Risk-Sensitive Optimal Execution via a Conditional Value-at-Risk Objective^{*}

Seungki Min KAIST skmin@kaist.ac.kr Ciamac C. Moallemi Columbia University ciamac@gsb.columbia.edu

Costis Maglaras Columbia University c.maglaras@gsb.columbia.edu Current Revision: January 2022

PRELIMINARY VERSION — DO NOT DISTRIBUTE

Abstract

We consider a liquidation problem in which a risk-averse trader tries to liquidate a fixed quantity of an asset in the presence of market impact and random price fluctuations. When deciding the liquidation strategy, the trader encounters a trade-off between the transaction costs incurred due to market impact and the volatility risk of holding the position. Our formulation begins with a continuous-time and infinite horizon variation of the seminal model of Almgren and Chriss (2000), but we define as the objective the conditional value-at-risk (CVaR) of the implementation shortfall, and allow for dynamic (adaptive) trading strategies. In this setting, remarkably, we are able to derive closed-form expressions for the optimal liquidation strategy and its value function.

Our results yield a number of important practical insights. We are able to quantify the benefit of adaptive policies over optimized static (pre-committed) policies. The relevant improvement depends only on the level of risk aversion, and grows without bound as the trader becomes more risk neutral. For moderate levels of risk aversion, the optimal dynamic policy outperforms the optimal static policy by 5–15%, and outperforms the optimal volume weighted average price (VWAP) policy by 15–25%. This improvement is achieved through dynamic policies that exhibit "aggressiveness-in-the-money": trading is accelerated when price movements are favorable (to minimize risk), and is slowed when price movements are unfavorable (to minimize transaction costs). Overall, the optimal dynamic policies exhibit much better performance in the right tail of worst outcomes, relative to optimal static policies.

From a mathematical perspective, our analysis exploits the dual representation of CVaR to convert the problem to a continuous-time, zero sum dynamic game. In this setting, we leverage the idea of the state-space augmentation, recently applied to certain discrete-time Markov decision processes with a CVaR objective. We obtain a partial differential equation describing the optimal value function, which is separable and a special instance of the Emden–Fowler equation. This leads to a closed-form solution. As our problem is a special case of a continuous-time linear-quadratic-Gaussian control problem with a CVaR objective, these results may be interesting in broader settings.

^{*}The authors wish to thank Carlos Gomez Gascon and Paul Glasserman for helpful discussions.

1. Introduction

Optimal execution is a problem of significant importance for algorithmic traders in modern financial markets. Here, a trader must decide an optimal trading strategy to buy (or sell) a fixed quantity of an asset. Typically, there is a trade-off between trading quickly, which minimizes the risk of adverse price movements, and trading slowly, which minimizes transaction costs. In their seminal paper, Almgren and Chriss (2000) framed this liquidation problem as a mean-variance optimization problem, that optimizes a combination of the average cost (i.e., the expected implementation shortfall) and the variability of the cost (i.e., the variance of the implementation shortfall). In that setting, they explicitly derive the optimal liquidation schedule, which is parameterized by the risk-aversion level of the trader. Their framework and suggested solution have become standards in this area, serving as a starting point for more complicated models and trading strategies, both in the academic literature and among practitioners.

However, the Almgren and Chriss (2000) analysis is restricted to static strategies: they only considered deterministic schedules under which the trader pre-commits to a fixed trading schedule, i.e., algorithms that do not adapt to changing market conditions such as the price of the asset. This restriction makes the analysis considerably more straightforward and tractable, and allows for closed-form solutions. This leaves open possible additional benefit from dynamic, adaptive trading strategies, however. Indeed, most practitioners utilize adaptive strategies. This is often done through ad hoc or heuristic means, such as model predictive control (MPC), where the trader periodically resolves for a new deterministic policy as time evolves.

Indeed, it has been an important objective to incorporate dynamic strategies into the Almgren-Chriss framework, but in a more principled way. Some practitioners such as Kissell and Malamut (2005) have suggested a series of heuristics that are price-adaptive in a particular way: they liquidate more aggressively when the price moves in a favorable direction, and liquidate more slowly when the price moves in an adverse direction. Lorenz and Almgren (2007) observed that this "aggressiveness in-the-money" (AIM) behavior can strictly improve on the optimal deterministic strategy in the mean-variance criterion. In a subsequent work, Lorenz and Almgren (2011) develop a dynamic programming technique by which approximate solutions can be numerically obtained, and show that such approximate solutions exhibit AIM, despite the lack of an analytic solution. See also Forsyth (2011) for the continuous-time version of this analysis.

Another stream of work (including the present paper) introduces alternative risk criteria other than the mean-variance objective, so as to formulate the problem into a more tractable form. For example, Schied and Schöneborn (2009) formulate the problem as an expected utility maximization problem, and derive a Hamilton-Jacobi-Bellman (HJB) equation that characterizes the optimal adaptive strategy. They find that the optimal strategy is either aggressive- or passive-in-themoney, respectively, if and only if the utility function displays increasing or decreasing risk aversion, respectively, but an analytic solution is not available. Gatheral and Schied (2011) propose an alternative risk criterion that utilizes the time-averaged risk exposure to the price change driven by the geometric Brownian motion (more precisely, the risk term is formulated as the time integration of the position value process, i.e., the product of the position process and the price process), and explicitly solve for the optimal strategy that is shown to exhibit the AIM behavior. However, their proposed risk criteria is ad hoc, and may encourage, for example, negative positions to reduce risk. Hence, the resulting optimal policies are not "liquidate-only": they may trade in both directions (buying and selling) in excess of what is strictly necessary to reduce the position to zero. Forsyth et al. (2012) investigate the use of the quadratic variation of the position value process as a risk criteria, and observe that the classic static solution of Almgren and Chriss (2000) is again optimal when the price process is driven by the arithmetic Brownian motion. The authors also consider the geometric Brownian motion, under which the optimal solution is numerically determined and shown to be price-adaptive, and they report its AIM behavior through numerical examples. Lin et al. (2015) introduces a composite dynamic coherent risk measures and derive the optimal solution that is tractable but static. One can also consider an entropic risk measure introduced such as that of Glasserman and Xu (2013), but it can be shown that the resulting strategy is also not priceadaptive.

To summarize, the prior work in this area either (i) yields only approximate numerical solutions; (ii) incorporates only ad hoc risk criteria; (iii) yields policies that are not liquidate-only and trade in both directions; or (iv) illustrates no benefit from dynamic policies over static policies. In contrast, our work simultaneously addresses all of these issues, and, to our knowledge, is the first paper to do so.

