A Global Time-Consistent CVaR Hedging Strategy

Shima Jalili* Patrice Gaillardetz†

Abstract

We study the evaluation of financial derivatives through a discrete-time hedging approach. We propose a global multistage hedging framework based on Conditional Value-at-Risk that also ensures time consistency. Our strategy builds on the temporal decomposition of coherent risk measures introduced by Pflug and Pichler [2016]. To efficiently solve the resulting hedging problem, we employ backward dynamic programming algorithms enhanced with parametric techniques to reduce computational complexity and mitigate the curse of dimensionality. The versatility of the proposed time-consistent strategy is illustrated through numerical experiments involving European call option.

Keywords: Discrete time; Time-Consistency; Global-Hedging; Conditional Value-at-Risk; Decomposition-Theorem,

*shima Jalili, Ph.D. e-mail: shima.jalili@mail.concordia.ca

†Patrice Gaillardetz, Professor at Concordia University, e-mail: patrice.gaillardetz@concordia.ca

1 Introduction

Risk minimization using coherent risk measures has become a cornerstone of modern portfolio management and efficient hedging strategies. The seminal work by Artzner et al. [1999] introduced the concept of coherent risk measures, highlighting the properties of subadditivity, translation invariance, positive homogeneity, and monotonicity as essential for consistent risk assessments. Among these measures, Conditional Value-at-Risk (CVaR) has emerged as a particularly influential choice to capture tail risk more effectively than traditional Value-at-Risk (VaR).

The application of these measures is especially critical for derivative pricing and hedging in incomplete markets, where perfect replication is not possible. In discrete-time settings, the groundwork for risk-minimization was laid by Föllmer and Sondermann [1985], who introduced hedging strategies that minimize the squared hedging error. This quadratic approach was further developed through recursive computation methods Föllmer and Schweizer [1988] and a comprehensive solution to the global mean-variance hedging problem Schweizer [1995]. Building on these ideas, subsequent research has shifted focus from variance to tail risk, employing coherent risk measures to better capture extreme losses. Xu [2006], for instance, uses convex risk measures to introduce the notion of risk efficiency. More recent work by Gaillardetz and Hachem [2022] utilizes coherent risk measures and stochastic programming to derive optimal hedging strategies under broader risk preferences.

The literature on dynamic risk measures is extensive, with many authors contributing important insights. For instance, Riedel [2004] analyzed dynamic coherent risk measures, underscoring the importance of reconciling immediate risk assessments with future information. In a related vein, Detlefsen and Scandolo [2005] investigated the conditional approach to convex risk measures, clarifying how nested relationships can be maintained. These theoretical works reinforce the foundational principles of time consistency that are central to multi-period risk management.

In multi-period settings, analysts have increasingly recognized that measuring risk solely at a single future date may lead to re-optimization pitfalls as information unfolds. This awareness has spurred a rich literature on dynamic or time-consistent risk measures, where the objective is to ensure that risk preferences at one stage do not contradict those formed at subsequent stages. The work of Shapiro [2012] provides a clear treatment of time consistency by examining how multi-stage stochastic optimization can be carried out when the risk measure respects a recursive property. Shapiro [2012] framework focuses on discrete-time, multi-step decision processes, showing that if a dynamic risk measure is properly nested or exhibits certain monotonicity relationships, then backward-induction algorithms will produce stable, non-contradictory solutions over time.

A parallel development has come from researcher Pflug and Pichler [2016], who have studied how a global coherent risk measure, for example, a final-horizon CVaR can be decomposed into stage-wise conditional components in a way that inherently preserves time consistency. Their "temporal decomposition" theorem ensures that one can break down a terminal risk functional into its intermediate conditional versions, yet recombine them without losing any information or altering the original preference ordering. This construction is especially relevant for multi-stage hedging problems where decision-makers want both a terminal loss criterion (such as CVaR at the end of an investment horizon) and stage-wise guidance for overbalancing or adjusting hedges. By guaranteeing a match between local (conditional) and global (unconditional) perspectives, Pflug and Pichler [2016]'s framework removes the risk of dynamic inconsistency, which can otherwise arise if naive conditioning leads to re-rankings of outcomes at intermediate stages.

In a more applied context, Godin [2016] has examined how to minimize CVaR in a global dynamic hedging problem, leveraging the foundational optimization framework of Rockafellar et al. [2000]. This "global" approach treats path of trading decisions up to maturity as a single problem, with the final CVaR capturing the tail risk of terminal losses. While this method effectively solves for a final-horizon objective at a specific level of risk preference, its dynamic strategy is governed only by basic self-financing and admissibility constraints, as a result, a strategy that is optimal from an initial viewpoint can lead the hedger into intermediate positions that become unacceptable as information unfolds. Hence, this exposes a key weakness of a purely global objective: it is susceptible to time-inconsistency, where initial risk preferences conflict with rational decisions at later stages.

This paper directly addresses this vulnerability by introducing a hedging framework that is time-consistent by design. We leverage the temporal decomposition theorem of Pflug and Pichler [2016] to construct a global CVaR objective from a recursive composition of conditional risk measures such that governed by self-financing and time-consistency constraints. This methodology imposes a coherent risk-management structure at every stage of the hedging horizon, ensuring that investors' preferences remain consistent as information evolves. The result is a dynamic hedging strategy where each intermediate decision is inherently aligned with the ultimate risk-minimization goal. Our approach eliminates the risk of adopting unacceptable intermediate positions and provides a theoretically sound and practically robust foundation for multi-period hedging under CVaR.

This paper is divided as follows. Section 2 describes the market in which the current work takes place and defines the hedging problem. In Section 3 details Coherent risk measure and its properties. In Section 4, we discuss the time-consistency and temporal decomposition theorem. In Section 5, we give a brief description of the global CVaR presented in Godin [2016], and then we apply the temporal decomposition theorem to obtain a type of global time-consistent. Section 6 illustrates a numerical example involving hedging of European call options. We conclude in the last section.

2 Financial Framework

This section outlines the mathematical framework for the discrete-time financial market and presents the theoretical setup for hedging. We analyze an arbitrage-free financial market without transaction costs. Let the time horizon in years be segmented into T successive dates $t = \{0, 1, 2, \ldots, T\}$. Let $(\Omega, \mathcal{F} = (\mathcal{F}_t)_{0 \le t \le T}, \mathbb{P})$, be the probability space where Ω is the set of all finite state space, $\mathcal{F}_t$ is a discrete filtration describing the information available at time t, and $\mathbb{P}$ is the physical measure. Let $S_t = (S_t^{(0)}, S_t^{(1)}, \ldots, S_t^{(d)})$ denote the asset prices process at time $t = 0, 1, \ldots, T$, where $S_t^{(y)}$ denotes the price of the asset y at time t for $y = 0, 1, \ldots, d$.

2.1 Trading Strategy

Hedging involves using an investment portfolio to offset the risk associated with a financial position. Hedging's main goal is to reduce or completely eliminate risk exposures by creating cash flows that are equal in magnitude but opposite in direction to those of the existing exposure. So, what follows are some definitions in the theory of hedging from the mathematical finance literature.

