Dynamic Portfolio Optimization under CVaR Constraints

Anran Hu[*], Silvana M. Pesenti[†], Xiaofei Shi [‡]

Abstract

We study continuous-time dynamic portfolio optimization under a Conditional Value-at-Risk (CVaR) constraint on the investor's terminal loss. For a general class of convex trading objectives, we exploit the auxiliary-threshold representation of CVaR to establish the existence of an optimal strategy and strong duality without requiring market completeness. These results motivate a dual-based nested bisection–golden-search algorithm over the threshold and Lagrangian multiplier, where the inner iterations reduce to standard unconstrained stochastic control problems. We prove that the resulting strategies converge to the optimal control as the number of iterations tends to infinity. Numerical experiments recover the Merton policy when the risk constraint is nonbinding. When the constraint is binding, the optimal strategy becomes state dependent: the investor reduces risky exposure following adverse outcomes but preserves, and near maturity may increase, exposure following favorable outcomes. Thus, a terminal CVaR constraint produces an asymmetric reallocation across states rather than uniform de-risking. Nontraded endowment risk amplifies the conservative adjustment, whereas price impact lowers desired positions and adjustment speeds.

Keywords: Dynamic portfolio optimization; Conditional Value-at-Risk; Constrained stochastic control; Lagrangian duality; Incomplete markets

JEL Classification: G11; C61; C63

1 Introduction

Financial investors seek to generate returns while managing the risks associated with their investment strategies. Merton's foundational analyses of continuous-time portfolio selection and the efficient portfolio frontier established canonical benchmarks for this trade-off [23, 24]. Subsequent work has extended portfolio choice in several directions, including market models with price impact (e.g. [36, 11]), trading costs (e.g. [9, 26]), robust and risk-aware decision criteria (e.g. [19]), and benchmark-relative portfolio optimization (e.g. [27]).

Downside-risk bounds are particularly relevant in portfolio selection when regulatory mandates or trading-desk requirements impose direct limits on losses in adverse states. Unlike conventional risk aversion, such bounds target the downside tail directly. Conditional Value-at-Risk (CVaR), also known as Expected Shortfall, is a downside-tail risk measure that averages losses in a prescribed tail of the loss distribution [2, 34]. In static portfolio optimization, CVaR may either be minimized as the risk objective [33] or imposed as a constraint on an otherwise return-oriented objective [21, 3]. Related discrete-time multiperiod formulations include CVaR-based risk control in [7] and dynamic mean–CVaR portfolio selection with a focus on time consistency in [35]. In the continuous-time

[*]Columbia University, Department of Industrial Engineering and Operations Research; ah4277@columbia.edu.
[†]University of Toronto, Department of Statistical Sciences; silvana.pesenti@utoronto.ca.
[‡]University of Toronto, Department of Statistical Sciences; xf.shi@utoronto.ca.

setting, one strand of the literature studies the cases where the problem can be recast as the choice of a replicable terminal payoff, due to additional assumptions such as market completeness [16, 14]. Another strand considers self-financing portfolio choice under dynamically re-evaluated risk constraints and exploits specialized growth-optimal or CRRA structures [32, 25].

A common way is to incorporate downside risk through a soft penalty in the investor's objective. Although convenient, such a formulation requires the investor to specify an exogenous penalty coefficient, and an arbitrary choice of this coefficient does not ensure compliance with a prescribed risk budget. In many institutional, regulatory, and wealth-management applications, however, the relevant mandate is an explicit upper bound on downside risk. Motivated by this consideration, we treat the terminal CVaR bound as an exogenous risk-management requirement rather than a preference penalty incorporated into the performance objective. Hard-constraint formulations also arise more broadly in stochastic control, including problems with expectation and terminal-law constraints (e.g. [30, 10]), and in mean-field games with state or other feasibility constraints (e.g. [6, 18]).

In this paper, we study a general continuous-time dynamic portfolio optimization problem subject to a hard CVaR constraint on terminal loss. Our framework does not require self-financing wealth dynamics or market completeness and accommodates cumulative external cash flows and nontraded endowment risk. In this general setting, the problem cannot be reduced to a static choice of terminal payoff: nontraded risks need not be replicable, while general running costs make the entire state–control path relevant. In the frictional extensions, the current position becomes an additional state variable and trading speed becomes the control, so the timing and speed of portfolio adjustment also affect optimality. The optimal strategy must therefore be characterized dynamically rather than recovered solely from an optimal terminal payoff.

To analyze this general portfolio optimization problem under hard CVaR constraints, we use the Rockafellar–Uryasev representation of CVaR to obtain a convex primal–dual formulation over the adapted trading strategy and an auxiliary scalar threshold. This formulation preserves the dynamic nature of the portfolio problem while yielding a modular solution approach based on standard unconstrained stochastic control problems. It provides the basis for both our theoretical analysis and the computational method developed below.

Contributions. Our contribution to this line of research is threefold.

First, we provide a convex-analytic treatment of continuous-time portfolio optimization under a hard terminal CVaR constraint in a market that may be incomplete. The investor minimizes a general convex expected running and terminal cost over admissible adapted trading strategies. The framework covers, among other examples, linear–quadratic and CARA specifications, as well as extended-valued CRRA and logarithmic criteria whenever the corresponding finite-feasibility condition is satisfied. We show that the CVaR threshold can be restricted to a common compact interval and establish existence of a primal optimizer. Under strict convexity, the optimal trading strategy is unique. We further establish existence of Lagrangian minimizers for every multiplier. Under a Slater condition, we prove strong duality, existence and an explicit bound for an optimal multiplier, and recovery of a primal optimizer from the dual problem. These results do not require market completeness.

Second, we exploit the resulting max–min dual structure to develop a modular numerical method. For each fixed Lagrange multiplier and CVaR threshold, the innermost problem is a standard unconstrained stochastic control problem, allowing existing dynamic-programming, PDE, or stochastic-control solvers to be used as an oracle. An inner golden-section search selects the threshold, while an outer bisection adjusts the multiplier using the CVaR constraint residual. We quantify the error of the inner search and prove convergence of the combined procedure when the inner accuracy is increased appropriately relative to the outer bisection. In particular, the constraint residual converges to zero and the objective value converges to the primal optimum. Under strict convexity, the computed controls converge weakly to the unique optimal strategy; strong convexity strengthens this conclusion to strong convergence in $L^2_{\mathbb{F}}$.

Third, we use numerical experiments to examine the effect of a binding terminal CVaR constraint in four environments: a frictionless complete-market benchmark, an incomplete market with nontraded endowment risk, a model with quadratic trading-rate regularization, and a model with square-root price impact. The first three settings remain within the convex structure motivating our analysis, while the square-root specification provides a robustness experiment beyond the setting covered by our convergence theory. Three empirical patterns emerge across the reported calibrations. First, the CVaR constraint reshapes the terminal-wealth distribution asymmetrically, with the adjustment concentrated in the downside tail rather than taking the form of a uniform contraction. This illustrates the distinction from conventional risk aversion: CVaR targets adverse tail outcomes directly rather than penalizing risk throughout the distribution. Second, the response combines an initially conservative adjustment with state-dependent feedback. Although the constraint is imposed only at maturity, when it binds, current wealth becomes a relevant state variable for the exposure policy: the computed strategies reduce risky exposure following adverse outcomes while maintaining greater participation following favorable outcomes. Third, in the reported zero-interest, no-positive-inflow experiments, the unconditional loss-CVaR diagnostic remains below the terminal risk limit at every reported intermediate date across all four environments. Although the constraint is imposed only at maturity, the behavior of the unconditional interim CVaR diagnostic suggests that, in these experiments, the terminal constraint may also discipline risk taking earlier in the investment horizon.

More broadly, our work is connected to the literature on portfolio optimization under alternative downside-risk and distributional constraints. Related formulations impose VaR constraints [4, 31, 8] or the closely related Capital-at-Risk constraint [12], formulate portfolio choice in terms of wealth quantiles [17], employ dynamic multivariate risk measures [13], or impose utility-based shortfall-risk constraints [15]. Our analysis also relates to benchmark-relative portfolio formulations that restrict terminal payoffs through divergence constraints, including Bregman–Wasserstein constraints [29] and their asymmetric α-Bregman–Wasserstein extension [28].

The remainder of the paper is organized as follows. Section 2 introduces the market, endowment, objective, and terminal CVaR constraint. Section 3 develops the primal and dual formulations. Section 4 presents the nested search algorithm and its convergence analysis. Section 5 studies the complete- and incomplete-market benchmarks, and Section 6 adds quadratic trading rate regularization and nonlinear price impact. Section 7 concludes.

2 Market Model and CVaR-Constrained Portfolio Problem

This section formulates the dynamic portfolio management problem studied in the paper. We first introduce a general continuous-time market with traded assets, cumulative endowment, and potentially nontraded background risk. Portfolio strategies are evaluated through a running and terminal cost criterion and are required to satisfy a CVaR constraint on terminal losses. This formulation accommodates incomplete markets and provides the setting for the reformulation and algorithm developed in the subsequent sections.

We then specialize the general model to a one-stock linear–quadratic Black–Scholes benchmark with terminal CVaR constraints. The benchmark is analytically transparent: in the absence of a binding CVaR constraint, the optimal risky-asset holding reduces to the classical Merton dollar exposure, while an active constraint changes the feedback strategy to protect the lower tail of terminal wealth. This structure also provides a common framework for the numerical experiments with endowment risk and trading frictions.

We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t\in[0,T]}, \mathbb{P})$ satisfying the usual conditions, where $[0,T]$ denotes the investment horizon. Let $\left\{B_t = (B_t^{(1)}, B_t^{(2)}, \dots, B_t^{(d)})^\top\right\}_{t\in[0,T]}$, be a standard d-dimensional Brownian motion adapted to $\mathbb{F}$, and let $\{B_t^\perp\}_{t\in[0,T]}$ be a 1-dim Brownian motion on $\mathbb{F}$ that is independent of $\{B_t\}_{t\in[0,T]}$.

The market contains one risk-free asset and $m \le d$ risky assets. The risk-free asset earns a constant interest rate $r \ge 0$. The risky asset price vector $S_t = (S_t^{(1)}, S_t^{(2)}, \dots, S_t^{(m)})^\top$ satisfies

$$dS_t = rS_t dt + \operatorname{diag}\{S_t\}\left(\mu(t,S_t)dt + \sigma(t,S_t)\, dB_t\right), \qquad S_0 \in (0,\infty)^m. \tag{2.1}$$

Here the $\mu : [0,T]\times\mathbb{R}^m \to \mathbb{R}^m$ is the vector of excess dollar returns, and $\sigma : [0,T]\times\mathbb{R}^m \to \mathbb{R}^{m\times d}$ is the volatility matrix. Since m may be smaller than d, the risky assets need not span all Brownian shocks.

The investor receives a cumulative endowment ζ with dynamics

$$d\zeta_t = db_t + \beta_t\, dB_t + \beta_t^\perp dB_t^\perp. \tag{2.2}$$

The finite-variation process b captures deterministic or predictable cash flows, including lump-sum payments, while the predictable row vector $\beta_t \in \mathbb{R}^{1\times d}$ captures the endowment's exposure to Brownian shocks. Only the component of β_t lying in the span of the traded volatility matrix can be hedged by the risky assets. In contrast, the remaining component $\beta_t^\perp \in \mathbb{R}$, driven by the independent Brownian motion $B^\perp$, is nontraded background risk and is a source of market incompleteness in the model.

We assume that the filtered probability space supports a process S satisfying (2.1), and we fix this weak solution throughout the paper. We impose the following regularity conditions on the exogenous coefficients and endowment processes:

Assumption 2.1 (Regularity of exogenous processes)**.** The functions μ and σ are jointly continuous. In addition, μ and σ are uniformly bounded, and for every $(t,s)\in[0,T]\times\mathbb{R}^m$, $\sigma(t,s)\sigma(t,s)^\top$ is positive definite. The process b is predictable and of finite variation, with $b_0 = 0$, and β and $\beta^\perp$ are predictable. Moreover, there exists $p > 1$ such that

$$\mathbb{E}\left[|b|_T^{2p} + \left(\int_0^T \|\beta_t\|^2 + (\beta_t^\perp)^2\, dt\right)^p\right] < \infty.$$

Let $\phi = \left\{\phi_t = \left(\phi_t^{(1)}, \phi_t^{(2)}, \dots, \phi_t^{(m)}\right)^\top\right\}_{t\in[0,T]}$ denote the portfolio of the investor, where $\phi_t^{(i)}$ represents the number of shares held in the (i)–th risky asset. Available to the investor are trading strategies that are $\mathbb{R}^m$-valued predictable processes $\varphi = \{\varphi_t := \operatorname{diag}\{S_t\}\phi_t\}_{t\in[0,T]}$, where $\varphi_t^{(i)} = \phi_t^{(i)} S_t^{(i)}$ denotes the dollar risky exposure in the (i)–th risky asset. The remaining wealth is invested in the risk-free asset, with the controlled wealth process $W^\varphi = \{W_t^\varphi\}_{t\in[0,T]}$ follows

$$\begin{aligned} dW_t^\varphi &= r\left(W_t^\varphi - \phi_t^\top S_t\right)dt + \phi_t^\top dS_t + d\zeta_t \\ &= \left(rW_t^\varphi + \varphi_t^\top \mu(t,S_t)\right)dt + db_t + \left(\varphi_t^\top\sigma(t,S_t) + \beta_t\right)dB_t + \beta_t^\perp dB_t^\perp, \qquad W_0^\varphi = w_0\,, \end{aligned} \tag{2.3}$$

where $w_0 > 0$ denotes the initial wealth of the investor.

Next, we denote by $\mathcal{A}$ the set of admissible strategies as follows:

Assumption 2.2. The admissible strategy set is

$$\mathcal{A} := \left\{ \varphi \in L^2_{\mathbb{F}}([0,T] \times \Omega; \mathbb{R}^m) : \varphi_t \in A, \quad dt \otimes d\mathbb{P}\text{-a.e.} \right\}.$$

where the action set $A \subset \mathbb{R}^m$ is nonempty, compact, and convex.

The investor evaluates a strategy through the cost functional

$$J(\varphi) := \mathbb{E}\left[\int_0^T f(t, W_t^\varphi, \varphi_t)\, dt + g(W_T^\varphi) \right], \tag{2.4}$$

where f is the running cost function and g is the terminal cost function. As the investor aims to minimize the cost functional (2.4), the utility maximization framework fits within our setting by treating costs as negative rewards.

In addition to the cost functional, the investor requires a risk constraint on a function $\ell : \mathbb{R} \to \mathbb{R}$ applied to the terminal portfolio wealth. The risk constraint is given by the Conditional Value-at-Risk (CVaR) [1]. The CVaR at a confidence level $\alpha \in (0,1)$ for the loss random variable $\ell(W_T)$ is defined as

$$\mathrm{CVaR}_\alpha(\ell(W_T)) = \frac{1}{1-\alpha} \int_\alpha^1 \mathrm{VaR}_u(\ell(W_T))du, \tag{2.5}$$

where the Value-at-Risk at a confidence level of $\alpha \in (0,1)$ is given by

$$\mathrm{VaR}_\alpha(\ell(W_T)) = \inf\{x \in \mathbb{R} : \mathbb{P}(\ell(W_T) \le x) \ge \alpha\}. \tag{2.6}$$

Given a confidence level $\alpha \in (0,1)$ and a risk limit $c \in \mathbb{R}$, the investor requires a strategy that satisfies

$$\mathrm{CVaR}_\alpha\big(\ell(W_T^\varphi)\big) \le c. \tag{2.7}$$

For the choice $\ell(w) = -w$, the CVaR constraint limits the average severity of the worst $(1-\alpha)$ fraction of terminal losses to c. Summarizing, the investor's CVaR-constrained portfolio problem is therefore

$$\begin{aligned} \inf_{\varphi \in \mathcal{A}} \quad & J(\varphi) \qquad\qquad (2.8) \\ \text{subject to} \quad & \mathrm{CVaR}_\alpha\big(\ell(W_T^\varphi)\big) \le c, \\ & (S_t, W_t^\varphi) \text{ satisfy (2.1) and (2.3).} \end{aligned}$$

Throughout the paper, we impose the following regularity conditions on the objective and terminal loss functions.

Assumption 2.3 (Regularity of cost functionals)**.** The functions f, g, ℓ satisfy the following conditions.

(i) The function $f : [0,T] \times \mathbb{R} \times A \to (-\infty, +\infty]$ is jointly Borel measurable and, for almost every $t \in [0,T]$, the mapping $(w,a) \mapsto f(t,w,a)$ is proper, convex, and lower semicontinuous. Moreover, there exists $\underline{f} \in L^1(0,T)$ such that $f(t,w,a) \ge \underline{f}(t)$ for almost every t and all $(w,a) \in \mathbb{R} \times A$.