In this paper, we consider the conditional value-at-risk (CVaR; also known as average value-atrisk, tail conditional expectation, or expected shortfall) as an objective. CVaR is a risk measure that quantifies the tail risk. Given a quantile value $q \in (0, 1]$ and a random variable that represents the cost, the CVaR value at level q is defined as the average of the worst q-fraction of the outcomes, i.e., the tail average beyond the q^{th} quantile of the cost distribution. Smaller values of q focus on performance in the worst cases, i.e., the trader is more risk averse, while q = 1, on the other hand, corresponds to a risk neutral trader. Starting from the pioneering work of Artzner et al. (1999), CVaR has received much attention for its intuitive definition and for nice mathematical properties as a convex and coherent risk measure.

From the point of view of using CVaR as an optimization objective, the (static) optimization of the CVaR value for a single period can be done efficiently, by utilizing its primal variational representation (Rockafellar and Uryasev, 2002). In the multi-period setting of a Markov decision problem (MDP), however, the dynamic optimization of the CVaR value of the total cost is challenging. To begin, this objective is time-inconsistent. Moreover, as observed by Artzner et al. (2007) and Shapiro (2009), the optimal action at a point in time may not be completely determined by the current state of the MDP, but may depend on the entire history, and therefore the conventional dynamic programming techniques may not work.

More recent work has adopted the idea of state augmentation to overcome this issue: by in-

troducing an extra state variable, an optimal policy can be sufficiently characterized as a Markov process defined on this augmented state space. Broadly speaking, these studies develop CVaR MDP frameworks using two kinds of state augmentation. The first kind introduces an extra state variable that represents the running cost, and derives the dynamic programming principle from the primal variational representation of CVaR, i.e., $\text{CVaR}_q[C] = \min_{c \in \mathbb{R}} \{c + \frac{1}{q}\mathbb{E}[(C-c)^+]\}$ (see Rockafellar and Uryasev, 2002). This state augmentation scheme is adopted in, for example, Bäuerle and Ott (2011); Huang and Guo (2016); Miller and Yang (2017); Chow et al. (2018); Backhoff-Veraguas and Tangpi (2020). The second kind of state augmentation introduces an extra state variable that represents the quantile value, and derives the dynamic programming principle from the dual variational representation of CVaR, i.e., $\text{CVaR}_q[C] = \sup_{\mathbb{Q} \in \mathbb{Q} \ll \mathbb{R}, \frac{d\mathbb{Q}}{d\mathbb{R}} \leq q^{-1}} \mathbb{E}_{\mathbb{Q}}[C]$ (see Artzner et al., 1999). The work of Pflug and Pichler (2016); Chow et al. (2015); Chapman et al. (2018); Li et al. (2020) belongs to this category.

This paper. In this paper, we seek an adaptive liquidation strategy that minimizes the CVaR value of the implementation shortfall in a continuous-time, infinite-horizon variation of the classical Almgren-Chriss framework. We adopt the second kind of state augmentation described above. More specifically, we consider an augmented state space represented as (X_t, Q_t) , where $X_t \in \mathbb{R}$ is the current remaining position size of the trader, and $Q_t \in [0, 1]$ is a quantile value that represents the current level of risk neutrality. We observe that the dynamic optimization of the CVaR objective can be represented as a (continuous-time) stochastic game between the trader who controls the position process X_t and the adversary who controls the quantile process Q_t . By analyzing the Nash equilibrium of this game, we can identify the minimal CVaR value that the trader can achieve, and specify the trader's optimal policy and the adversary's optimal policy, which are formulated as time-stationary Markov policies on (X_t, Q_t) .

Practical contributions. Using our approach, we can express the optimal dynamic liquidation strategy in an analytic closed-form. This allows us to characterize properties of the optimal strategy, and yields a number of insights consistent with how optimal execution algorithms are widely employed in practice.

First, we show that the optimal strategy is liquidate-only, i.e., trades in only one direction and it keeps liquidating until it completes the execution. This is intuitive since, in practice, traders typically do not want to increase their positions or establish new positions during the liquidation process. We also observe that the optimal trading strategy becomes more aggressive and trades more quickly as the level of risk aversion increases. This is also intuitive since price risk can be minimized by trading quickly.

Second, and more interestingly, we show that the optimal trading strategy is dynamic and responds to stochastic fluctuations in the price process in a way that accelerates trading (to minimize price risk) when price movements are favorable, and slows trading (to minimize transaction costs) when price movements are unfavorable — in other words, it exhibits aggressiveness-in-the-money (AIM). The intuition behind the optimality of AIM is as follows:

Recall that the CVaR measure is the tail average beyond the q^{th} quantile — the CVaR objective is not impacted by cost realizations below this threshold but only requires the average cost above this threshold to be minimized. As an extreme case, imagine a situation when the price movements have been so favorable that the trader has earned a large profit from the holding positions and the current cumulative cost is far below the quantile threshold. The CVaR objective of a trader in this scenario is not impacted by the cost (so long as it remains below the threshold) and the trader should thus be willing to pay a large, deterministic transaction cost (up to the gap between the current cumulative cost and the threshold) so as to complete the execution quickly and minimize the risk of the total cost exceeding the threshold. In the opposite case, imagine a scenario where the current cumulative cost is far above the threshold. Since the total cost is very likely to ultimately exceed the threshold, the CVaR objective of the trader becomes the expected cost and is approximately risk neutral. The trader is thus encouraged to slow down trading so as to minimize (deterministic) transaction costs thereafter. As illustrated in these two extreme cases, the optimal strategy responds asymmetrically to price changes, in the way that it becomes more aggressive in an adverse situation, due to the intrinsic asymmetry of the CVaR measure.

Our stochastic game interpretation of this dynamic CVaR optimization problem provides an alternative (and more formal) justification of AIM. On the augmented state space, the adversary's quantile process Q_t represents the current level of risk neutrality (starting at value q), and is coupled with the price process. In particular, the adversary's optimal quantile process Q_t^* represents the probability that the current sample path is among the q-fraction of worst outcomes for the trader, and hence, Q_t^* increases when the price moves in an unfavorable direction. When Q_t^* increases, consequently, the trader becomes more risk-neutral and is encouraged to trade slowly, which is consistent with AIM.

Third, we are also able to quantify the benefit of dynamic trading by comparing the optimal dynamic policy to two benchmark static policies: the optimal static policy and an optimized version of a policy that trades at a constant rate. The later policy is meant to be representative of volume weighted average price (VWAP) polices that are popular in practice. The relative improvement of the optimal dynamic policy depends on the problem parameters only through the risk aversion q, and can be characterized analytically. For moderate levels of risk aversion, the optimal dynamic policy outperforms the optimal static policy by 5–15%, and outperforms the optimal dynamic policy by 15–25%. Moreover, relative improvement is increasing in q and grows without bound as $q \nearrow 1$, i.e., as the investor becomes more risk neutral. Since most traders using optimal execution algorithms are large investors trading a small portion of their overall portfolio over a short time horizon, their utility functions are nearly linear from the perspective of the optimal execution problem, hence the nearly risk neutral regime ($q \approx 1$), where the relative benefits of dynamic trading are greatest, is also the most practically relevant regime.