A trading strategy $\theta = (\theta_t)_{0 \le t \le T} = (\theta_t^{(0)}, \theta_t^{(1)}, \ldots, \theta_t^{(d)})_{0 \le t \le T}$, is an adapted stochastic process, where $\theta_t^{(0)}$, $\theta_t^{(1:d)} = (\theta_t^{(1)}, \ldots, \theta_t^{(d)})$ respectively denotes the portion invested in the risk-free asset and invested in the risky assets during the time interval $[t, t+1)$. The value of the portfolio is $\mathcal{V}_t = \sum_{y=0}^{d} \theta_t^{(y)} S_t^{(y)}$. Let the risk-free asset be $S_t^{(0)} = e^{rt}$, where r is the annualized risk-free rate. Strategy θ is self-financing if at time t, once the new price process $S_t = (S_t^{(0)}, \ldots, S_t^{(d)})$ is quoted, the investor adjusts their portfolio positions from θ_t to θ_{t+1} without cash infusions into or withdrawals from the portfolio; i.e, for $t = 0, 1, \ldots, T-2$:

$$\theta_t S_{t+1}^\top = \theta_{t+1} S_{t+1}^\top, \tag{2.1}$$

where $\top$ denotes the transpose operator. One of the characterization of self-financing is to write $\mathcal{V}_t$ for all $t = 1, 2, \ldots, T$ as:

$$\mathcal{V}_t = e^{rt}\mathcal{V}_0 + \sum_{k=0}^{t-1} \theta_k (S_{k+1} - S_k e^r)^\top e^{r(t-k-1)}. \tag{2.2}$$

3 Coherent Risk Measure

In incomplete market, eliminating the risk with a perfect hedging is not guaranteed, and partial hedging will be adopted and some residual risk at expiration will be tolerated. In this situation, an approach in the optimal hedging portfolio is to minimize the convex risk measure. The theory of risk measures in mathematical finance has become mainstream. Coherent risk measures were first introduced and discussed by Artzner et al. [1999]. Let $\mathcal{X}$ be a set of random variables and $X^{(1)}, X^{(2)} \in \mathcal{X}$. A risk measure $\rho : \mathcal{X} \to \mathbb{R}$ is coherent if it satisfies the following properties:

(i) Monotonicity: if $X^{(1)} \leq X^{(2)}$ then $\rho(X^{(1)}) \leq \rho(X^{(2)})$,

(i) Translation Invariance: $\rho(X^{(1)} + c) = \rho(X^{(1)}) - c, \quad \forall c \in \mathbb{R}$,

(iii) Convexity: $\rho(\lambda X^{(1)} + (1-\lambda)X^{(2)}) \leq \lambda\rho(X^{(1)}) + (1-\lambda)\rho(X^{(2)}), \quad \forall \lambda \in [0, 1]$.

Furthermore, due to the translation invariance property of ρ, optimal hedging strategies remain unchanged regardless of the initial portfolio value. Therefore, for optimal hedging strategies $\theta^{*(1:d)}$ we have

$$\begin{aligned} \theta^* &= \arg\min_\theta [\rho(\Phi(S_T^{(1)}) - (e^{rT}\mathcal{V}_0 + \sum_{k=0}^{T-1} \theta_k (S_{k+1} - S_k e^r)^\top e^{r(T-k-1)}))], \\ &= \arg\min_\theta [\rho(\Phi(S_T^{(1)}) - \sum_{k=0}^{T-1} \theta_k (S_{k+1} - S_k e^r)^\top e^{r(T-k-1)}]. \end{aligned} \tag{3.3}$$

In this study, we focus on the Conditional Value-at-Risk (CVaR) which is formally defined for $\alpha \in (0, 1)$. Rockafellar et al. [2000] propose the definition of CVaR as as the solution of an optimization problem

$$\mathrm{CVaR}_\alpha(X) = \min_\pi (\pi + \frac{1}{1-\alpha} E[(X-\pi)^+]), \tag{3.4}$$

where $(x)^+ = \max\{0, x\}$, and π is associated with the value-at-risk at level α

$$\text{VaR}_\alpha(X) = \min_x P(X \le x) \ge \alpha. \tag{3.5}$$

(3.4) shows a key structural property of CVaR which is belonging to the class of polyhedral risk measures, making it a polyhedral function (piecewise linear). As a result, this structure is computationally advantageous because it allows the convex optimization problem to be precisely reformulated as a tractable linear program. Moreover, Artzner et al. [1999] and Föllmer and Schied [2011] show that the duality representation of any convex coherent risk measure can be described as the maximum expected random variables X over a set of dual variables. Therefore, the dual representation proposed by Pflug and Pichler [2016]. At confidence level α is

$$\text{CVaR}_\alpha(X) = \sup_{Z \in \mathcal{Z}} \mathbb{E}[ZX], \tag{3.6}$$

where $\mathcal{Z}$ is the set of adapted processes in dual space in which CVaR admits its Fenchel–Moreau (conjugate) representation; see Pflug and Pichler [2016]. Furthermore, $\mathbb{E}(Z) = 1$, satisfying the additional constraint $0 \le (1-\alpha)Z \le 1$.

4 Time-consistency and Temporal Decomposition of Risk Measures

Time-consistency has become a critical concept in modern risk management, particularly in dynamic settings where decisions and evaluations evolve over time. In continuous time, Delbaen [2006] first investigated time-consistent coherent risk measures. Later, Cheridito et al. [2006] extended this analysis to the convex setting within a discrete-time framework. Jobert and Rogers [2008] demonstrated that time-consistent valuations can be constructed through backward induction of one-period static risk measures. In their work, they highlighted how static risk measures can be glued together to provide dynamic and time-consistent assessments.

The core idea behind time consistency is that once a partial decision or preference ordering is fixed at an earlier time, it should remain valid at all subsequent times in the absence of contradictory new information. In Bielecki et al. [2017], a dynamic convex risk measure $(\rho(X|\mathcal{F}_t))_{0\le t\le T}$ is defined as a time-consistent risk measure if the following conditions hold.

(i) $\rho(X^{(1)}|\mathcal{F}_{t+1}) \ge \rho(X^{(2)}|\mathcal{F}_{t+1}) \Longrightarrow \rho(X^{(1)}|\mathcal{F}_t) \ge \rho(X^{(2)}|\mathcal{F}_t), \quad t = 0, 1, 2, \dots, T-1,$ (4.7)

(ii) $\rho(X^{(1)}|\mathcal{F}_t) = \rho((-\rho(X^{(1)}|\mathcal{F}_{t+s}))|\mathcal{F}_t), \quad \forall \quad t, s \geq 0,$ (4.8)

where (4.8) considered a recursive property of the time-consistent risk measure. These properties provide a stable and logical way to measure risk over time. As a result, decisions based on that risk measurement such as portfolio optimization, capital allocation, or insurance reserve management are also consistent and reliable. When combined with backward algorithms (e.g., dynamic programming), they enable consistent policies over an entire decision horizon.

A time-consistent strategy, as discussed in Boda and Filar [2006] adheres to the principle of optimality from dynamic programming (see Bellman and Dreyfus [1962]).

