S-shaped Utility Maximization with VaR Constraint and Partial Information

Dongmei Zhu*, Ashley Davey†, Harry Zheng‡

Abstract

We study S-shaped utility maximisation with VaR constraint and unobservable drift coefficient. Using the Bayesian filter, the concavification principle, and the change of measure, we give a semi-closed integral representation for the dual value function and find a critical wealth level that determines if the constrained problem admits a unique optimal solution and Lagrange multiplier or is infeasible. We also propose three algorithms (Lagrange, simulation, deep neural network) to solve the problem and compare their performances with numerical examples.

Communicated by Nizar Touzi.
Keywords. S-shaped utility maximization, VaR constraint, Partial information, Bayesian filter, Dual control, Physics informed neural network.
2020 Mathematics Subject Classification. 93E20, 93E11, 91G80, 90C46, 49M29

1 Introduction

Optimal portfolio via expected utility maximization has been extensively studied, see Pham [19] for expositions. The S-shaped utility has drawn particularly great attention since the ground-breaking work of Kahneman and Tversky[17] on the prospect theory. Carpenter [6] is the first in solving S-shaped utility maximization with the concavification principle in a complete market. There are many papers in the literature on the subject, for example, Berkelaar et al. [4] incorporate prospect theory and derive closed-form solutions for optimal portfolio choice under loss aversion; Jin and Zhou [16] study a general continuous-time behavioural portfolio selection model with S-shaped utility and probability distortion; Ingersoll and Jin [15] discuss realization utility with reference-dependent preferences.

S-shaped utility maximization may lead to extreme loss due to its risk-seeking nature in the loss region. To mitigate this risk, one may incorporate some risk measures into the model. The most common one is the value at risk (VaR), defined as the maximum portfolio loss that may occur during a given period with a pre-set confidence level, which satisfies the regulatory requirements. There are also many papers in the literature on the subject, for example, Basak and Shapiro [2] embed the VaR into a utility maximization framework and study its implication for optimal portfolio policies. Yiu [27] imposes the VaR as a dynamic constraint and derive the optimal constrained portfolio allocation by dynamic programming technique. Chen et al. [7] focus on a utility maximization

*School of Economics and Management, Southeast University, Nangjing, China. zhudongmei@seu.edu.cn. Supported in part by the National Social Science Fund of China (Grant No.24CJY063).

†Department of Mathematics, Imperial College, London SW7 2BZ, UK. AshleyDavey@live.co.uk.

‡Department of Mathematics, Imperial College, London SW7 2BZ, UK. h.zheng@imperial.ac.uk. Supported in part by EPSRC (UK) grant (EP/V008331/1).

problem under multiple VaR-type constraints and their effects. Bensoussan et al. [3] discuss a Merton problem with an additional variance of terminal wealth term in objective function, leading to a time-inconsistent problem. Dong and Zheng [12] study S-shaped utility maximization with a VaR constraint and solve the problem by the dual control method.

Aforementioned papers assume fully observable models with deterministic coefficients or observable random coefficients. In real financial markets investors often can only observe partial information of risky assets, not full information needed for valuation and optimization, for example, stock price processes but not stock growth rates, or equity values but not firm values, etc. To circumvent these issues, the filtering theory (see Bain and Crisan [1]) is normally used to extract the information of unobservable random parameters with observable information, see Detemple [11]; Karatzas and Xue [18] for introduction of this field in asset pricing and utility maximization. There are three typical models for unobservable random parameters, including linear diffusion models, leading to the Kalman filter, see Brendle [5]; finite state Markov chain models, leading to the Wonham filter, see Rieder and Buerle [22]; Sass [23], and random vector models with prior distribution, leading to the Bayesian filter, see De Franco et al.[10]; Ekström and Vaicenavicius [13].

In this paper we extend the work of Dong and Zheng [12] to models with partial information. Specifically, we assume the drift coefficient of the risky asset is an unobservable random variable with some prior distribution. Using the Bayesian filter, we can transform the model into an equivalent fully observable one, which results in an additional filtered state process. The dual control approach in Dong and Zheng [12] no longer applies as the joint distribution of the dual and filtered state processes is unknown, which is in sharp contrast to Dong and Zheng [12] where only the distribution of the dual state process is needed and is known to be lognormally distributed, so one can easily compute the dual value function or its integral representation, then find the primal value function and optimal control with the primal-dual relation. With an additional filtered state process, the distribution of the dual process depends on that of the filtered state process and is in general unknown, which makes difficult to express the dual value function in semi-closed integral form.

To overcome the difficulty, we use a measure change to reduce the dimension of the dual state variables by one when the prior distribution of unobservable drift coefficient is a discrete distribution with two states and then characterize the dual value function with a semi-closed integral representation. We find a critical wealth level that determines if the S-shaped utility maximization with VaR constraint and partial information admits a unique optimal solution and Lagrange multiplier or is infeasible, that is, the VaR constraint is not satisfied for any admissible control strategies. We give a constructive proof of our main result, Theorem 3.1, which leads to an exact algorithm, called Lagrange algorithm, to solve the problem numerically. We also propose two other algorithms to solve the dual problem numerically, one is Monte Carlo simulation as both dual and filtered state processes can be easily simulated, albeit their joint distribution is unknown, the other is the Physics-Informed Neural Network (PINN) method (see Raissi et al. [20]) that approximates the dual value function with a neural network and uses the dual HJB equation as a loss function. Deep learning has been used for solving HJB equations for various stochastic control problems without additional constraints, see for example (Davey and Zheng [8]; Han and Jentzen [14], Wang et al. [25]). We extend the scope of these papers with an input parameter representing a Lagrange multiplier to solve control problems with additional VaR constraint.

The rest of the paper is organized as follows. In Section 2 we formulate the S-shaped utility maximization problem with VaR constraint and unobservable drift coefficient and discuss the Bayesian filter and the dual formulation. In Section 3 we state the main result of the paper, Theorem 3.1, that shows there is a critical wealth level for the existence of optimal solution and characterize the optimal terminal wealth and the corresponding Lagrange multiplier. In Section 4 we propose three algorithms (Lagrange, simulation, PINN) for solving the problem. In Section 5 we provide numerical examples with our algorithms. Section 6 concludes.

2 Model and Equivalent Problem

Let $(\Omega, \mathbb{F}, \mathcal{F}, \mathbb{P})$ be a filtered probability space, where $\mathbb{P}$ is the probability measure and the filtration $\mathbb{F} = \{\mathcal{F}_t, t \in [0, T]\}$ satisfies the usual conditions. The market consists of one riskless asset S_0 and one risky asset S, satisfying, for $0 \le t \le T$,

$$\begin{aligned} dS_0(t) &= rS_0(t)dt, \\ dS(t) &= \mu S(t)dt + \sigma S(t)dW(t), \end{aligned}$$

where $\{W(t), t \in [0, T]\}$ is a standard Brownian motion, adapted to the filtration $\mathbb{F}$, r and σ are positive constants, and μ is a $\mathcal{F}_0$ measurable random variable. We assume μ and W are unobservable and independent of each other. Let $\mathbb{F}^S = \{\mathcal{F}_t^S, t \in [0, T]\}$ be the natural filtration generated by the risky asset S, augmented with all $\mathbb{P}$-null sets in $\mathcal{F}$. The filtration $\mathbb{F}^S$ is observable and $\mathbb{F}^S \subset \mathbb{F}$. We further assume that random variable μ takes two values μ^h and μ^l with probability p and $1 - p$. To avoid triviality, we assume $\mu^l < \mu^h$ and $p \in (0, 1)$. There is only one risky asset in the market for simplicity, which can be easily generalized to multiple risky assets with correlated Brownian motions.

Let $\pi(t)$ be the proportion of wealth invested in the risky asset at time t, then the wealth process X satisfies the following stochastic differential equation (SDE):

$$dX(t) = X(t)(r + \pi(t)(\mu - r))dt + X(t)\pi(t)\sigma dW(t),\ X(0) = x_0, \quad 0 \le t \le T, \tag{2.1}$$

where π is $\mathcal{F}^S$-progressively measurable and satisfies $\mathbb{E}\left[\int_0^T |\pi(t)|^2 dt\right] < \infty$, called an *admissible control*. The set of all admissible controls is denoted by $\mathcal{A}$. We consider a general S-shaped utility function given by

$$U(x) = \begin{cases} -\infty,\ x < 0, \\ -U_2(\theta - x),\ 0 \le x < \theta, \\ U_1(x - \theta),\ x \ge \theta, \end{cases}$$

where U_1, U_2 are strictly increasing, strictly concave, continuously differentiable with $U_1(0) = U_2(0) = 0$, and θ is a positive constant. Additionally, $\lim_{x\to+\infty} U_1(x) = +\infty$, $\lim_{x\to+\infty} U_1'(x) = 0$, $\lim_{x\to+\infty} \frac{xU_1'(x)}{U_1(x)} < 1$ and $\lim_{x\to 0^+} U_i'(x) = +\infty$, for $i = 1, 2$. In what follows, I_i denotes the inverse function of U_i' for $i = 1, 2$. U is convex for $0 \le x \le \theta$ and concave for $x \ge \theta$, indicating the behavioral change from risk seeking to risk averse at a reference point θ. Our problem is to maximize the expected utility of terminal wealth with a quantile constraint:

$$\begin{cases} \sup_{\pi\in\mathcal{A}} \mathbb{E}[U(X^\pi(T)))], \\ \text{s.t. } X^\pi(t) \text{ satisfies (2.1)}, \\ \mathbb{P}(X^\pi(T) \ge L) \ge 1 - \varepsilon, \end{cases} \tag{2.2}$$

where $0 \le \varepsilon \le 1$ is a constant given in advance. In this paper, we assume $L < \theta$. The case $L \ge \theta$ can be similarly discussed.