Last, numerical experiments yield insight to some useful features of the dynamic optimal policy. Namely, this policy effectively controls the worst outcomes in the (right) tail of the cost distribution, resulting in a cost distribution is very distinguished from the ones induced by deterministic policies. This feature may appeal to the practitioners who are not particularly interested in optimizing the CVaR value. For example, we observe that the class of CVaR-optimal dynamic policies achieves better median performance and better tail probabilities than the class of deterministic policies. See Figure 10 and the accompanying discussion for further details.

Mathematical contributions. We introduce a novel and technically sound approach to developing the CVaR MDP framework in the continuous-time setting. To sketch our approach briefly, we first introduce a scaled version of CVaR that allows us to avoid the ambiguity of CVaR in the corner case (i.e., when q = 0) and inherently induces the concavity of the objective (Proposition 1). We then utilize the martingale representation theorem, by which we can rewrite the CVaR objective as a maximization problem for an adversary who controls the quantile process Q_t against the decision maker (Theorem 1), and the problem can be translated into a continuous-time stochastic game between the decision maker and the adversary who are competing over the expected value of the risk-adjusted outcome (Theorem 2). After this step, the objective now decomposes over time, hence we can develop a CVaR dynamic programming principle (Theorem 3). This naturally leads to the Hamilton-Jacobi-Bellman (HJB) partial differential equations that characterize the optimal solution. The HJB equation is separable across the state variables, and hence admits a closed form solution for the value function.

From the HJB equation, we are able to define candidate optimal control policies for the trader and adversary. Unfortunately, we are not able to show that these policies as feasible. Instead, we construct a series of feasible policies that converge pointwise to the candidate optimal policy and in value to the optimal value function (see Theorem 6 and the accompanying discussion).

Our approach leverages the idea of state augmentation that incorporates the quantile value as an extra state variable, which has been suggested and utilized in prior work (Pflug and Pichler, 2016; Chow et al., 2015; Chapman et al., 2018; Li et al., 2020), but in exclusively discrete-time settings. To our knowledge, our work is the first to apply this in continuous-time, which introduces new technical challenges. We clearly take advantage of the continuous-time setting: the martingale representation theorem, which is intrinsically continuous-time, is fundamental in our analysis as it allows us to parameterize the adversary's control policy with a real-valued stochastic process that can be tractably optimized. Moreover, the choice of continuous-time is critical in this application, since the optimization of an analogous discrete-time model would involve a Bellman equation that is not separable, and hence is unlikely to admit an analytic solution (see the discussion at the end of §3.1).

As our problem is a special case of a continuous-time linear-quadratic-Gaussian control problem with a CVaR objective, these technical results may be interesting in broader settings, and feed into the broader emerging literature of continuous-time differential games.

Organization of paper. In §2, we introduce the notation and formally describe the model and

the problem. In §3, we develop the CVaR MDP framework for which we sequentially introduce a martingale representation of the CVaR objective, the game-theoretic representation of the problem, and the Markov policies defined on the augmented state space. In §4, we derive the HJB equation, identify its solution, and characterize the optimal liquidation strategy. In §5, we compare the optimal adaptive strategy with two deterministic strategies: the optimal deterministic strategy and the optimized VWAP strategy. In §6, we provide simulation results that illustrate the optimal strategy. In the appendix, we provide the proofs that are deferred from §2–§5.

2. Problem

We consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$, where \mathbb{F} is a natural filtration of a Brownian motion $(W_t)_{t\geq 0}$ that satisfies the usual conditions. We denote by $\mathcal{L}^p(\Omega, \mathcal{F}, \mathbb{P})$ (or simply by \mathcal{L}^p) the set of \mathcal{F} -measurable random variables $X : \Omega \to \mathbb{R}$ such that $\mathbb{E}|X|^p < \infty$. Given a sequence of random variables $(X_n)_{n\in\mathbb{N}}$, it is said that $X_n \xrightarrow{\mathcal{L}^p} X$ if $\lim_{n\to\infty} \mathbb{E}|X_n - X|^p = 0$. The time index set is denoted by $\mathbb{T} \triangleq [0, \infty)$. We also define \mathcal{P} to be the set of progressively measurable stochastic processes in this filtered probability space. We denote $\mathbb{R}_+ \triangleq [0, \infty)$ and $\mathbb{R}_- \triangleq (-\infty, 0]$.

2.1. Model

We consider a continuous-time and infinite-horizon version of the setting of Almgren and Chriss (2000). We postulate a trader who wants to liquidate $x \in \mathbb{X}$ units of an asset over an *infinite-time horizon*.¹ Here, x can be negative if the trader wants to acquire the asset. We define $\mathbb{X} \triangleq [-M, M] \subset \mathbb{R}$, where $M \ge 0$ is an arbitrary large number² that represents the largest possible position size that the trader is allowed to own.

Liquidation strategy. The trader's liquidation policy is represented with a real-valued continuoustime stochastic process $\pi \triangleq (\pi_t)_{t\geq 0}$, where $\pi_t \in \mathbb{R}$ specifies the *liquidation rate* at time t. Given an initial position size x and a liquidation strategy π , the trader's position process $X^{x,\pi} \triangleq (X_t^{x,\pi})_{t\geq 0}$ is determined by

$$X_t^{x,\pi} = x - \int_{s=0}^t \pi_s ds;$$

i.e., the trader liquidates π_t units of the asset per unit time (or acquires $-\pi_t$ units of the asset per unit time if $\pi_t < 0$). While deferring a formal statement to the end of this subsection, we will restrict our attention to the policies under which the trader's position varies continuously over time (i.e., involves no impulse trades) and vanishes eventually (i.e., converges to zero as t goes to

¹The choice of an infinite-horizon setting is important since as it will allow an analytical solution of the problem. In particular, in an infinite-horizon setting, the optimal policy is stationary over time, and hence time can be eliminated as a state variable. As we will see in subsequent sections, this results in a simpler Hamilton-Jacobi-Bellman (HJB) equation that can be solved analytically. Beyond this, in some sense the infinite-horizon setting is more elegant from a modeling perspective since it endogenizes the effective time horizon of the trader dynamically as a function of risk aversion. That said, the optimal policy for a finite-horizon variation of our setting could be obtained via a numerical solution of the HJB equation, and we expect the qualitative insights to be no different.

 $^{^{2}}$ We require that the position size be bounded in order to resolve technical difficulties (see Remark 1).

infinity).