(ii) The function $g : \mathbb{R} \to (-\infty, +\infty]$ is proper, convex, and lower semicontinuous.

(iii) The terminal loss $\ell : \mathbb{R} \to \mathbb{R}$ is convex and lower semicontinuous. Moreover, for the same $p > 1$ as in Assumption 2.1, there exists $L_\ell > 0$ such that

$$|\ell(w)| \leq L_\ell(1 + |w|^p), \qquad w \in \mathbb{R}.$$

Remark 2.4. Together with Assumptions 2.1 and 2.2, Assumption 2.3iii ensures that $\ell(W_T^\varphi)$ is integrable for every $\varphi \in \mathcal{A}$. Indeed, the standard moment estimate for the wealth equation, together with the boundedness of μ and σ and the compactness of A, gives $\mathbb{E}\big[|W_T^\varphi|^{2p}\big] < \infty$ for every $\varphi \in \mathcal{A}$. Assumption 2.3 iii then implies $\mathbb{E}\big[|\ell(W_T^\varphi)|^2\big] < \infty$, and hence $\ell(W_T^\varphi)$ is integrable. A uniform square-integrability estimate is established in Proposition 3.2.

Remark 2.5. Assumption 2.3ii ensures that there exist $a_g, b_g \in \mathbb{R}$ such that $g(w) \geq a_g w + b_g$, for all $w \in \mathbb{R}$. Together with Assumption 2.3i on f and the uniform first-moment bound on admissible wealth processes, this yields $\inf_{\varphi \in \mathcal{A}} J(\varphi) > -\infty$. We henceforth fix $J_{\text{low}} \in \mathbb{R}$ such that $J(\varphi) \geq J_{\text{low}}$ for any $\varphi \in \mathcal{A}$.

Allowing the objective costs to take the value $+\infty$ permits domain restrictions in utility-based criteria. In particular, the framework accommodates CARA terminal costs, as well as extended-valued CRRA and logarithmic costs whenever a finite-cost feasible strategy exists.

For the existence of a primal optimizer, we require the feasibility of the optimization problem.

Assumption 2.6 (Feasibility)**.** There exists at least one admissible strategy $\varphi \in \mathcal{A}$ such that

$$J(\varphi) < \infty \qquad \text{and} \qquad \mathrm{CVaR}_\alpha(\ell(W_T^\varphi)) \leq c.$$

A simple sufficient condition for the feasibility assumption, is that $0 \in A$, the no-trade wealth

$$W_T^0 = e^{rT} w_0 + \int_0^T e^{r(T-t)}\, d\zeta_t$$

already satisfies $\mathrm{CVaR}_\alpha(\ell(W_T^0)) \leq c$, and the corresponding objective $J(0)$ is finite.

The formulation in Section 2 is deliberately stated at a general level. It accommodates alternative objective specifications, including CARA terminal costs and, whenever the effective-domain and finite-feasibility requirements are satisfied, CRRA and logarithmic costs, as well as richer market dynamics and additional state variables to include consumption control or market frictions. The reformulation and algorithm developed in Sections 3 and 4 are driven primarily by the convex structure of the CVaR representation and do not depend on the particular objective function or market specification.

3 Constrained Convex Optimization Reformulation

In this section, we reformulate the risk-constrained stochastic control problem (2.8) as a convex constrained optimization problem. The main device is the auxiliary-variable representation of CVaR, which converts the terminal risk constraint into a convex constraint in both the control and an additional scalar threshold variable. This formulation then allows us to introduce a Lagrangian dual problem and establish strong duality under an additional Slater-type condition.

3.1 Reformulation and Primal Problem

By Remark 2.4, $\ell(W_T^\varphi)$ is integrable for every $\varphi \in \mathcal{A}$. Hence, its CVaR at level $\alpha \in (0,1)$ admits the representation introduced in [33]:

$$\mathrm{CVaR}_\alpha(\ell(W_T^\varphi)) = \inf_{\eta\in\mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha}\mathbb{E}\big[(\ell(W_T^\varphi) - \eta)^+\big] \right\}. \tag{3.1}$$

The infimum is attained, and its minimizers are the generalized α-quantiles of $\ell(W_T^\varphi)$. Further, for each $(\varphi, \eta) \in \mathcal{A} \times \mathbb{R}$, define

$$C(\varphi,\eta) := \eta + \frac{1}{1-\alpha}\mathbb{E}\big[(\ell(W_T^\varphi) - \eta)^+\big] - c. \tag{3.2}$$

Then the constraint $\mathrm{CVaR}_\alpha(\ell(W_T^\varphi)) \le c$ is equivalent to the existence of $\eta \in \mathbb{R}$ such that $C(\varphi,\eta) \le 0$. Thus the portfolio optimization problem can be written as

$$p^\star := \inf_{\varphi\in\mathcal{A},\ \eta\in\mathbb{R}} J(\varphi) \quad \text{subject to} \quad C(\varphi,\eta) \le 0. \tag{P}$$

Note that the scalar variable η is not a trading decision. It is an auxiliary threshold variable introduced by the CVaR representation. This formulation separates the portfolio control φ from the scalar risk threshold η, while preserving convexity.

We now present the main analytical properties of the primal formulation. These properties will be used later to derive the Lagrangian dual problem and the dual-based numerical algorithm.

Proposition 3.1 (Convexity of the primal problem)**.** *Under Assumptions 2.2 and 2.3, the extended-valued functional $J : \mathcal{A} \to (-\infty, +\infty]$ defined in (2.4) is convex in φ, and the finite-valued mapping $C : \mathcal{A} \times \mathbb{R} \to \mathbb{R}$ defined in (3.2) is jointly convex in (φ, η), and the primal problem is a convex optimization problem.*

Proof. For any $\varphi^1, \varphi^2 \in \mathcal{A}$ and $\theta \in [0,1]$, define $\varphi^\theta := \theta\varphi^1 + (1-\theta)\varphi^2$. By the linearity of the state equation and the uniqueness of solutions,

$$W_t^{\varphi^\theta} = \theta W_t^{\varphi^1} + (1-\theta) W_t^{\varphi^2}, \qquad 0 \le t \le T.$$

Since $f(t,\cdot,\cdot)$ is jointly convex and g is convex, we have

$$f(t, W_t^{\varphi^\theta}, \varphi_t^\theta) \le \theta f(t, W_t^{\varphi^1}, \varphi_t^1) + (1-\theta) f(t, W_t^{\varphi^2}, \varphi_t^2),$$

and similarly $g(W_T^{\varphi^\theta}) \le \theta g(W_T^{\varphi^1}) + (1-\theta) g(W_T^{\varphi^2})$. Taking expectations gives $J(\varphi^\theta) \le \theta J(\varphi^1) + (1-\theta)J(\varphi^2)$. Therefore, the objective function $J(\varphi)$ is convex in the control φ on $\mathcal{A}$.

It remains to show that C is convex. Let $\eta^\theta := \theta\eta^1 + (1-\theta)\eta^2$. Since ℓ is convex and W_T^φ is affine in φ,

$$\ell(W_T^{\varphi^\theta}) \le \theta \ell(W_T^{\varphi^1}) + (1-\theta)\ell(W_T^{\varphi^2}).$$

Therefore,

$$\ell(W_T^{\varphi^\theta}) - \eta^\theta \le \theta(\ell(W_T^{\varphi^1}) - \eta^1) + (1-\theta)(\ell(W_T^{\varphi^2}) - \eta^2).$$

Because $x \mapsto x_+$ is convex and nondecreasing, it follows that

$$(\ell(W_T^{\varphi^\theta}) - \eta^\theta)^+ \le \theta(\ell(W_T^{\varphi^1}) - \eta^1)^+ + (1-\theta)(\ell(W_T^{\varphi^2}) - \eta^2)^+.$$

Taking expectations and adding the affine term η shows that $C(\varphi,\eta)$ is jointly convex. Hence the feasible set $\{(\varphi,\eta) : \varphi \in \mathcal{A},\ C(\varphi,\eta) \le 0\}$ is convex, and the primal problem is convex. $\square$

We next establish a uniform square-integrability estimate for the terminal losses and use it to restrict the auxiliary threshold to a common compact interval that does not depend on the constraint level c.

Proposition 3.2 (Compactification of the CVaR threshold)**.** *Under Assumptions 2.1, 2.2, and 2.3, there exists $M_X < \infty$ such that $\sup_{\varphi\in\mathcal{A}} \mathbb{E}\left[|\ell(W_T^\varphi)|^2\right] \le M_X^2$. Let $E = \left[-\frac{M_X}{\sqrt{\alpha}}, \frac{M_X}{\sqrt{1-\alpha}}\right]$, then for every $\varphi \in \mathcal{A}$,*

$$\inf_{\eta\in\mathbb{R}} \left\{\eta + \frac{1}{1-\alpha}\mathbb{E}[(\ell(W_T^\varphi) - \eta)^+]\right\} = \inf_{\eta\in E} \left\{\eta + \frac{1}{1-\alpha}\mathbb{E}[(\ell(W_T^\varphi) - \eta)^+]\right\}.$$

Consequently, (P) *is equivalent to*

$$p^\star = \inf_{\varphi\in\mathcal{A},\ \eta\in E} J(\varphi) \quad \text{subject to} \quad C(\varphi,\eta) \le 0. \tag{P_E}$$

Proof. By the standard moment estimate for SDEs with coefficients of linear growth, see [22] or [20, Theorem 2.1], for $p \ge 1$ from Assumption 2.1, there exists $C_p > 0$, independent of $\varphi \in \mathcal{A}$, such that

$$\sup_{\varphi\in\mathcal{A}} \mathbb{E}[|W_T^\varphi|^{2p}] \le C_p(1 + |w_0|^{2p}) < \infty.$$

Since $|\ell(w)| \le K(1+|w|^p)$, we have $|\ell(W_T^\varphi)|^2 \le 2K^2\left(1 + |W_T^\varphi|^{2p}\right)$, and therefore $\sup_{\varphi\in\mathcal{A}} \mathbb{E}[|\ell(W_T^\varphi)|^2] < \infty$. Hence, there exists $M_X > 0$ that only depends on the compact set A and exogenous market coefficients, such that $\sup_{\varphi\in\mathcal{A}} \mathbb{E}[|\ell(W_T^\varphi)|^2] \le M_X^2$.

Now fix $\varphi \in \mathcal{A}$ and write $X = \ell(W_T^\varphi)$. The minimizers of $\eta \mapsto \eta + \frac{1}{1-\alpha}\mathbb{E}[(X-\eta)^+]$ are the α-quantiles of X, namely the points η satisfying

$$\mathbb{P}(X < \eta) \le \alpha \le \mathbb{P}(X \le \eta).$$

If $\eta < -\frac{M_X}{\sqrt{\alpha}}$, then

$$\mathbb{P}(X \le \eta) \le \mathbb{P}(|X| \ge |\eta|) \le \frac{\mathbb{E}[|X|^2]}{\eta^2} < \alpha.$$

Thus η cannot be an α-quantile. Similarly, if $\eta > \frac{M_X}{\sqrt{1-\alpha}}$, then

$$\mathbb{P}(X \ge \eta) \le \mathbb{P}(|X| \ge \eta) \le \frac{\mathbb{E}[|X|^2]}{\eta^2} < 1-\alpha.$$

Hence $\mathbb{P}(X < \eta) > \alpha$, so η cannot be an α-quantile. Therefore every minimizer lies in E, and the infimum over $\mathbb{R}$ equals the infimum over E.

The equivalence between (P) and (P_E) follows directly. □

Proposition 3.3 (Existence of a primal optimizer)**.** *Assume Assumptions 2.1, 2.2, 2.3, and 2.6 hold. Then the compactified primal problem*

$$p^\star = \inf_{\varphi\in\mathcal{A},\ \eta\in E} J(\varphi) \quad \text{subject to} \quad C(\varphi,\eta) \le 0$$

admits an optimal solution $(\varphi^\star, \eta^\star) \in \mathcal{A}\times E$. Consequently, the original primal problem over $\mathcal{A}\times\mathbb{R}$ also admits an optimal solution. In particular, one may choose an optimal solution with $\eta^\star \in E$.

Proof. By Assumption 2.6 and Proposition 3.2, the feasible set of the compactified problem is nonempty. Define

$$\mathbb{A} := \{(\varphi, \eta) \in \mathcal{A} \times E : C(\varphi, \eta) \leq 0\}.$$

Since A is compact, the set $\mathcal{A}$ is bounded in $L^2_{\mathbb{F}}([0,T] \times \Omega; \mathbb{R}^m)$. Moreover, $\mathcal{A}$ is convex by the convexity of A, and it is closed: if $\varphi^n \to \varphi$ in $L^2_{\mathbb{F}}$, then, up to a subsequence, $\varphi^n_t \to \varphi_t$ $dt \otimes d\mathbb{P}$-a.e.; since A is closed, $\varphi_t \in A$ a.e. Hence $\varphi \in \mathcal{A}$. Thus $\mathcal{A}$ is closed, bounded, and convex in the Hilbert space $L^2_{\mathbb{F}}$, and is therefore weakly compact. Since E is compact, $\mathcal{A} \times E$ is weakly compact.

It remains to show that $\mathbb{A}$ is weakly closed. By Proposition 3.1, C is convex. Moreover, by the linearity of the controlled wealth dynamics (2.3), we can see that for $\varphi^1, \varphi^2 \in \mathcal{A}$,

$$de^{-rt}(W^{\varphi^1}_t - W^{\varphi^2}_t) = e^{-rt} \left(\varphi^1_t - \varphi^2_t\right)^\top (\mu(t, S_t)dt + \sigma(t, S_t)dB_t),$$

which implies

$$\begin{aligned}
\mathbb{E}\left[(W^{\varphi^1}_t - W^{\varphi^2}_t)^2\right] &\leq 2\mathbb{E}\left[\int_0^t e^{2r(t-u)} \left(t\left|\left(\varphi^1_u - \varphi^2_u\right)^\top \mu(u, S_u)\right|^2 + \left\|\left(\varphi^1_u - \varphi^2_u\right)^\top \sigma(u, S_u)\right\|^2\right) du\right] \\
&\leq M\mathbb{E}\left[\int_0^t \left\|\varphi^1_u - \varphi^2_u\right\|^2 \, du\right], \qquad (3.3)
\end{aligned}$$

where M only depends on r, T, the set A, and the uniform bounds of μ and σ. The above stability estimate (3.3) for the wealth equation, together with the lower semicontinuity of ℓ and Fatou's lemma, implies that C is lower semicontinuous in the strong topology of $L^2_{\mathbb{F}} \times \mathbb{R}$. Since a convex strongly lower semicontinuous functional is weakly lower semicontinuous, C is weakly lower semicontinuous. Therefore the sublevel set of C

$$\mathbb{A} = \{(\varphi, \eta) \in \mathcal{A} \times E : C(\varphi, \eta) \leq 0\}$$

is weakly closed. Hence $\mathbb{A} \subset \mathcal{A} \times E$ is weakly compact.

Let $(\varphi^n, \eta^n) \subset \mathbb{A}$ be a minimizing sequence, so that $J(\varphi^n) \to p^\star$. By weak compactness of $\mathbb{A}$, there exist a subsequence, still denoted by (φ^n, η^n), and a pair $(\varphi^\star, \eta^\star) \in \mathbb{A}$ such that

$$\varphi^n \rightharpoonup \varphi^\star \quad \text{in } L^2_{\mathbb{F}}, \qquad \eta^n \to \eta^\star.$$

We next verify that J is weakly lower semicontinuous. Suppose first that $\varphi^n \to \varphi$ strongly in $L^2_{\mathbb{F}}$. By (3.3),

$$W^{\varphi^n} \to W^\varphi \quad \text{in } L^2([0,T] \times \Omega), \qquad W^{\varphi^n}_T \to W^\varphi_T \quad \text{in } L^2(\Omega).$$

Passing to a subsequence, the corresponding convergences hold almost everywhere.

Since $f(t, w, a) \geq \underline{f}(t)$ with $\underline{f} \in L^1(0,T)$, the function $f - \underline{f}$ is nonnegative. The lower semicontinuity of f and Fatou's lemma therefore yield

$$\mathbb{E}\left[\int_0^T f(t, W^\varphi_t, \varphi_t)\, dt\right] \leq \liminf_{n\to\infty} \mathbb{E}\left[\int_0^T f(t, W^{\varphi^n}_t, \varphi^n_t)\, dt\right].$$

Moreover, since g is proper, convex, and lower semicontinuous, there exist $a_g, b_g \in \mathbb{R}$ such that $g(w) \geq a_g w + b_g$, for all $w \in \mathbb{R}$. Applying Fatou's lemma to the nonnegative function $g(w) - a_g w - b_g$, and using the L^1 convergence of $W^{\varphi^n}_T$, gives

$$\mathbb{E}[g(W^\varphi_T)] \leq \liminf_{n\to\infty} \mathbb{E}[g(W^{\varphi^n}_T)].$$

Thus J is strongly lower semicontinuous.