2.1 Filtering and Equivalent Formulation

We now transform the primal problem (2.2) into an equivalent completely observable problem. Denote the filter estimate of μ by

$$\hat{\mu}(t) = \mathbb{E}[\mu|\mathcal{F}_t^S]$$

and the innovation process $\hat{W}$ by

$$\hat{W}(t) := \sigma^{-1} \int_0^t (\mu - \hat{\mu}(s))ds + W(t), \quad t \in [0,T].$$

Then $\hat{W}$ is a $(\mathbb{P}, \mathbb{F}^S)$-Brownian motion (see Sass[23]). We can rewrite equivalently the asset price process S as

$$dS(t) = \hat{\mu}(t)S(t)dt + \sigma S(t)d\hat{W}(t),$$

and the wealth process X as

$$dX(t) = X(t)(r + \pi(t)(\hat{\mu}(t) - r))dt + X(t)\pi(t)\sigma d\hat{W}(t). \tag{2.3}$$

The filtered drift process $\hat{\mu}$ satisfies the SDE:

$$d\hat{\mu}(t) = \psi(\hat{\mu}(t))d\hat{W}(t), \tag{2.4}$$

where $\psi(u) = \sigma^{-1}(u - \mu^l)(\mu^h - u)$, and $\hat{\mu}(0) = \mathbb{E}[\mu] = p\mu^h + (1-p)\mu^l \in (\mu^l, \mu^h)$.

Remark 2.1. *The unobservable drift μ is a two-point random variable, which results in a closed-form formula for function ψ, critically important in introducing the likelihood ratio process and the dimension reduction procedure, see Section 3. In general, ψ does not have a closed-form formula except in the normal or two-point distribution case. We have to use other methods to deal with it and we leave it for future research.*

We now have a fully observed control problem with state processes $X, \hat{\mu}$ and the problem (2.2) is equivalent to the following problem:

$$\begin{cases} \sup_{\pi \in \mathcal{A}} \mathbb{E}[U(X^\pi(T))], \\ \text{s.t. } X^\pi(t) \text{ satisfies (2.3)}, \\ \hat{\mu}(t) \text{ satisfies (2.4)}, \\ \mathbb{P}(X^\pi(T) \geq L) \geq 1 - \varepsilon, \end{cases} \tag{2.5}$$

We can solve problem (2.5) in two steps: First, solve an unconstrained problem:

$$\begin{cases} \sup_{\pi \in \mathcal{A}} \mathbb{E}[U_\lambda(X^\pi(T))], \\ \text{s.t. } X^\pi(t) \text{ satisfies (2.3)}, \\ \hat{\mu}(t) \text{ satisfies (2.4)}. \end{cases} \tag{2.6}$$

where

$$U_\lambda(x) := U(x) + \lambda \mathbb{I}_{\{x \geq L\}}$$

and $\lambda \geq 0$ is a Lagrange multiplier to be determined. Second, find λ^* such that the quantile constraint and the complementary slackness condition are satisfied:

$$\begin{cases} \mathbb{P}(X^{\pi^*(\lambda),\lambda}(T) \geq L) \geq 1 - \varepsilon, \\ \lambda(\mathbb{P}(X^{\pi^*(\lambda),\lambda}(T) \geq L) - 1 + \varepsilon) = 0. \end{cases} \tag{2.7}$$

The relation of problems (2.5) and (2.6) is discussed in [12, Lemma 2.3] that shows if there exists a nonnegative constant λ^* such that $X^{\pi^*,\lambda^*}(T)$ solves problem (2.6) and satisfies condition (2.7), then it also solves problem (2.5).

2.2 Concavified Utility and Dual Problem

The utility function U_λ in (2.6) is discontinuous at $x = L$ as well as nonconcave. We can use the concavification technique (see Carpenter [6]) to solve the unconstrained problem (2.6) as both state processes X and $\hat{\mu}$ are driven by the same Brownian motion $\hat{W}$ and we have a complete market model. The concave envelope of U is given by

$$U^c(x) = \begin{cases} -\infty, & x < 0, \\ c_z x - U_2(\theta), & 0 \le x < z, \\ U_1(x-\theta), & x \ge z, \end{cases} \tag{2.8}$$

where

$$c_x = U_1'(x-\theta), \ x > \theta$$

and $z > \theta$ is the unique solution to the equation

$$U_1(x-\theta) + U_2(\theta) - xU_1'(x-\theta) = 0. \tag{2.9}$$

For a fixed $\lambda \ge 0$, denote by U_λ^c the concave envelope of U_λ and V_λ^c the dual function of U_λ^c, defined by

$$V_\lambda^c(y) = \sup_{x\ge 0}\{U_\lambda^c(x) - xy\}, \quad y > 0,$$

and $x^{*,\lambda}(y)$ the maximizer of $V_\lambda^c(y)$. We can characterize U_λ^c and $x^{*,\lambda}(y)$ as follows:

Proposition 2.2. *(Dong and Zheng [12]) Let*

$$k_\lambda := \frac{U_2(\theta) - U_2(\theta - L) + \lambda}{L},$$

and $\tilde{z} \in (\theta, z)$ is the unique solution to the equation

$$U_1(x-\theta) + U_2(\theta - L) - (x-L)U_1'(x-\theta) = 0. \tag{2.10}$$

(1) If $k_\lambda > c_{\tilde{z}}$, then

$$U_\lambda^c(x) = \begin{cases} -\infty, & x < 0, \\ k_\lambda x - U_2(\theta), & 0 \le x < L, \\ c_{\tilde{z}}(x-L) - U_2(\theta - L) + \lambda, & L \le x < \tilde{z}, \\ U_1(x-\theta) + \lambda, & x \ge \tilde{z}, \end{cases}$$

and

$$x^{*,\lambda}(y) = \begin{cases} \theta + I_1(y), & y < c_{\tilde{z}}, \\ L, & c_{\tilde{z}} \le y < k_\lambda, \\ 0, & y \ge k_\lambda. \end{cases} \tag{2.11}$$

(2) If $k_\lambda \le c_{\tilde{z}}$, then

$$U_\lambda^c(x) = \begin{cases} -\infty, & x < 0, \\ c_{z_0} x - U_2(\theta), & 0 \le x < \tilde{z}_0, \\ U_1(x-\theta) + \lambda, & x \ge \tilde{z}_0, \end{cases}$$

and

$$x^{*,\lambda}(y) = \begin{cases} \theta + I_1(y), & y < c_{\tilde{z}_0}, \\ 0, & y \ge c_{\tilde{z}_0}, \end{cases} \tag{2.12}$$

where $\tilde{z}_0 \in [\tilde{z}, z]$ is the unique solution to the equation

$$U_1(x-\theta) + U_2(\theta) + \lambda - xU_1'(x-\theta) = 0. \tag{2.13}$$

If $\lambda = 0$, then $\tilde{z}_0 = z$, $\tilde{U}_\lambda^c(x)$ and $x^{\lambda,0}(y)$ are given by (2.8) and (2.12) respectively.

Figure 1 shows an example of utilities U_λ and their concavified counterparts U_λ^c.

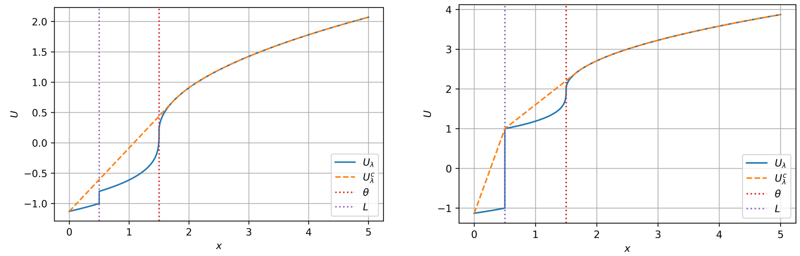

(a) One line segment concavified utility

(b) Two line segment concavified utility

Figure 1: Utility U_λ and concavified utility U_λ^c: (a) $k_\lambda \le U_1'(\tilde{z}-\theta)$ and (b) $k_\lambda > U_1'(\tilde{z}-\theta)$.

Now we consider the auxiliary stochastic control problem with fixed λ:

$$\begin{cases} \sup_{\pi\in\mathcal{A}} \mathbb{E}[U_\lambda^c(X^\pi(T))], \\ \text{s.t. } X^\pi(t) \text{ satisfies (2.3)}, \\ \hat{\mu}(t) \text{ satisfies (2.4)}. \end{cases} \tag{2.14}$$

Denote the value function of (2.14) by

$$u_\lambda^c(t,x,\hat{\mu}) := \sup_{\pi\in\mathcal{A}} \mathbb{E}[U_\lambda^c(X^\pi(T))|X^\pi(t)=x,\hat{\mu}(t)=\hat{\mu}] \tag{2.15}$$

and the constraint probability function h by

$$h_\lambda(t,x,\hat{\mu}) := \mathbb{E}[\mathbb{1}_{X^{\pi^*,\lambda}(T)\ge L}|X^\pi(t)=x,\hat{\mu}(t)=\hat{\mu}], \tag{2.16}$$

where $X^{\pi^*,\lambda}$ is the optimal state process of (2.14). The concavification principle states that problems (2.6) and (2.14) are equivalent (see Reichlin [21, Theorem 5.1]) and the optimal solution for (2.14) is the same as that for (2.6). The HJB equation for problem (2.14) is given by

$$\frac{\partial u_\lambda^c}{\partial t} + \sup_\pi \left((rx + x\pi(\hat{\mu}-r))\frac{\partial u_\lambda^c}{\partial x} + \frac{1}{2}x^2\pi^2\sigma^2\frac{\partial^2 u_\lambda^c}{\partial x^2} + \frac{1}{2}\psi^2\frac{\partial^2 u_\lambda^c}{\partial \hat{\mu}^2} + x\pi\sigma\psi\frac{\partial^2 u_\lambda^c}{\partial\hat{\mu}\partial x}\right) = 0$$

with the terminal condition $u_\lambda^c(T,x,\hat{\mu}) = U_\lambda^c(x)$. This is a nonlinear PDE with two state variables and is in general difficult to solve. We may use the dual method to solve it. The dual state process Y is strictly positive and satisfies the following SDE:

$$dY(t) = -Y(t)rdt - Y(t)\sigma^{-1}(\hat{\mu}(t)-r)d\hat{W}(t), \quad Y(0)=y_0. \tag{2.17}$$