Liquidation cost. Following the framework of Almgren and Chriss (2000), we define the *cost process* (or loss process) $C^{x,\pi} \triangleq (C_t^{x,\pi})_{t\geq 0}$ as

$$C_t^{x,\pi} \triangleq \int_{s=0}^t \frac{1}{2} \eta \pi_s^2 ds - \int_{s=0}^t \sigma X_s^{x,\pi} dW_s.$$
(1)

The first term $\int_{s=0}^{t} \frac{1}{2} \eta \pi_s^2 ds$ represents the (cumulative) transaction cost incurred by the temporary price impact, where the coefficient $\eta > 0$ reflects the illiquidity of the asset. The second term $\int_{s=0}^{t} \sigma X_s^{x,\pi} dW_s$ represents the (cumulative) loss incurred by the random fluctuation of the market price, where $\sigma > 0$ is the volatility of the price process and W_t is the standard Brownian motion. The total cost (i.e., total implementation shortfall) $C_{\infty}^{x,\pi} \triangleq \lim_{t\to\infty} C_t^{x,\pi}$ is a random variable of interest that we want to minimize via a CVaR objective.

Note that we do not consider permanent price impact in our formulation. Within the continuoustime Almgren–Chriss framework, it is well known that the contribution of permanent linear price impact to the implementation shortfall is path-independent; i.e., it does not depend on the liquidation strategy as long as the strategy clears all the positions eventually.³ Therefore, without loss of generality, we can ignore the presence of permanent impact since it does not make any difference in the trader's decision making. See Almgren and Chriss (2000) for details.

Admissible strategies. We now formally define the set of admissible policies $\Pi(x)$ as

$$\Pi(x) \triangleq \left\{ \pi : \mathbb{T} \times \Omega \to \mathbb{R} \middle| \begin{array}{l} \pi \in \mathcal{P}, \\ \mathbb{E}\left[\left(\int_{t=0}^{\infty} \pi_t^2 dt \right)^2 \right] < \infty, \ \mathbb{E}\left[\int_{t=0}^{\infty} |X_t^{x,\pi}|^2 dt \right] < \infty, \\ X_t^{x,\pi} \in \mathbb{X}, \ \forall t \ge 0 \end{array} \right\}.$$
(2)

An admissible policy $\pi \in \Pi(x)$ can be dynamic and so it can adjust the trading rate adaptively to the price changes. By this definition, impulse trades are not allowed by the constraint $\mathbb{E}\left[\left(\int_{t=0}^{\infty} \pi_t^2 dt\right)^2\right] < \infty$, and a non-vanishing position is also not allowed by the constraint $\mathbb{E}\left[\int_{t=0}^{\infty} |X_t^{x,\pi}|^2 dt\right] < \infty$. These conditions are to guarantee that the limit $C_{\infty}^{x,\pi} \triangleq \lim_{t\to\infty} C_t^{x,\pi}$ is an integrable random variable: individual terms in (1) converge in \mathcal{L}^2 as $t \to \infty$, and therefore $C_t^{x,\pi} \stackrel{\mathcal{L}^2}{\to} C_{\infty}^{x,\pi}$ for some $C_{\infty}^{x,\pi} \in \mathcal{L}^2$.

Remark 1. The last condition $X_t^{x,\pi} \in \mathbb{X} \triangleq [-M, M]$ is seemingly restrictive, but enforcing boundedness of the position path is merely a technical condition necessary for our later analysis. In particular, we will see that under the optimal policy, the position size will be monotonically decreasing towards zero over time. Therefore, we need only choose M so that the initial position size is feasible, i.e., $|x| \leq M$. Beyond this, the choice of M will not be a binding constraint: it will neither influence the optimal policy nor the optimal value function.

³Suppose that liquidating the asset at rate v permanently shifts the market price at rate $-\lambda v$, i.e., permanent linear price impact. Given a feasible policy π , the associated transaction costs can be expressed $\int_{t=0}^{\infty} \lambda \pi_t X_t^{x,\pi} dt$, which is always equal to $\frac{1}{2}\lambda x^2$, independent of the policy, given that $\lim_{t\to\infty} X_t^{x,\pi} = 0$.

2.2. Scaled Conditional Value-at-Risk

The conditional value-at-risk (CVaR) at a quantile level $q \in (0, 1]$ is a mapping from $\mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P})$ to \mathbb{R} . While there exist several definitions of CVaR in the literature, we consider the one based on its dual representation as a coherent risk measure (Artzner et al., 1999; Shapiro, 2009): given a random variable $C \in \mathcal{L}^1$,

$$\operatorname{CVaR}_{q}[C] \triangleq \sup_{Q^{\dagger} \in \mathcal{Q}^{\dagger}(q)} \mathbb{E}\left[CQ^{\dagger}\right] \quad \text{where} \quad \mathcal{Q}^{\dagger}(q) \triangleq \left\{Q^{\dagger} \in \mathcal{L}^{\infty}(\Omega, \mathcal{F}, \mathbb{P}), \ 0 \le Q^{\dagger} \le \frac{1}{q}, \ \mathbb{E}Q^{\dagger} = 1\right\},$$
(3)

i.e., it is defined as a maximization over a random variable Q^{\dagger} that has bounded support [0, 1/q]and an expected value of one (or equivalently, a maximization over a probability measure such that it is absolutely continuous with respect to \mathbb{P} and its Radon–Nikodyn derivative with respect to \mathbb{P} is upper bounded by 1/q almost surely).

We introduce a scaled version of the CVaR measure, which surprisingly simplifies our analysis.

Definition 1 (Scaled Conditional Value-at-Risk). Given a random variable $C \in \mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P})$ and a quantile $q \in [0, 1]$, the scaled conditional value-at-risk (S-CVaR) at level q is

$$S-CVaR_q[C] \triangleq \sup_{Q \in \mathcal{Q}(q)} \mathbb{E}[CQ], \qquad (4)$$

where $\mathcal{Q}(q)$ represents the risk envelope:

$$\mathcal{Q}(q) \triangleq \{ Q \in \mathcal{L}^{\infty}(\Omega, \mathcal{F}, \mathbb{P}) \mid 0 \le Q \le 1, \ \mathbb{E}Q = q \}.$$

The S-CVaR measure is obtained by simply scaling the risk envelope of CVaR by q. As a result, S-CVaR_q[C] is also given by $q \times \text{CVaR}_q[C]$ for $q \neq 0$. Nevertheless, it naturally incorporates the case q = 0 into its definition, for which CVaR is not well defined.⁴ It further has the following useful properties:

Proposition 1 (Properties of S-CVaR). For any random variable $C \in \mathcal{L}^1$, S-CVaR_q[C] satisfies the following properties:

- (i) S-CVaR_q[C] = $q \times \text{CVaR}_q[C]$, for any $q \in (0, 1]$.
- (*ii*) S-CVaR₀[C] = 0 and S-CVaR₁[C] = $\mathbb{E}C$.
- (iii) $|S-CVaR_q[C]| \leq \mathbb{E}|C|$ and $S-CVaR_q[C] \geq q\mathbb{E}C$ for any $q \in [0,1]$.
- (iv) The mapping $q \mapsto \text{S-CVaR}_q[C]$ is concave and continuous on [0, 1].