Definition 4.1. *A hedging strategy $(\theta_t^*)_{0\leq t\leq T-1}$ considered time-consistent if it exhibits the following properties:*

(i) if the strategy θ_t^* at each time step t, $t = 0, \ldots, T-1$ is chosen by

$$\theta_t^* \in \underset{\theta_t, \theta^*_{t+1:T-1}}{\operatorname{argmin}} \rho\big(X|\mathcal{F}_t\big)\,, \tag{4.9}$$

then the strategy θ_t^* will be the optimal strategy in the problem

$$\min_\theta \rho(X), \tag{4.10}$$

(ii) if the strategy

$$\theta^* \in \arg\min_\theta \rho\big(X\big), \tag{4.11}$$

it also satisfies

$$\theta^*_{t:T-1} \in \underset{\theta_{t:T-1}}{\operatorname{argmin}} \rho\big(X|\mathcal{F}_t\big)\,, \qquad \forall\, t = 1, \ldots, T-1. \tag{4.12}$$

In the class of multistage risk measures, to obtain time consistency, nested conditional risk measures were proposed by Ruszczyński and Shapiro [2006], Shapiro [2009] and Ruszczyński [2010]. Another approach is based on a risk measure that is dynamically adapted according to available information, as proposed by Pflug and Pichler [2014]. The temporal decomposition theorem (see Theorem 21 of Pflug and Pichler [2014]) shows how the risk measure, conditional on $\mathcal{F}_t$, can be reassembled to represent the risk measure at the initial time, and needs to be adjusted based on the information revealed, and the decision has to be changed over time to meet the optimal decision. The temporal decomposition theorem for the risk measure ρ states that

$$\rho(X) = \sup_{Z_t\in\mathcal{Z}} \mathbb{E}\left[Z_t \rho\left(X \mid \mathcal{F}_t\right)\right], \tag{4.13}$$

where random variables Z_t are adapted to the information, satisfying $\mathbb{E}(Z_t|\mathcal{F}_t) = 1$. Moreover, letting $\mathcal{F}_t \subset \mathcal{F}_s \subset \mathcal{F}$, the risk conditional functional obeys the nested decomposition and has the following recursive representation

$$\rho(X \mid \mathcal{F}_t) = \sup_{Z_s \in \mathcal{Z}} \mathbb{E}[Z_s \rho_{Z_s}(X \mid \mathcal{F}_s) \mid \mathcal{F}_t], \tag{4.14}$$

where Z_s is $\mathcal{F}_s$ -measurable and ρ_{Z_t} is the conditional risk function obtained from a basic risk function ρ through its conditional version given $\mathcal{F}_t$, so

$$\rho_{Z_t}(X|\mathcal{F}_t) = \sup_{Z' \in \mathcal{Z}} \left\{ \mathbb{E}\left(XZ' \mid \mathcal{F}_t\right) \right\}, \tag{4.15}$$

where $\mathbb{E}\left(Z' \mid \mathcal{F}_t\right) = 1$, and $\rho^*(Z_t Z') = 0$ which ρ^* is the conjugate of ρ which is defined as $\rho^*(Z) = \sup_{X \in \mathcal{X}} \mathbb{E}(XZ) - \rho(X)$. (4.13) and (4.14) capture the temporal or backward decomposition of a dynamic risk measure. Hence, in a time-consistent framework, we can build up the unconditional risk via an iterated supremum over conditional risks. Therefore according to (4.13), we can break down the terminal risk measure into conditional or stage-wise components at each stage in a multi-period model. Pflug and Pichler [2016] states the decomposition theorem for risk measure CVaR. So, for all $\alpha \in [0, 1]$

$$\mathrm{CVaR}_\alpha(X) = \sup_{Z_t \in \mathcal{Z}} \mathbb{E}\left[Z_t \cdot \mathrm{CVaR}_{1-(1-\alpha)Z_t}(X|\mathcal{F}_t)\right], \tag{4.16}$$

where the supremum is over all random variables Z_t which are $\mathcal{F}_t$-measurable, and satisfy $\mathbb{E}(Z_t) = 1$ and $0 \le (1-\alpha)Z_t \le 1$. Moreover, $\mathrm{CVaR}_{1-(1-\alpha)Z_t}(X|\mathcal{F}_t)$ is the conditional CVaR at the random risk level $1 - (1-\alpha)Z_t$. The constraints in the representation (4.16) are the same as for the unconditional dual representation (3.6), except that it depends not only on X and $\mathcal{F}$, but also on the dual random variable Z, which is additionally adapted to the new available information. Hence, according to (4.14) the nested decomposition of CVaR leading to a time consistent formulation:

$$\mathrm{CVaR}(X|\mathcal{F}_t) = \sup_{Z_s \in \mathcal{Z}} \mathbb{E}[Z_s \mathrm{CVaR}_{1-(1-\alpha)Z_s}(X \mid \mathcal{F}_s) \mid \mathcal{F}_t], \tag{4.17}$$

for $t < s$.

5 Optimal Global CVaR

5.1 The Global CVaR

Godin [2016] addresses the problem of constructing a dynamic hedging strategy that minimizes the CVaR of the terminal hedging error in an incomplete market with both transaction costs and non-Gaussian asset returns. The global CVaR_α optimization is equivalent to

$$\min_{\theta\in\Theta^{(1)}} e^{-rT}\mathrm{CVaR}_\alpha((\Phi(S_T^{(1)})-\mathcal{V}_T)), \tag{5.18}$$

where Φ is the option payoff, and $\Theta^{(1)}$ is the set of all self-financing strategies where $\mathcal{V}_T$ is given by (2.2). The problem stated as minimizing (5.18) is tackled by leveraging a key characterization of CVaR. Using the equivalent expression (3.4), the CVaR of a random variable X can be written as:

$$\mathrm{CVaR}_\alpha(X) = \min_\pi \mathbb{E}[f_{\pi,\alpha}^{(CVaR)}(X)] = \min_\pi \mathbb{E}[\pi + \frac{1}{1-\alpha}(X-\pi)\mathbb{I}_{\{X>\pi\}}]. \tag{5.19}$$

This reformulation, developed by Rockafellar et al. [2000] and also applied studies such as Boda and Filar [2006], transforms the complex task of minimizing CVaR into a more tractable joint minimization problem as follows:

$$\begin{aligned}\min_{\theta\in\Theta^{(1)}} \mathrm{CVaR}_\alpha(X) &= \min_{\theta\in\Theta^{(1)}}\min_\pi \mathbb{E}\left[f_{\pi,\alpha}^{(\mathrm{CVaR})}(X)\right]\\ &= \min_\pi \min_{\theta\in\Theta^{(1)}} \mathbb{E}\left[f_{\pi,\alpha}^{(\mathrm{CVaR})}(X)\right].\end{aligned} \tag{5.20}$$

François et al. [2012] show that the minimization of the expected auxiliary function for a fixed π can be solved using dynamic programming. The Bellman equation provides a backward-recursive scheme to compute the value function at each time step t, and simultaneously identifies the optimal trade that achieves this minimum. At $t=0$, this process yields the strategy that minimizes $\mathbb{E}[f_{\pi,\alpha}^{(CVaR)}(\Phi(S_T^{(1)})-\mathcal{V}_T)]$ for the chosen π. The minimum CVaR is attained at minimal value π^*. The persistence of optimality over time depends on the model specification such as asset dynamics, payoff structure, admissible trades, and thus is not guaranteed in general. Therefore, the time-consistency concept focuses on keeping the optimal strategy optimal over time which is different from the definition used in Shapiro [2012] and Kupper and Schachermayer [2009], which focus instead on a structural property of the risk measure itself.