Since there is no control in SDE (2.17), the dual problem is reduced to a simple evaluation of expectation of dual function at $Y(T)$, that is,

$$\begin{cases} \mathbb{E}[V_\lambda^c(Y(T))], \\ \text{s.t. } Y(t) \text{ satisfies (2.17)}, \\ \hat{\mu}(t) \text{ satisfies (2.4)}. \end{cases} \tag{2.18}$$

The dual value function is defined by

$$v_\lambda^c(t,y,\hat{\mu}) := \mathbb{E}[V_\lambda^c(Y(T))|Y(t)=y,\hat{\mu}(t)=\hat{\mu}] \tag{2.19}$$

and the dual constraint function by

$$g_\lambda(t,y,\hat{\mu}) := \mathbb{E}\left[\mathbb{1}_{x^{*,\lambda}(Y(T))\geq L}\middle| Y(t)=y, \hat{\mu}(t)=\hat{\mu}\right], \tag{2.20}$$

where $\mathbb{1}_S$ is an indicator that equals 1 if S happens and 0 otherwise. By Feynman-Kac formula, we have

$$\frac{\partial v^c_\lambda}{\partial t} - ry\frac{\partial v^c_\lambda}{\partial y} + \frac{1}{2}y^2\sigma^{-2}(\hat{\mu}-r)^2\frac{\partial^2 v^c_\lambda}{\partial y^2} + \frac{1}{2}\psi^2\frac{\partial^2 v^c_\lambda}{\partial \hat{\mu}^2} - y\sigma^{-1}(\hat{\mu}-r)\psi\frac{\partial^2 v^c_\lambda}{\partial \mu \partial y} = 0 \tag{2.21}$$

with the terminal condition $v^c_\lambda(T,y,\hat{\mu}) = V^c_\lambda(y)$. The optimal terminal wealth $X^{*,\lambda}(T)$ for primal problem (2.6) is given by Proposition (2.2), that is, $X^{*,\lambda}(T) = x^{*,\lambda}(Y(T))$ and y_0 is determined by the binding budget constraint $\mathbb{E}[X^{*,\lambda}(T)Y(T)] = x_0y_0$. The optimal wealth process $X^{*,\lambda}$ is determined by $X^{*,\lambda}(t) = \mathbb{E}[X^{*,\lambda}(T)Y(T)/Y(t)|\mathcal{F}^S_t]$ for $t\in[0,T]$, the optimal control process $\pi^{*,\lambda}$ by the martingale representation theorem, and the optimal value function u^c_λ by $u^c_\lambda(t,x,\hat{\mu}) = \inf_{y>0}(v^c_\lambda(t,y,\hat{\mu}) + xy)$.

3 Main Results

To express $\mathbb{E}[V^c_\lambda(Y(T))]$ explicitly in terms of integral representation, we need to know the joint distribution of Y and $\hat{\mu}$, which is unknown. To address this, we employ a measure change technique (Xing et al. [26]). We introduce a new process W^Q by

$$dW^Q(t) := \frac{\hat{\mu}(t)-\mu^l}{\sigma}dt + d\hat{W}(t),$$

and a new probability measure by

$$\left.\frac{dQ}{dP}\right|_{\mathcal{F}^S_T} := \exp\left(-\frac{1}{2}\int_0^T\left(\frac{\hat{\mu}(u)-\mu^l}{\sigma}\right)^2 du - \int_0^T \frac{\hat{\mu}(u)-\mu^l}{\sigma}d\hat{W}(u)\right).$$

By Girsanov's theorem, W^Q is a standard Brownian motion under the new probability measure Q. Under this new measure Q, the wealth process X satisfies

$$dX(t) = X(t)(r+\theta_l\sigma\pi(t))dt + X(t)\sigma\pi(t)dW^Q(t),$$

where $\theta_l = \frac{\mu^l - r}{\sigma}$, and the corresponding dual process

$$d\mathcal{Y}(t) = \mathcal{Y}(t)\left(-rdt - \theta_l dW^Q(t)\right).$$

Let

$$\Phi(t) := \frac{\hat{\mu}(t)-\mu^l}{\mu^h - \hat{\mu}(t)}, \quad t\in[0,T].$$

Note that $\Phi(t)$ is well defined as $\hat{\mu}(t)\in(\mu^l,\mu^h)$ a.s. for all $t\in[0,T]$ due to $\hat{\mu}(0)\in(\mu^l,\mu^h)$ (see Décamps et al. [9, Lemma 3.1]). Then $\Phi(t)$ satisfies the SDE:

$$d\Phi(t) = \Theta\Phi(t)dW^Q(t)$$

with $\Phi(0) = \phi := \frac{\hat{\mu}(0)-\mu^l}{\mu^h-\hat{\mu}(0)}$, where $\Theta = \frac{\mu^h-\mu^l}{\sigma}$. In addition, let

$$F(t) := \frac{1+\Phi(t)}{1+\phi}, \tag{3.1}$$

then $F(t)$ satisfies the SDE:

$$dF(t) = \sigma^{-1}(\hat{\mu}(t) - \mu^l)F(t)dW^Q(t)$$

with $F(0) = 1$, and

$$F(t) = \left.\frac{dP}{dQ}\right|_{\mathcal{F}_t^S}.$$

Applying Ito's lemma, we deduce that

$$d(F(t)Y(t)) = -rF(t)Y(t)dt - F(t)Y(t)\theta_l dW^Q(t)$$

with $F(0)Y(0) = y_0$, which yields that $\mathcal{Y}(t) = F(t)Y(t)$ under measure Q. We also get that

$$\Phi(t) = \phi \exp\left\{\Theta W^Q(t) - \frac{1}{2}\Theta^2 t\right\},$$

and

$$\mathcal{Y}(t) = y_0 \exp\left\{-\theta_l W^Q(t) - \left(r + \frac{1}{2}\theta_l^2\right) t\right\}.$$

Using the above observations, the dual value function in (2.18) becomes

$$\mathbb{E}[V_\lambda^c(Y(T))] = \mathbb{E}^Q[F(T)V_\lambda^c(\mathcal{Y}(T)/F(T))] = \mathbb{E}^Q\left[\frac{1+\Phi(T)}{1+\phi} V_\lambda^c\left(\frac{1+\phi}{1+\Phi(T)}\mathcal{Y}(T)\right)\right].$$

Since Φ and $\mathcal{Y}$ can be expressed in terms of Q-Brownian motion W^Q, we can write out the integral representation of the dual value function (2.19) and the constraint function (2.20) explicitly, that is,

$$v_\lambda^c(t, y, \hat{\mu}) = \int_{\mathbb{R}} \Psi(t, x, \hat{\mu})\, V_\lambda^c\left(\frac{y \exp\left\{-\theta_l x - \left(r + \frac{1}{2}\theta_l^2\right)(T-t)\right\}}{\Psi(t, x, \hat{\mu})}\right) p(t, x)dx, \tag{3.2}$$

$$g_\lambda(t, y, \hat{\mu}) = \int_{\mathbb{R}} \Psi(t, x, \hat{\mu})\, \mathbb{1}_{x^{*,\lambda}\left(\frac{y \exp\left\{-\theta_l x - \left(r + \frac{1}{2}\theta_l^2\right)(T-t)\right\}}{\Psi(t, x, \hat{\mu})}\right) \geq L} p(t, x)dx, \tag{3.3}$$

where

$$\Psi(t, x, \hat{\mu}) := \frac{1 + \phi \exp\left\{\Theta x - \frac{1}{2}\Theta^2(T-t)\right\}}{1+\phi},$$

$$p(t, x) := \frac{1}{\sqrt{2\pi(T-t)}} \exp\left\{-\frac{x^2}{2(T-t)}\right\}.$$

Condition (2.7) can be written as

$$\begin{cases} \mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T) \geq L\}}] \geq 1 - \varepsilon, \\ \lambda(\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T) \geq L\}}] - 1 + \varepsilon) = 0. \end{cases} \tag{3.4}$$

For a given $\lambda \geq 0$, the optimal terminal wealth is given by

$$X^{\pi^*,\lambda}(T) = x^{*,\lambda}(Y(T)), \tag{3.5}$$

where $x^{*,\lambda}$ is given by (2.11) or (2.12), depending on the value λ,

$$Y(T) = \frac{\mathcal{Y}(T)}{F(T)} = y_0(1+\phi)H(T),$$

and

$$H(T) := \frac{\exp\{-\theta_l W^Q(T) - (r + \frac{1}{2}\theta_l^2)T\}}{1 + \phi \exp\left\{\Theta W^Q(T) - \frac{1}{2}\Theta^2 T\right\}}. \tag{3.6}$$

Additionally, y_0 is determined by the binding budget constraint $\mathbb{E}[X^{*,\lambda}(T)Y(T)] = x_0 y_0$, that is,

$$\mathbb{E}^Q[F(T)X^{\pi^*,\lambda}(T)(1+\phi)H(T)] = x_0. \tag{3.7}$$

Combining (3.4) and (3.7), we can derive solutions (y_0, λ^*). We next state the main theorem on the existence and uniqueness of the Lagrange multiplier and the feasibility condition.

Theorem 3.1. *Let $x_0 > 0$ and H_ε^* be the unique solution of the equation*

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le H_\varepsilon^*\}}] = 1 - \varepsilon, \tag{3.8}$$

where $F(T)$ and $H(T)$ are given in (3.1) and (3.6). Denote by

$$\hat{x}_\varepsilon := \mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}(1+\phi)H(T)]. \tag{3.9}$$

Then the following results hold.