⁴When q = 0, $\text{CVaR}_0[C]$ is typically defined as the essential supremum of $C \in \mathcal{L}^1$, which can be infinite if the loss distribution has an unbounded support. We have defined the S-CVaR measure using the dual representation of CVaR so that we can effectively avoid the ambiguity at q = 0.

(v) Suppose that C is a continuous random variable whose distribution is atomless. Then,

S-CVaR_q[C] =
$$\mathbb{E}\left[C\mathbb{I}_{\left\{C \ge F_C^{-1}(1-q)\right\}}\right]$$

where $F_C^{-1}(\cdot)$ is the inverse distribution function of C, and $\operatorname{CVaR}_q[C] = \mathbb{E}\left[C \left| C \ge F_C^{-1}(1-q)\right]\right]$.

The proof can be found in Appendix B. Properties (i)–(iii) provide basic characterizations of S-CVaR. The property (v) provides an interpretation of S-CVaR as a truncated average as opposed to the interpretation of CVaR as a conditional average. We particularly highlight property (iv) that shows the concavity and continuity of the S-CVaR value with respect to q, which is a crucial property that will be exploited in our analysis.

One can interpret the definition (4) as a maximization problem for an adversary. This adversary selects a set of scenarios so as to maximize the average cost within the selected scenario, given a constraint that the total measure of the selected scenarios should be q. Informally,⁵ the optimized random variable Q^* is an indicator random variable such that $Q^*(\omega) = 1$ if the scenario ω is among the worst q-fraction of the scenarios, and $Q^*(\omega) = 0$ otherwise.

2.3. Risk-sensitive Execution with a CVaR Objective

We now introduce the CVaR risk criterion into the setting described in §2.1. In particular, we seek an adaptive strategy that minimizes the CVaR value of the implementation shortfall, given an initial position $x \in \mathbb{X}$ and a target quantile $q \in (0, 1]$. Without loss of generality, we formulate this optimization problem via an S-CVaR objective and define the value function $V : \mathbb{X} \times [0, 1] \to \mathbb{R}$ as

$$V(x,q) \triangleq \inf_{\pi \in \Pi(x)} \operatorname{S-CVaR}_q \left[C_{\infty}^{x,\pi} \right].$$
(*)

Note that the above formulation includes the case q = 0. By Proposition 1, the value function V(x,q) is well-defined at q = 0, and the minimal CVaR value is simply given by V(x,q)/q for any $q \neq 0$. We aim to identify the optimal value function V(x,q) as well as its corresponding optimal liquidation strategy π^* .

Recall that the objective S-CVaR_q[$C_{\infty}^{x,\pi}$] concerns the worst q-fraction of outcomes. When q = 1, the problem reduces to a risk-neutral liquidation problem. When q takes a smaller value, the problem is equivalent to considering a more risk-averse trader who concerns a smaller fraction of worst cases, being wary of more extreme cases. We anticipate that the trader uses this quantile value $q \in [0, 1]$ as an input to our algorithm so as to control the level of risk-aversion that he wants to achieve. In practice, we do not expect the traders to use an extremely small quantile value such as q = 0.01 or q = 0.05: since they encounter this sort of liquidation task often, possibly on a daily basis, it would be too conservative for them to optimize their performance in the worst 1% or 5%

⁵When the loss distribution has an atom at the q^{th} quantile, the extremal random variable $Q^{\star}(\omega)$ may take a fractional value.

of cases at a cost of sacrificing their performance in the normal 99% or 95% of cases.

3. CVaR Dynamic Programming Principle

Using the definition of the S-CVaR measure (4), the risk-sensitive optimal execution problem (*) can be formulated as

$$V(x,q) = \inf_{\pi \in \Pi(x)} \sup_{Q \in \mathcal{Q}(q)} \mathbb{E} \left[C_{\infty}^{x,\pi} Q \right].$$

As discussed in §2.2, we can think of an adversary who optimizes a random variable $Q \in Q(q)$ so as to select the worst q-fraction of sample paths against the trader who employs a liquidation policy $\pi \in \Pi(x)$. In this section, we reformulate the adversary's optimization problem as an optimization over a real-valued continuous-time stochastic process rather than a random variable, and interpret the risk-sensitive optimal control problem as a continuous-time stochastic game between the trader and the adversary. To this end, we develop a continuous-time dynamic programming principle by exploiting the recursive structure of this game.

3.1. Martingale Representation of CVaR Objective

We consider an arbitrary random variable $C \in \mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P})$ and derive an alternative representation of S-CVaR_q[C]. The results in this subsection are valid not only in the context of the liquidation problem, but also in any filtered probability space generated by a Brownian motion.

We define the *adversary's policy* as a real-valued continuous-time stochastic process $\gamma \triangleq (\gamma_t)_{t \ge 0}$, which determines the *adversary's quantile process* $Q^{q,\gamma} \triangleq (Q_t^{q,\gamma})_{t \ge 0}$:

$$Q_t^{q,\gamma} = q + \int_{s=0}^t \gamma_s \, dW_s,$$

where $W = (W_t)_{t \ge 0}$ is the Brownian motion that drives the random price fluctuation. We sometimes call γ_t the quantile diffusion rate by analogy to the liquidation rate π_t .

The set of admissible adversary's policies is defined as

$$\Gamma(q) \triangleq \left\{ \gamma : \mathbb{T} \times \Omega \to \mathbb{R} \mid \gamma \in \mathcal{P}, \ 0 \le Q_t^{q,\gamma} \le 1, \forall t \ge 0 \right\}.$$
(5)

Given an admissible adversary's policy $\gamma \in \Gamma(q)$, its corresponding quantile process $Q^{q,\gamma}$ is a (local) martingale starting at $q \in [0, 1]$ whose diffusion term is governed by γ . In particular, the quantile process is required to take values within [0, 1] and, as a result, it has the following properties (the proof is provided in Appendix C):

Proposition 2 (Properties of the adversary's quantile process Q). For any $\gamma \in \Gamma(q)$,

(i) $(Q_t^{q,\gamma})_{t\geq 0}$ is a continuous and bounded martingale taking values in [0,1], and hence $\mathbb{E}[Q_{\tau}^{q,\gamma}] = q$ for any stopping time τ .

- (ii) $Q^{q,\gamma}_{\infty} \triangleq \lim_{t \to \infty} Q^{q,\gamma}_t$ exists in [0,1] almost surely, and also $\mathbb{E}[Q^{q,\gamma}_{\infty}] = q$.
- (iii) Once $Q_t^{q,\gamma}$ hits 0 or 1, it never escapes thereafter.