Since J is convex by Proposition 3.1, it is weakly lower semicontinuous. Hence

$$J(\varphi^\star) \le \liminf_{n\to\infty} J(\varphi^n) = p^\star.$$

Since $(\varphi^\star, \eta^\star) \in \mathbb{A}$ is feasible, the reverse inequality $p^\star \le J(\varphi^\star)$ follows from the definition of $p^\star$. Therefore $J(\varphi^\star) = p^\star$, and $(\varphi^\star, \eta^\star)$ is a primal optimizer.

Finally, Proposition 3.2 shows that restricting η to E does not change the value of the primal problem. Hence the same pair is also an optimal solution of the original primal problem over $\mathcal{A} \times \mathbb{R}$. □

Finally, we establish the uniqueness of the optimal control, under strictly convexity assumption.

Assumption 3.4. The objective functional J is strictly convex on $\mathcal{A}$; that is, for any distinct $\varphi, \psi \in \mathcal{A}$ satisfying $J(\varphi) < \infty$ and $J(\psi) < \infty$, and any $\theta \in (0,1)$,

$$J\big(\theta\varphi + (1-\theta)\psi\big) < \theta J(\varphi) + (1-\theta)J(\psi).$$

In the following, we provide some explicit conditions on the problem parameters that can guarantee the above the strict convexity of J.

Proposition 3.5. *Suppose that Assumptions 2.1, 2.2, and 2.3 hold. Then the objective functional J is strictly convex on $\mathcal{A}$, if one of the following conditions holds.*

(i) For almost every $t \in [0,T]$, the function $(w,a) \mapsto f(t,w,a)$ is strictly convex on its effective domain, in the sense that, for every $\theta \in (0,1)$,

$$f\big(t, \theta w + (1-\theta)w', \theta a + (1-\theta)a'\big) < \theta f(t,w,a) + (1-\theta) f(t,w',a')$$

whenever $(w,a) \neq (w',a')$ and both $f(t,w,a)$ and $f(t,w',a')$ are finite.

(ii) There exists $\underline{\nu} > 0$ such that $\sigma(t,s)\sigma(t,s)^\top \succeq \underline{\nu} I_m$, for all $(t,s) \in [0,T]\times\mathbb{R}^m$, and the function g is strictly convex on its effective domain. More precisely, for every $\theta \in (0,1)$,

$$g\big(\theta w + (1-\theta)w'\big) < \theta g(w) + (1-\theta)g(w')$$

whenever $w \neq w'$ and both $g(w)$ and $g(w')$ are finite.

Condition ii covers several standard expected-utility objectives. In particular, when $f \equiv 0$, it applies to the negative exponential (CARA) utility

$$g(w) = e^{-\gamma w}, \qquad \gamma > 0,$$

as well as to the lower-semicontinuous extended-valued negative CRRA cost. For $\gamma > 0$, define

$$U_\gamma(w) = \begin{cases} \dfrac{w^{1-\gamma} - 1}{1-\gamma}, & \gamma \neq 1, \\ \log w, & \gamma = 1, \end{cases} \qquad w > 0,$$

and define the negative CRRA cost g_γ as the lower-semicontinuous extension of $-U_\gamma$ to $\mathbb{R}$, namely

$$g_\gamma(w) = \begin{cases} -U_\gamma(w), & w > 0, \\ \lim_{x\downarrow 0} -U_\gamma(x), & w = 0, \\ +\infty, & w < 0. \end{cases}$$

For every $\gamma > 0$, g_γ is proper, convex, lower semicontinuous, and strictly convex on its effective domain. The limiting case at $\gamma = 1$ is the negative logarithmic utility $g(w) = -\log w$.

Proof. Let $\varphi, \psi \in \mathcal{A}$ be distinct in $L^2_{\mathbb{F}}$ and satisfy $J(\varphi), J(\psi) < \infty$. Let $\theta \in (0,1)$, and define $\varphi^\theta := \theta\varphi + (1-\theta)\psi$. Since $\mathcal{A}$ is convex, $\varphi^\theta \in \mathcal{A}$. Moreover, the wealth equation is affine in the control, so

$$W^{\varphi^\theta} = \theta W^\varphi + (1-\theta) W^\psi.$$

Since $\varphi \neq \psi$ in $L^2_{\mathbb{F}}$, the set $\{(t,\omega) : \varphi_t(\omega) \neq \psi_t(\omega)\}$ has positive $dt \otimes d\mathbb{P}$ measure. On this set, $(W^\varphi_t, \varphi_t) \neq (W^\psi_t, \psi_t)$. If Condition (i) holds, then

$$\mathbb{E}\int_0^T f\big(t, W_t^{\varphi^\theta}, \varphi_t^\theta\big)\,dt < \theta\mathbb{E}\int_0^T f(t, W_t^\varphi, \varphi_t)\,dt + (1-\theta)\mathbb{E}\int_0^T f(t, W_t^\psi, \psi_t)\,dt.$$

By the convexity of g,

$$\mathbb{E}\big[g(W_T^{\varphi^\theta})\big] \leq \theta\mathbb{E}\big[g(W_T^\varphi)\big] + (1-\theta)\mathbb{E}\big[g(W_T^\psi)\big].$$

Combining the two inequalities gives $J(\varphi^\theta) < \theta J(\varphi) + (1-\theta)J(\psi)$.

If Condition ii holds, let $\varphi, \psi \in \mathcal{A}$ be distinct in $L^2_{\mathbb{F}}$ and satisfy $J(\varphi), J(\psi) < \infty$. Define $u_t := \varphi_t - \psi_t$, $X_t := W_t^\varphi - W_t^\psi$. Since the endowment terms cancel, the wealth difference satisfies

$$dX_t = \Big(rX_t + u_t^\top \mu(t, S_t)\Big)\,dt + u_t^\top \sigma(t, S_t)\,dB_t, \qquad X_0 = 0.$$

Let $\overline{\mu} := \sup_{(t,s)\in[0,T]\times\mathbb{R}^m} |\mu(t,s)| < \infty$. Applying Itô's formula to $|X_t|^2$ and taking expectations gives

$$\frac{d}{dt}\mathbb{E}[|X_t|^2] = 2r\mathbb{E}[|X_t|^2] + 2\mathbb{E}\Big[X_t u_t^\top \mu(t, S_t)\Big] + \mathbb{E}\Big[u_t^\top \sigma(t, S_t)\sigma(t, S_t)^\top u_t\Big].$$

By Young's inequality,

$$2X_t u_t^\top \mu(t, S_t) \geq -\frac{\underline{\nu}}{2}|u_t|^2 - \frac{2\overline{\mu}^2}{\underline{\nu}}|X_t|^2.$$

Together with $u_t^\top \sigma(t, S_t)\sigma(t, S_t)^\top u_t \geq \underline{\nu}|u_t|^2$, and $r \geq 0$, this yields

$$\frac{d}{dt}\mathbb{E}[|X_t|^2] \geq -K\mathbb{E}[|X_t|^2] + \frac{\underline{\nu}}{2}\mathbb{E}[|u_t|^2], \qquad K := \frac{2\overline{\mu}^2}{\underline{\nu}}.$$

Since $X_0 = 0$, multiplying by e^{Kt} and integrating gives

$$\mathbb{E}\Big[|W_T^\varphi - W_T^\psi|^2\Big] \geq \frac{\underline{\nu}}{2}\int_0^T e^{-K(T-t)}\mathbb{E}[|\varphi_t - \psi_t|^2]\,dt \geq \frac{\underline{\nu}}{2}e^{-KT}\|\varphi - \psi\|^2_{L^2_{\mathbb{F}}}. \tag{3.4}$$

Since $\varphi \neq \psi$ in $L^2_{\mathbb{F}}$, it follows that $\mathbb{P}\Big(W_T^\varphi \neq W_T^\psi\Big) > 0$. Hence, by the strict convexity of g,

$$\mathbb{E}\big[g(W_T^{\varphi^\theta})\big] < \theta\mathbb{E}\big[g(W_T^\varphi)\big] + (1-\theta)\mathbb{E}\big[g(W_T^\psi)\big].$$

By the convexity of f,

$$\mathbb{E}\int_0^T f\big(t, W_t^{\varphi^\theta}, \varphi_t^\theta\big)\,dt \leq \theta\mathbb{E}\int_0^T f(t, W_t^\varphi, \varphi_t)\,dt + (1-\theta)\mathbb{E}\int_0^T f(t, W_t^\psi, \psi_t)\,dt.$$

Combining the two inequalities also implies $J(\varphi^\theta) < \theta J(\varphi) + (1-\theta)J(\psi)$. Thus, J is strictly convex on $\mathcal{A}$. $\square$

Proposition 3.6 (Uniqueness of the optimal control)**.** *Suppose that Assumptions 2.1, 2.2, 2.3, 2.6, and 3.4 hold. Then the primal optimal control $\varphi^\star$ is unique.*

Proof. Let (φ^1, η^1) and (φ^2, η^2) be two primal optimizers. For $\theta \in (0, 1)$, define

$$\varphi^\theta := \theta\varphi^1 + (1-\theta)\varphi^2, \qquad \eta^\theta := \theta\eta^1 + (1-\theta)\eta^2.$$

By convexity of the feasible set, $(\varphi^\theta, \eta^\theta)$ is feasible. If $\varphi^1 \neq \varphi^2$, by strict convexity of J,

$$J(\varphi^\theta) < \theta J(\varphi^1) + (1-\theta)J(\varphi^2) = p^\star.$$

This contradicts the optimality of (φ^1, η^1) and (φ^2, η^2). Therefore $\varphi^1 = \varphi^2$. □

Remark 3.7 (Uniqueness of the CVaR threshold)**.** The auxiliary variable $\eta^\star$ is not unique in general, since the objective J does not depend on η. For a fixed optimal control $\varphi^\star$, the set of admissible thresholds is

$$\{\eta \in E : H_{\varphi^\star}(\eta) \leq c\}, \qquad H_{\varphi^\star}(\eta) := \eta + \frac{1}{1-\alpha}\mathbb{E}\big[(\ell(W_T^{\varphi^\star}) - \eta)^+\big].$$

Thus $\eta^\star$ is unique only if this set is a singleton. In particular, if the CVaR constraint is active and $c = \mathrm{CVaR}_\alpha(\ell(W_T^{\varphi^\star}))$, then $\eta^\star$ must be a minimizer of $H_{\varphi^\star}$. In this case, uniqueness of $\eta^\star$ is equivalent to uniqueness of the α-quantile of the terminal loss

$$X^\star := \ell(W_T^{\varphi^\star}).$$

A sufficient condition is that the distribution function of $X^\star$ is strictly increasing in a neighborhood of its α-quantile. If the constraint is slack, or if the quantile set contains an interval, then $\eta^\star$ need not be unique.

3.2 Dual formulation

We introduce the Lagrangian dual problem associated with the compactified primal problem (P_E).

For $(\varphi, \eta) \in \mathcal{A} \times E$ and $\lambda \geq 0$, define the Lagrangian

$$\begin{aligned}\mathcal{L}(\varphi, \eta, \lambda) :&= J(\varphi) + \lambda C(\varphi, \eta)\\ &= J(\varphi) + \lambda\left[\eta + \frac{1}{1-\alpha}\mathbb{E}\big[(\ell(W_T^\varphi) - \eta)^+\big] - c\right].\end{aligned}$$

The corresponding dual function is

$$q(\lambda) := \inf_{\varphi\in\mathcal{A},\ \eta\in E} \mathcal{L}(\varphi, \eta, \lambda), \qquad \lambda \geq 0, \tag{3.5}$$

and the dual problem is

$$d^\star := \sup_{\lambda\geq 0} q(\lambda). \tag{D}$$

More explicitly,

$$d^* = \sup_{\lambda\geq 0}\inf_{\varphi\in\mathcal{A},\ \eta\in E}\left\{J(\varphi) + \lambda\left(\eta + \frac{1}{1-\alpha}\mathbb{E}\left[\left(\ell(W_T^\varphi) - \eta\right)^+\right] - c\right)\right\}.$$

Remark 3.8. By Proposition 3.2, restricting η to E does not change the primal problem. It also does not change the dual formulation. Indeed, for fixed φ and $\lambda > 0$, minimizing $\mathcal{L}(\varphi, \eta, \lambda)$ over η is equivalent to minimizing

$$\eta + \frac{1}{1-\alpha}\mathbb{E}\big[(\ell(W_T^\varphi) - \eta)^+\big]$$

over η. Proposition 3.2 shows that this minimum is attained in E. When $\lambda = 0$, the Lagrangian is independent of η. Hence

$$\inf_{\varphi\in\mathcal{A},\ \eta\in\mathbb{R}} \mathcal{L}(\varphi, \eta, \lambda) = \inf_{\varphi\in\mathcal{A},\ \eta\in E} \mathcal{L}(\varphi, \eta, \lambda), \qquad \lambda \geq 0.$$

We now establish the property of the dual formulation (D) and its connection with the primal formulation (P_E).

Proposition 3.9 (Existence of Lagrange minimizers)**.** *Under Assumptions 2.1, 2.2, 2.3, and 2.6, for every $\lambda \geq 0$, the inner problem of* (D), $\inf_{\varphi\in\mathcal{A},\ \eta\in E} \mathcal{L}(\varphi, \eta, \lambda)$, *admits a minimizer.*

Suppose, in addition, that Assumption 3.4 holds. Then, for every $\lambda \geq 0$, the Lagrangian is strictly convex in its control component on its effective domain: for any $(\varphi^1, \eta^1), (\varphi^2, \eta^2) \in \mathcal{A} \times E$ with $\varphi^1 \neq \varphi^2$ in $L^2_{\mathbb{F}}$ and $J(\varphi^1), J(\varphi^2) < \infty$, and any $\theta \in (0, 1)$,

$$\mathcal{L}\big(\theta\varphi^1 + (1-\theta)\varphi^2, \theta\eta^1 + (1-\theta)\eta^2, \lambda\big) < \theta\mathcal{L}(\varphi^1, \eta^1, \lambda) + (1-\theta)\mathcal{L}(\varphi^2, \eta^2, \lambda).$$

Consequently, all Lagrangian minimizers at a fixed λ have the same control component.

Proof. For fixed $\lambda \geq 0$, the Lagrangian is given by

$$\mathcal{L}(\varphi, \eta, \lambda) = J(\varphi) + \lambda C(\varphi, \eta).$$

By the compactification result, the minimization is over $\mathcal{A} \times E$. As in the proof of Proposition 3.3, $\mathcal{A}$ is weakly compact and E is compact. Moreover, J and C are weakly lower semicontinuous, and therefore $\mathcal{L}(\cdot, \cdot, \lambda)$ is weakly lower semicontinuous on $\mathcal{A} \times E$. Hence the infimum is attained.

If Assumption 3.4 holds, J is strictly convex on $\mathcal{A}$, while C is jointly convex in (φ, η). Hence, for every $\lambda \geq 0$, the Lagrangian

$$\mathcal{L}(\varphi, \eta, \lambda) = J(\varphi) + \lambda C(\varphi, \eta)$$

is strictly convex in its control component. Therefore, two Lagrangian minimizers at the same multiplier cannot have different control components, by an argument similar to the proof of Proposition 3.6. □

We next present a weak duality result, which states that the dual value is not larger than the primal value.

Proposition 3.10 (Weak duality)**.** *Under Assumptions 2.1, 2.2, 2.3, and 2.6, for every $\lambda \geq 0$, we have $q(\lambda) \leq p^*$. Consequently, $d^* \leq p^*$.*

Proof. Let (φ, η) be primal feasible. Then $C(\varphi, \eta) \leq 0$. Since $\lambda \geq 0$,

$$\mathcal{L}(\varphi, \eta, \lambda) = J(\varphi) + \lambda C(\varphi, \eta) \leq J(\varphi).$$

Taking the infimum over feasible (φ, η) gives $q(\lambda) \leq p^*$. Taking the supremum over $\lambda \geq 0$ gives $d^* \leq p^*$. □

To establish strong duality, we require the following Slater's condition. This condition strengthens Assumption 2.6 to strict feasibility, which is standard in the optimization literature on guarantee strong duality (see *e.g.*, [5]).