5.2 A Global Time-Consistent CVaR

In Godin [2016], the global CVaR problem (5.18) is solved without any control beyond enforcing the self-financing constraint and the requirements that hedging positions be predictable and adhere to pre-defined admissibility criteria, rather than being governed by additional, explicit controls on the dynamic trading path. Consequently, the hedger can end up in positions that might be unacceptable later, despite being globally optimal from the initial viewpoint. However, time consistency guarantees that investors' preferences implied by a dynamic risk measure remain consistent over time, so that the evolution of information does not affect the original risk assessment or the optimal strategy. Herein, the aim is to solve the global CVaR at confidence level α when the hedging strategy is also time-consistent, that is satisfied in (4.7) and (4.8). Therefore, the optimization problem becomes

$$\min_{\theta\in\Theta^{(2)}} e^{-rT}\mathrm{CVaR}_\alpha((\Phi(S_T^{(1)})-\mathcal{V}_T)), \tag{5.21}$$

where $\Theta^{(2)}$ is the set of all self-financing strategies that satisfy the time-consistent condition. This time-consistent constraint requires that the optimal policy obeys the dynamic programming principle evaluating decisions backward for global optimality and forward to ensure that each subproblem remains consistent as new information unfolds. Consequently, decisions must be made in a nested, stage-by-stage structure rather than all at once. Hence, a multistage model with nested risk measures can guarantee that each decision uses exactly the information available at that point. Therefore, we apply the backward dynamic recursion to obtain optimal decisions. It begins at the maturity date and works backward to ensure that the decisions made at each step are optimal given the future outcomes and are consistent with the overall objective. The solution to such a multi-stage optimization problem, where decisions are made sequentially over time, is typically found using dynamic programming, which is based on relationships between optimal value functions (or cost-to-go function). Let $J_t(\alpha,\xi_t)$ be the backward cost-to-go function where the hedging portfolio value ξ_t and CVaR_α are the state variables. The backward process is initialized by the solution for the final time interval, from $T-1$ to T. The time-consistency condition requires that the strategy θ_{T-1} minimizes the CVaR of the final hedging error. Therefore, for all risk-aversion level $\alpha\in[0,1]$, the cost-to-go function is (5.21) is

$$J_{T-1}(\alpha,\xi_{T-1}) = \min_{\theta_{T-1}} e^{-r}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)})-\xi_T|\mathcal{F}_{T-1}), \tag{5.22}$$

where θ_{T-1} satisfies the self-financing condition $\xi_{T-1}=\theta_{T-1}S_{T-1}^\top$ and state equation $\xi_T=\xi_{T-1}+\theta_{T-1}(S_T-S_{T-1})^\top$. By the definition of the time-consistent condition in (4.11), the strategy θ_{T-1} belongs to the set of admissible strategies $\Theta^{(2)}$ since it minimizes the conditional risk in the final period.

For the time interval t to $t+1$ $(t=0,\dots,T-2)$, enforcing the time-consistent conditions for θ leads to

$$J_t(\alpha,\xi_t) = \min_{\theta_t,\theta_{t+1:T-1}\in\Theta^{(2)}} e^{-r(T-t)}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)})-\mathcal{V}_T|\mathcal{F}_t), \tag{5.23}$$

under the self-financing constraint $\xi_t = \theta_t S_t^\top$ and state equation $\xi_{t+1} = \xi_t + \theta_t(S_{t+1} - S_t)^\top$. Applying the temporal decomposition (4.17), we have

$$\begin{aligned} J_t(\alpha, \xi_t) &= \min_{\theta_t} e^{-r(T-t)} \sup_{Z_{t+1}\in\mathcal{Z}, \theta_{t+1:T-1}\in\Theta^{(2)}} \mathbb{E}[Z_{t+1}\mathrm{CVaR}_{1-(1-\alpha)Z_{t+1}}(\Phi(S_T^{(1)}) - \mathcal{V}_T|\mathcal{F}_{t+1})|\mathcal{F}_t] \\ &= \min_{\theta_t} e^{-r} \sup_{Z_{t+1}\in\mathcal{Z}} \mathbb{E}[Z_{t+1}J_{t+1}((1-(1-\alpha)Z_{t+1}), \xi_{t+1})|\mathcal{F}_t], \end{aligned} \tag{5.24}$$

where $E[Z_{t+1}|\mathcal{F}_t] = 1$ and $0 \le (1-\alpha)Z_{t+1} \le 1$.

Notably, we work with a restricted admissible set of strategies, denoted by $\Theta^{(1)}$, which are $\mathcal{F}_t$-adapted and self-financing. This is a more constrained set than the classical "for all policies" in Definition 4.1. Under these constraints, solving the problem backward in time yields a policy that is time-consistent and belongs to $\Theta^{(2)}$.

Proposition 1. *The cost-to-go function for the interval t and $t+1$ is $J_t(\alpha, \xi_t)$ can be written as an affine linear equation*

$$J_t(\alpha, \xi_t) = -\xi_t + \beta_{t,\alpha}, \tag{5.25}$$

for $t = 0, ..., T-1$ where $\beta_{t,\alpha}$ is the optimal intercept term.

Proof. The optimal hedging strategy is determined by minimizing CVaR for a given ξ_t. Hence we have

$$\begin{aligned} J_t(\alpha, \xi_t) &= \min_{\theta_{t:T-1}} e^{-r(T-t)}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)}) - \mathcal{V}_T|\mathcal{F}_t) \\ &= \min_{\theta_{t:T-1}} e^{-r(T-t)}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)}) - \sum_{k=t}^{T-1}\theta_k(S_{k+1} - S_k e^r)^\top e^{r(T-t-k)} - e^{r(T-t)}\xi_t|\mathcal{F}_t) \\ &= \min_{\theta_{t:T-1}} e^{-r(T-t)}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)}) - \sum_{k=t}^{T-1}\theta_k(S_{k+1} - S_k e^r)^\top e^{r(T-t-k)}|\mathcal{F}_t) - \xi_t, \end{aligned} \tag{5.26}$$

which is induced by transition invariance property. Setting

$$\beta_{T-1,\alpha} = \min_{\theta_{t:T-1}} e^{-r(T-t)}\mathrm{CVaR}_\alpha(\Phi(S_T^{(1)}) - \sum_{k=t}^{T-1}\theta_i(S_{k+1} - S_k e^r)^\top e^{r(T-t-k)}|\mathcal{F}_t), \tag{5.27}$$

leads to (5.25). □

Hence, the time-consistent problem is solved using linear programming the last period from $T-1$ to T, to determine the intercept $\beta_{T-1,\alpha}$ of the cost-to-go function $J_{T-1}(\alpha, \xi_{T-1})$. Then we propagate the result to the initial time to obtain the optimal hedging strategy using with non-linear programming along with interpolation and extrapolation for cost-to-go function. Therefore, the optimization reduces to determining only the intercept term β. Therefore, (5.25) significantly simplifies the dimensionality while still preserving optimality across all initial wealth levels.

Herein, we stress out that, due to the time inconsistency of the global CVaR, the problem solved by our algorithm is not equivalent to the global one solved by Godin [2016]; however, it can be considered as an alternative framework that demonstrates the behavior of the hedger followed in practice, as well as an approximation of the global problem.

6 Numerical Experiment

In this section, the performance of hedging procedures obtained by solving a global hedging problem (5.21) within a recombining lattice framework, considering two underlying asset price models: a binomial model and a jump-diffusion model. To investigate the impact on hedging, we will compare the distribution of the discounted value of hedging errors

$$\mathcal{V}_0 + e^{-rT}(\Phi(S_T^{(1)}) - \mathcal{V}_T),$$

incurred by hedging strategies. We consider one-year, at-the-money European call options with $\Phi(S_T^{(1)}) = \max(0, S_T^{(1)} - K)$. The option is hedged using the risk-free asset and the underlying risky asset ($d = 1$). To assess hedging performance, we conduct 25,000 simulations of the underlying asset price paths. The time-consistent hedging strategy used in these simulations is derived numerically by implementing a backward induction to solve the non-linear dynamic programming equation recursively. The principal challenge at each stage of this recursion is the evaluation of the expectation term in Equation (5.24). The supremum is obtained using a stochastic sampling approach. Specifically, the hit-and-run MCMC algorithm (see Smith [1984]) generates a large, representative sample from the linearly-constrained set $\mathcal{Z}$, transforming the supremum into a computationally feasible maximum over the sample. The algorithm first employs interpolation and extrapolation to approximate the future cost-to-go function J_{t+1}. Subsequently, a non-linear optimization is performed to find the optimal hedging strategy that minimizes the present value of hedging errors. This entire procedure is repeated for all nodes at all time steps to construct the complete strategy. By computing the mean, standard deviation of errors, and also VaR and CVaR at $\alpha = 95\%,\ 99\%$ to assess the tail performance.