By the martingale representation theorem, any random variable Q in the risk envelope $\mathcal{Q}(q)$ can be represented as the limit of a quantile process $Q^{q,\gamma}$ for some $\gamma \in \Gamma(q)$, and vice versa:

Lemma 1. For any $q \in [0,1]$, $\mathcal{Q}(q) = \{Q^{q,\gamma}_{\infty} | \gamma \in \Gamma(q)\}$.

Proof. Consider an arbitrary $\tilde{Q} \in \mathcal{Q}(q)$, and Doob martingale $(Q_t)_{t\geq 0}$ generated by \tilde{Q} , i.e., $Q_t \triangleq \mathbb{E}[\tilde{Q}|\mathcal{F}_t]$ for each t (we have $\lim_{t\to\infty} Q_t = \tilde{Q}$ and $Q_0 = \mathbb{E}\tilde{Q} = q$). By the martingale representation theorem (Protter, 2015, Thm. 43 in Chap. IV), there exists a predictable process γ such that $Q_t = Q_0 + \int_{s=0}^t \gamma_s dW_s$. Since $Q_t = \mathbb{E}[\tilde{Q}|\mathcal{F}_t] \in [0, 1]$ for any t, we have an admissible adversary's policy $\gamma \in \Gamma(q)$ and therefore $\tilde{Q} \in \{Q_{\infty}^{q,\gamma} | \gamma \in \Gamma(q)\}$ and $\mathcal{Q}(q) \subseteq \{Q_{\infty}^{q,\gamma} | \gamma \in \Gamma(q)\}$.

Now consider an arbitrary $\gamma \in \Gamma(q)$ and let $\tilde{Q} \triangleq Q_{\infty}^{q,\gamma}$. Trivially, $\mathbb{E}\tilde{Q} = Q_0 = q$ and $\tilde{Q} \in [0,1]$. Therefore, $Q_{\infty}^{q,\gamma} \in \mathcal{Q}(q)$, and hence $\mathcal{Q}(q) \supseteq \{Q_{\infty}^{q,\gamma} | \gamma \in \Gamma(q)\}$.

This alternative representation of the risk envelope Q(q) immediately leads to the following representation of the S-CVaR value, under which the adversary optimizes over the set of stochastic processes instead of the set of random variables:

Theorem 1 (Martingale representation of the CVaR objective). For any given random variable $C \in \mathcal{L}^1(\Omega, \mathcal{F}, \mathbb{P})$ and $q \in [0, 1]$, we have

$$S-CVaR_q[C] = \sup_{\gamma \in \Gamma(q)} \mathbb{E}\left[CQ_{\infty}^{q,\gamma}\right].$$
(6)

Proof. The result immediately follows from Lemma 1 and the definition of S-CVaR (4).

To better understand this adversary's optimization problem, consider a discrete-time setting with two periods, six sample paths and q = 1/2, as illustrated in Figure 1. The adversary is asked to assign the quantile values on the individual nodes so as to maximize the objective, $\mathbb{E}[CQ_2]$, given a constraint that the resulting quantile process (Q_0, Q_1, Q_2) needs to be a martingale.



Figure 1: An illustration of the discrete-time version of the adversary's martingale in a two-period setting with $q = \frac{1}{2}$. The underlying probability space is represented with a tree, in which each node represents the current state of the system, and each branch represents a possible transition from one state to another. We assume that six sample paths are equally likely to be realized. The label next to each terminal node represents the total cost C incurred along each sample path. The value in each node represents the value of the (optimal) adversary's martingale process at each state.

It is easy to see that the adversary's optimal solution is to select the worst q-fraction of sample paths, i.e., to assign the quantile value 1 to the q-fraction of terminal nodes with the largest realized cost and assign the quantile value 0 to the rest of terminal nodes. The quantile values of the non-terminal nodes are sequentially determined via backward induction (from t = 2 to t = 0) by averaging the quantile values of subsequent nodes. Here, the quantile value at each node means how likely the current sample path ends at one of the terminal nodes selected by the adversary.⁶

We can alternatively interpret this optimization problem as a sequential decision-making problem that the adversary assigns the quantile values in a forward direction (from t = 0 to t = 2). It begins with assigning the quantile value q to the root node. Starting from the root node, the adversary is asked to allocate quantile values to the subsequent nodes given a budget constraint that the average of the quantile values assigned to these nodes should equal the quantile value at the current node. This is repeated until all terminal nodes get assigned their quantile values. Here, the quantile value at each node means how much fraction of sample paths can be selected thereafter, and the budget constraint guarantees the resulting quantile process be a martingale.

The latter interpretation is closed related to Theorem 1. Indeed, for example, can consider a binomial tree model as a discrete-time approximation of the underlying Brownian motion, in which every node has two branches representing the events of price moving up or down by a small amount. In this binomial approximation, the quantile value allocation problem that the adversary solves at each node involves only one decision variable: given the current quantile value Q_t , the allocation

⁶Consider the optimal martingale Q^* that solves (6), and recall that the random variable $Q^*_{\infty}(\omega)$ indicates whether the sample path ω is among the worst *q*-fraction of scenarios. As a Doob martingale, the quantile process $Q^*_t = \mathbb{E}[Q^*_{\infty}|\mathcal{F}_t]$ is its running estizmate at time *t*, i.e., the likelihood that the current sample path will end up with one of the worst scenarios.

of next quantile values is of the form $\{(1 + \theta_t)Q_t, (1 - \theta_t)Q_t\}$, where θ_t is the real-valued decision variable. Here, determining the value of θ_t is effectively equivalent to determining the diffusion rate of Q_t , which is in fact the adversary's policy γ_t in our formulation.

We highlight that this martingale representation greatly simplifies the adversary's optimization problem and is only available in a continuous-time setting. In the discrete-time setting investigated in Pflug and Pichler (2016); Chow et al. (2015); Chapman et al. (2018); Li et al. (2020), for example, the adversary in each period needs to solve a much more complicated optimization that may involve many decision variables (if there are n possible random outcomes at a state, the quantile value allocation problem involves n - 1 real-valued decision variables). It may not easy in such cases to solve the adversary's optimization problem even numerically.

3.2. Risk-sensitive Liquidation as a Continuous-time Stochastic Game

We now return to the liquidation problem, and define an *outcome function* J as a function of the trader's policy $\pi \in \Pi(x)$ and the adversary's policy $\gamma \in \Gamma(q)$ at each $x \in \mathbb{X}$ and $q \in [0, 1]$:

$$J(\pi, \gamma; x, q) \triangleq \mathbb{E}\left[C_{\infty}^{x, \pi} Q_{\infty}^{q, \gamma}\right].$$

By Theorem 1, the value function (*) can be formulated as

$$V(x,q) = \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} J(\pi,\gamma;x,q).$$

The following theorem characterizes this value function as an equilibrium outcome of a continuoustime stochastic game between the trader and the adversary.