Assumption 3.11 (Slater condition)**.** There exist $\bar{\varphi} \in \mathcal{A}$ and $\delta > 0$ such that

$$J(\bar{\varphi}) < \infty, \qquad \mathrm{CVaR}_\alpha\big(\ell(W_T^{\bar{\varphi}})\big) \le c - \delta.$$

Remark 3.12. Under Assumption 3.11, Proposition 3.2 ensures that choosing $\bar{\eta}$ that minimizes $C(\bar{\varphi}, \eta)$, guarantees that

$$C(\bar{\varphi}, \bar{\eta}) = \mathrm{CVaR}_\alpha(\ell(W_T^{\bar{\varphi}})) - c \le -\delta.$$

Now we are ready to state the strong duality result.

Theorem 3.13 (Strong duality)**.** *Under Assumptions 2.1, 2.2, 2.3, and 3.11, the primal problem and dual problem have the same optimal value: $p^* = d^*$. Moreover, the dual problem admits an optimal multiplier $\lambda^* \ge 0$.*

Proof. Define the perturbation value function

$$v(u) := \inf_{\varphi \in \mathcal{A},\ \eta \in E} \{J(\varphi) : C(\varphi, \eta) \le u\}.$$

Then $v(0) = p^*$. Since J and C are convex, v is convex. Assumption 3.11 implies $v(u) < +\infty$ for all $u > -\delta$ because $(\bar{\varphi}, \bar{\eta})$ is feasible for such u. By the discussion in Remark 2.5, $v(u) > -\infty$. Hence v is finite on an open interval containing 0. As a finite convex function on an open interval, v is continuous at 0 and $\partial v(0) \neq \emptyset$.

Let $\xi \in \partial v(0)$. Since relaxing the constraint can only decrease the value, v is nonincreasing, and therefore $\xi \le 0$. Define $\lambda^* := -\xi \ge 0$. The subgradient inequality gives, for all u,

$$v(u) \ge v(0) - \lambda^* u.$$

For any $(\varphi, \eta) \in \mathcal{A} \times E$, taking $u = C(\varphi, \eta)$ yields

$$J(\varphi) \ge v(C(\varphi, \eta)) \ge p^* - \lambda^* C(\varphi, \eta),$$

where the first inequality holds because (φ, η) is feasible for the perturbed constraint level $u = C(\varphi, \eta)$. Therefore

$$J(\varphi) + \lambda^* C(\varphi, \eta) \ge p^*$$

for all (φ, η). Taking the infimum over $(\varphi, \eta) \in \mathcal{A} \times E$ gives $q(\lambda^*) \ge p^*$. By Proposition 3.10, $q(\lambda^*) \le p^*$. Hence $q(\lambda^*) = p^*$, so $d^* = p^*$ and λ^* is a dual optimizer. □

Similar to Proposition 3.2, we provide a bound for the multiplier λ, which is useful in the numerical algorithm.

Proposition 3.14. *Under Assumptions 2.1,2.2, 2.3 and 3.11, let $\bar{\varphi} \in \mathcal{A}$ be chosen as in Assumption 3.11, and let J_{low} be a lower bound of J on $\mathcal{A}$ from Remark 2.5. Define*

$$\Lambda := \frac{J(\bar{\varphi}) - J_{\mathrm{low}}}{\delta} \in [0, \infty).$$

Then every dual optimizer λ^* satisfies $0 \le \lambda^* \le \Lambda$.

Proof. The lower bound $\lambda^* \geq 0$ follows from dual feasibility. By Theorem 3.13, $q(\lambda^*) = p^*$. Since $J(\varphi) \geq J_{\text{low}}$ for all $\varphi \in \mathcal{A}$, we have $p^* \geq J_{\text{low}}$. On the other hand, by definition of the dual function, choosing $\bar{\varphi}, \bar{\eta}$ as in Remark 3.12,

$$q(\lambda^*) = \inf_{\varphi \in \mathcal{A}, \eta \in E} \mathcal{L}(\varphi, \eta, \lambda^*) \leq \mathcal{L}(\bar{\varphi}, \bar{\eta}, \lambda^*).$$

Therefore,

$$q(\lambda^*) \leq J(\bar{\varphi}) + \lambda^* C(\bar{\varphi}, \bar{\eta}) \leq J(\bar{\varphi}) - \lambda^* \delta.$$

Combining the inequalities gives

$$J_{\text{low}} \leq p^* = q(\lambda^*) \leq J(\bar{\varphi}) - \lambda^* \delta.$$

Hence

$$\lambda^* \leq \frac{J(\bar{\varphi}) - J_{\text{low}}}{\delta} = \Lambda.$$

This proves the claim. ☐

The following corollary explains how a primal optimizer can be recovered once a dual optimizer has been found. It also shows that, under strict convexity, the control component recovered from any Lagrangian minimizer is the unique primal optimal control.

Corollary 3.15 (Recovery of a primal optimizer)**.** *Suppose Assumptions 2.1,2.2, 2.3 and 3.11 hold. Let $\lambda^\star$ be a dual optimizer. Suppose*

$$(\varphi^\star, \eta^\star) \in \arg \min_{\varphi \in \mathcal{A}, \eta \in E} \mathcal{L}(\varphi, \eta, \lambda^\star)$$

and

$$C(\varphi^\star, \eta^\star) \leq 0, \qquad \lambda^\star C(\varphi^\star, \eta^\star) = 0.$$

Then $(\varphi^\star, \eta^\star)$ is a primal optimizer.

If in addition Assumption 3.4 also holds, then any Lagrangian minimizer

$$(\widehat{\varphi}, \widehat{\eta}) \in \arg \min_{\varphi \in \mathcal{A}, \eta \in E} \mathcal{L}(\varphi, \eta, \lambda^\star)$$

recovers the unique primal optimal control.

Proof. Since $(\varphi^\star, \eta^\star)$ minimizes the Lagrangian at $\lambda^\star$, $\mathcal{L}(\varphi^\star, \eta^\star, \lambda^\star) = q(\lambda^\star)$. By strong duality, $q(\lambda^\star) = p^\star$. Using complementary slackness,

$$J(\varphi^\star) = J(\varphi^\star) + \lambda^\star C(\varphi^\star, \eta^\star) = \mathcal{L}(\varphi^\star, \eta^\star, \lambda^\star) = p^\star.$$

Since $C(\varphi^\star, \eta^\star) \leq 0$, the pair is primal feasible and attains the primal value. Hence it is a primal optimizer.

It remains to prove the recovery statement. Let $(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}}) \in \mathcal{A} \times E$ be a primal optimizer, whose existence follows from Proposition 3.3. Since $(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}})$ is feasible,

$$\mathcal{L}(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}}, \lambda^\star) = J(\varphi^{\mathrm{p}}) + \lambda^\star C(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}}) \leq J(\varphi^{\mathrm{p}}) = p^\star.$$

On the other hand, since $\lambda^\star$ is dual optimal, strong duality gives $q(\lambda^\star) = p^\star$, and by definition of q,

$$q(\lambda^\star) \leq \mathcal{L}(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}}, \lambda^\star).$$

Therefore $\mathcal{L}(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}}, \lambda^\star) = q(\lambda^\star)$, so $(\varphi^{\mathrm{p}}, \eta^{\mathrm{p}})$ is a Lagrangian minimizer at $\lambda^\star$.

By Proposition 3.9, under the strict convexity of J, the control component of the Lagrangian minimizer at $\lambda^\star$ is unique. Hence any Lagrangian minimizer $(\widehat{\varphi}, \widehat{\eta})$ at $\lambda^\star$ satisfies $\widehat{\varphi} = \varphi^{\mathrm{p}}$. Thus $\widehat{\varphi}$ is the unique primal optimal control. ☐

4 Algorithm and Analysis

4.1 Nested Bisection-Golden Search Method

In this section, motivated by the dual formulation in Section 3.2, we propose a nested bisection method for the portfolio optimization problem with CVaR constraint, by solving the compactified dual problem. By Propositions 3.2 and 3.14, it is enough to solve the dual problem

$$d^\star = \sup_{\lambda \in [0,\Lambda]} q(\lambda), \qquad q(\lambda) = \inf_{\varphi \in \mathcal{A}, \eta \in E} \mathcal{L}(\varphi, \eta, \lambda).$$

The proposed algorithm consists of a control oracle for fixed (λ, η), an inner golden-section search over η, and an outer bisection over λ.

Throughout this section, Assumptions 2.1–2.3, 3.11, and 3.4 are in force. Assumption 3.4 ensures that both J and $\mathcal{L}(\cdot, \eta, \lambda)$ are strictly convex in the control component. Consequently, every fixed-(λ, η) control problem and every joint inner problem $q(\lambda)$ have a unique control component.

Control oracle. For fixed (λ, η), define the modified terminal cost

$$\Psi_{\lambda,\eta}(w) := g(w) + \frac{\lambda}{1-\alpha}(\ell(w) - \eta)^+.$$

Then the fixed-(λ, η) subproblem is

$$\inf_{\varphi \in \mathcal{A}} \mathcal{L}(\varphi, \eta, \lambda) = \inf_{\varphi \in \mathcal{A}} \mathbb{E}\left[\int_0^T f(t, W_t^\varphi, \varphi_t)\, dt + \Psi_{\lambda,\eta}(W_T^\varphi)\right] + \lambda(\eta - c). \tag{4.1}$$

The last term $\lambda(\eta - c)$ is independent of φ, so the control oracle only needs to solve a standard stochastic control problem with terminal cost $\Psi_{\lambda,\eta}$. We denote its output by

$$(\varphi_{\lambda,\eta}, V_\lambda(\eta)) := \mathsf{Control}(\lambda, \eta),$$

where $\varphi_{\lambda,\eta} = \arg\min_{\varphi \in \mathcal{A}} \mathcal{L}(\varphi, \eta, \lambda)$ is the unique solution of (4.1) under Assumption 3.4, and $V_\lambda(\eta) = \mathcal{L}(\varphi_{\lambda,\eta}, \eta, \lambda) = \inf_{\varphi \in \mathcal{A}} \mathcal{L}(\varphi, \eta, \lambda)$ is the corresponding optimal value.

Inner golden-section search over the CVaR threshold. We next construct a method for evaluating the dual function q defined in (3.5). For any fixed $\lambda > 0$,

$$\begin{aligned} q(\lambda) &= \inf_{\varphi \in \mathcal{A},\ \eta \in E} \mathcal{L}(\varphi, \eta, \lambda) \\ &= \inf_{\eta \in E} \inf_{\varphi \in \mathcal{A}} \mathcal{L}(\varphi, \eta, \lambda) = \inf_{\eta \in E} V_\lambda(\eta). \end{aligned}$$

Thus, evaluating $q(\lambda)$ reduces to a one-dimensional minimization of V_λ over the compact interval E.

Proposition 4.1 (Properties of the inner value function)**.** *Suppose that Assumptions 2.1, 2.2, 2.3, and 3.11 hold. For every $\lambda \in [0, \Lambda]$, the function V_λ is finite and convex on E. Moreover,*

$$|V_\lambda(\eta) - V_\lambda(\eta')| \le \lambda \kappa_\alpha |\eta - \eta'|, \qquad \eta, \eta' \in E,$$

where $\kappa_\alpha := \max\left\{1, \frac{\alpha}{1-\alpha}\right\}$. Consequently, $\mathcal{E}_\lambda := \arg\min_{\eta \in E} V_\lambda(\eta)$ is a nonempty closed interval.

Proof. Let $\bar{\varphi}$ be the finite-cost strategy in Assumption 3.11. Since $C(\bar{\varphi}, \eta)$ is finite, $V_\lambda(\eta) \le \mathcal{L}(\bar{\varphi}, \eta, \lambda)$. Moreover, by Remark 2.5 and the nonnegativity of the hinge term,

$$\mathcal{L}(\varphi, \eta, \lambda) \ge J_{\text{low}} + \lambda(\eta - c), \qquad \varphi \in \mathcal{A}.$$

Hence $V_\lambda(\eta) > -\infty$. For fixed λ, the affine dependence of W^φ on φ, the convexity of f, g, and ℓ, and the convexity of $\mathcal{A}$ imply that $(\varphi, \eta) \mapsto \mathcal{L}(\varphi, \eta, \lambda)$ is jointly convex. Partial minimization over φ therefore shows that V_λ is convex.

For fixed $x \in \mathbb{R}$, define $r_x(\eta) := \eta + \frac{1}{1-\alpha}(x - \eta)^+$. Every subgradient of r_x has an absolute value bounded by $\kappa_\alpha := \max\left\{1, \frac{\alpha}{1-\alpha}\right\}$, and hence $|r_x(\eta) - r_x(\eta')| \le \kappa_\alpha |\eta - \eta'|$. It follows that, uniformly over $\varphi \in \mathcal{A}$,

$$\big|\mathcal{L}(\varphi, \eta, \lambda) - \mathcal{L}(\varphi, \eta', \lambda)\big| \le \lambda \kappa_\alpha |\eta - \eta'|.$$

Taking the infimum over φ in both directions gives

$$|V_\lambda(\eta) - V_\lambda(\eta')| \le \lambda \kappa_\alpha |\eta - \eta'|.$$

Thus V_λ is continuous. Since E is compact, $\mathcal{E}_\lambda$ is nonempty and closed, and its convexity follows from the convexity of V_λ. □

Since V_λ is convex, it decreases before its minimizer set and increases after it. This allows us to locate a minimizer using only function values. Starting from an interval $[\eta_{\text{low}}, \eta_{\text{high}}] \subseteq E$ containing $\mathcal{E}_\lambda$, define

$$\rho := \frac{\sqrt{5} - 1}{2}, \qquad \eta_1 = \eta_{\text{high}} - \rho(\eta_{\text{high}} - \eta_{\text{low}}), \qquad \eta_2 = \eta_{\text{low}} + \rho(\eta_{\text{high}} - \eta_{\text{low}}).$$

We evaluate V_λ at these two interior points η_1, η_2 using the control oracle

$$(\varphi_{\lambda,\eta}, V_\lambda(\eta)) = \mathsf{Control}(\lambda, \eta).$$

If $V_\lambda(\eta_1) \le V_\lambda(\eta_2)$, convexity implies that the interval $[\eta_{\text{low}}, \eta_2]$ contains a minimizer, so the right-hand portion may be discarded. Otherwise, the interval $[\eta_1, \eta_{\text{high}}]$ contains a minimizer, and the left-hand portion may be discarded. The points are chosen according to the golden ratio so that one previous function evaluation can be reused after each interval reduction. Thus, after the initial two evaluations, each iteration requires only one new call to the control oracle. The algorithm is summarized as follows.

Algorithm 1: EtaGoldenSearch(λ)

Data: Multiplier $\lambda > 0$, interval $E = [\eta_{\text{low}}, \eta_{\text{high}}]$, number of reductions $N_\eta \geq 1$
Result: Approximate inner solution $(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda)$
$\rho \leftarrow (\sqrt{5} - 1)/2$;
$\eta_1 \leftarrow \eta_{\text{high}} - \rho(\eta_{\text{high}} - \eta_{\text{low}})$;
$\eta_2 \leftarrow \eta_{\text{low}} + \rho(\eta_{\text{high}} - \eta_{\text{low}})$;
$(\varphi_i, v_i) \leftarrow \mathsf{Control}(\lambda, \eta_i)$, $i = 1, 2$;
for $j = 1, \ldots, N_\eta$ **do**
  if $v_1 \leq v_2$ **then**
    $\eta_{\text{high}} \leftarrow \eta_2$;
    $(\eta_2, \varphi_2, v_2) \leftarrow (\eta_1, \varphi_1, v_1)$;
    $\eta_1 \leftarrow \eta_{\text{high}} - \rho(\eta_{\text{high}} - \eta_{\text{low}})$;
    $(\varphi_1, v_1) \leftarrow \mathsf{Control}(\lambda, \eta_1)$;
  else
    $\eta_{\text{low}} \leftarrow \eta_1$;
    $(\eta_1, \varphi_1, v_1) \leftarrow (\eta_2, \varphi_2, v_2)$;
    $\eta_2 \leftarrow \eta_{\text{low}} + \rho(\eta_{\text{high}} - \eta_{\text{low}})$;
    $(\varphi_2, v_2) \leftarrow \mathsf{Control}(\lambda, \eta_2)$;
  end
end
return (φ_1, η_1) *if* $v_1 \leq v_2$*; otherwise* (φ_2, η_2);

Outer bisection over the Lagrange multiplier. The inner golden-section search via Algorithm 1 approximately evaluates $q(\lambda)$ for each fixed multiplier λ. We now maximize the concave dual function q over $[0, \Lambda]$. The key observation is that its derivative is given by the constraint residual, which we define in the following two cases of λ:

For every $\lambda > 0$, let

$$(\varphi_\lambda, \eta_\lambda) \in \arg \min_{\varphi \in \mathcal{A}, \eta \in E} \mathcal{L}(\varphi, \eta, \lambda),$$

and define the residual

$$R(\lambda) := C(\varphi_\lambda, \eta_\lambda).$$

Under Assumption 3.4, the control component φ_λ is unique. Although the optimizing threshold η_λ may not be unique, the value of the residual is unique. Indeed, if η_λ and $\widetilde{\eta}_\lambda$ are both optimal thresholds for the same control, then

$$J(\varphi_\lambda) + \lambda C(\varphi_\lambda, \eta_\lambda) = J(\varphi_\lambda) + \lambda C(\varphi_\lambda, \widetilde{\eta}_\lambda).$$

Since $\lambda > 0$, $C(\varphi_\lambda, \eta_\lambda) = C(\varphi_\lambda, \widetilde{\eta}_\lambda)$. Thus, $R(\lambda)$ is well defined.