1. *If $x_0 > \hat{x}_\varepsilon$, then there exists a unique $\lambda^* \ge 0$ such that $X^{\pi^*,\lambda^*}(T)$ in (3.5) is the optimal solution to problem (2.5).*

2. *If $x_0 = \hat{x}_\varepsilon$, then there is only one solution $X^{\pi^*,\lambda^*}(T) = L\mathbb{I}_{H(T)<H_\varepsilon^*}$ a.s..*

3. *If $x_0 < \hat{x}_\varepsilon$, then problem (2.5) is infeasible, that is, condition (2.7) is not satisfied.*

Proof. We first show that there exists a unique solution H_ε^* to equation (3.8). Define $f(x) := \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le x\}}]$ for $x \ge 0$. Then f is continuous and strictly increasing, $f(0) = 0 < 1-\varepsilon$ and $f(\infty) = \mathbb{E}^Q[F(T)] = \mathbb{E}^Q[1+\Phi(T)]/(1+\phi) = (1+\Phi(0))/(1+\phi) = 1 > 1-\varepsilon$ as Φ is a Q-martingale. Therefore, there exsits a unique $H_\varepsilon^* > 0$ such that $f(H_\varepsilon^*) = 1-\varepsilon$.

We now discuss the case $x_0 > \hat{x}_\varepsilon$. For a fixed $\lambda \ge 0$, the optimal terminal wealth $X^{\pi^*,\lambda}(T)$ is given by (3.5). If we can find a solution (y_0, λ^*) to equations (3.4) and (3.7), then $X^{\pi^*,\lambda^*}(T)$ is the optimal solution to problem (2.5).

Case I: $H_\varepsilon^* \le \frac{c_z}{y_0(1+\phi)}$. If we choose $\lambda = 0$, then $X^{\pi^*,\lambda}(T)$ is given by

$$X^{\pi^*,\lambda}(T) = (\theta + I_1(y_0(1+\phi)H(T)))\mathbb{I}_{\{H(T)\le\frac{c_z}{y_0(1+\phi)}\}}. \tag{3.10}$$

We have

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T)\ge L\}}] \;=\; \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le\frac{c_z}{y_0(1+\phi)}\}}] \ge \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le H_\varepsilon^*\}}] = 1-\varepsilon.$$

The quantile constraint (3.4) is satisfied with $\lambda^* = 0$. We next show that there is a unique solution y_0 to equation (3.7). Denote by $f(y_0) := \mathbb{E}^Q[F(T)X^{\pi^*,\lambda}(T)(1+\phi)H(T)]$. We can check that f is continuous, strictly decreasing, and $\lim_{y_0\to 0^+} f(y_0) = \infty, \lim_{y_0\to\infty} f(y_0) = 0 < x_0$, then there exists a unique y_0 satisfying (3.7).

Case II: $\frac{c_z}{y_0(1+\phi)} < H_\varepsilon^* \le \frac{c_{\tilde z}}{y_0(1+\phi)}$. In this case, if $\lambda^* = 0$, then

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T)\ge L\}}] = \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le \frac{c_z}{y_0(1+\phi)}\}}] < \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\le H_\varepsilon^*\}}] = 1-\varepsilon,$$

which implies $\lambda^* = 0$ is impossible as (3.4) is not satisfied. We must have $\lambda^* > 0$ and the quantile constraint is binding, that is,

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T)\ge L\}}] = 1-\varepsilon. \tag{3.11}$$

Since $X^{\pi^*,\lambda}(T) = x^{*,\lambda}(Y(T))$ and $x^{*,\lambda}$ is given by (2.11) or (2.12), we next discuss its form. If $k_\lambda > c_{\tilde z}$, then

$$X^{\pi^*,\lambda}(T) = (\theta + I_1(y_0(1+\phi)H(T)))\mathbb{I}_{\{H(T)<\frac{c_{\tilde z}}{y_0(1+\phi)}\}} + L\mathbb{I}_{\{\frac{c_{\tilde z}}{y_0(1+\phi)}\le H(T)<\frac{k_\lambda}{y_0(1+\phi)}\}}, \tag{3.12}$$

and

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T)\ge L\}}] = \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)<\frac{k_\lambda}{y_0(1+\phi)}\}}] > \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}] = 1-\varepsilon$$

as $k_\lambda > c_{\tilde z} \ge y_0(1+\phi)H_\varepsilon^*$, which is a contradiction to (3.11). We must have $k_\lambda \le c_{\tilde z}$, then

$$X^{\pi^*,\lambda}(T) = (\theta + I_1(y_0(1+\phi)H(T)))\mathbb{I}_{\{H(T)<\frac{c_{\tilde z_0}}{y_0(1+\phi)}\}}. \tag{3.13}$$

To ensure (3.11) holds, we have $c_{\tilde z_0} = y_0(1+\phi)H_\varepsilon^*$ and from (2.13) we define

$$\lambda_1^*(y_0) := \tilde z_0 U_1'(\tilde z_0 - \theta) - U_1(\tilde z_0 - \theta) - U_2(\theta). \tag{3.14}$$

To show $\lambda_1^*(y_0) > 0$, define $g(x) := xU_1'(x-\theta) - U_1(x-\theta) - U_2(\theta)$. Since g is strictly decreasing and $\tilde z_0 < z$, we have $\lambda_1^*(y_0) = g(\tilde z_0) > g(z) = 0$. The proof of existence and uniqueness of y_0 is similar to that in case I.

Case III: $H_\varepsilon^* > \frac{c_{\tilde z}}{y_0(1+\phi)}$. In this case, similar to Case II, we have $\lambda^* > 0$ and $X^{\pi^*,\lambda}(T)$ is given by (3.12) or (3.13). If $X^{\pi^*,\lambda}(T)$ were given by (3.13), we would have

$$\mathbb{E}^Q[F(T)\mathbb{I}_{\{X^{\pi^*,\lambda}(T)\ge L\}}] = \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)<\frac{c_{\tilde z_0}}{y_0(1+\phi)}\}}] < \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}] = 1-\varepsilon$$

as $c_{\tilde z_0} \le c_{\tilde z} < H_\varepsilon^* y_0(1+\phi)$, which is a contradiction to (3.11). We must have $k_\lambda > c_{\tilde z}$ and $X^{\pi^*,\lambda}(T)$ is given by (3.12). To ensure (3.11) holds, we have $k_\lambda = y_0(1+\phi)H_\varepsilon^*$ and by the expression of k_λ we define

$$\lambda_2^*(y_0) = k_\lambda L + U_2(\theta - L) - U_2(\theta). \tag{3.15}$$

Then

$$\lambda_2^*(y_0) > LU_1'(\tilde z - \theta) + U_2(\theta - L) - U_2(\theta) = \tilde z U_1'(\tilde z - \theta) - U_1(\tilde z - \theta) - U_2(\theta) = g(\tilde z) > g(z) = 0.$$

The second equation holds as $U_1(\tilde z - \theta) + U_2(\theta - L) - (\tilde z - L)U_1'(\tilde z - \theta) = 0$.

To show the existance and uniqueness of y_0, define $f(y_0) := \mathbb{E}^Q[F(T)X^{\pi^*,\lambda^*}(T)(1+\phi)H(T)]$. We can check that f is continuous, strictly decreasing, $\lim_{y_0\to 0^+} f(y_0) = \infty$ and $\lim_{y_0\to\infty} f(y_0) = \mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}(1+\phi)H(T)] = \hat x_\varepsilon < x_0$, which shows there exists a unique y_0 to equation (3.7).

We next discuss the case $x_0 = \hat x_\varepsilon$. Combining (3.7), we have

$$\mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}H(T)] = \mathbb{E}^Q[F(T)X^{\pi^*,\lambda^*}(T)H(T)]. \tag{3.16}$$

If $H_\varepsilon^* \le \frac{c_z}{y_0(1+\phi)}$, then $\lambda^* = 0$ and $X^{\pi^*}(T)$ is given by (3.10). We have

$$\mathbb{E}^Q[F(T)X^{\pi^*,\lambda^*}(T)H(T)] \ge \mathbb{E}^Q[F(T)z\mathbb{I}_{H(T)<\frac{c_z}{y_0(1+\phi)}}H(T)] > \mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}H(T)],$$

which is a contradiction to (3.16). Similarly, if $\frac{c_z}{y_0(1+\phi)} < H_\varepsilon^* \le \frac{c_{\tilde{z}}}{y_0(1+\phi)}$, then $X^{\pi^*,\lambda}(T)$ is given by (3.13) with $c_{\tilde{z}_0} = y_0(1+\phi)H_\varepsilon^*$, this gives that

$$\mathbb{E}^Q[F(T)X^{\pi^*,\lambda^*}(T)H(T)] \ge \mathbb{E}^Q[F(T)\tilde{z}_0\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}H(T)] > \mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}H(T)],$$

which is again a contradiction to (3.16). If $H_\varepsilon^* > \frac{c_{\tilde{z}}}{y_0(1+\phi)}$, then $X^{\pi^*,\lambda}(T)$ is given by (3.12) with $k_\lambda = y_0(1+\phi)H_\varepsilon^*$, this gives that

$$\mathbb{E}^Q[F(T)X^{\pi^*,\lambda^*}(T)H(T)] \ge \mathbb{E}^Q[F(T)L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}H(T)].$$

To ensure (3.16) holds, we must have $y_0 = \infty$, then $k_\lambda = \infty$, $\lambda = \infty$, and $X^{\pi^*,\lambda^*}(T) = L\mathbb{I}_{H(T)<H_\varepsilon^*}$.