Theorem 2 (CVaR optimization as a continuous-time stochastic game). The value function V(x,q) is the outcome at the Nash equilibrium of the zero-sum game in which the trader and the adversary compete over the outcome J,

$$V(x,q) = \inf_{\pi \in \Pi(x)} \max_{\gamma \in \Gamma(q)} J(\pi,\gamma;x,q) = \max_{\gamma \in \Gamma(q)} \inf_{\pi \in \Pi(x)} J(\pi,\gamma;x,q),$$

where an optimal solution $\gamma \in \Gamma(q)$ must exist for each maximization.

Theorem 2 states that the minimax solution equals the maximin solution. This means that the value function is, as a saddle point, the equilibrium outcome at which each player simultaneously plays the best response against the other player's strategy. This may not always hold true for a general class of risk-sensitive control problems: the convexity of the outcome function with respect to the trader's policy and the convexity of the policy space are required in our proof (Appendix C.1) in order to apply Sion's minimax theorem (Sion, 1958).

The following remark provides an alternative interpretation of this game based on the Girsanov theorem.

Remark 2. Recall that the dual representation of the CVaR objective (3) involves a multiplication with a random variable Q^{\dagger} , which corresponds to an absolutely continuous change of measure with Radon-Nikodym derivative Q^{\dagger} . In terms of the adversary's quantile process $Q^{q,\gamma}$, this change of measure can by written as $Q^{\dagger} \triangleq Q_{\infty}^{q,\gamma}/q$. By the Girsanov theorem, the price process has a drift under this new measure, and is given by

$$dP = \sigma d\tilde{W}_t + \sigma \frac{\gamma_t}{Q_t^{q,\gamma}} dt,$$

where \tilde{W} is a Brownian motion under the new measure. In other words, the trader's optimization problem given the adversary's policy γ is can also be viewed as a risk-neutral liquidation problem under altered price dynamics in the presence of adversary who can introduce drift into the price process.

3.3. CVaR Dynamic Programming Principle

In this subsection, we develop a continuous-time dynamic programming principle. We first state a proposition that characterizes a temporal structure of the game.

Proposition 3 (Time decomposition). Fix $x \in \mathbb{X}$ and $q \in [0,1]$. For any trader's policy $\pi \in \Pi(x)$, adversary's policy $\gamma \in \Gamma(q)$, and stopping time τ , we have

$$J(\pi,\gamma;x,q) = \mathbb{E}\left[C_{\tau}^{x,\pi}Q_{\tau}^{q,\gamma} + \mathbb{E}\left[\left(C_{\infty}^{x,\pi} - C_{\tau}^{x,\pi}\right)Q_{\infty}^{q,\gamma}\right|\mathcal{F}_{\tau}\right]\right].$$

Proof. Observe that $C_{\tau}^{x,\pi}$ is \mathcal{F}_{τ} -measurable and $\mathbb{E}[Q_{\infty}^{q,\gamma}|\mathcal{F}_{\tau}] = Q_{\tau}^{q,\gamma}$ since $Q^{q,\gamma}$ is a martingale. Utilizing the tower property, we obtain $J(\pi,\gamma;x,q) = \mathbb{E}\left[C_{\infty}^{x,\pi}Q_{\infty}^{q,\gamma}\right] = \mathbb{E}\left[C_{\infty}^{x,\pi}Q_{\infty}^{q,\gamma} - (C_{\infty}^{x,\pi} - C_{\tau}^{x,\pi})Q_{\infty}^{q,\gamma}\right] = \mathbb{E}\left[\mathbb{E}(C_{\tau}^{x,\pi}Q_{\infty}^{q,\gamma}|\mathcal{F}_{\tau}) + \mathbb{E}((C_{\infty}^{x,\pi} - C_{\tau}^{x,\pi})Q_{\infty}^{q,\gamma}|\mathcal{F}_{\tau})\right] = \mathbb{E}\left[C_{\tau}^{x,\pi}Q_{\tau}^{q,\gamma} + \mathbb{E}((C_{\infty}^{x,\pi} - C_{\tau}^{x,\pi})Q_{\infty}^{q,\gamma}|\mathcal{F}_{\tau})\right].$

Note that $C_{\infty}^{x,\pi} - C_{\tau}^{x,\pi}$ represents the cost realized after time τ . Proposition 3 states that the final outcome can be decomposed into two terms: one term describes the subgame before time τ , and the other term describes the subgame after time τ .

Observe that, in the subgame after time τ , the trader is liquidating $X_{\tau}^{x,\pi}$ shares and the adversary is selecting a $Q_{\tau}^{q,\gamma}$ -fraction of the future scenarios realized thereafter. This time decomposition naturally leads to the following dynamic programming principle:

Theorem 3 (CVaR dynamic programming principle). For any $x \in \mathbb{X}$, $q \in [0,1]$, and a stopping time τ with $\mathbb{E}[\tau] < \infty$, we have

$$V(x,q) = \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{x,\pi}_{\tau} Q^{q,\gamma}_{\tau} + V \left(X^{x,\pi}_{\tau}, Q^{q,\gamma}_{\tau} \right) \right].$$
(7)

Theorem 3 provides the optimality principle in the form of Bellman's equation: at the equilibrium, the outcome after time τ can be sufficiently described by the subgame equilibrium $V(X_{\tau}^{x,\pi}, Q_{\tau}^{q,\gamma})$.

The trader is minimizing the (risk-adjusted) cost up to time τ in a consideration of his future state X_{τ} , while the adversary is simultaneously maximizing the (risk-adjusted) cost up to time τ in a consideration of his future state Q_{τ} , and the subgame starts at those future states. Like Theorem 2, Theorem 3 relies on the saddle-point characterization of the equilibrium; i.e., it does not matter which player commits his policy first in the subgame. The formal proof can be found in Appendix C.3.

3.4. (X, Q)-Markov Policies

Theorem 3 implies that the *augmented state space* (X_t, Q_t) is sufficient to describe the remaining subgame at time t. Therefore, if a policy is reasonable, its action at time t $(\pi_t \text{ or } \gamma_t)$ should be determined by the current position size X_t and the current quantile value Q_t . To formalize this idea, we introduce time-stationary Markov policies running on this augmented state space:

Definition 2 ((X,Q)-Markov policies). We say that a trader's policy π is an (X,Q)-Markov policy coupled with γ if

$$\pi_t(\omega) = f\left(X_t^{x,\pi}(\omega), Q_t^{q,\gamma}(\omega)\right), \quad \forall t, \omega$$
(8)

for some measurable function $f : \mathbb{R} \times [0, 1] \to \mathbb{R}$.

