non-self-financing counterparts, especially at lower initial portfolio values. However, they show rapid CVaR reductions with moderate increases in initial cost, often stabilizing or even slightly increasing CVaR at higher f_0, suggesting high sensitivity to lower investment that tapers off as more capital is used to set up the portfolio. The local stochastic programming strategies, while displaying less extreme CVaR reductions, demonstrate more stable cost-to-risk relationships across all norms, indicating cost efficiency that makes it advantageous for controlled-risk scenarios without high initial costs. This stability can be attractive for hedgers who prioritize consistency and lower cost sensitivity. Given that the lowest CVaR for the stochastic programming strategies in Table 5.4 is around $f_0 \leq 15\%$ when calibrated, the strategies provide strong evidence for comparable risk management properties that can rival that of super-replication in terms of robustness. For instance, the SF-SP-$\ell_{1,1}$ strategy obtains a minimum tail risk of 15.15% for $f_0 = 13\%$ whereas the super replicating strategies in Table 5.1 for $\Delta = 6$ show CVaR values around 15.02% for a corresponding 15.05% initial portfolio value.

The analysis of Table 5.4 also reveals important distinctions in risk and cost management across the self-financing and non-self-financing stochastic programming strategies for different norms. Global stochastic programming, particularly SF-SP-$\ell_{1,1}$, are cost-sensitive, as they quickly reduce CVaR with moderate increases in f_0. This efficiency makes them suitable for scenarios requiring rapid risk reduction but with budget constraints. Non-self-financing stochastic programming strategies show consistent CVaR reduction, especially in SP-$\ell_{\infty,1}$, where increased f_0 provides limited but steady risk management. This makes local stochastic programming strategies attractive for stable hedging portfolios where cost sensitivity needs control. In summary, SF-SP strategies offer flexibility with higher initial risk but quick improvements as f_0 rises, while local SP strategies provide steady, reliable CVaR reductions at higher initial costs, with the $\ell_{1,1}$ norm being the most cost-effective for risk management across both strategy types.

Conclusion

In this thesis, we sought to develop cost-effective and robust hedging strategies that minimize risk exposure for contingent claims. The research is motivated by the need to address the limitations of traditional hedging strategies, particularly in the presence of market incompleteness and transaction costs. We begin by developing a model framework based on a discrete process with a finite number of realizations. The structure is defined as an event tree, which is convenient for stochastic and deterministic dynamic optimization. The tree consists of nodes representing discrete process realizations at each period and branches indicating transitions between nodes. Our framework assumes a recombining tree to reduce computational complexity. We introduce the $\ell_{p,q}$ norm as a convex risk measure, with modifications to include asymmetric parameters to expand the flexibility and robustness when hedging. The choice of risk measure and norm depends on the investor's risk tolerance and the desired level of conservatism.

For our proposed models, the super-replicating strategy ensures that the hedge portfolio dominates the claim at each node, making it a conservative approach. We provide a detailed optimization problem to find the least-cost super-replicating portfolio. The strategy is shown to be effective but expensive, as it requires the portfolio to cover potential losses fully. On the other hand, the $\ell_{p,q}$ norm as constraint strategy allows for some degree of positive losses by controlling the loss function with a threshold parameter. Different norms, such as $\ell_{1,q}$, $\ell_{2,q}$, and $\ell_{\infty,q}$, are used to limit the losses. The optimization problem is formulated to minimize the hedge portfolio value subject to the loss constraint, providing a more cost-effective alternative to super-replication.

In the $\ell_{p,q}$ norm as objective strategy, the objective function directly minimizes the norm. The goal is to obtain a hedging strategy that minimizes the $\ell_{p,q}$ norm of the losses. The optimization problem includes constraints to control conservatism and ensure feasibility. The portfolio as state variable strategy introduces state variables representing future portfolio values and minimizes the sum of local losses and future cost-to-go functions. Using state variables in the optimization process enables dynamic adjustments to the hedging strategy, enhancing its robustness and adaptability. We propose three variations: stochastic programming, dynamic coherent risk, and barrier on future risk. Each variation aims to balance current and future risks, providing a dynamic and robust hedging strategy.