We finally discuss the case $x_0 < \hat{x}_\varepsilon$. Suppose solutions exist. If $H_\varepsilon^* \le \frac{c_z}{y_0(1+\phi)}$, then $\lambda^* = 0$ and $X^{\pi^*}(T)$ is given by (3.10). We have

$$\mathbb{E}^Q[X^{\pi^*,\lambda}(T)\mathcal{Y}(T)] \ge \mathbb{E}^Q[L\mathbb{I}_{\{H(T)<H_\varepsilon^*\}}y_0(1+\phi)H(T)F(T)] = \hat{x}_\varepsilon y_0 > x_0y_0,$$

which is a contradiction. Similarly, if $\frac{c_z}{y_0(1+\phi)} < H_\varepsilon^* \le \frac{c_{\tilde{z}}}{y_0(1+\phi)}$ or $H_\varepsilon^* > \frac{c_{\tilde{z}}}{y_0(1+\phi)}$, then $X^{\pi^*,\lambda}(T)$ is given by (3.13) with $c_{\tilde{z}_0} = y_0(1+\phi)H_\varepsilon^*$ or (3.12) with $k_\lambda = y_0(1+\phi)H_\varepsilon^*$, respectively, we would again have $\mathbb{E}^Q[X^{\pi^*,\lambda}(T)\mathcal{Y}(T)] > x_0y_0$, which is a contradiction. We conclude that there is no feasible solution if $x_0 < \hat{x}_\varepsilon$. □

Remark 3.2. *If $\varepsilon = 0$, which requires $X^\pi(T) \ge L$ a.s., then $H_\varepsilon^* = +\infty$, $\hat{x}_\varepsilon = (1+\phi)\mathbb{E}^Q[LF(T)H(T)]$ and*

$$X^{\pi^*,\lambda^*}(T) = (\theta + I_1(Y(T)))\mathbb{I}_{\{Y(T)<c_{\tilde{z}}\}} + L\mathbb{I}_{\{Y(T)\ge c_{\tilde{z}}\}}.$$

If $\varepsilon = 1$, which removes the quantile constraint, then $H_\varepsilon^* = 0$, $\hat{x}_\varepsilon = 0$ and

$$X^{\pi^*,\lambda^*}(T) = (\theta + I_1(Y(T)))\mathbb{I}_{\{Y(T)\le c_z\}}.$$

Remark 3.3. *Chen et al. [7] discuss risk management with multiple VaR constraints and solve the problem with the backward approach, that is, first solve the last period problem with only terminal VaR constraint, then solve the second last period problem with intertemporal indirect value function and intertemporal VaR constraint, and continue until cover the first period. Their approach critically depends on the distribution of risky asset price process being known (a geometric Brownian motion), which provides an integral representation of the intertemporal indirect value function. The same idea in theory is applicable to this paper. Due to unobservable drift, the distribution of risky asset price process is unknown after filtering, however, thanks to the change of measure, the dual value function v_λ^c has an integral representation (3.2) and the Lagrange multiplier λ^* can be determined by a variation of Theorem 3.1 for the last period and the intertemporal indirect value function $u_{\lambda^*}^c$ is given by the conjugate function of $v_{\lambda^*}^c$. This approach is nevertheless highly complex. To highlight the main ideas and approaches of the paper, we only consider a terminal VaR constraint. There may be some other more streamlined ways of dealing with intertemporal VaR and other constraints, which is worth further reserch. We thank the referee for this interesting query.*

4 Algorithms

In this section we discuss three algorithms to solve problem (2.6) with conditions (2.7). The first one is the exact algorithm based on Theorem 3.1, the second the dual simulation method based on (2.18), and the third the PINN method to solve the dual HJB equation (2.21).

4.1 Lagrange Algorithm

The proof of Theorem 3.1 is constructive, and from this we propose a numerical algorithm to solve the problem. We note that the values of y_0 are different in different $X^{\pi^*,\lambda^*}(T)$, which are denoted by y_0^1, y_0^2, y_0^3 in (3.10), (3.13), and (3.12) respectively. Now we give an algorithm to derive the optimal terminal wealth $X^{\pi^*,\lambda^*}(T)$ and the corresponding Lagrange multiplier λ^*. The exact algorithm is the following.

Step 0 Input initial wealth x_0, initial drift estimate $\hat{\mu}(0)$, confidence level ε, then compute z (2.9), $\tilde{z}$ (2.10), H_ε^* (3.8), and $\hat{x}_\varepsilon$ (3.9). If $x_0 < \hat{x}_\varepsilon$, then the problem is infeasible, stop. If $x_0 = \hat{x}_\varepsilon$, then set $X^{\pi^*,\lambda^*}(T) = L\mathbb{I}_{H(T)<H_\varepsilon^*}$, stop. Otherwise, go to Step 1.

Step 1 Compute y_0^3 from the budget constraint (3.7) using $X^{\pi^*,\lambda}(T)$ from (3.12). If $y_0^3 > \frac{c_{\tilde{z}}}{H_\varepsilon^*(1+\phi)}$, then set $X^{\pi^*,\lambda^*}(T) = X^{\pi^*,\lambda}(T)$ and $\lambda^* = \lambda_2(y_0^3)$ from (3.15), stop. Otherwise, go to Step 2.

Step 2 Compute y_0^2 from (3.7) using $X^{\pi^*,\lambda}(T)$ from (3.13). If $y_0^2 > \frac{c_z}{H_\varepsilon^*(1+\phi)}$, then set $X^{\pi^*,\lambda^*}(T) = X^{\pi^*,\lambda}(T)$ and $\lambda^* = \lambda_1(y_0^2)$ from (3.14), stop. Otherwise, go to Step 3.

Step 3 Set $X^{\pi^*,\lambda^*}(T)$ from (3.10) and $\lambda^* = 0$, stop.

Since λ^* can be expressed as a function of y_0 that is related to ε, we write the optimal terminal wealth as $X^{\pi^*,\varepsilon}(T)$ instead of $X^{\pi^*,\lambda^*}(T)$ and characterize its form in terms of ε. Specifically, denote by $\varepsilon^* = 1 - \mathbb{P}(H(T) \leq \frac{c_z}{y_0(1+\phi)})$ and $\varepsilon_* = 1 - \mathbb{P}(H(T) \leq \frac{c_{\tilde{z}}}{y_0(1+\phi)})$. Since $\tilde{z} < z$, we have $c_{\tilde{z}} > c_z$ and $\varepsilon_* \leq \varepsilon^*$. If $\varepsilon \geq \varepsilon^*$, then $\mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq\frac{c_z}{y_0(1+\phi)}\}}] \geq \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq H_\varepsilon^*\}}] = 1-\varepsilon$, we have $H_\varepsilon^* \leq \frac{c_z}{y_0(1+\phi)}$ and $X^{\pi^*,\varepsilon}(T)$ is given by (3.10). If $\varepsilon_* \leq \varepsilon < \varepsilon^*$, then $\mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq\frac{c_{\tilde{z}}}{y_0(1+\phi)}\}}] \geq \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq H_\varepsilon^*\}}] > \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq\frac{c_z}{y_0(1+\phi)}\}}]$, we have $H_\varepsilon^* \in (\frac{c_z}{y_0(1+\phi)}, \frac{c_{\tilde{z}}}{y_0(1+\phi)}]$ and $X^{\pi^*,\varepsilon}(T)$ is given by (3.13). If $\varepsilon < \varepsilon_*$, then $\mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq H_\varepsilon^*\}}] > \mathbb{E}^Q[F(T)\mathbb{I}_{\{H(T)\leq\frac{c_{\tilde{z}}}{y_0(1+\phi)}\}}]$, we have $H_\varepsilon^* > \frac{c_{\tilde{z}}}{y_0(1+\phi)}$ and $X^{\pi^*,\varepsilon}(T)$ is given by (3.12).

4.2 Dual Simulation Algorithm

First we consider running Monte Carlo simulations for the dual problem (2.18). Due to the unconstrained, complete market nature of the problem, the dual problem is reduced to evaluation of an expectation, so Monte Carlo methods are well suited. The only difficulty is in converting to the primal problem, we need to find the dual start parameter y^* associated to x_0, minimizing $v_\lambda^c(0, y, \hat{\mu}(0)) + x_0 y$, or equivalently solving $\partial_y v_\lambda^c(0, y, \hat{\mu}(0)) + x_0 = 0$. Combined with the need to find the Lagrange multiplier λ^* leads to a coupled optimization problem to solve. Optimization over y is very easy, as we can evaluate $\partial_y v_\lambda^c(0, y, \hat{\mu}(0))$ again using simulation. We therefore find y^* for a range of λ values, then can find the right λ to solve the constrained problem.

Fix a sample size $M \in \mathbb{N}$ and discretization size $N \in \mathbb{N}$. Let $h = \frac{T}{N}$ be the step size of the corresponding discretization on $[0, T]$. Define the following Monte Carlo function

$$\mathrm{MC}^d(y, \lambda) = \frac{1}{M} \sum_{i=1}^{M} V_\lambda^c \left(y \zeta_N^i \right), \tag{4.1}$$

where

$$\begin{aligned} \zeta_{n+1}^i &= \zeta_n^i \left(1 - hr - \sqrt{h} \sigma^{-1} (\hat{\mu}_n^i - r) Z_n^i \right), \\ \hat{\mu}_{n+1}^i &= \hat{\mu}_n^i + \sqrt{h} \psi \left(nh, \hat{\mu}_n^i \right) Z_n^i, \end{aligned} \tag{4.2}$$

for $i = 1, \dots, M$ and $n = 0, \dots, N-1$, $\zeta_0^i = 1$, $\hat{\mu}_0^i = \hat{\mu}(0)$, and Z_n^i are independent standard normal random variables. Given fixed λ, the optimal dual start $\hat{y}$ should (approximately, given sufficiently large M) minimize the function $y \mapsto \mathrm{MC}^d(y, \lambda) + x_0 y$ for $y > 0$. We therefore define the gradient descent update, for $k \in \mathbb{N}$

$$\begin{aligned} y_{k+1} &= y_k - \delta (\partial_y \mathrm{MC}^d(y_k, \lambda) + x_0) \\ &= y_k - \delta \left(\frac{1}{M} \sum_{i=1}^{M} \zeta_N^i (V_\lambda^c)' \left(y_k \zeta_N^i \right) + x_0 \right) \end{aligned}$$

for some learning rate $\delta > 0$ and initial $y_0 > 0$. Assuming convergence of the algorithm to some $y^* = \lim_{k \to \infty} y_k$, we output the primal Monte Carlo simulation

$$\begin{aligned} u_\lambda^c(t, x_0, \hat{\mu}) &= \frac{1}{M} \sum_{i=1}^{M} U_\lambda^c \left(x^{*,\lambda} \left(y^* \zeta_N^i \right) \right), \\ h_\lambda(t, x_0, \hat{\mu}) &= \frac{1}{M} \sum_{i=1}^{M} \mathbb{1}_{x^{*,\lambda} \left(y^* \zeta_N^i \right) \geq L}. \end{aligned}$$

This method applies for only one value of λ. However, the algorithm can be run in parallel for multiple values $(\lambda_j)_{j=1}^J \subset [0, \infty)$. We can then interpolate between these points to produce u and h as functions of λ.