For $\lambda = 0$, we define (φ_0, η_0) as follows. Solving the unconstrained problem yields φ_0. Next, we choose

$$\eta_0 \in \arg \min_{\eta \in E} \left\{ \eta + \frac{1}{1 - \alpha} \mathbb{E}[(\ell(W_T^{\varphi_0}) - \eta)^+] \right\},$$

which yields the residual

$$R(0) = C(\varphi_0, \eta_0) = \mathrm{CVaR}_\alpha(\ell(W_T^{\varphi_0})) - c.$$

If $R(0) \leq 0$, then the unconstrained optimizer is feasible and the optimal multiplier is $\lambda^\star = 0$. Otherwise, the constraint is active, and we search for a root of R on $(0, \Lambda]$.

Proposition 4.2 (Dual residual). *Suppose that Assumptions 2.1, 2.2, 2.3, 3.4, and 3.11 hold. Then R is nonincreasing and continuous on $[0,\infty)$, and therefore uniformly continuous on $[0,\Lambda]$. Moreover, q is differentiable on $(0,\infty)$, with $q'(\lambda) = R(\lambda)$ for all $\lambda > 0$.*

Proof. By Proposition 3.2, there exists $M_C < \infty$ such that

$$|C(\varphi,\eta)| \le M_C, \qquad (\varphi,\eta) \in \mathcal{A} \times E.$$

Together with the finite-cost strategy provided by Assumption 3.11 and the lower bound on J, this implies that q is finite. Moreover,

$$|q(\lambda_2) - q(\lambda_1)| \le M_C|\lambda_2 - \lambda_1|, \qquad \lambda_1, \lambda_2 \ge 0,$$

so q is continuous.

Let $0 \le \lambda_1 < \lambda_2$, and choose a Lagrangian minimizer at each multiplier, using (φ_0, η_0) when $\lambda_1 = 0$. Optimality gives

$$q(\lambda_2) \le q(\lambda_1) + (\lambda_2 - \lambda_1)R(\lambda_1)$$

and

$$q(\lambda_1) \le q(\lambda_2) - (\lambda_2 - \lambda_1)R(\lambda_2).$$

Therefore,

$$R(\lambda_2) \le \frac{q(\lambda_2) - q(\lambda_1)}{\lambda_2 - \lambda_1} \le R(\lambda_1). \tag{4.2}$$

In particular, R is nonincreasing.

Next, we show the continuity. Let $\lambda_n \to \lambda$ monotonically and choose

$$(\varphi_n, \eta_n) \in \operatorname{argmin}_{\varphi\in\mathcal{A},\ \eta\in E}\mathcal{L}(\varphi,\eta,\lambda_n).$$

By weak compactness, along a subsequence, $\varphi_n \rightharpoonup \bar\varphi$ and $\eta_n \to \bar\eta$. The weak lower semicontinuity of the Lagrangian, the boundedness of C, and the continuity of q imply that

$$(\bar\varphi, \bar\eta) \in \operatorname{argmin}_{\varphi\in\mathcal{A},\ \eta\in E}\mathcal{L}(\varphi,\eta,\lambda).$$

Suppose first that $\lambda_n \downarrow \lambda$ and write $R(\lambda_n) \to r$. Monotonicity gives $r \le R(\lambda)$. By the weak lower semicontinuity of C, $C(\bar\varphi,\bar\eta) \le r$. If $\lambda > 0$, then $C(\bar\varphi,\bar\eta) = R(\lambda)$. If $\lambda = 0$, the definition of η_0 gives $R(0) \le C(\bar\varphi,\bar\eta)$. In either case, $R(\lambda) \le r$, and hence $r = R(\lambda)$. Thus R is right-continuous, including at zero.

Now suppose that $\lambda_n \uparrow \lambda$ with $\lambda > 0$, and write $R(\lambda_n) \to r$. Monotonicity gives $r \ge R(\lambda)$. Moreover,

$$J(\varphi_n) = q(\lambda_n) - \lambda_n R(\lambda_n).$$

Using the weak lower semicontinuity of J gives

$$q(\lambda) - \lambda R(\lambda) = J(\bar\varphi) \le \liminf_{n\to\infty} J(\varphi_n) = q(\lambda) - \lambda r.$$

Since $\lambda > 0$, this implies $r \le R(\lambda)$. Therefore $r = R(\lambda)$, proving left continuity.

Finally, for $\lambda > 0$, applying (4.2) to λ and $\lambda + h$ yields

$$R(\lambda + h) \le \frac{q(\lambda+h) - q(\lambda)}{h} \le R(\lambda).$$

Letting $h \downarrow 0$ and using right continuity gives $q'_+(\lambda) = R(\lambda)$. Similarly,

$$R(\lambda) \le \frac{q(\lambda) - q(\lambda - h)}{h} \le R(\lambda - h),$$

and left continuity gives $q'_-(\lambda) = R(\lambda)$. Thus q is differentiable and $q'(\lambda) = R(\lambda)$ for $\lambda > 0$. $\square$

In the active-constraint case, $\lambda^\star \neq 0$. Define the dual optimal set $\mathcal{Z} := \{\lambda \in [0, \Lambda] : R(\lambda) = 0\}$. By Theorem 3.13, Proposition 3.14, and the differentiability of q, the set Z is nonempty and coincides with the set of dual optimizers. Since R is nonincreasing, $R(\lambda) > 0$ means that the multiplier is too small, so the search must move to the right. Similarly, $R(\lambda) < 0$ means that the multiplier is too large, so the search must move to the left. Thus, the root of R can be located by bisection.

For a queried multiplier λ, Algorithm 1 returns $(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) = \mathsf{EtaGoldenSearch}(\lambda)$. We use the corresponding residual $\widehat{R}(\lambda) := C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda)$ in the outer bisection. Starting from $[0, \Lambda]$, the method evaluates $\widehat{R}$ at the midpoint of the current interval. If $\widehat{R}(\lambda) > 0$, it retains the right half; if $\widehat{R}(\lambda) < 0$, it retains the left half. The main algorithm is shown below.

Algorithm 2: Nested bisection–golden search algorithm

```
Data: Multiplier bound Λ, maximum iteration N_λ ≥ 1
Result: Approximate primal-dual solution (φ̂, η̂, λ̂)
Solve the unconstrained problem and compute R(0);
if R(0) ≤ 0 then
    return (φ_0, η_0, 0);
end
λ_low ← 0, λ_high ← Λ;
for k = 0, 1, ..., N_λ − 1 do
    λ_mid ← (λ_low + λ_high)/2;
    (φ_mid, η_mid) ← EtaGoldenSearch(λ_mid);
    R_mid ← C(φ_mid, η_mid);
    if R_mid = 0 then
        return (φ_mid, η_mid, λ_mid);
    end
    if R_mid > 0 then
        λ_low ← λ_mid;
    else
        λ_high ← λ_mid;
    end
end
return (φ_mid, η_mid, λ_mid);
```

4.2 Convergence Analysis

We analyze the performance of the main algorithm, under exact evaluations of the control oracle. We first study the convergence of the inner golden-search over the CVaR threshold (Algorithm 1) and then establish the convergence of Algorithm 2.

For a nonempty set $S \subset \mathbb{R}$, we write

$$\operatorname{dist}(\eta, S) := \inf_{\eta' \in S} |\eta - \eta'|.$$

Proposition 4.3 (Convergence of the inner golden-section search)**.** *Let $\rho := \frac{\sqrt{5}-1}{2}$. Suppose that Assumptions 2.1, 2.2, 2.3, and 3.11 hold, and that the control and value used by Algorithm 1 are exact. After N_η interval reductions, Algorithm 1 returns $(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda)$ satisfying*

$$\operatorname{dist}\big(\widehat{\eta}_\lambda, \mathcal{E}_\lambda\big) \leq |E| \rho^{N_\eta},$$

and

$$0 \leq V_\lambda(\widehat{\eta}_\lambda) - q(\lambda) \leq \lambda \kappa_\alpha |E| \rho^{N_\eta}.$$

Proof. By Proposition 4.1, V_λ is convex and $\mathcal{E}_\lambda$ is a nonempty closed interval.

Consider a current search interval $[\eta_{\text{low}}, \eta_{\text{high}}]$ that intersects $\mathcal{E}_\lambda$, and let $\eta_1 < \eta_2$ be the two golden-section points. If $V_\lambda(\eta_1) \le V_\lambda(\eta_2)$, convexity implies that a minimizer exists in $[\eta_{\text{low}}, \eta_2]$, so the algorithm may discard $(\eta_2, \eta_{\text{high}}]$. Similarly, if $V_\lambda(\eta_1) > V_\lambda(\eta_2)$, a minimizer exists in $[\eta_1, \eta_{\text{high}}]$. Thus every update retains an interval intersecting $\mathcal{E}_\lambda$.

Each golden-section update reduces the interval length by the factor ρ. Therefore, after N_η reductions, the remaining interval has length $|E|\rho^{N_\eta}$ and still intersects $\mathcal{E}_\lambda$. Since $\widehat{\eta}_\lambda$ lies in this interval,

$$\operatorname{dist}\big(\widehat{\eta}_\lambda, \mathcal{E}_\lambda\big) \le |E|\rho^{N_\eta}.$$

Choose $\eta_\lambda^\star \in \mathcal{E}_\lambda$ such that $|\widehat{\eta}_\lambda - \eta_\lambda^\star| = \operatorname{dist}\big(\widehat{\eta}_\lambda, \mathcal{E}_\lambda\big)$. Since $q(\lambda) = \min_{\eta \in E} V_\lambda(\eta) = V_\lambda(\eta_\lambda^\star)$, the Lipschitz estimate in Proposition 4.1 yields

$$0 \le V_\lambda(\widehat{\eta}_\lambda) - q(\lambda) \le \lambda\kappa_\alpha |\widehat{\eta}_\lambda - \eta_\lambda^\star| \le \lambda\kappa_\alpha |E|\rho^{N_\eta}.$$

□

Theorem 4.4. *Suppose that Assumptions 2.1, 2.2, 2.3, 3.11, and 3.4 hold, and that the control, value, and residual evaluations used by the algorithm are exact. Consider the active-constraint case* $R(0) > 0$, *and let* $(\widehat{\varphi}, \widehat{\eta}, \widehat{\lambda})$ *be the output of the Algorithm 2 with* N_λ *outer iterations and* N_η *inner iterations.*

Let $\rho := \frac{\sqrt{5}-1}{2}$. *If* $N_\lambda \to \infty$ *and* $2^{N_\lambda}\rho^{N_\eta} \to 0$, *then*

$$\operatorname{dist}(\widehat{\lambda}, \mathcal{Z}) \to 0, \qquad C(\widehat{\varphi}, \widehat{\eta}) \to 0, \qquad J(\widehat{\varphi}) \to p^\star.$$

Moreover,

$$\widehat{\varphi} \rightharpoonup \varphi^\star \qquad \text{weakly in } L^2_{\mathbb{F}},$$

where $\varphi^\star$ *is the unique primal optimal control.*

Proof. For every multiplier λ queried by the outer bisection, let $(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda)$ denote the output of the inner bisection and the associated exact control oracle. Then by Proposition 4.3,

$$0 \le \mathcal{L}(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda, \lambda) - q(\lambda) \le \lambda\kappa_\alpha |E|\rho^{N_\eta} \le \Lambda\kappa_\alpha |E|\rho^{N_\eta}.$$

For any $\mu \in [0, \Lambda]$, the definition of q and the affine dependence of the Lagrangian on the multiplier give

$$\begin{aligned} q(\mu) &\le \mathcal{L}(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda, \mu) \\ &= \mathcal{L}(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda, \lambda) + (\mu - \lambda)C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) \\ &\le q(\lambda) + \Lambda\kappa_\alpha |E|\rho^{N_\eta} + (\mu - \lambda)C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda). \end{aligned}$$

Among all multipliers queried during the N_λ outer iterations and at the final midpoint, the smallest possible distance from the endpoints of $[0, \Lambda]$ is $\Lambda 2^{-(N_\lambda+1)}$. Hence every queried multiplier λ satisfies

$$\lambda \pm \Lambda 2^{-(N_\lambda+1)} \in [0, \Lambda].$$

Taking $\mu = \lambda + \Lambda 2^{-(N_\lambda+1)}$ in the preceding inequality gives

$$C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) \ge \frac{q\big(\lambda + \Lambda 2^{-(N_\lambda+1)}\big) - q(\lambda)}{\Lambda 2^{-(N_\lambda+1)}} - 2\kappa_\alpha |E| 2^{N_\lambda}\rho^{N_\eta}.$$

Since q is concave and $q' = R$,

$$\frac{q\big(\lambda + \Lambda 2^{-(N_\lambda+1)}\big) - q(\lambda)}{\Lambda 2^{-(N_\lambda+1)}} \geq R\Big(\lambda + \Lambda 2^{-(N_\lambda+1)}\Big).$$

Similarly, taking $\mu = \lambda - \Lambda 2^{-(N_\lambda+1)}$ and using concavity yields

$$C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) \leq R\Big(\lambda - \Lambda 2^{-(N_\lambda+1)}\Big) + 2\kappa_\alpha |E| 2^{N_\lambda} \rho^{N_\eta}.$$

Combining the two inequalities proves

$$\begin{aligned} R\Big(\lambda + \Lambda 2^{-(N_\lambda+1)}\Big) - 2\kappa_\alpha |E| 2^{N_\lambda} \rho^{N_\eta} &\leq C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) \\ &\leq R\Big(\lambda - \Lambda 2^{-(N_\lambda+1)}\Big) + 2\kappa_\alpha |E| 2^{N_\lambda} \rho^{N_\eta}. \end{aligned}$$

Consequently,

$$|C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) - R(\lambda)| \leq \sup_{\substack{\lambda_1, \lambda_2 \in [0,\Lambda] \\ |\lambda_1 - \lambda_2| \leq \Lambda 2^{-(N_\lambda+1)}}} |R(\lambda_1) - R(\lambda_2)| + 2\kappa_\alpha |E| 2^{N_\lambda} \rho^{N_\eta}.$$

By Proposition 4.2, with $R(0) := q'_+(0) = \min_{\eta \in E} C(\varphi_0, \eta)$ and $R(\Lambda) := q'_-(\Lambda)$, the residual R is continuous on $[0, \Lambda]$. Hence it is uniformly continuous. Therefore, if $N_\lambda \to \infty$ and $2^{N_\lambda} \rho^{N_\eta} \to 0$, then

$$\sup_\lambda |C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda) - R(\lambda)| \to 0, \tag{4.3}$$

where the supremum is over the multipliers queried by the outer bisection.

We next prove that $\operatorname{dist}(\widehat{\lambda}, \mathcal{Z}) \to 0$. Fix $\varepsilon > 0$. Since R is continuous and $\mathcal{Z} = \{\lambda \in [0, \Lambda] : R(\lambda) = 0\}$, compactness implies that

$$\inf_{\substack{\lambda \in [0,\Lambda] \\ \operatorname{dist}(\lambda, \mathcal{Z}) \geq \varepsilon}} |R(\lambda)| > 0,$$

provided that the set over which the infimum is taken is nonempty. By (4.3), for sufficiently large N_λ and N_η, $C(\widehat{\varphi}_\lambda, \widehat{\eta}_\lambda)$ and $R(\lambda)$ have the same sign at every queried midpoint satisfying $\operatorname{dist}(\lambda, \mathcal{Z}) \geq \varepsilon$.

We claim that, throughout the outer bisection, the current bracket intersects the closed ε-neighborhood of $\mathcal{Z}$. This is true for the initial bracket $[0, \Lambda]$. Suppose it holds for the current bracket, and let λ be its midpoint. If $\operatorname{dist}(\lambda, \mathcal{Z}) \leq \varepsilon$, then either half selected by the algorithm contains λ as an endpoint, and hence still intersects the ε-neighborhood of $\mathcal{Z}$.