6.1 Binomial Model

In this section, we assume that the price process has $(2m+1)$ distinct outcomes at each time step, and $S_t^{(1)}$ at time t, the time $t+h$ price at node j is governed by $d^j u^{2m-j} S_t^{(1)}$ for $j = 0, \cdots, 2m$, and the increasing and decreasing factors u and d are assumed such that the price process converges to the lognormal distribution underlying a continuous geometric Brownian motion with annual drift μ, volatility σ, and with n dates per year with evenly time interval of length $h = \frac{1}{n}$. In other words, $u = e^{\sigma\sqrt{h}}$ and $d = u^{-1}$. Let $p_{i+j,t+h|i,t}$ be the conditional probability to reach state $(i+j, t+h)$ from (i,t). Hence, at period t, the price process is distributed according to a binomial distribution given by:

$$p_{i+j,t+h|i,t} = \binom{2m}{j} \left(\frac{e^{\mu h/2m} - d}{u-d} \right)^{2m-j} \left(1 - \frac{e^{\mu h/2m} - d}{u-d} \right)^{j}, \tag{6.28}$$

for $j = 0, 1, \ldots, 2m$, $i = 0, 1, \ldots, \frac{t(2m)}{h}$, and $t = 0, h, 2h, \ldots, T-h$. For this model, the parameter values are derived through a moment-matching equivalent from the Merton Jump Diffusion model's parameters, consistent with those presented in Coleman et al. [2007]. These are shown in Table 1.

r	μ	σ	K
0.03	0.075	0.190	1

Table 1: Parameters of Binomial Model.

Figure 1 provides hedging error histograms for both hedging strategies within the binomial model. The global and the time-consistent strategies are obtained assuming $\alpha = 95\%$, 12 trading dates ($h = 1/12$), and 11 branches ($m = 5$). The legends show the initial portfolio value, the expected value, the standard deviation, $\mathrm{VaR}_{99\%}$, and the $\mathrm{CVaR}_{99\%}$ of the hedging errors. While both strategies yield a nearly identical expected hedging error, their risk metrics are fundamentally different. The TC approach is significantly more effective at limiting the average magnitude of the worst-case losses. It achieves a $\mathrm{CVaR}_{99\%}$ of 15.19%, which is notably lower than the global strategy's 16.44%. The lower value for the TC strategy means it provides better protection against extreme, costly errors. This superior tail risk management stems from its period-by-period optimization, which offers better control and prevents large deviations from accumulating over time. Furthermore, the TC strategy exhibits a lower standard deviation of hedging errors (2.15% vs. 2.37%), indicating more predictable and stable outcomes. This is visually confirmed in the Figure 1. The TC distribution

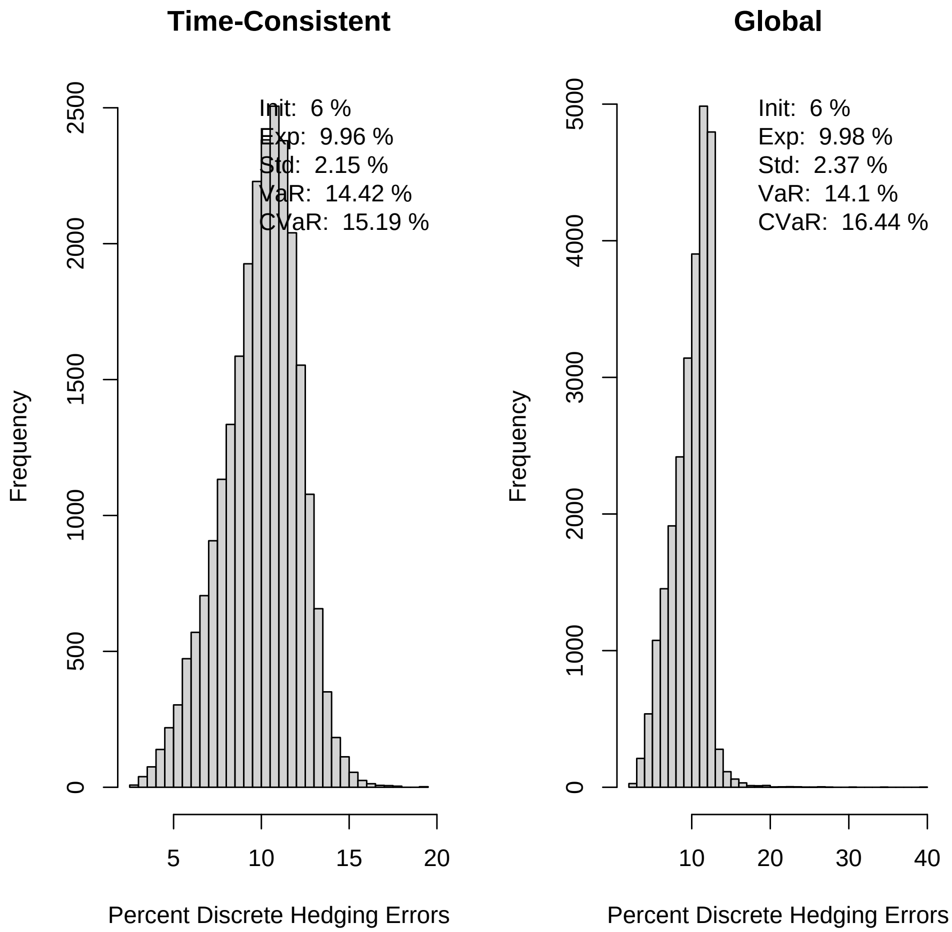

Figure 1: Distribution of the present value of the hedging errors under Binomial model for European call option.

is compact, with its errors almost entirely contained within a narrow range of 4% to 20%. In the contrast, the global distribution has a much wider range, with a long tail extending 40%. This visual difference leads to a critical insight: the global strategy presents a lower $\mathrm{VaR}_{99\%}$ of 14.1%, but a significantly higher $\mathrm{CVaR}_{99\%}$. This combination is the classic signature of a distribution with a fat tail.

6.2 Jump diffusion Model

In this subsection, we explain how to model a price process governed by jump diffusion process. Random jumps J_t are assumed to have a log-normal distribution, that is, $\ln(J_t) \overset{\text{i.i.d}}{\sim} N(\mu_J, \gamma_J^2)$, where μ_J

is the mean jump size and γ_J^2 is the variance of the jump size. Hence, the relative jump amplitude of S_t is log-normally distributed with the expected jump size $\kappa = \mathbb{E}^{\mathbb{P}}(J_t - 1) = e^{\mu_J + \frac{1}{2}(\gamma_J^2)} - 1$. According to Amin [1993], the discretized form of the continuous process of the jump diffusion model, in state j and time t, $t = 0, h, 2h \dots, T$, is constructed as follows:

$$S_{t,j}^{(1)} = S_0^{(1)} e^{(\alpha t + j\sigma\sqrt{h})}, \tag{6.29}$$

for $j = 0, \pm 1, \dots, \pm 2(mn)$, where

$$\alpha = \mu - \frac{1}{2}\sigma^2 - \lambda\kappa, \tag{6.30}$$

is the drift of the stock price, where λ is the jump intensity, and $j\sigma\sqrt{h}$ represents up or down movements of the diffusion process. The natural logarithmic scale of (6.29) is :

$$\ln(S_{t,j}^{(1)}) = \ln(S_0^{(1)}) + (\alpha t + j\sigma\sqrt{h}). \tag{6.31}$$

Restricted to a fixed time t, the state space is a grid such that consecutive points are spaced $\sigma\sqrt{h}$ apart.