Remark 4.1. *The training range may intersect with the infeasible region, points $x, \hat{\mu}$ at which the constrained problem (2.5) starting at $x, \hat{\mu}$ has no solution for some $\varepsilon \in [0, 1]$. However, the unconstrained problem (2.14) is well defined and has a solution for any start points $x \geq 0, \hat{\mu} \in [\mu^l, \mu^h]$ and multiplier $\lambda \geq 0$. This is also the case in the next algorithm.*

4.3 Dual PINN Algorithm

By Feynmann Kac, the dual value function (2.19) satisfies a PDE (2.21) which can be solved numerically. In particular, due to the simple nature of the dual problem, the dual PDE is linear, suggesting high effectiveness of numerical methods. The PINN method assumes a bounded state space, and samples points in this space. The dual value functions is defined to be a neural network, to which the PDE operator can be applied. This term is combined with the error at terminal time, to produce a loss function that can be minimized over the parameters of the neural network.

Typical convergence analysis for the PINN method requires PDEs with uniform Lipschitz continuity of the terminal function Shin et al. [24]. For the HJB equation satisfied by the dual value function, the terminal function is only a Lipschitz continuous function of y away from $y = 0$. Therefore, even in the simple case of the dual problem, our PDEs fall out of the provably convergent class of PDEs. However, we can still attempt this method and compare to the simulation approach. Just like the dual simulation method, we need to find the optimal dual start after finding the dual value function, in conjunction with the Lagrange multiplier.

First we define a neural network. Let $\mathcal{N} \in \mathbb{N}$ be fixed, this is the number of "nodes" in our network. Let $f\colon \mathbb{R}^d \times \Xi^{d,\mathcal{N}} \to \mathbb{R}$ for some input dimension $d \in \mathbb{N}$, where $\Xi^{d,\mathcal{N}} = \mathbb{R}^{d\times\mathcal{N}} \times \mathbb{R}^{\mathcal{N}\times\mathcal{N}} \times \mathbb{R}^{\mathcal{N}\times 1} \times \mathbb{R}^{\mathcal{N}} \times \mathbb{R}^{\mathcal{N}} \times \mathbb{R}$ is the parameter space. For $X \in \mathbb{R}^d$ and $\Theta = (A_1, A_2, A_3, b_1, b_2, b_3) \in \Xi^{d,\mathcal{N}}$ we define the neural network f as

$$f(X;\Theta) = L_3 \circ \eta \circ L_2 \circ \eta \circ L_1(X), \qquad X \in \mathbb{R}^d$$

where $L_i(X) = A_i X + b_i$ are linear maps (A is the "weight", b is the "bias") for $i = 1, 2, 3$ and η is a non-linear "activation" function, applying the function $x \mapsto \tanh(x)$ element-wise. Each output of $\eta \circ L_i$, and the input X are referred to as a layer. We have the input layer, the output layer $f(X;\Theta)$, and the two so-called "hidden layers" in between. In the sequel we will take $d = 4$.

The dual PINN method aims to find the value of $v^c_\lambda(t, y, \hat{\mu})$ and for $t \in [0, T]$, $y \in \mathcal{Y} \subset [0, \infty)$, $\hat{\mu} \in \mathcal{M} \subset [\mu^l, \mu^h]$, and $\lambda \in \Lambda \subset [0, \infty)$. The function $v(t, y, \hat{\mu}, \lambda) := v^c_\lambda(t, y, \hat{\mu})$ solves the PDE (2.21). We can solve the HJB equation using the PINN method to find the function v^c. We can then solve the optimality condition for $y^* := y^*(\lambda)$ at time 0 for a range of λ. We can then evaluate

$$g_\lambda(0, y^*(\lambda), \hat{\mu}(0)) = \frac{1}{M}\sum_{i=1}^{M} \mathbb{1}_{x^{*,\lambda}(y^*(\lambda)\zeta^i_N) \geq L}$$

using Monte Carlo simulation (4.2).

We initialize a neural network function $v(t, y, \hat{\mu}, \lambda; \Theta^v)$ depending on some parameters Θ^v. We define the following loss functionals for any twice differentiable $v\colon [0, T] \times \mathcal{Y} \times \mathcal{M} \times \Lambda \to \mathbb{R}$.

$$\begin{aligned}
\mathbb{L}(v) = \sum_{i=1}^{K_c} \Big| &\partial_t v(t^c_i, y^c_i, \mu^c_i, \lambda^c_i) - r y^c_i \partial_y v(t^c_i, y^c_i, \mu^c_i, \lambda^c_i) + \frac{1}{2}(y^c_i)^2 \left|\sigma^{-1}(\mu^c_i - r)\right|^2 \partial_{yy} v(t^c_i, y^c_i, \mu^c_i, \lambda^c_i) \\
&+ \frac{1}{2}\psi(t^c_i, \mu^c_i)^2 \partial_{\mu\mu} v(t^c_i, y^c_i, \mu^c_i, \lambda^c_i) - \sigma^{-1} y^c_i(\mu^c_i - r\mathbb{1})\psi(t^c_i, \mu^c_i)\partial_{y\mu} v(t^c_i, y^c_i, \mu^c_i, \lambda^c_i)\Big|^2, \\
\mathbb{L}_V(v) = \sum_{i=1}^{K_b} &\left|v(T, y^b_i, \mu^b_i, \lambda^b_i) - V^c_{\lambda^b_i}(y^b_i)\right|^2,
\end{aligned}$$

where $t^c_i, y^c_i, \mu^c_i, \lambda^c_i$ for $i = 1, \ldots, K_c$ are generated uniformly from $[0, T] \times \mathcal{Y} \times \mathcal{M} \times \Lambda$ ("collocation points") and $y^b_i, \mu^b_i, \lambda^b_i$ for $i = 1, \ldots, K_b$ are generated uniformly from $\mathcal{Y} \times \mathcal{M} \times \Lambda$ ("boundary points"). We minimize the sum of these functions using gradient descent, starting at some arbitrary Θ^v_0, evaluated element wise along Θ

$$\Theta^v_{k+1} = \Theta^v_k - \delta \partial_\Theta \left[\mathbb{L}(v(\cdot;\Theta^v_k)) + \mathbb{L}_V(v(\cdot;\Theta^v_k))\right],$$

for $k \in \mathbb{N}$ and some learning rate $\delta > 0$. Assuming convergence, let Θ^v be the converged parameter set. For the dual problem, we need to find the optimal dual start y^* as well as the optimal Lagrange multiplier λ^* for fixed x_0 and $\hat{\mu}$, which solve the coupled optimization problem

$$\begin{aligned}
y^* &= \arg\min_{y>0} \left\{v^c_{\lambda^*}(0, y, \hat{\mu}) + x_0 y\right\}, \\
\lambda^* &= \arg\min_{\lambda \geq 0} \left\{\lambda \left|g^c_\lambda(0, y^*, \hat{\mu}) - (1-\varepsilon)\right| + \max\left(0, (1-\varepsilon) - g^c_\lambda(0, y^*, \hat{\mu})\right)\right\}.
\end{aligned} \tag{4.3}$$

To solve this problem, we first find y^* solving $\partial_y v(t, y^*, \hat{\mu}, 0; \Theta^v) + x_0 = 0$. We can then test if $g_0(t, y^*, \hat{\mu}) \geq (1-\varepsilon)$, in which case $\lambda^* = 0$ and we are done. Otherwise, we are searching for $y^*, \lambda^* > 0$ such that $\partial_y v(t, y^*, \hat{\mu}, \lambda^*; \Theta^v) + x_0 = 0$ and $g_{\lambda^*}(0, y^*, \hat{\mu}) - (1-\varepsilon) = 0$. We can use existing numerical optimization algorithms that take the functions $v(t, y, \hat{\mu}, \lambda; \Theta^v)$ and $g_\lambda(t, y, \hat{\mu})$ and solve

$$|\partial_y v(0, y^*, \hat{\mu}, \lambda^*; \Theta^v) + x_0|^2 + |g_{\lambda^*}(0, y^*, \hat{\mu}) - (1-\varepsilon)|^2 = 0,$$

over the trained ranged $y, \lambda \in \mathcal{Y} \times \Lambda$. If we have either $y^* \in \partial\mathcal{Y}$ or $\lambda^* \in \partial\Lambda\backslash\{0\}$ on the boundary on the training region, then it is likely we have not found the true optimizers and we should increase the training range and repeat training.

If we are only interested in finding y^* and the corresponding value and constraint for fixed $t, x, \hat{\mu}, \lambda$, we only need to solve $\partial_y v(t, y^*, \hat{\mu}, \lambda; \Theta^v) + x = 0$. This can again be solved numerically, and we then output

$$u^c_\lambda(t, x, \hat{\mu}) = v(t, y^*, \hat{\mu}, \lambda; \Theta^v) + xy,$$
$$h_\lambda(t, x, \hat{\mu}) = g_\lambda(t, y^*, \hat{\mu}).$$

In practise, we find a $y^*(\lambda)$ for each λ, and then find λ^* such that $(\lambda^*, y^*(\lambda^*))$ satisfies (4.3). We do this because we have an explicit representation for the neural network v and it's derivative, so can easily optimize using it, unlike g_λ which is evaluated via simulation. Also, once we have $y^*(\lambda)$, we simulate the optimal terminal wealth via $X_N^{\pi^*,\lambda} = x^{*,\lambda}(y^*(\lambda)\zeta_N)$ using (4.2).