Otherwise, $\operatorname{dist}(\lambda, \mathcal{Z}) > \varepsilon$, so the approximate residual has the same sign as $R(\lambda)$. Since R is nonincreasing, its zero set $\mathcal{Z}$ is an interval. If λ lies to the left of the ε-neighborhood of $\mathcal{Z}$, then $R(\lambda) > 0$, and the algorithm retains the right half of the bracket, which still intersects that neighborhood. Similarly, if λ lies to the right of the ε-neighborhood, then $R(\lambda) < 0$, and the algorithm retains the left half, which again intersects the neighborhood. The claim therefore follows by induction.

Consequently, the final bracket contains some point $\widetilde{\lambda}$ such that $\operatorname{dist}(\widetilde{\lambda}, \mathcal{Z}) \leq \varepsilon$. Since $\widehat{\lambda}$ also lies in the final bracket, whose length is at most $\Lambda 2^{-N_\lambda}$,

$$\operatorname{dist}(\widehat{\lambda}, \mathcal{Z}) \leq |\widehat{\lambda} - \widetilde{\lambda}| + \operatorname{dist}(\widetilde{\lambda}, \mathcal{Z}) \leq \Lambda 2^{-N_\lambda} + \varepsilon.$$

Taking the limit superior and then letting $\varepsilon \downarrow 0$ gives $\operatorname{dist}(\widehat{\lambda}, \mathcal{Z}) \to 0$.

Because R is continuous and vanishes on $\mathcal{Z}$, $R(\widehat{\lambda}) \to 0$. Applying (4.3) at the final multiplier yields

$$|C(\widehat{\varphi}, \widehat{\eta})| \leq \left| C(\widehat{\varphi}, \widehat{\eta}) - R(\widehat{\lambda}) \right| + |R(\widehat{\lambda})| \to 0.$$

Applying Proposition 4.3 at $\widehat{\lambda}$ gives

$$\mathcal{L}(\widehat{\varphi}, \widehat{\eta}, \widehat{\lambda}) \leq q(\widehat{\lambda}) + \Lambda\kappa_\alpha |E| \rho^{N_\eta} \leq p^\star + \Lambda\kappa_\alpha |E| \rho^{N_\eta}.$$

Therefore,

$$J(\widehat{\varphi}) - p^\star \leq \Lambda\kappa_\alpha |E| \rho^{N_\eta} + \widehat{\lambda} \left| C(\widehat{\varphi}, \widehat{\eta}) \right|.$$

On the other hand, for any $\lambda^\star \in \mathcal{Z}$, strong duality implies

$$p^\star = q(\lambda^\star) \leq J(\widehat{\varphi}) + \lambda^\star C(\widehat{\varphi}, \widehat{\eta}).$$

Since $\widehat{\lambda}, \lambda^\star \in [0, \Lambda]$, we conclude that

$$|J(\widehat{\varphi}) - p^\star| \leq \Lambda\kappa_\alpha |E| \rho^{N_\eta} + \Lambda \left| C(\widehat{\varphi}, \widehat{\eta}) \right| \to 0.$$

It remains to prove the convergence of the controls. Since $\mathcal{A}$ is weakly compact in $L^2_{\mathbb{F}}([0, T] \times \Omega; \mathbb{R}^m)$ and E is compact, every subsequence of $(\widehat{\varphi}, \widehat{\eta})$ admits a further subsequence, still denoted by $(\widehat{\varphi}, \widehat{\eta})$, such that

$$\widehat{\varphi} \rightharpoonup \overline{\varphi} \quad \text{weakly in } L^2_{\mathbb{F}}, \qquad \widehat{\eta} \to \overline{\eta} \in E.$$

By the weak lower semicontinuity of J and C, $J(\overline{\varphi}) \leq \liminf J(\widehat{\varphi}) = p^\star$, and $C(\overline{\varphi}, \overline{\eta}) \leq \liminf C(\widehat{\varphi}, \widehat{\eta}) = 0$. Thus $(\overline{\varphi}, \overline{\eta})$ is primal feasible and attains the primal value. Hence $\overline{\varphi}$ is a primal optimal control. By Assumption 3.4, the primal optimal control is unique, and therefore $\overline{\varphi} = \varphi^\star$. Thus every weakly convergent subsequence of $\widehat{\varphi}$ has limit $\varphi^\star$. Since $\mathcal{A}$ is weakly compact, it follows that the entire sequence satisfies

$$\widehat{\varphi} \rightharpoonup \varphi^\star \qquad \text{weakly in } L^2_{\mathbb{F}}.$$

□

The strict-convexity condition in Theorem 4.4 guarantees the weak convergence. Strong convergence follows under the following stronger condition.

Assumption 4.5. J is strongly convex on its effective domain with respect to the $L^2_{\mathbb{F}}$ norm. That is, there exists a constant $\kappa > 0$ such that for any $\varphi, \psi \in \mathcal{A}$ satisfying $J(\varphi) < \infty$, $J(\psi) < \infty$, and $\theta \in (0, 1)$,

$$J\big(\theta\varphi + (1-\theta)\psi\big) \leq \theta J(\varphi) + (1-\theta) J(\psi) - \frac{\kappa}{2}\theta(1-\theta)\|\varphi - \psi\|^2_{L^2_{\mathbb{F}}}.$$

Similar to Proposition 3.5, we also provide some explicit conditions to guarantee the strongly convex property of J.

Proposition 4.6. *Suppose that Assumptions 2.1, 2.2, and 2.3 hold. Then the objective functional J is strongly convex on its effective domain, if one of the following conditions holds.*

(i) The function $(w, a) \mapsto f(t, w, a)$ is strongly convex. That is, there exist a constant $\kappa_f > 0$ such that, for almost every $t \in [0, T]$, every $(w, a), (w', a') \in \mathbb{R} \times A$, and every $\theta \in (0, 1)$,

$$\begin{aligned} f\big(t, \theta w + (1-\theta)w', \theta a + (1-\theta)a'\big) &\leq \theta f(t, w, a) + (1-\theta) f(t, w', a') \\ &\quad - \frac{\kappa_f}{2}\theta(1-\theta)\left(|w - w'|^2 + |a - a'|^2\right). \end{aligned}$$

(ii) There exists $\underline{\nu} > 0$ such that $\sigma(t,s)\sigma(t,s)^\top \succeq \underline{\nu} I_m$, for all $(t,s) \in [0,T] \times \mathbb{R}^m$, and the function g is strongly convex on $\mathbb{R}$. More precisely, there exists a constant $\kappa_g > 0$ such that for every $w, w' \in \mathbb{R}$ and every $\theta \in (0,1)$,

$$g\big(\theta w + (1-\theta)w'\big) \le \theta g(w) + (1-\theta) g(w') - \frac{\kappa_g}{2}\theta(1-\theta)\left(|w-w'|^2\right).$$

Proof. Let $\varphi, \psi \in \mathcal{A}$ satisfy $J(\varphi), J(\psi) < \infty$, and let $\theta \in (0,1)$. Set $\varphi^\theta := \theta\varphi + (1-\theta)\psi$. As in the proof of Proposition 3.5, the affine wealth dynamics imply $W^{\varphi^\theta} = \theta W^\varphi + (1-\theta)W^\psi$.

Suppose first that Condition i holds. By the strong convexity of f and the convexity of g,

$$\begin{aligned} J(\varphi^\theta) &\le \theta J(\varphi) + (1-\theta)J(\psi) \\ &\quad - \frac{\kappa_f}{2}\theta(1-\theta)\mathbb{E}\int_0^T \left(|W_t^\varphi - W_t^\psi|^2 + |\varphi_t - \psi_t|^2\right) dt. \end{aligned}$$

Dropping the first nonnegative term gives

$$J(\varphi^\theta) \le \theta J(\varphi) + (1-\theta)J(\psi) - \frac{\kappa_f}{2}\theta(1-\theta)\|\varphi-\psi\|^2_{L^2_{\mathbb{F}}}.$$

Hence J is κ_f-strongly convex.

Now suppose that Condition ii holds. Combining (3.4), the strong convexity of g and the convexity of the running-cost term therefore implies

$$J(\varphi^\theta) \le \theta J(\varphi) + (1-\theta)J(\psi) - \frac{c\kappa_g}{2}\theta(1-\theta)\|\varphi-\psi\|^2_{L^2_{\mathbb{F}}},$$

where $c := \frac{\nu}{2}e^{-KT}$. Thus J is $c\kappa_g$-strongly convex. □

Corollary 4.7 (Strong convergence of the computed controls). *Suppose that Assumptions 2.1, 2.2, 2.3, 3.11, and 4.5 hold. Then*

$$\widehat{\varphi} \to \varphi^\star \qquad \textit{strongly in } L^2_{\mathbb{F}}.$$

Proof. Let $\lambda^\star \in \mathcal{Z}$. By strong duality and primal recovery, there exists $\eta^\star \in E$ such that $(\varphi^\star, \eta^\star)$ minimizes $\mathcal{L}(\cdot,\cdot,\lambda^\star)$. Since C is jointly convex, $\mathcal{L}(\cdot,\cdot,\lambda^\star)$ is κ-strongly convex in its control component. Consequently,

$$\frac{\kappa}{2}\|\widehat{\varphi} - \varphi^\star\|^2_{L^2_{\mathbb{F}}} \le \mathcal{L}(\widehat{\varphi}, \widehat{\eta}, \lambda^\star) - q(\lambda^\star).$$

Using $q(\lambda^\star) = p^\star$, the right-hand side equals

$$J(\widehat{\varphi}) - p^\star + \lambda^\star C(\widehat{\varphi}, \widehat{\eta}),$$

which converges to zero by Theorem 4.4. Hence

$$\widehat{\varphi} \to \varphi^\star \qquad \text{strongly in } L^2_{\mathbb{F}}.$$

□

5 Numerical Experiments

This section evaluates Algorithm 2 in a scalar linear–quadratic Black–Scholes model. This benchmark isolates the dynamic effect of a terminal CVaR constraint, first in a complete market, and then in the presence of non-traded risks.

Let the traded stock be driven by a Brownian motion B, and let $B^{\perp}$ be an independent Brownian motion driving the nontraded risk exposure. With $r = 0$,

$$dS_t = S_t(\mu\, dt + \sigma\, dB_t), \qquad \mu > 0, \quad \sigma > 0, \quad S_0 > 0, \tag{5.1}$$

and the endowment dynamic satisfies $d\zeta_t = \beta^{\perp} dB_t^{\perp}$. With φ_t denoting dollar risky exposure, wealth evolves as

$$dW_t^{\varphi} = \varphi_t \mu\, dt + \varphi_t \sigma\, dB_t + \beta^{\perp} dB_t^{\perp}, \qquad W_0^{\varphi} = w_0. \tag{5.2}$$

The CVaR-constrained portfolio management problem is therefore:

$$\begin{aligned} \inf_{\varphi \in \mathcal{A}} \quad & J(\varphi) = \mathbb{E}\left[\int_0^T \left\{\frac{\gamma}{2}\left[(\sigma\varphi_t)^2 + (\beta^{\perp})^2\right] - \mu\varphi_t\right\} dt\right] \\ \text{subject to} \quad & \mathrm{CVaR}_{\alpha}(-W_T^{\varphi}) \le c. \end{aligned} \tag{5.3}$$

For $c < 0$, the constraint requires average terminal wealth in the worst $(1-\alpha)$ fraction of outcomes to be at least $-c$; it is not a pathwise wealth floor. When the CVaR constraint is nonbinding, pointwise minimization gives the Merton dollar exposure

$$\bar{\varphi}_t = \frac{\mu}{\gamma\sigma^2}. \tag{5.4}$$

The independent endowment shock adds variance but does not change (5.4). When the CVaR constraint is binding, however, the investor may reduce traded exposure to offset its effect on the lower tail.

Table 1 reports the common calibration. The nonbinding and binding CVaR limits are $c = -0.86$ and $c = -0.94$, respectively. For fixed (λ, η), the inner dynamic program computes a Markov feedback control on a discretized state grid. We optimize η by golden-section search and use an outer bisection in λ to enforce the CVaR constraint. Code and further implementation details are available at https://github.com/xf-shi/Dynamic-Portfolio-under-CVaR.

Table 1: Common parameters used in the numerical experiments.

Parameter	Value
Horizon	$T = 1$
Risk-free rate	$r = 0.00$
Excess return	$\mu = 0.08$
Volatility	$\sigma = 0.20$
Initial stock price	$S_0 = 1$
Risk aversion	$\gamma = 5$
Initial wealth	$w_0 = 1$
Merton exposure	$\bar{\varphi}_t = 0.40$
CVaR confidence level	$\alpha = 0.95$
Nonbinding CVaR limit	$c = -0.86$
Binding CVaR limit	$c = -0.94$

5.1 Complete Market Benchmark

For the complete-market benchmark, we suppress the orthogonal factor $B^\perp$ and take the filtration to be generated by the traded Brownian motion B. Equivalently, there is no nontraded endowment risk and the single traded risky asset spans the single source of market uncertainty. In the complete market, the nontraded risk exposure is $\beta^\perp = 0$. Table 2 shows that the constraint is nonbinding at $c = -0.8600$: the Merton strategy has terminal loss CVaR $-0.8671 < -0.8600$ and $\widehat{\lambda} = 0$. In the binding case with $c = -0.9400$, the calibrated multiplier is 0.0720. Average risky exposure falls from 0.4000 to 0.2981, whereas expected terminal wealth declines by only about 0.8%.

Case	c	$\widehat{\lambda}$	$\mathrm{CVaR}_\alpha(-W_T^\star)$	$\mathbb{E}[W_T^\star]$	average $\varphi_t^\star$
Nonbinding	-0.8600	0	-0.8671	1.0322	0.4000
Binding	-0.9400	0.0720	-0.9406	1.0242	0.2981

Table 2: Summary statistics for the complete-market benchmark in the nonbinding and binding cases of the terminal CVaR constraint.

Figure 1 shows that the policy for the binding case is not a constant rescaling of the Merton exposure. It returns toward 0.400 along the selected favorable path, which ends at $W_T^\star \approx 1.311$, but falls to about 0.068 late in the selected unfavorable path, which ends at $W_T^\star \approx 0.947$. Thus the binding constraint induces state-dependent de-risking while retaining participation in favorable states.

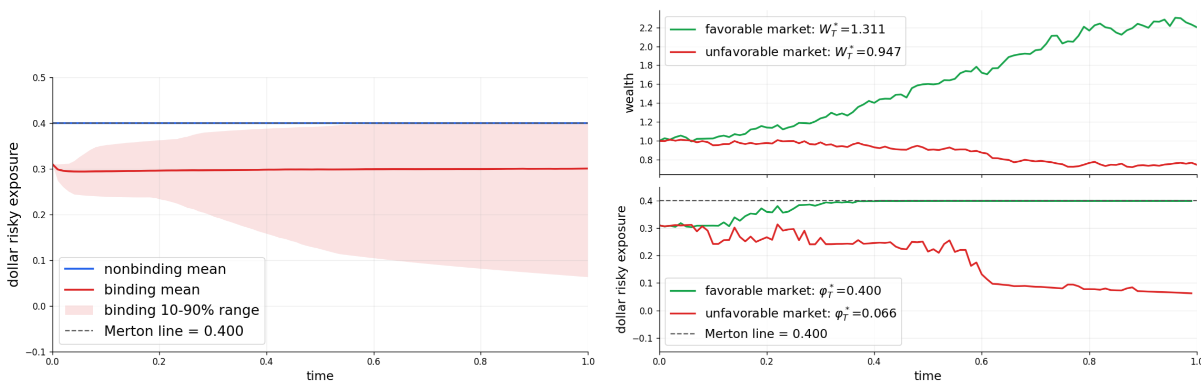

Figure 1: Optimal dollar risky exposure in the complete market. Left: mean exposure in the binding case and its 10–90% simulation band, with the Merton exposure from the nonbinding case. Right: selected favorable (green) and unfavorable (red) wealth paths and their associated exposures.

The policy for the binding case primarily compresses the lower tail of terminal wealth (Figure 2). The right panel also reports the unconditional diagnostic $\mathrm{CVaR}_\alpha(-W_t^\star)$ for $t < T$. In this zero-rate, no-inflow calibration it remains below the numerical level c throughout the horizon. This observation is not a dynamic constraint and need not persist under other calibrations.