Let l represent the set of all local and non-local states, for a non-local event, then under $\mathbb{P}$ measure the probability mass function assigned to each non-adjacent node is given by:

$$p(l) = N(\alpha h + (l + \frac{1}{2})\sigma\sqrt{h}) - N(\alpha h + (l - \frac{1}{2})\sigma\sqrt{h}), \qquad l \notin \{-1, 0, 1\}, \tag{6.32}$$

where $N(.)$ is the standard normal cumulative distribution function. The probabilities $p(l)$ correspond to a discretization of the continuous-time jump distribution. The remaining probability to be assigned lies in the interval $[\alpha h - (1 + \frac{1}{2})\sigma\sqrt{h}, \alpha h + (1 + \frac{1}{2})\sigma\sqrt{h}]$ for the points $l \in \{-1, 0, 1\}$. We assign this entire probability mass to the point corresponding to $l = 0$ and assume that $p(l) = 0$ for $l = \pm 1$. Therefore,

$$p(0) = N(\alpha h + (\frac{1}{2})\sigma\sqrt{h}) - N(\alpha h + (-\frac{1}{2})\sigma\sqrt{h}). \tag{6.33}$$

The probability of a jump occurrence (a rare event) in the discrete approximation at any period is λh, and multiple jumps cannot occur a single time interval. Therefore, the probability that the process does not jump is $1 - \lambda h$. Consequently, the transition probability mass can be assigned at each node due to jumps and local price changes and is given by

$$
\begin{aligned}
&\Pr(\ln(S^{(1)}_{t+h,j+1}) - \ln(S^{(1)}_{t,j}) = \alpha h + \sigma\sqrt{h}) = q(1-\lambda h),\\
&\Pr(\ln(S^{(1)}_{t+h,j-1}) - \ln(S^{(1)}_{t,j}) = \alpha h - \sigma\sqrt{h}) = (1-q)(1-\lambda h),\\
&\Pr(\ln(S^{(1)}_{t+h,l+j}) - \ln(S^{(1)}_{t,j}) = \alpha h + l\sigma\sqrt{h}) = p(l)\lambda h,
\end{aligned}
\tag{6.34}
$$

where $l \in \{-nm, \ldots, -2, 0, 2, \ldots, nm\}$. To finalize the discrete formulation and construct the lattice, we must now define the value of q. According to the assumption of deliverability presented in Amin [1993], the jump-size distribution under the martingale probability measure is the same as under the physical probability measure. In Amin [1993], the mean produced by the jump component is adjusted, i.e., Amin [1993] solves the probability with $qu + (1-q)d = \frac{e^{rh} - \lambda h E^{\mathbb{P}}[J_t]}{1-\lambda h}$, then by setting $u = e^{\alpha h + \sigma\sqrt{h}}$, and $d = e^{\alpha h - \sigma\sqrt{h}}$, q follows as:

$$
q = \frac{\frac{e^{rh} - \lambda h(\kappa+1)}{(1-\lambda h)} - e^{\alpha h - \sigma\sqrt{h}}}{e^{\alpha h + \sigma\sqrt{h}} - e^{\alpha h - \sigma\sqrt{h}}}. \tag{6.35}
$$

The parameter values for the jump model dynamics are the same as in the ones from Coleman et al. [2007], and are presented in Table 2. Figure 2 provides the hedging error distributions when the process follows the jump model under the same parameters ($\alpha = 95\%$, $h = 1/12$, $m = 5$). Like the binomial model, both strategies yield a nearly identical expected hedging error of approximately 9.8%. It is intuitively, hedging is inherently riskier when jumps are possible, and result in higher $\text{VaR}_{99\%}$ and $\text{CVaR}_{99\%}$ with respect to the binomial model for both approaches. Moreover, the range of errors is much broader for the global approach, with extreme outcomes (very negative and very positive errors) that rarely occur in the TC setting that vary between 5% to 30%. The TC approach, by contrast, produces a tighter, more symmetric, and more stable error distribution, avoiding extreme scenarios with $\text{CVaR}_{99\%} = 17.82\%$. For the global approach, the hedging error distribution is highly leptokurtic which is often described as a spike and fat tails shape. Therefore, the lower $\text{VaR}_{99\%} = 12.9\%$ can indeed be a product of the tall peak and the concentration of the outcomes in a very narrow range. The range of errors from below -20% to over 60% represents the long tails produced by the global approach. These tails, though low in frequency, represent rare but catastrophic hedging failures. They are caused by the strategy's inability to handle sudden, large market jumps. From Figures 1 and 2, we can conclude that the wider spread and longer tails of the hedging errors demonstrate that the global method produces a heavier-tailed error distribution, indicating a greater likelihood of extreme deviations. In contrast, the Time-Consistent method yields a portfolio with tighter error dispersion, making it more robust and efficient in managing hedging risk.

r	μ	σ	λ	μ_J	γ_J
0.03	0.08	0.2	1	0.02	0.1

Table 2: Parameters of the Merton jump-diffusion model.

6.3 Performance and Sensitivity Analysis

Sensitivity to number of trades n and branches m

The use of the lattice structure makes it necessary to study the sensitivity of different hedging statistics when the number of rebalancing and outcomes varies. Table 3 gives a sensitivity analysis of the performance of the time-consistent hedging strategy for different values of h and m based on the binomial model. It reports key performance metrics including expected value, standard deviation, VaR, and CVaR evaluated at two distinct confidence levels: the optimization level of 95% and the more extreme tail level of 99%. The most significant improvement in hedging performance is achieved by increasing the rebalancing frequency. For example, in the most uncertain market scenario ($m = 8$), doubling the number of trading dates h from 1/8 to 1/16 decreases the $\text{CVaR}_{99\%}$ by more than 2% (from 16.805% to 14.615%). This demonstrates that more frequent adjustments are powerfully effective at mitigating risk, especially when market complexity is high. Conversely, the negative impact of rising market uncertainty at each time step, which is least severe when hedging is already frequent. When trading often $h = 1/16$, doubling the number of market outcomes from $m = 4$ to $m = 8$ leads to a far smaller increase in $\text{CVaR}_{99\%}$ of just 0.25% (from 14.362% to 14.615%). Moreover, the expected value of hedging cost is approximately constant across all levels. In the subsequent numerical experiments, we fix the parameter values at $h = 1/12$ and $m = 5$.