5 Numerical Examples

In this section we solve the constrained problem (2.6) with all methods. For the algorithms that solve the unconstrained problem, given fixed $\varepsilon > 0$, we numerically solve the problem (2.6) in two steps. Firstly, we set $\lambda = 0$ and find the values of $u^c_0(t, x, \hat{\mu})$ and $h_0(t, x, \hat{\mu})$ given by (2.15) and (2.16) respectively. If $h_0(t, x, \hat{\mu}) \geq 1 - \varepsilon$ then we output $u^c_0(t, x, \hat{\mu})$ and $\lambda^* = 0$ and we are done. Otherwise, we find $u^c_\lambda(t, x, \hat{\mu})$ and $h_\lambda(t, x, \hat{\mu})$ for a sufficiently large range of $\lambda > 0$. We then find λ^* such that $h_{\lambda^*}(t, x, \hat{\mu}) = 1 - \varepsilon$.

5.1 Data

Unless otherwise mentioned, we take $T = 1.0$, $d = 1$, $\theta = 1.5$, $L = 0.9$, $x_0 = 1.0$, $r = 0.05$, $\sigma = 0.2$, $\lambda = 0.2$. We use the utility $U_1(x) = \sqrt{x}$, $U_2(x) = x^{0.3}$. With this configuration, problem (2.5) admits a unique solution for any $\varepsilon \in [0, 1]$ by Theorem 3.1. In this setting we have

$$\tilde{z}_0 = \theta + \left(\sqrt{(U_2(\theta) + \lambda)^2 + \theta} - (U_2(\theta) + \lambda)\right)^2,$$
$$\tilde{z} = \theta + \left(\sqrt{U_2(\theta - L)^2 + (\theta - L)} - U_2(\theta - L)\right)^2.$$

In all algorithms, we use grids $\mathcal{X} = [0.2, 2.0]$, $\mathcal{Y} = [0.2, 2.0]$, $\Lambda = [0.0, 2.5]$, and $\mathcal{M} = [0.03, 0.1]$. We use a grid of $(\lambda_j)_{j=1}^J \subset \Lambda$ equally spaced points in Λ with $J = 51$ (we choose 50 subintervals over the region, then J accounts for the midpoints, including both end points). For neural networks, we use a neural network structure with 2 hidden layers, with tanh activation function. For the primal simulation method the network layers have 10 hidden nodes, and for the dual PINN method they have 100 nodes. For dual simulation we take $N = 100$, $M = 100000$, and run the algorithm for 200 steps with $\delta = 0.1$. We run the PINN algorithm generating 2000 and

ε	Method	λ^*	y^*	u	$u^c_{\lambda^*}$	$\mathbb{P}(X^{\pi^*,\lambda^*}(T) = L)$	$\mathbb{P}(X^{\pi^*,\lambda^*}(T) = 0)$
0	Dual sim	1.65	1.883	-0.561	1.086	0.746	0.003
0	Dual pinn	1.7	1.911	-0.601	1.13	0.782	0
0	Lagrange	1.659	1.885	-0.564	1.095	0.751	0
0.1	Dual sim	1.453	1.795	-0.411	0.898	0.5	0.102
0.1	Dual pinn	1.464	1.806	-0.435	0.91	0.52	0.1
0.1	Lagrange	1.452	1.794	-0.411	0.896	0.5	0.1
0.35	Dual sim	0.478	1.214	-0.095	0.216	0	0.35
0.35	Dual pinn	0.614	1.295	-0.114	0.31	0	0.35
0.35	Lagrange	0.483	1.216	-0.095	0.219	0	0.35
1	Dual sim	0	0.946	-0.086	-0.085	0	0.395
1	Dual pinn	0	0.94	-0.046	-0.084	0	0.374
1	Lagrange	0	0.945	-0.085	-0.085	0	0.395

Table 1: Various statistics for different values of ε, applied at $t = 0$, $x_0 = 1.0$ and $\hat{\mu}(0) = 0.07$.

200 collocation and boundary points respectively, and use $\delta = 0.01$. We run until either 100000 iteration steps, or the loss function is below 0.00005. The subsequent constrained optimization is performed using `Scipy`, using the neural network function as input into the equation. We solve the discrete distribution problem, taking $\psi(u) = \sigma^{-1}(u - \mu^l)(\mu^h - u)\mathbb{1}_{\mu^l \leq u \leq \mu^h}$ to facilitate comparison with the exact algorithm. The neural networks are implemented using `Tensorflow` and parameters are optimized using the `ADAM` algorithm. Code to implement the algorithms can be found in `https://github.com/Ashley-Davey/ML-For-Quantile`.

5.2 Numerics

Table 1 displays some statistics for specific values of ε. By inverting the primal or dual constraint functions, we compute the value and other information for a single value of ε. We give the values of λ^* and y^*, along with the concavified problem value $u^c_{\lambda^*}(0, x_0, \hat{\mu}(0))$ (2.15), the "true value" $u(0, x_0, \hat{\mu}(0)) := \mathbb{E}[U(X^{\pi^*,\lambda^*}(T))|X^{\pi^*,\lambda^*}(0) = x, \hat{\mu}(0) = \hat{\mu}]$ as a Monte Carlo evaluation of the utility at the outputted optimal state $X^{\pi^*,\lambda^*}(T)$ of the Lagrange problem with optimal Lagrange multiplier, and the probability that this state matches the lower limit and 0 exactly.

Figure 2 shows the results for the constrained problem for all methods. For each method, we find the value and constraint function for a range of λ, and take the (right) inverse the graph of $\lambda \mapsto h_\lambda(0, x_0, \hat{\mu}(0))$ to get a mapping $\varepsilon \mapsto \lambda^*(\varepsilon)$ for $\varepsilon \leq \varepsilon_0 := 1 - \mathbb{P}(X^{\pi^*,0}(T) \geq L)$. For $\varepsilon > \varepsilon_0$, the solution of the constrained problem is equal to the solution at $\varepsilon = \varepsilon_0$ as the quantile constraint is non-binding. Figure 2 (a) shows the graph of ε against Lagrange multiplier λ^*, (b) the corresponding primal value $\mathbb{E}[U^c_\lambda(X^{\pi^*,\lambda}(T))]$, and (c) the constraint probability $\mathbb{P}(X^{\pi^*,\lambda}(T) \geq L)$. We run the algorithms 10 times and take an average. Figure 2 (b) is made up of two sections, the kink point is the point at which the concavified utility moves from one line segment to two, and the structure changes much faster in terms of ε. The simulation and PINN values appear to agree, with a slight gap between the two problems accounting for numerical error. Figure 2 (c) verifies that in the binding region $\varepsilon \in [0, \varepsilon_0]$ we have $\varepsilon = 1 - \mathbb{P}(X^{\pi^*,\lambda^*}(T) \geq L)$, and after this point we have $\mathbb{P}(X^{\pi^*,\lambda^*}(T) \geq L) = 1 - \varepsilon_0$ and $\lambda^* = 0$. The differences in the graph correspond to the differences in each algorithms' approximation of ε_0, which is where each graph of Figure 2 (a) hits the x-axis, or the value of $\mathbb{P}(X^{\pi^*,\lambda^*}(T) = 0)$ when $\varepsilon = 1$ in Table 1.

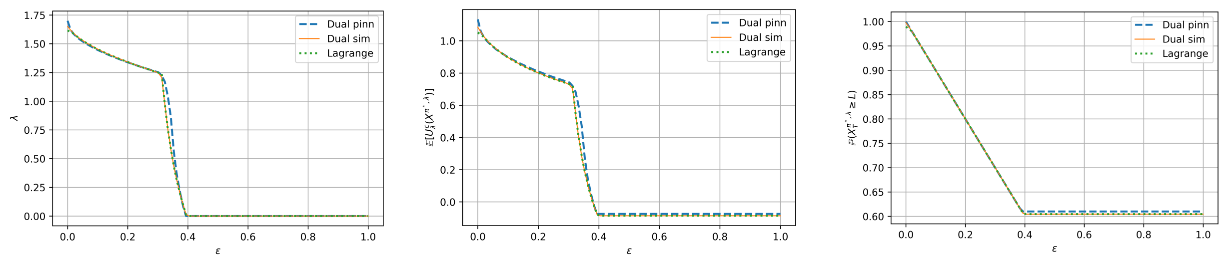

(a) ε against λ^*. (b) ε against $\mathbb{E}[U_\lambda^c(X^{\pi^*,\lambda^*}(T))]$. (c) ε against $\mathbb{P}(X^{\pi^*,\lambda^*}(T) \geq L)$

Figure 2: Numerical results for the concavified problem.

Figure 3 shows the distribution of the optimal terminal wealth $X^{\pi^*,\lambda}(T)$ at terminal time when $\lambda \in \{0, 1.5, 2.5\}$. For the dual PINN and dual simulation methods, we generate this graph by simulating the dual state process and applying the function $x^{*,\lambda}$ given in Proposition 2.2. Where there is an atom at $x = 0$ and $x = L$, we separate the distributions of each algorithm to make them clearer, but they all still refer to the same atom. We see the concavification principle applies, with the terminal state taking values at points x where $U_\lambda(x) = U_\lambda^c(x)$. The continuous section of the distributions have an exponentially decreasing tail in all cases, indicated by the linear segment of the log-scaled graphs.

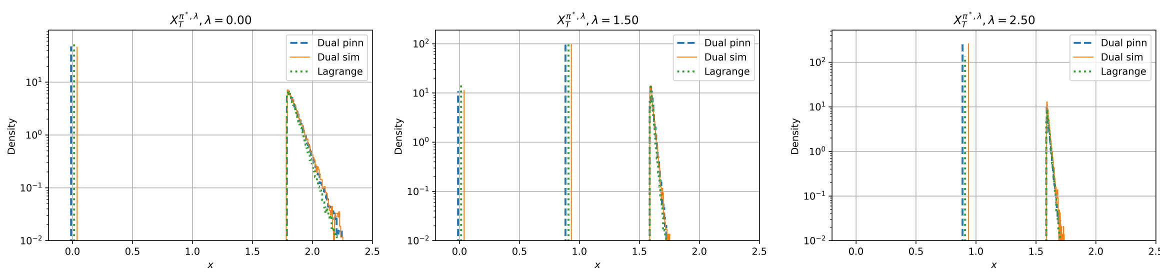

(a) $\lambda = 0$. (b) $\lambda = 1.5$. (c) $\lambda = 2.5$

Figure 3: Distribution of $X^{\pi^*,\lambda}(T)$ for three different values of λ, log scaled.