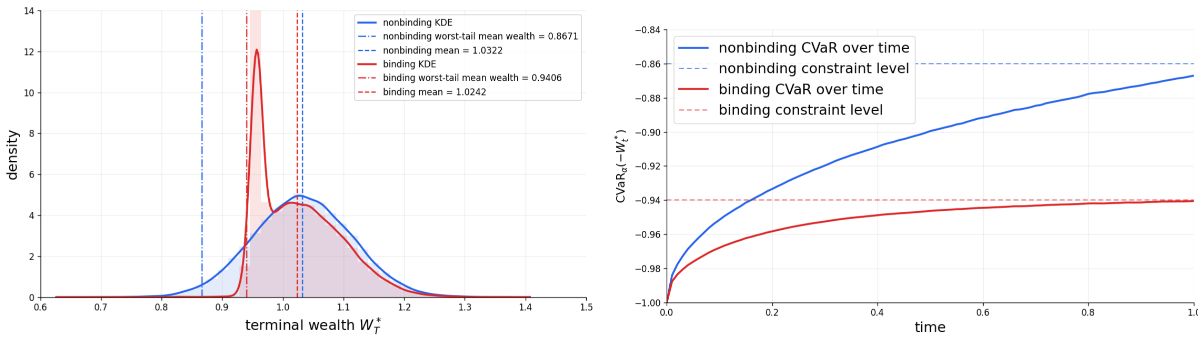

Figure 2: Wealth and lower-tail risk in the complete market. Left: terminal wealth distributions in the nonbinding and binding cases. Right: $\mathrm{CVaR}_\alpha(-W_t^\star)$ as a diagnostic over time, with the terminal constraint levels.

This lower-tail reshaping is qualitatively related to complete-market mean-risk and quantile models [14, 16, 17]. Here, however, the quadratic variation penalty regularizes the solution: the binding constraint leads to smooth de-risking that depends on the state rather than to a digital or lottery-like terminal payoff.

5.2 Incomplete Market

We set the nontraded risk exposure to be $\beta^\perp = 0.02$. In this incomplete market setting, the algorithm of Section 4 remains applicable.

In the nonbinding case, nontraded risk leaves the Merton exposure unchanged but moves terminal loss CVaR from -0.8671 to -0.8620. In the binding case with $c = -0.9400$, the calibrated multiplier is 0.1400 and average risky exposure falls to 0.2549, which is 14.5% below the complete-market value in the binding case and 36.3% below the Merton exposure. Expected terminal wealth falls by about 1.1% relative to the policy for the nonbinding case.

Case	c	$\widehat{\lambda}$	$\mathrm{CVaR}_\alpha(-W_T^\star)$	$\mathbb{E}[W_T^\star]$	average $\varphi_t^\star$
Nonbinding	-0.86	0.0000	-0.8620	1.0322	0.4000
Binding	-0.94	0.1400	-0.9405	1.0206	0.2549

Table 3: Summary statistics for the incomplete market in the nonbinding and binding cases of the terminal CVaR constraint.

Figure 3 displays the resulting feedback response in the binding case. Along the selected favorable path, wealth ends at approximately 1.237 and exposure returns to the Merton level after $t = 0.6$. Along the selected unfavorable path, wealth ends near 0.956 and average exposure after $t = 0.6$ falls to approximately 0.043. The investor therefore offsets unhedgeable tail risk indirectly by reducing traded risk in unfavorable states.

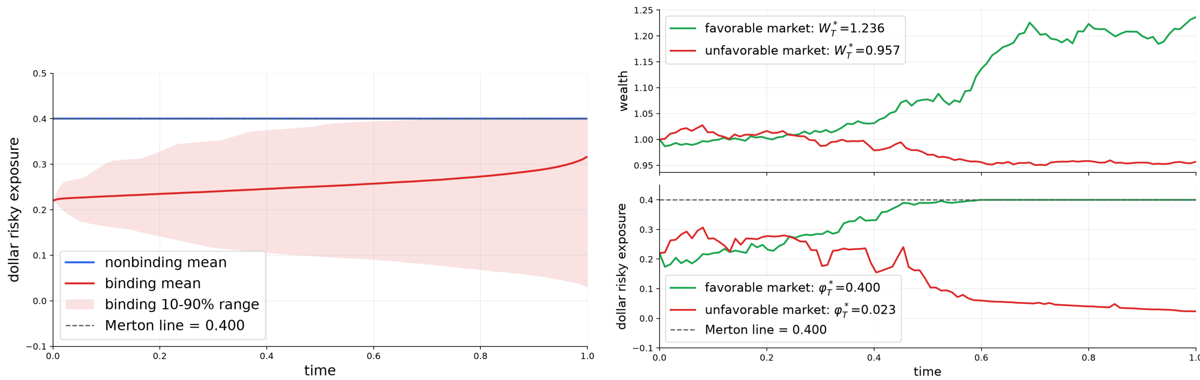

Figure 3: Optimal dollar risky exposure in the incomplete market. Left: mean exposure in the binding case and its 10–90% simulation band, with the Merton exposure from the nonbinding case. Right: selected favorable (green) and unfavorable (red) wealth paths and their associated exposures.

The distributional comparison in Figure 4 confirms that the policy for the binding case removes mass from severe low-wealth outcomes. The right-panel trajectories are interim diagnostics only; the optimization imposes the CVaR restriction at T, not at intermediate dates.

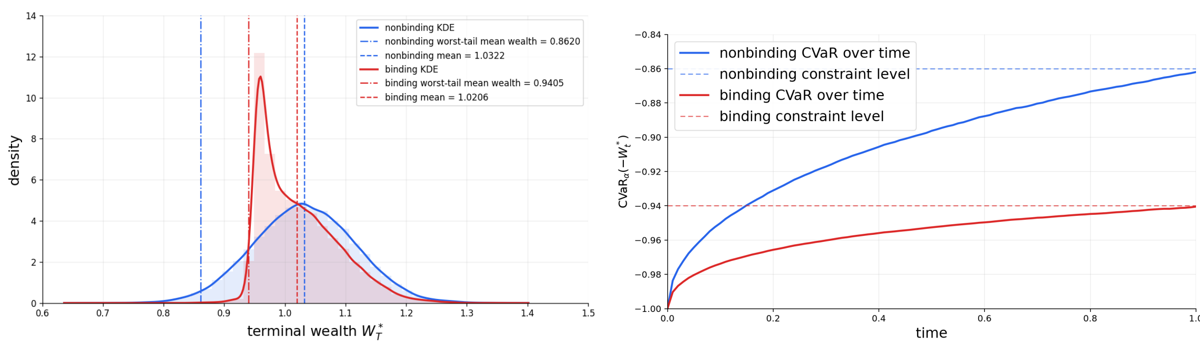

Figure 4: Wealth and lower-tail risk in the incomplete market. Left: terminal wealth distributions in the nonbinding and binding cases. Right: $\mathrm{CVaR}_\alpha(-W_t^\star)$ as a diagnostic over time, with the terminal constraint levels.

6 Portfolio Adjustment and Price Impact

With the presence of trading regularization or existence of price impact, the dollar risky amount becomes a state variable. We keep φ_t to represent the dollar risky amount, and the control $\dot{\varphi}$ is referred to as the signed dollar trading rate, where the dynamic between the dollar risky amount φ and the control $\dot{\varphi}$ remains to be linear.

$$\begin{aligned} d\varphi_t &= \dot{\phi}_t S_t dt + \phi_t dS_t \\ &= \dot{\varphi}_t dt + \varphi_t \left(\mu dt + \sigma dB_t\right), \qquad \varphi_0 = 0, \end{aligned} \tag{6.1}$$

where recall that ϕ_t stands for the total shares hold by the investor, and $\dot{\phi}_t$ represents her signed trading rate in shares. Although the results in Sections 2–4 are stated for a scalar wealth state and direct portfolio controls, the same arguments extend directly to finite-dimensional affine controlled state dynamics under the corresponding convexity, compactness, and integrability conditions. We use this extension in Section 6.1.

The control $\dot{\phi}$ is used in both specifications below, but the economic interpretations and theoretical properties of the two specifications differ. In Section 6.1, the quadratic term is an objective regularizer that smooths portfolio adjustment; it is not an execution cost and is therefore not deducted from wealth. The augmented state dynamics remain affine in the trading-rate control, and the quadratic regularizer preserves the convex structure underlying the analysis in Sections 3 and 4. In Section 6.2, by contrast, price impact changes the execution price and the resulting cost is deducted directly from wealth. The resulting wealth dynamics depend nonlinearly on the trading-rate control through $|\dot{\phi}_t|^{3/2}$ and therefore fall outside the affine-state setting of Sections 3 and 4. We use the square-root specification as a numerical robustness experiment beyond the setting covered by our convergence theory.

6.1 Quadratic Trading Rate Regularization

We first add a quadratic penalty on the signed dollar trading rate to smooth portfolio adjustment:

$$J(\dot{\varphi}) = \mathbb{E}\left[\int_0^T \left(\frac{\gamma}{2}\sigma^2\varphi_t^2 + \frac{\Lambda}{2}\dot{\varphi}_t^2 - \mu\varphi_t\right) dt\right]. \tag{6.2}$$

Since the regularization only shows up in the target functional, the wealth dynamics remain the same as (5.2). The CVaR-constrained control problem is

$$\begin{aligned} &\inf_{\dot{\varphi}\in\mathcal{A}} \quad J(\dot{\varphi}) \\ &\text{subject to} \quad \mathrm{CVaR}_\alpha(-W_T) \le c. \end{aligned} \tag{6.3}$$

The term $\Lambda\dot{\varphi}_t^2/2$ is a regularization term in the performance criterion, which penalizes abrupt changes in the risky position and produces smoother trading policies. We use the calibration from Table 1 and set $\Lambda = 0.01$; larger values of Λ place greater weight on smooth portfolio adjustment.

Table 4 shows that the CVaR constraint is nonbinding at $c = -0.86$. In the binding case, average exposure falls from 0.3104 to 0.2415 and $\widehat{\lambda}$ rises to 0.06001. Expected terminal wealth declines from 1.0247 to 1.0192, while terminal loss CVaR reaches -0.9413. The policy for the binding case uses the calibrated pair $(\lambda, \eta) = (0.06001, -0.947499)$.

Case	c	$\widehat{\lambda}$	$\mathrm{CVaR}_\alpha(-W_T^\star)$	$\mathbb{E}[W_T^\star]$	average $\varphi_t^\star$
Nonbinding	-0.86	0.0000	-0.8971	1.0247	0.3104
Binding	-0.94	0.06001	-0.9413	1.0192	0.2415

Table 4: Summary statistics with quadratic trading-rate regularization in the nonbinding and binding cases of the terminal CVaR constraint.

Figure 5 shows that both policies trade most rapidly near the initial date and that their mean trading rates approach zero near maturity. The policy for the binding case trades less aggressively and maintains a lower exposure profile. The exposure bands remain nondegenerate in the nonbinding case because $\dot{\varphi}$ contains the stock-price diffusion in (6.1).

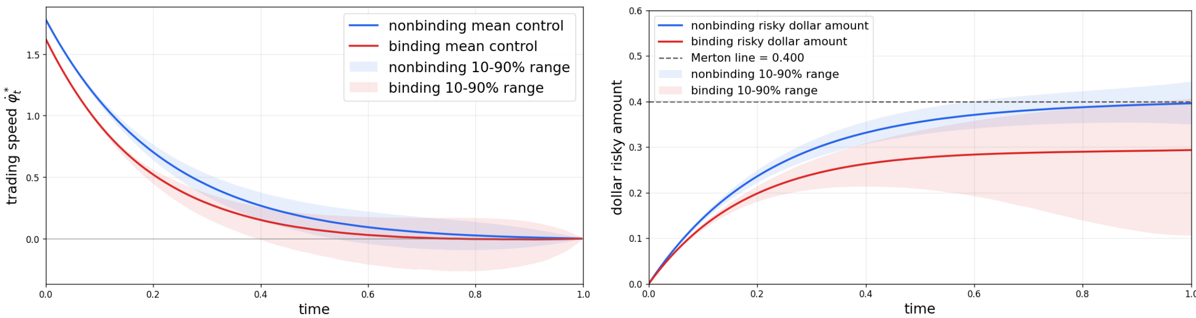

Figure 5: Signed dollar trading rate and dollar risky exposure under the quadratically regularized specification. Left: mean trading rates and 10–90% simulation bands. Right: associated exposure profiles and bands.

In the binding case, terminal wealth is more concentrated and has less downside mass (Figure 6). The right panel reports $\mathrm{CVaR}_\alpha(-W_t)$ as an interim diagnostic. The constraint is imposed at maturity.

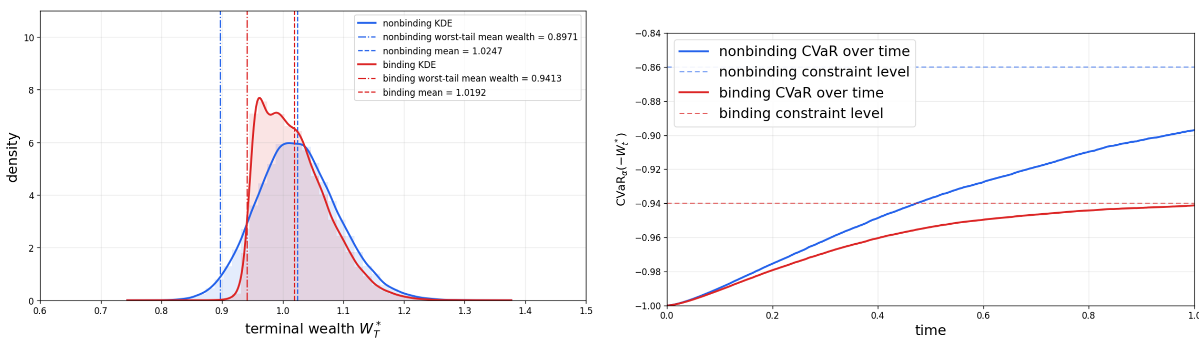

Figure 6: Wealth and lower-tail risk under quadratic trading-rate regularization. Left: terminal-wealth distributions. Right: out-of-sample $\mathrm{CVaR}_\alpha(-W_t)$ diagnostics and terminal constraint levels.

The selected paths in Figure 7 illustrate the feedback mechanism under the binding constraint. The favorable path ends at wealth 1.085 and approaches the Merton exposure, whereas the unfavorable path ends near 0.942 and reduces exposure to approximately 0.064 late in the horizon. The trading rates along both paths decrease toward zero as maturity approaches. In the nonbinding case, the trading rate converges to zero as maturity approaches because further repositioning offers no remaining expected-return benefit while still incurring the quadratic trading-rate penalty. In the binding case, the last reported preterminal rate remains slightly nonzero because the terminal-loss term associated with the CVaR constraint creates an additional incentive to adjust exposure over the final decision interval.

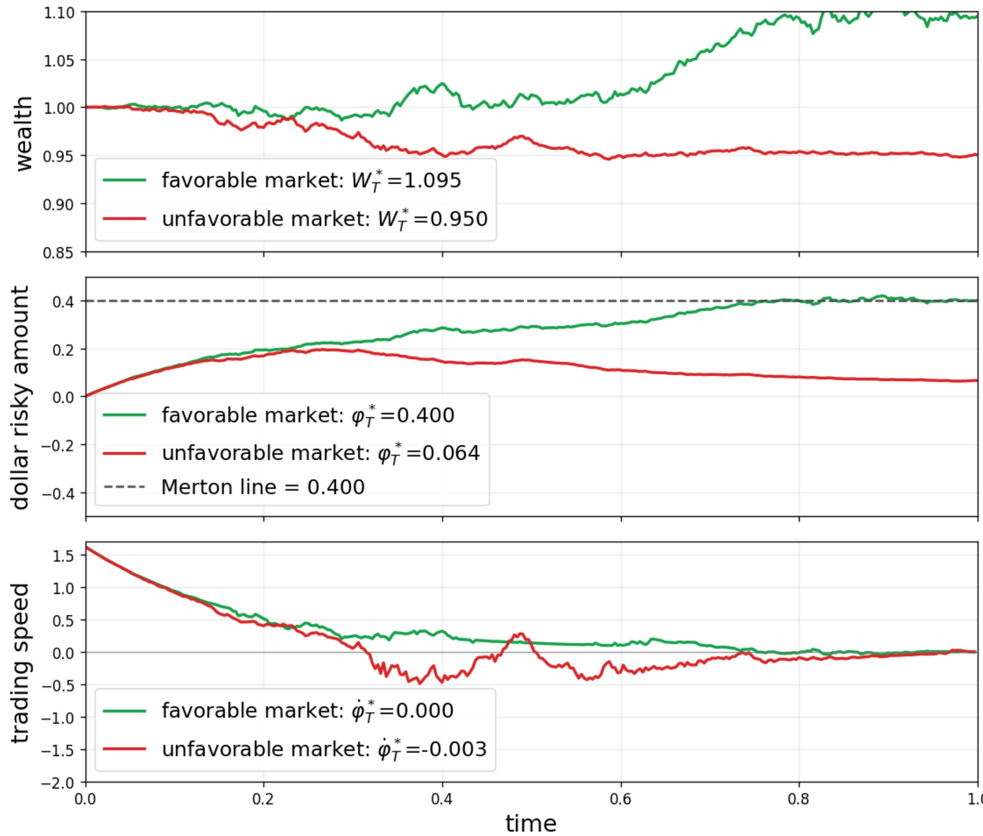

Figure 7: Wealth, dollar risky exposure, and signed dollar trading rate along selected favorable (green) and unfavorable (red) realizations in the binding case.