Finally, we introduce another approach to self-financing strategies, where the portfolio is adjusted over time without additional capital injections. A self-financing super-replicating strategy is developed, ensuring the portfolio value supersedes the contingent claim at maturity. The optimization problem is formulated to find the least-cost self-financing strategy. We extend the concept of self-financing to stochastic programming strategies which is the same as dynamic coherent risk under self-financing conditions and compare the robustness of the self-financing strategies to their local counterpart and ℓ_2 or quadratic hedging strategy with equal transition probabilities.

To assess the performance of our models, we apply the hedging strategies to European call options. Our financial model assumes a uni-dimensional discrete lattice where the stock process evolves. Detailed numerical examples are provided to compare the performance of different hedging strategies. Key metrics such as initial portfolio value, expected hedging error, and CVaR are analyzed. The results show that strategies with controlled losses offer a cost-effective alternative to super-replication while maintaining similar tail risk profiles. The proposed strategies also exhibit robust performance under various market conditions, with consistent expected hedging errors and CVaR values.

By comparing the performance of local and global hedging strategies, our results indicate that they offer marginal improvements in expected hedging errors and tail risk, making them a viable alternative in certain market conditions without transaction costs. Our results also affirm the suitability of minimizing the mean square error under self-financing conditions. We observe that the widely used self-financing quadratic hedging technique offers a cost-efficient alternative while maintaining reasonable tail risk levels, making it attractive for cost-conscious hedgers.

In summary, each hedging strategy displays unique strengths. Both local and global super-replicating strategies are reliable and conservative. Norm as constraint strategies offer robust cost-effective alternatives to super-replication whereas the norm as objective strategies may appeal to hedgers interested in lower cost of capital. Self-financing quadratic hedging strategy is cost-effective but riskier, and portfolio as state variable strategies under specific norms excel with increased granularity, presenting a flexible and dynamic solution for managing risk and cost in volatile markets. This versatility enables hedgers to select strategies based on investment objectives, capital constraints, and risk tolerance, fostering a nuanced approach to portfolio optimization and risk management.

Recommendations for Future Work

The findings of this thesis have significant implications for the design and implementation of hedging strategies in incomplete financial markets. Introducing the self-financing strategy further enhances the practical applicability of the proposed methods. Future research on semi-robust risk-minimizing strategies could extend our proposed strategies to American-style products, considering the interactions between different stopping times and their impact on hedging performance. Also, applying the proposed strategies to real-world financial products and markets could validate their effectiveness and identify potential areas for improvement. However, developing more efficient algorithms for solving optimization problems, particularly for large-scale portfolios, could enhance the practical applicability of the strategies. Lastly, investigating the impact of different behavioural factors such as utility functions on the choice of hedging strategies could provide valuable insights for personalized risk-minimizing solutions. The semi-robust risk-minimizing strategies proposed are model-independent and adaptable to different discrete filtrations and model frameworks with possible extensions to non-linear optimization techniques.

Bibliography

Beatrice Acciaio and Irina Penner. Dynamic risk measures. *Advanced Mathematical Methods for Finance*, pages 1–34, 2011.

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

Chaithanya Bandi and Dimitris Bertsimas. Robust option pricing. *European Journal of Operational Research*, 239(3):842–853, 2014.

Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations Research*, 52(1): 35–53, 2004.

Dimitris Bertsimas, Leonid Kogan, and Andrew W Lo. Hedging derivative securities and incomplete markets: An ϵ-arbitrage approach. *Operations Research*, 49(3):372–397, 2001.

Phelim P. Boyle. A lattice framework for option pricing with two state variables. *Journal of Financial and Quantitative Analysis*, 23(1):1–12, 1988.

Phelim P. Boyle and David Emanuel. Options on the general mean. *University of British Columbia, Canada, Working paper*, 1980.

Alexandre Carbonneau and Frédéric Godin. Equal risk pricing of derivatives with deep hedging. *Quantitative Finance*, 21(4):593–608, 2021.

Sebastián Ceria and Robert A Stubbs. Incorporating estimation errors into portfolio selection: Robust portfolio construction. *Journal of Asset Management*, 7(2):109–127, 2006.

Guan-Yu Chen, Ken Palmer, and Yuan-Chung Sheu. The least cost super replicating portfolio in the boyle–vorst model with transaction costs. *International Journal of Theoretical and Applied Finance*, 11(01):55–85, 2008.