5.3 Problem Feasibility

For fixed $\varepsilon = 0.2$, we may compute the value of $\hat{x}_\varepsilon$ defined in Theorem 3.1 as $\hat{x}_\varepsilon \approx 0.66$. This is the minimum wealth needed to achieve $\mathbb{P}(X^{\pi^*,\lambda}(T) \geq L) = 0.8$ for some $\lambda \geq 0$. Figure 4 (a) plots the constraint function against λ for feasible $x = 0.73$, infeasible $x = 0.6$ and the threshold $x = \hat{x}_\varepsilon \approx 0.66$. We see that $\hat{x}_\varepsilon$ is exactly the point such that $\lim_{\lambda\to\infty} \mathbb{P}(X^{\pi^*,\lambda}(T) \geq L) = 1 - \varepsilon$, where $X^{\pi^*,\lambda}(T)$ maximizes the value function u_λ^c. For $x < \hat{x}$ the constraint function is lower than $1 - \varepsilon$ for all λ, so there is no solution to the constrained problem. With the Lagrange multiplier method, the objective function of the optimization problem is given by, for fixed $\varepsilon \in [0, 1]$

$$\begin{aligned} J(\pi, \lambda; t, x, \hat{\mu}) &:= \mathbb{E}\left[U(X^\pi(T)) + \lambda\left(\mathbb{1}_{X^\pi(T)\geq L} - (1-\varepsilon)\right)\middle| X^\pi(t) = x, \hat{\mu}(t) = \hat{\mu}\right] \\ &= \mathbb{E}\left[U_\lambda^c(X^\pi(T)) | X^\pi(t) = x, \hat{\mu}(t) = \hat{\mu}\right] - \lambda(1-\varepsilon). \end{aligned}$$

Prior to this, the final term $-\lambda(1-\varepsilon)$ has been omitted as it does not play a role in the optimization of J over π. Define the corresponding (unconstrained) full value function

$$\bar{u}_\lambda^c(t, x, \hat{\mu}) = u_\lambda^c(t, x, \hat{\mu}) - \lambda(1-\varepsilon),$$

where the optimal π^* depending on λ has been found. We use the dual simulation algorithm to plot this function at an infeasible $x = 0.6$, feasible $x = 0.8$, and the transition point $x = \hat{x} \approx 0.66$ at $t = 0$. This plot is given in Figure 4 (b). If a solution exists, then it is a saddle point of the function $\lambda \mapsto \bar{u}^c_\lambda(0, x, \hat{\mu})$. In the feasible region there is a unique saddle point at the solution, in the infeasible region there is no solution, and in the midpoint there is a limiting saddle point at $\lambda^* = \infty$ which corresponds to the only feasible wealth $X^{\pi^*,\infty}(T) := L\mathbb{1}_{H_T < H^*_\varepsilon}$.

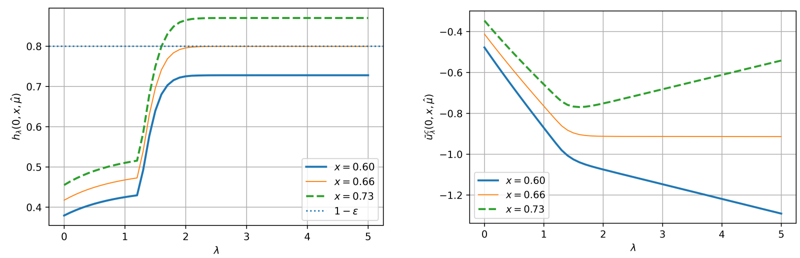

(a) λ against constraint probability for 3 values of x.

(b) λ against full value function for 3 values of x.

Figure 4: Results for dual simulation and discrete algorithm in both feasible and infeasible regions.

6 Conclusions

In this paper we solve S-shaped utility maximization incorporating both partial information and VaR constraint. We convert the original unobservable model into an equivalent fully observable one with an additional filtered state process. We then solve the problem in two steps, first, solve an unconstrained problem with the concavification principle and the dual method, and second, find the Lagrange multiplier and the initial dual state value for the constrained problem. We use a change of measure approach to overcoming the difficulty of the unknown joint distribution of the dual and filtered state processes and characterize the dual value function in a semi-closed integral form. We identify a critical wealth level that makes the constrained problem admits a unique optimal solution or is infeasible. We also propose three algorithms (Lagrange, simulation, deep neural network) to numerically solve the problem and compare their performances with numerical examples. There remain many open problems. For example, if unobservable drift follows a general prior distribution, not necessarily a Bernoulli distribution, the current change of measure approach no longer works, how can we solve such a problem? We leave this and other open problems for future research.

Acknowledgement. The authors are grateful to anonymous reviewers whose comments and suggestions have helped to improve the paper of the previous version.

References

[1] Bain A, Crisan D (2009) Fundamentals of stochastic filtering. Springer, New York

[2] Basak S, Shapiro A (2001) Value-at-risk-based risk management: optimal policies and asset prices. The Review of Financial Studies 14:371-405. https://doi.org/10.1093/rfs/14.2.371.

[3] Bensoussan A, Hoe S, Kim J, Yan Z (2022) A risk extended version of merton's optimal consumption and portfolio selection. Operations Research 70:815-829. https://doi.org/10.1287/opre.2021.2197.

[4] Berkelaar AB, Kouwenberg R, Post T (2004) Optimal portfolio choice under loss aversion. The Review of Economics and Statistics 86:973-987. https://doi.org/10.1162/0034653043125167.

[5] Brendle S (2006) Portfolio selection under incomplete information. Stochastic Processes and their Applications 116:701-723. https://doi.org/10.1016/j.spa.2005.11.010.

[6] Carpenter J N (2000) Does option compensation increase managerial risk appetite?. The Journal of Finance 55:2311-2331. https://doi.org/10.1111/0022-1082.00288.

[7] Chen A, Nguyen T, Stadje M (2018) Risk management with multiple VaR constraints. Mathematical Methods of Operations Research 88:297-337. https://doi.org/10.1007/s00186-018-0637-1.

[8] Davey A, Zheng H (2022) Deep learning for constrained utility maximisation. Methodology and Computing in Applied Probability 24:661-692. https://doi.org/10.1007/s11009-021-09912-3.

[9] Décamps JP, Mariotti T, Villeneuve S (2005) Investment timing under incomplete information. Mathematics of Operations Research 30:472–500. https://doi.org/10.1287/moor.1040.0132.

[10] De Franco C, Nicolle J, Pham H (2019) Bayesian learning for the Markowitz portfolio selection problem. International Journal of Theoretical and Applied Finance 22:1950037. https://doi.org/10.1142/S0219024919500377.

[11] Detemple JB (1986) Asset pricing in a production economy with incomplete information. The Journal of Finance 41:383-391. https://doi.org/10.1111/j.1540-6261.1986.tb05043.x.

[12] Dong Y, Zheng H (2020) Optimal investment with S-shaped utility and trading and value at risk constraints: an application to defined contribution pension plan. European Journal of Operational Research 281:341-356. https://doi.org/10.1016/j.ejor.2019.08.034.

[13] Ekström E, Vaicenavicius J (2016) Optimal liquidation of an asset under drift uncertainty. SIAM Journal on Financial Mathematics 7:357-381. https://doi.org/10.1137/15M1033265.

[14] Han J, Jentzen A, E W (2018) Solving high-dimensional partial differential equations using deep learning. Proceedings of the National Academy of Sciences 115:8505-8510. https://doi.org/10.1073/pnas.1718942115.

[15] Ingersoll JE, Jin LJ (2013) Realization utility with reference-dependent preferences. The Review of Financial Studies 26:723-767. https://doi.org/10.1093/rfs/hhs116.

[16] Jin H, Zhou XY (2008) Behavioral portfolio selection in continuous time. Mathematical Finance 18:385-426. https://doi.org/10.1111/j.1467-9965.2008.00339.x.

[17] Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47:363-391. http://dx.doi.org/10.2307/1914185.

[18] Karatzas I, Xue XX (1991) A note on utility maximization under partial observations 1. Mathematical Finance 1:57-70. https://doi.org/10.1111/j.1467-9965.1991.tb00009.x.

[19] Pham H (2009) Continuous-time stochastic control and optimization with financial applications. Springer, New York.

[20] Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378:686-707. https://doi.org/10.1016/j.jcp.2018.10.045.

[21] Reichlin C (2013) Utility maximization with a given pricing measure when the utility is not necessarily concave. Mathematics and Financial Economics 7:531-556. https://doi.org/10.1007/s11579-013-0093-x.

[22] Rieder U, Bäuerle N (2005) Portfolio optimization with unobservable Markov-modulated drift process. Journal of Applied Probability 42:362-378. https://doi.org/10.1239/jap/1118777176.

[23] Sass J (2007) Utility maximization with convex constraints and partial information. Acta Applicandae Mathematicae 97:221-238. https://doi.org/10.1007/s10440-007-9124-z.

[24] Shin Y, Darbon J, Karniadakis GE (2020) On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type pdes. arXiv:2004.01806.

[25] Wang C, Li S, He D, Wang L (2022) Is L^2 physics informed loss always suitable for training physics informed neural network?. Advances in Neural Information Processing Systems 35:8278-8290.

[26] Xing J, Ma J, Zheng H (2025) A simple integral equation approach for optimal investment stopping problems with partial information. Mathematics of Operations Research, published online. https://doi.org/10.1287/moor.2023.0268

[27] Yiu KFC (2004) Optimal portfolios under a value-at-risk constraint. Journal of Economic Dynamics and Control 28:1317-1334. https://doi.org/10.1016/S0165-1889(03)00116-7.

Data Availability Statement. Data sharing not applicable - no new data generated, as the article describes entirely theoretical research.