Sensitivity to different risk levels α

Figure 3 demonstrates the initial value of objective functions for different thresholds ($\alpha \in (0.5, 0.99)$) for both hedging strategies under the jump and binomial models. It is evident that the approaches do not yield convergent results, and the discrepancy widens in the tail under the jump model. When comparing optimization methods, the global CVaR curves display varying steepness, whereas

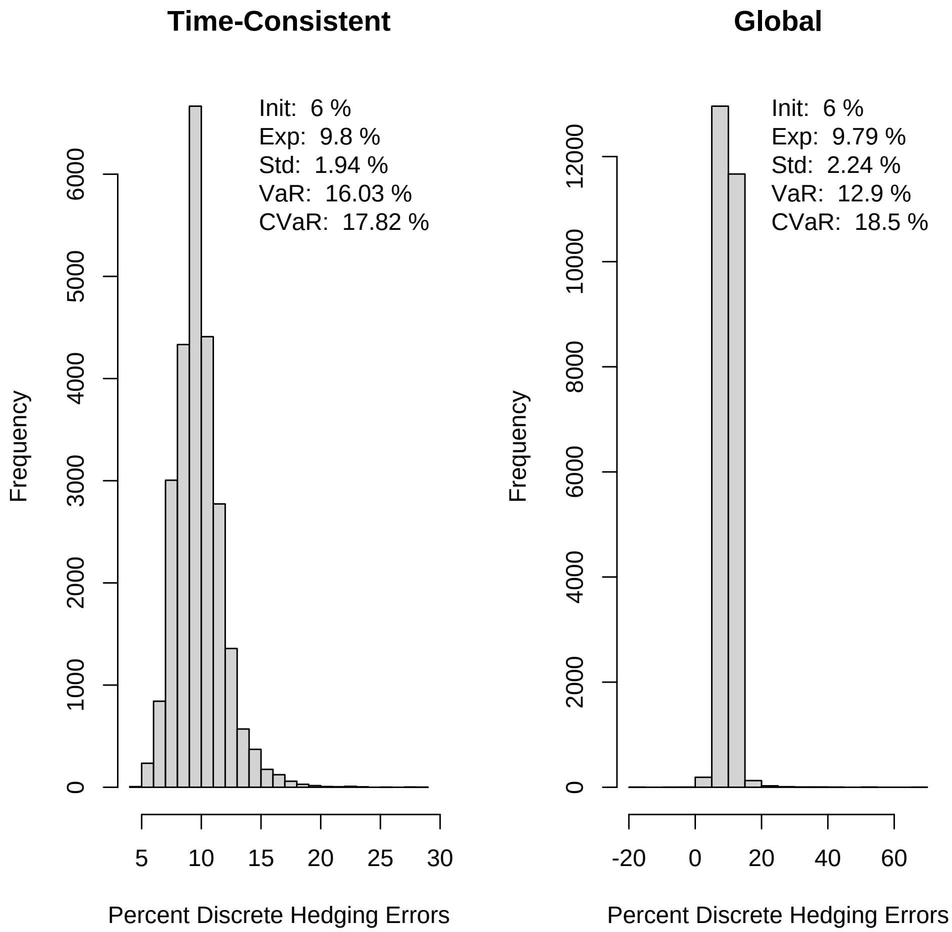

Figure 2: Distribution of the present value of the hedging errors under Jump diffusion model for European call option.

the TC slope changes smoothly across all confidence levels, indicating a more stable and reliable response to shifts in α. By construction, the global approach seeks the strategy that minimizes the hedging error CVaR at a chosen threshold. Consequently, the global strategy will, by definition, yield a lower hedging error CVaR than the time-consistent approach at the evaluation threshold. However, the TC approach produces a higher hedging error CVaR because it delivers a more robust hedge: it accounts for risk along each potential path and incorporates rare high-risk events, even at lower thresholds. In contrast, the global method may overlook these events due to their initially negligible probabilities.

Although the global strategy yields a lower hedging error CVaR, it may result in higher ob-

h	m	Exp.	std.	$\text{VaR}_{95\%}$	$\text{CVaR}_{95\%}$	$\text{VaR}_{99\%}$	$\text{CVaR}_{99\%}$
1/8	4	9.937	2.571	13.771	14.722	15.279	16.241
	5	9.964	2.576	13.868	14.891	15.492	16.546
	8	9.953	2.619	13.946	15.033	15.726	16.805
1/12	4	9.955	2.162	13.111	13.859	14.328	15.056
	5	9.956	2.151	13.103	13.903	14.416	15.192
	8	9.947	2.164	13.204	14.082	14.585	15.516
1/16	4	9.945	1.903	12.702	13.358	13.775	14.362
	5	9.996	1.924	12.748	13.465	13.910	14.612
	8	9.966	1.946	12.801	13.514	13.926	14.615

Table 3: Hedging errors statistics for the time-consistent hedging strategy for different values of trading dates ($h = 1/8, 1/12, 1/16$) and outcomes ($m = 4, 5, 8$).

jective function values. This occurs when the threshold is set to a higher level, such as 99%, the calculation admits previously negligible rare scenarios like large market jumps into its risk set. Because these events are so severe, their potential losses dominate and result in a very high objective function value. Consequently, achieving this high-threshold optimum requires a prohibitively expensive hedge, as significant capital is needed to cover these extreme potential deficiencies.

As the figure 3 shows this significant difference in higher threshold and considerable for the jump model.

6.4 Sensitivity to Option Moneyness

We perform sensitivity analysis of the models to different strike prices of the European call option. Tables 4 and 5 demonstrate how the hedging performance vary for both strategies under two models with different strike prices based on $\alpha = 95\%$. We consider the out-of-the-money (OTM) ($K = 1.05S_0^{(1)}$), at-the-money (ATM) ($K = S_0^{(1)}$), and in-the-money (ITM) ($K = 0.95S_0^{(1)}$) European options. In both Tables 4 and 5, the expected hedging errors for the time-consistent and global strategies are nearly identical across moneyness levels. However, their risk profiles differ substantially. The TC strategy consistently delivers a lower standard deviation, indicating greater

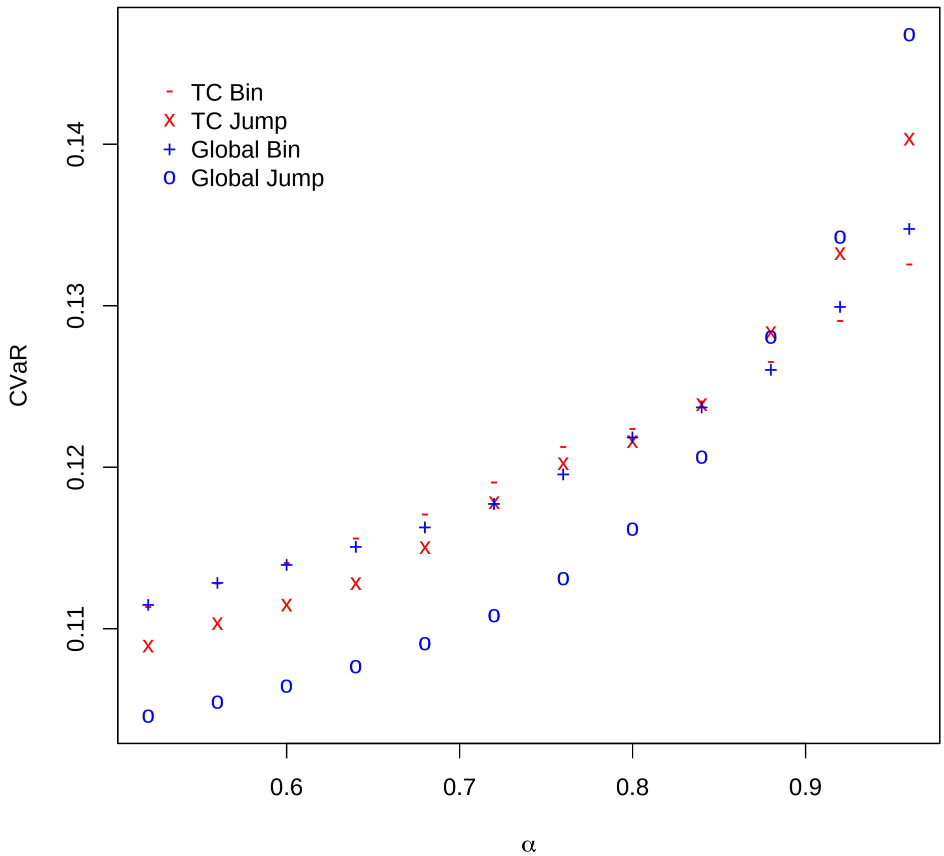

Figure 3: Objective functions for different values of risk measure thresholds α.

stability, and more importantly, it outperforms in mitigating extreme tail risk events. As discussed in Subsection 6.3, the global approach yields a lower CVaR at its 95% optimization threshold across all moneyness levels. However, the higher threshold 99% reveals that the global CVaR becomes significantly higher than that of the TC strategy. This demonstrates that the global method's optimality at a specific threshold comes at the cost of neglecting more extreme tail events. In effect, the global strategy's cost of risk at this extreme threshold becomes so high that it would make the derivative prohibitively expensive to hedge. At this point, the global strategy effectively reduces to a super-hedging strategy, which is well known to be costly.