Xinliang Chen, Griselda Deelstra, Jan Dhaene, Daniël Linders, and Michèle Vanmaele. On an optimization problem related to static super-replicating strategies. *Journal of Computational and Applied Mathematics*, 278:213–230, 2015.

Thomas F. Coleman, Dmitriy Levchenkov, and Yuying Li. Discrete hedging of american-type options using local risk minimization. *Journal of Banking & Finance*, 31(11):3398–3419, 2007.

Jurie Conradie. Asymmetric norms, cones and partial orders. *Topology and its Applications*, 193:100–115, 2015.

John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of financial Economics*, 7(3):229–263, 1979.

Mark Herbert Ainsworth Davis and J Martin C Clark. A note on super-replicating strategies. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 347(1684):485–494, 1994.

Yan Dolinsky and Ariel Neufeld. Super-replication in fully incomplete markets. *Mathematical Finance*, 28(2):483–515, 2018.

Laurent El Ghaoui, Francois Oustry, and Hervé Lebret. Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, 9(1):33–52, 1998.

Hans Föllmer and Alexander Schied. *Stochastic Finance: An Introduction in Discrete Time.* Walter de Gruyter, 2011.

Hans Föllmer and Martin Schweizer. Hedging by sequential regression: An introduction to the mathematics of option trading. *Astin Bulletin*, 2(147-160), 1988.

Raquel J. Fonseca and Berç Rustem. Robust hedging strategies. *Computers & operations research*, 39(11):2528–2536, 2012.

Patrice Gaillardetz and Saeb Hachem. Dynamic hedging in incomplete markets using risk measures. *IMA Journal of Management Mathematics*, 33(2):345–367, 2022.

Patrice Gaillardetz and Mehran Moghtadai. Partial hedging for equity-linked products using risk-minimizing strategies. *North American Actuarial Journal*, 21(4):580–593, 2017.

Patrice Gaillardetz and Emmanuel Osei-Mireku. Worst-case valuation of equity-linked products using risk-minimizing strategies. *North American Actuarial Journal*, 26(1):64–81, 2022.

Lluis M. García-Raffi, Salvador Romaguera, and Enrique A. Sánchez-Pérez. Sequence spaces and asymmetric norms in the theory of computational complexity. *Mathematical and Computer Modelling*, 36(1-2):1–11, 2002.

Bartosz Jaroszkowski and Max Jensen. Valuation of european options under an uncertain market price of volatility risk. *Applied Mathematical Finance*, 29(3):213–226, 2022.

Damien Lamberton, Huyên Pham, and Martin Schweizer. Local risk-minimization under transaction costs. *Mathematics of Operations Research*, 23(3):585–612, 1998.

Andreas Martin Lisewski and Olivier Lichtarge. Untangling complex networks: risk minimization in financial markets through accessible spin glass ground states. *Physica A: Statistical Mechanics and its Applications*, 389(16):3250–3253, 2010.

Pengfei Liu, Wenguo Yang, and Tiande Guo. A discussion on the conservatism of robust linear optimization problems. *Optimization*, 65(8):1641–1650, 2016.

Somayyeh Lotfi and Stavros A. Zenios. Robust var and cvar optimization under joint ambiguity in distributions, means, and covariances. *European Journal of Operational Research*, 269(2):556–576, 2018.

Antje Mahayni. Effectiveness of hedging strategies under model misspecification and trading restrictions. *International Journal of Theoretical and Applied Finance*, 6(05):521–552, 2003.

Jiří Matoušek and Bernd Gärtner. *Understanding and using linear programming*, volume 1. Springer, 2007.

Thomas Møller. Risk-minimizing hedging strategies for unit-linked life insurance contracts. *ASTIN Bulletin: The Journal of the IAA*, 28(1):17–47, 1998.

Stanley Pliska. *Introduction to mathematical finance*. Blackwell publishers Oxford, 1997.

Bruno Rémillard and Sylvain Rubenthaler. Optimal hedging in discrete time. *Quantitative Finance*, 13(6):819–825, 2013.