6.2 Square-Root Price Impact

We retain the same market and initial position, but now let execution price respond to signed dollar trading rate according to the square-root specification. When executing a stock, the execution price S^{exe} is usually different from the market price S_t, i.e.

$$S_t^{\text{exe}} = S_t \left[1 + \frac{\Lambda_{3/2}}{3/2} \operatorname{sign}(\dot\varphi_t)\sqrt{|\dot\varphi_t|}\right]. \tag{6.4}$$

The placement of the brackets in (6.4) ensures that $S_t^{\text{exe}} = S_t$ when $\dot\varphi_t = 0$. Since $\dot\varphi_t = \dot\phi_t S_t$, execution costs reduce wealth as follows:

$$\begin{aligned} dW_t &= \phi_t\, dS_t - (S_t^{\text{exe}} - S_t)\dot\phi_t\, dt \\ &= \varphi_t(\mu\, dt + \sigma\, dB_t) - \frac{\Lambda_{3/2}}{3/2}|\dot\varphi_t|^{3/2}\, dt. \end{aligned} \tag{6.5}$$

The corresponding portfolio management problem is

$$\begin{aligned} \inf_{\dot\varphi\in\mathcal{A}} \quad & \mathbb{E}\left[\int_0^T \left(\frac{\gamma}{2}\sigma^2\varphi_t^2 + \frac{\Lambda_{3/2}}{3/2}|\dot\varphi_t|^{3/2} - \mu\varphi_t\right) dt\right] \\ \text{subject to} \quad & \mathrm{CVaR}_\alpha(-W_T) \le c. \end{aligned} \tag{6.6}$$

Although the cost $|\dot{\varphi}|^{3/2}$ is convex, its nonlinear appearance in the wealth dynamics places this specification outside the affine-state setting of Sections 2–4. The convergence guarantees established above therefore do not apply directly. We set $\Lambda_{3/2} = 0.0045$ and use the remaining parameters from Table 1.

In the nonbinding case, terminal loss CVaR is -0.8839 and the multiplier is zero. In the binding case, the policy reduces average exposure from 0.3500 to 0.2422 and mean terminal wealth from 1.0260 to 1.0180, as shown in Table 5.

Case	c	$\widehat{\lambda}$	$\mathrm{CVaR}_\alpha(-W_T^\star)$	$\mathbb{E}[W_T^\star]$	average $\varphi_t^\star$
Nonbinding	-0.86	0.0000	-0.8839	1.0260	0.3500
Binding	-0.94	0.07818	-0.9409	1.0180	0.2422

Table 5: Summary statistics for square-root price impact in the nonbinding and binding cases of the terminal CVaR constraint.

Figure 8 shows faster initial accumulation under the policy for the nonbinding case. Under the policy for the binding case, later trading rates can become negative in unfavorable states, indicating partial liquidation, while mean exposure stabilizes near 0.24. Diffusion in (6.1) generates widening exposure bands under both policies.

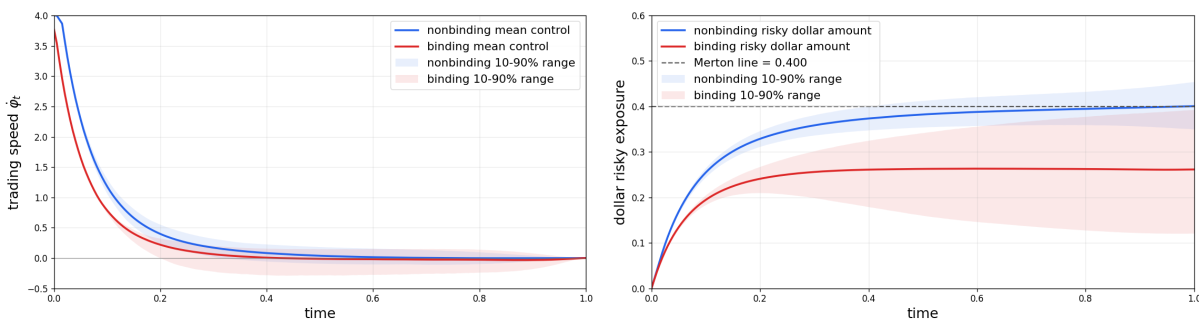

Figure 8: Trading and exposure under square-root price impact. Left: mean signed dollar trading rates and 10–90% simulation bands. Right: associated dollar risky exposures and bands, with the Merton reference.

The policy for the binding case compresses the terminal-wealth distribution and reduces extreme outcomes (Figure 9). Out-of-sample terminal loss CVaR is -0.9409, below the required level -0.9400.

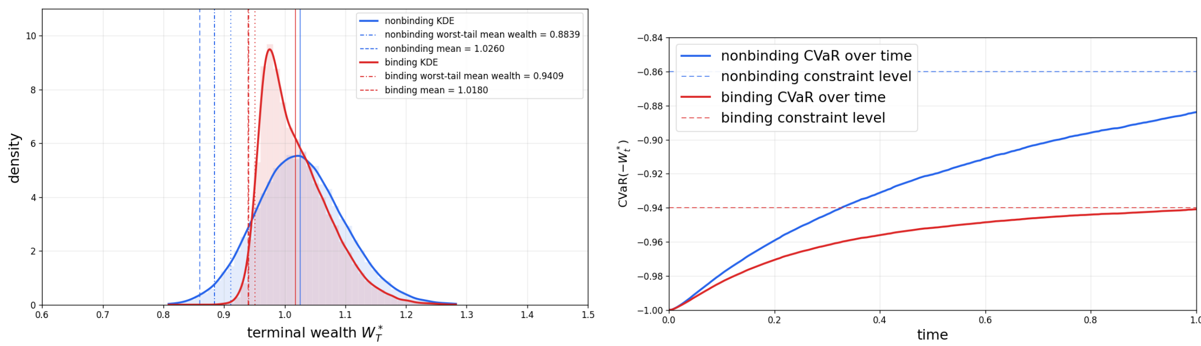

Figure 9: Wealth and tail risk under 3/2-power trading costs. Left: terminal wealth distributions. Right: out-of-sample loss-CVaR diagnostics and terminal constraint levels.

In Figure 10, the selected favorable path under the binding constraint ends at wealth 1.112 and returns toward the Merton exposure. The selected unfavorable path ends near 0.947 and reduces late-horizon exposure to approximately 0.097; its trading rate is negative over several intervals. Trading rates approach zero near maturity in both paths.

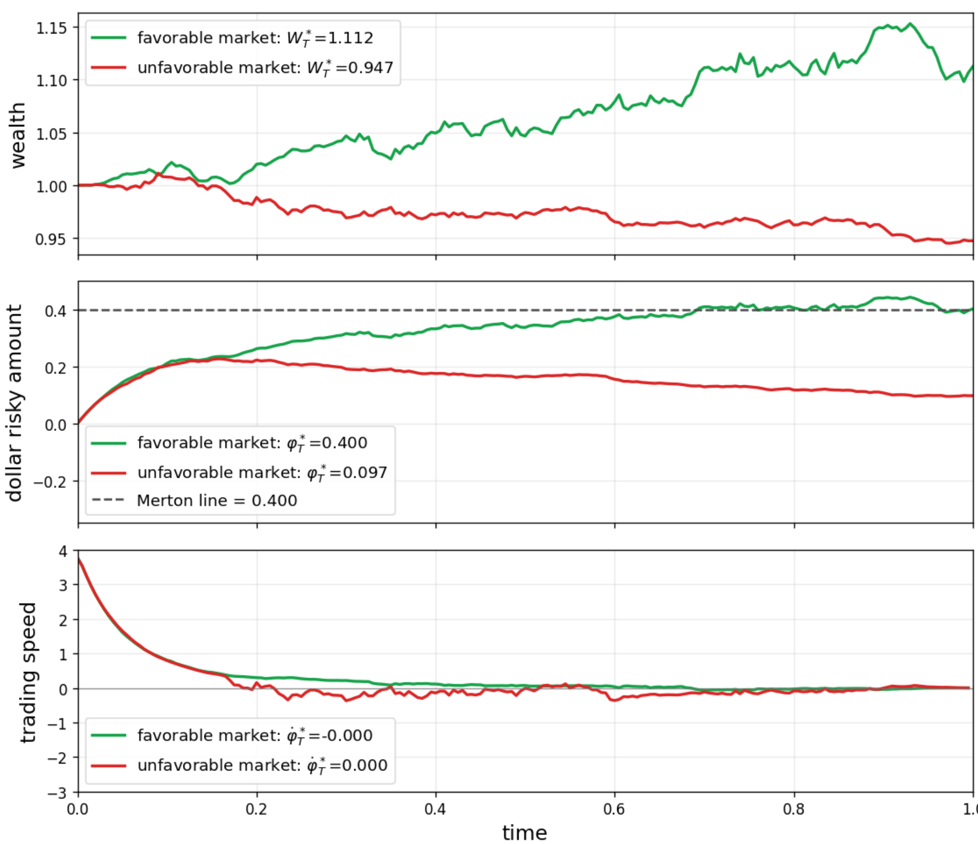

Figure 10: Wealth, dollar risky exposure, and signed dollar trading rate along selected favorable (green) and unfavorable (red) realizations in the binding case.

The 3/2-power specification produces a mean exposure in the binding case close to that under quadratic regularization but a different adjustment pattern: it permits faster initial accumulation and uses state-dependent sales when downside risk increases. Its held-out feasibility provides numerical evidence that the method remains effective in this nonlinear price-impact specification, although the convergence guarantees do not apply directly.

7 Conclusion

This paper studies a dynamic portfolio management problem under an explicit CVaR constraint on terminal loss. Using the auxiliary-threshold representation of CVaR, we obtain a convex formulation that accommodates general market settings, including incomplete market, nontraded endowment risks, and general convex trading objectives. We establish existence and strong duality and develop a modular numerical method that combines standard unconstrained stochastic-control problems with one-dimensional searches over the CVaR threshold and Lagrange multiplier, which is shown to converge globally.

The numerical experiments show that a binding terminal CVaR constraint does not induce uniform de-risking. Instead, the optimal policy reduces risky exposure following adverse outcomes while preserving participation in favorable states. Nontraded endowment risk strengthens this adjustment when the constraint binds, whereas trading frictions lower the desired exposure and slow portfolio adjustment. Across the environments considered, stronger lower-tail protection is achieved at a comparatively modest cost in expected terminal wealth. These findings illustrate how a terminal risk constraint can shape state-dependent portfolio decisions throughout the investment horizon.

Acknowledgments

SP and XS acknowledge financial support from the Natural Sciences and Engineering Research Council of Canada (RGPIN-2025-05847, RGPIN-2024-04569). AH thanks the Fields Institute for the FOCUS visitor support program and acknowledges financial support by InnoHK initiative, The Government of the HKSAR, and Laboratory for AI-Powered Financial Technologies.

References

[1] C. Acerbi and P. Simonetti. Portfolio optimization with spectral measures of risk. *arXiv preprint cond-mat/0203607*, 2002.

[2] C. Acerbi and D. Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, 2002.

[3] G. J. Alexander and A. M. Baptista. A comparison of VaR and CVaR constraints on portfolio selection with the mean-variance model. *Management Science*, 50(9):1261–1273, 2004.

[4] S. Basak and A. Shapiro. Value-at-risk-based risk management: Optimal policies and asset prices. *The Review of Financial Studies*, 14(2):371–405, 2001.

[5] S. Boyd and L. Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

[6] P. Cannarsa and R. Capuani. Existence and uniqueness for mean field games with state constraints. In *PDE models for multi-agent phenomena*, pages 49–71. Springer, 2018.

[7] Z. Chen. Multiperiod consumption and portfolio decisions under the multivariate GARCH model with transaction costs and CVaR-based risk control. *OR Spectrum*, 27(4):603–632, 2005.

[8] D. Cuoco, H. He, and S. Isaenko. Optimal dynamic trading strategies with risk limits. *Operations Research*, 56(2):358–368, 2008.

[9] M. Dai, Y. Jia, and H. Jin. Dynamic mean-variance portfolio selection with transaction costs. *Available at SSRN 4958481*, 2024.

[10] S. Daudin. Optimal control of diffusion processes with terminal constraint in law. *Journal of Optimization Theory and Applications*, 195(1):1–41, 2022.

[11] C. Edirisinghe, J. Chen, and J. Jeong. Optimal leveraged portfolio selection under quasi-elastic market impact. *Operations Research*, 71(5):1558–1576, 2023.

[12] S. Emmer, C. Klüppelberg, and R. Korn. Optimal portfolios with bounded capital at risk. *Mathematical Finance*, 11(4):365–384, 2001.

[13] Z. Feinstein and B. Rudloff. A recursive algorithm for multivariate risk measures and a set-valued Bellman's principle. *Journal of Global Optimization*, 68(1):47–69, 2017.

[14] J. Gao, K. Zhou, D. Li, and X. Cao. Dynamic mean-LPM and mean-CVaR portfolio optimization in continuous-time. *SIAM Journal on Control and Optimization*, 55(3):1377–1397, 2017.

[15] A. Gundel and S. Weber. Utility maximization under a shortfall risk constraint. *Journal of Mathematical Economics*, 44(11):1126–1151, 2008.

[16] X. D. He, H. Jin, and X. Y. Zhou. Dynamic portfolio choice when risk is measured by weighted VaR. *Mathematics of Operations Research*, 40(3):773–796, 2015.

[17] X. D. He and X. Y. Zhou. Portfolio choice via quantiles. *Mathematical Finance*, 21(2):203–231, 2011.

[18] A. Hu and Z. Lyu. Mean-field games with constraints. *arXiv preprint arXiv:2510.11843*, 2025.

[19] S. Jaimungal, S. M. Pesenti, Y. S. Wang, and H. Tatsat. Robust risk-aware reinforcement learning. *SIAM Journal on Financial Mathematics*, 13(1):213–226, 2022.

[20] Y.-H. Kim. On the p-th moment estimates for the solution of stochastic differential equations. *Journal of Inequalities and Applications*, 2014(1):395, 2014.

[21] P. Krokhmal, J. Palmquist, and S. Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *The Journal of Risk*, 4(2):11–27, 2002.

[22] X. Mao. *Stochastic differential equations and applications*. Elsevier, 2007.

[23] R. C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, Aug. 1969.

[24] R. C. Merton. An analytic derivation of the efficient portfolio frontier. *Journal of financial and quantitative analysis*, 7(4):1851–1872, 1972.

[25] S. Moreno-Bromberg, T. A. Pîrvu, and A. Réveillac. CRRA utility maximization under dynamic risk constraints. *Communications on Stochastic Analysis*, 7(2):203–225, 2013.

[26] J. Muhle-Karbe, J. Sefton, and X. Shi. Dynamic portfolio choice with intertemporal hedging and transaction costs. *Management Science*, dec 2025. Articles in Advance.

[27] S. M. Pesenti and S. Jaimungal. Portfolio optimization within a Wasserstein ball. *SIAM Journal on Financial Mathematics*, 14(4):1175–1214, 2023.

[28] S. M. Pesenti and T. Nguyen. Outperforming a benchmark with α-Bregman Wasserstein divergence. *arXiv preprint arXiv:2603.20580*, 2026.

[29] S. M. Pesenti, S. Vanduffel, Y. Yang, and J. Yao. Optimal payoff under Bregman-Wasserstein divergence constraints. *arXiv preprint arXiv:2411.18397*, 2024.

[30] L. Pfeiffer, X. Tan, and Y.-L. Zhou. Duality and approximation of stochastic optimal control problems under expectation constraints. *SIAM Journal on Control and Optimization*, 59(5):3231–3260, 2021.

[31] T. A. Pirvu. Portfolio optimization under the value-at-risk constraint. *Quantitative Finance*, 7(2):125–136, 2007.

[32] T. A. Pirvu and G. Žitković. Maximizing the growth rate under risk constraints. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 19(3):423–455, 2009.

[33] R. T. Rockafellar and S. Uryasev. Optimization of Conditional Value-at-Risk. *The Journal of Risk*, 2(3):21–41, 2000.

[34] R. T. Rockafellar and S. Uryasev. Conditional Value-at-Risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471, 2002.

[35] M. S. Strub, D. Li, X. Cui, and J. Gao. Discrete-time mean-CVaR portfolio selection and time-consistency induced term structure of the CVaR. *Journal of Economic Dynamics and Control*, 108:103751, 2019.

[36] G. Tsoukalas, J. Wang, and K. Giesecke. Dynamic portfolio execution. *Management Science*, 65(5):2015–2040, 2019.