In contrast, the TC strategy proves to be more robust, offering superior protection against these deeper risks regardless of the option's moneyness. For OTM options, the primary danger is a single large jump that abruptly makes the option valuable. Conversely, ITM options are most vulnerable to the significant, immediate loss from an adverse jump. The TC strategy's path-wise approach is better suited to handle these sudden, state-changing risks, underscoring its reliability for practical risk management.

Under the binomial model, the largest tail-risk reduction occurs for ITM options, where $\mathrm{CVaR}_{99\%}$ is lower by 1.56% (17.60% vs. 19.16%) compared to the global strategy. In contrast, under the Jump–Diffusion model, the most pronounced improvement, a substantial 1.79% at the 99% level, occurs for OTM options (15.56 vs. 17.35).

Strike	Strategy	Exp.	Std.	$\mathrm{VaR}_{95\%}$	$\mathrm{CVaR}_{95\%}$	$\mathrm{VaR}_{99\%}$	$\mathrm{CVaR}_{99\%}$
1	TC	9.96	2.15	13.10	13.90	14.42	15.19
	Global	9.98	2.37	12.46	13.57	14.10	16.44
1.05	TC	7.66	2.24	10.94	11.79	12.28	13.14
	Global	7.65	2.50	10.40	11.49	11.98	14.33
0.95	TC	12.76	1.93	15.61	16.37	16.86	17.60
	Global	12.77	2.07	14.76	16.15	16.78	19.16

Table 4: At-the-money, out-of-the-money, and in-the-money European options using the binomial model

7 Conclusions

This study develops and evaluates a global, multistage hedging framework based on CVaR that explicitly enforces time consistency. Our results confirm that a solution to the CVaR hedging problem exists which preserves optimality over time through a sequence of dynamically updated risk measures. Numerical experiments on European call options, under both binomial and jump-diffusion models, show that the proposed time-consistent hedging approach achieves notable risk reduction. In both market settings, the time-consistent strategy demonstrates superior control over extreme tail risks, consistently producing a lower CVaR at high confidence levels than the standard global optimization. This advantage is particularly pronounced in the jump-diffusion environment,

Strike	Strategy	Exp.	Std.	$\text{VaR}_{95\%}$	$\text{CVaR}_{95\%}$	$\text{VaR}_{99\%}$	$\text{CVaR}_{99\%}$
1	TC	9.79	1.95	13.21	14.93	15.99	17.74
	Global	9.79	2.24	12.46	13.68	12.90	18.50
1.05	TC	7.43	2.20	11.12	12.81	13.93	15.56
	Global	7.42	2.95	10.54	11.91	11.25	17.35
0.95	TC	12.55	1.70	15.63	17.27	18.23	20.00
	Global	12.58	2.13	14.66	16.07	15.56	21.57

Table 5: At-the-money, out-of-the-money, and in-the-money European options using the jump model.

where the likelihood of extreme outcomes is higher. The superior performance stems from the inherent coherence of the time-consistent approach; by ensuring that decisions remain aligned with the long-term objective at every stage, the framework avoids positions that could become vulnerable to rare, high-impact market shocks. This enhanced resilience confirms the practical value of our proposed method as a more robust and reliable hedging strategy for managing downside exposure under stress conditions.

References

Kaushik I Amin. Jump diffusion option valuation in discrete time. *The journal of finance*, 48(5): 1833–1863, 1993.

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

Richard Ernest Bellman and Stuart E Dreyfus. Applied dynamic programming. 1962.

Tomasz R Bielecki, Igor Cialenco, and Marcin Pitera. A survey of time consistency of dynamic risk measures and dynamic performance measures in discrete time: Lm-measure perspective. *Probability, Uncertainty and Quantitative Risk*, 2(1):3, 2017.

Kang Boda and Jerzy A Filar. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research*, 63:169–186, 2006.

Patrick Cheridito, Freddy Delbaen, and Michael Kupper. Dynamic monetary risk measures for bounded discrete-time processes. 2006.

Thomas F Coleman, Yohan Kim, Yuying Li, and M Patron. Robustly hedging variable annuities with guarantees under jump and volatility risks. *Journal of Risk and Insurance*, 74(2):347–376, 2007.

Freddy Delbaen. The structure of m–stable sets and in particular of the set of risk neutral measures. In *In memoriam Paul-André Meyer: Séminaire de Probabilités XXXIX*, pages 215–258. Springer, 2006.

Kai Detlefsen and Giacomo Scandolo. Conditional and dynamic convex risk measures. *Finance and stochastics*, 9:539–561, 2005.

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: The Journal of the IAA*, 18(2):147–160, 1988.

Hans Föllmer and Dieter Sondermann. Hedging of non-redundant contingent claims. 1985.

Pascal François, Geneviève Gauthier, and Frédéric Godin. Optimal hedging when the underlying asset follows a regime-switching markov process. In *25th Australasian Finance and Banking Conference*, 2012.

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

Frédéric Godin. Minimizing cvar in global dynamic hedging with transaction costs. *Quantitative Finance*, 16(3):461–475, 2016.

Arnaud Jobert and L Chris G Rogers. Valuations and dynamic convex risk measures. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 18(1): 1–22, 2008.

Michael Kupper and Walter Schachermayer. Representation results for law invariant time consistent functions. *Mathematics and Financial Economics*, 2(3):189–210, 2009.

Georg Ch Pflug and Alois Pichler. *Multistage stochastic optimization*, volume 1104. Springer, 2014.

Georg Ch Pflug and Alois Pichler. Time-consistent decisions and temporal decomposition of coherent risk functionals. *Mathematics of Operations Research*, 41(2):682–699, 2016.

Frank Riedel. Dynamic coherent risk measures. *Stochastic processes and their applications*, 112(2): 185–200, 2004.

R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

Andrzej Ruszczyński. Risk-averse dynamic programming for markov decision processes. *Mathematical programming*, 125:235–261, 2010.

Andrzej Ruszczyński and Alexander Shapiro. Conditional risk mappings. *Mathematics of operations research*, 31(3):544–561, 2006.

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

Alexander Shapiro. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, 37(3):143–147, 2009.

Alexander Shapiro. Time consistency of dynamic risk measures. *Operations Research Letters*, 40 (6):436–439, 2012.

Robert L Smith. Efficient monte carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 32(6):1296–1308, 1984.

Mingxin Xu. Risk measure pricing and hedging in incomplete markets. *Annals of Finance*, 2(1): 51–71, 2006.