Bruno Rémillard, Alexandre Hocquard, Hugues Langlois, and Nicolas Papageorgiou. Optimal hedging of american options in discrete time. In *Numerical Methods in Finance*, pages 145–170. Springer, 2012.

Frank Riedel. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112:185–200, 2004.

Ralph T. Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.

Ernst Roos and Dick den Hertog. Reducing conservatism in robust optimization. *INFORMS Journal on Computing*, 32(4):1109–1127, 2020.

Birgit Rudloff. Coherent hedging in incomplete markets. *Quantitative Finance*, 9(2):197–206, 2009.

Martin Schweizer. Variance-optimal hedging in discrete time. *Mathematics of Operations Research*, 20(1):1–32, 1995.

Martin Schweizer. Local risk-minimization for multidimensional assets and payment streams. *Banach Cent. Publ*, 83:213–229, 2008.

Maria Grazia Scutella and Raffaella Recchia. Robust portfolio asset allocation and risk measures. *Annals of Operations Research*, 204(1):145–169, 2013.

Halil M. Soner. Super-replication via dynamic programming and backward stochastic differential equations. *URL https://scholar.princeton.edu/sites/default/files/soner/files/super-replication-under-constraints.pdf*, 23/01/2023, 2008.

Seçil Sözüer and Aurélie C. Thiele. The state of robust optimization. *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 89–112, 2016.

Jeremy Staum. Incomplete markets. *Handbooks in Operations Research and Management Science*, 15:511–563, 2007.

Gerald L. Thompson, Fred M. Tonge, and Stanley Zionts. Techniques for removing nonbinding constraints and extraneous variables from linear programming problems. *Management Science*, 12(7):588–608, 1966.

Appendix A

Farka's Lemma

We introduce the variant of the popular *Farka's Lemma* used to establish the existence of a solution to convex linear optimization. The application of the lemma is used when we develop our self-financing strategies in Chapter 5.

Proposition A.1. ***(Farka's Lemma)*** *Let A be a real matrix with m rows and n columns, and let $b \in \mathbb{R}^m$ be a vector. Then exactly one of the following two possibilities occurs:*

$$1. \text{There exists a vector } x \in \mathbb{R}^n \text{ satisfying } Ax = b \text{ and } x \geq 0. \tag{A.1}$$

$$2. \text{There exists a vector } y \in \mathbb{R}^m \text{ such that } y^\top A \geq 0^\top \text{ and } y^\top b < 0. \tag{A.2}$$

Proof. We provide a brief proof of Farka's Lemma. We show that both (A.1) and (A.2) do not hold simultaneously. Note that $y^\top Ax = y^\top(Ax) = y^\top b < 0$ since by (A.1), $Ax = b$ and by (A.2) $y^\top b < 0$. But also $y^\top Ax = (y^\top A)x = (A^\top y)^\top x \geq 0$ since by (A.2) $A^\top y \geq 0$ and by (A.1) $x \geq 0$. This implies $y^\top Ax < 0$ and $y^\top Ax \geq 0$ at the same time. Hence, a contradiction and, thus, exactly one of the two equations is consistent. □

Several variants of Farka's Lemma answer questions about the feasibility of a linear equations and inequalities system. For example "When is there a non-negative solution for a system of linear equations?", "When is there a non-negative solution for a system of linear inequalities?" and lastly, "When is there ever a solution for a system of linear inequalities?". Concerning the latter question, Matoušek and Gärtner (2007) extends Farka's Lemma and proposes the following proposition.

Proposition A.2. *The system $Ax \leq b$ has a solution if every non-negative $y \in \mathbb{R}^m$ with $y^\top A = 0^\top$ also satisfies $y^\top b \geq 0$.*

Proof. Let $Ax \leq b$ for $x \in \mathbb{R}^n$. Then suppose $\exists\, y \geq 0$ such that $y^\top A = 0^\top$ but $y^\top b < 0$. Then $y^\top Ax = y^\top(Ax) \leq y^\top b < 0$ but $y^\top Ax = (y^\top A)x = 0$ hence a contradiction. Conversely, suppose $y \geq 0, y^\top A = 0^\top$ and $y^\top b \geq 0$. Then $y^\top Ax = (y^\top A)x = 0 \leq y^\top b$ and thus $Ax \leq b$. $\square$