Similarly, an adversary's policy γ is an (X, Q)-Markov policy coupled with π if

$$\gamma_t(\omega) = g\left(X_t^{x,\pi}(\omega), Q_t^{q,\gamma}(\omega)\right) \tag{9}$$

for some measurable function $g : \mathbb{R} \times [0, 1] \to \mathbb{R}$.

A policy pair (π, γ) is a mutually coupled (X, Q)-Markov policy pair if both (8) and (9) hold.

An (X, Q)-Markov policy is characterized by a function defined on the augmented state space. The function f or g specifies the liquidation rate or the quantile diffusion rate when the current position size is x and the current quantile level is q. Recall that we have defined a policy, π or γ , as a continuous-time stochastic process adapted to the Brownian motion, i.e., as a (progressively measurable) mapping $\pi : \mathbb{T} \times \Omega \to \mathbb{R}$. Strictly speaking, the function f or g does not completely determine one player's policy unless the other player's policy is specified. To avoid this ambiguity, when we describe an (X, Q)-Markov policy, we specify the other player's policy that is coupled with it.

To better understand, consider the policies π and γ that are mutually coupled (X, Q)-Markov policies induced by functions f and g. Under π and γ , the system is completely described by the coupled processes $(X_t, Q_t)_{t\geq 0}$ on the augmented state space, whose dynamics are given by the following stochastic differential equations:

$$dX_t = -f(X_t, Q_t) dt, \quad dQ_t = g(X_t, Q_t) dW_t,$$

with the initial states $X_0 = x$ and $Q_0 = q$. Even if γ is not an (X, Q)-Markov policy, we can still consider an (X, Q)-Markov policy π that is induced by f and coupled with γ , and then the position process will be given by $dX_t = -f(X_t, Q_t^{q,\gamma}) dt$.

Note also that the admissibility of an (X, Q)-Markov policy is not always guaranteed: it may fail to satisfy the admissible conditions given in (2) or (5), depending on the generating function f or g as well as the other player's policy coupled with it. See the discussions before and after Theorem 6.

4. Optimal Solution

In this section, we utilize the CVaR dynamic programming principle (Theorem 3) to derive a Hamilton–Jacobi–Bellman (HJB) equation for the risk-sensitive optimal execution problem, and identify the functional form of the value function and the optimal policies by solving this HJB equation. For all propositions/theorems stated in this section, we defer their proofs to Appendix D.

4.1. Minimal CVaR Cost

We first state Dynkin's formula that we can apply to the right-hand side of (7) in Theorem 3 so as to represent it as a time-integration.

Proposition 4 (Dynkin's formula). Consider a function $\hat{V} : \mathbb{X} \times [0,1] \to \mathbb{R}$ such that $\hat{V} \in \mathcal{C}^{1,2}(\mathbb{X} \times [0,1])$. For any $\pi \in \Pi(x)$, $\gamma \in \Gamma(q)$, and stopping time τ , we have

$$\mathbb{E}\left[C_{\tau}^{x,\pi}Q_{\tau}^{q,\gamma} + \widehat{V}\left(X_{\tau}^{x,\pi}, Q_{\tau}^{q,\gamma}\right)\right] = \widehat{V}\left(x,q\right) \\ + \mathbb{E}\left[\int_{t=0}^{\tau} \left\{\frac{\eta}{2}Q_{t}^{q,\gamma}\pi_{t}^{2} - \widehat{V}_{x}\left(X_{t}^{x,\pi}, Q_{t}^{q,\gamma}\right)\pi_{t}\right\}dt\right] \\ + \mathbb{E}\left[\int_{t=0}^{\tau} \left\{\frac{1}{2}\widehat{V}_{qq}\left(X_{t}^{x,\pi}, Q_{t}^{q,\gamma}\right)\gamma_{t}^{2} - \sigma X_{t}^{x,\pi}\gamma_{t}\right\}dt\right],$$

where $\widehat{V}_x(x,q) \triangleq \frac{\partial}{\partial x} \widehat{V}(x,q)$ and $\widehat{V}_{qq}(x,q) \triangleq \frac{\partial^2}{\partial q^2} \widehat{V}(x,q)$.

For the sake of argument, suppose that the value function V can be plugged into Proposition 4 in the place of \hat{V} . When considering an infinitesimal time interval (i.e., $\tau = dt$), we have

$$\begin{split} V(x,q) \stackrel{\text{Thm 3}}{=} & \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau}^{x,\pi} Q_{\tau}^{q,\gamma} + V \left(X_{\tau}^{x,\pi}, Q_{\tau}^{q,\gamma} \right) \right] \\ \stackrel{\text{Prop 4}}{=} & \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \left\{ V(x,q) + \left(\frac{\eta}{2} q \pi_0^2 - V_x \left(x, q \right) \pi_0 \right) dt + \left(\frac{1}{2} V_{qq} \left(x, q \right) \gamma_0^2 - \sigma x \gamma_0 \right) dt \right\} \\ &= V(x,q) + \inf_{\pi \in \Pi(x)} \left\{ \frac{\eta}{2} q \pi_0^2 - V_x \left(x, q \right) \pi_0 \right\} dt + \sup_{\gamma \in \Gamma(q)} \left\{ \frac{1}{2} V_{qq} \left(x, q \right) \gamma_0^2 - \sigma x \gamma_0 \right\} dt. \end{split}$$

Observe that the terms associated with the trader's policy π and the terms associated with the

adversary's policy γ can be separated. This informal argument suggests that the value function V has to satisfy the following HJB equation:

$$\min_{v \in \mathbb{R}} \left\{ \frac{\eta}{2} q v^2 - V_x(x, q) v \right\} + \max_{w \in \mathbb{R}} \left\{ \frac{1}{2} V_{qq}(x, q) w^2 - \sigma x w \right\} = 0.$$

In the following theorem, we make this argument more formally and identify the sufficient conditions to be the optimal value function.

Theorem 4 (Verification theorem). Consider a function $V^* : \mathbb{X} \times [0,1] \to \mathbb{R}_+$ satisfying

(i) $V^{\star} \in \mathcal{C}^{1,2}(\mathbb{X} \times (0,1))$, and, for any $x \in \mathbb{X}$ and $q \in (0,1)$, it satisfies

$$\min_{v \in \mathbb{R}} \left\{ \frac{\eta}{2} q v^2 - V_x^\star \left(x, q \right) v \right\} + \max_{w \in \mathbb{R}} \left\{ \frac{1}{2} V_{qq}^\star \left(x, q \right) w^2 - \sigma x w \right\} = 0, \tag{10}$$

where $V_x^{\star} \triangleq \frac{\partial V^{\star}}{\partial x}$ and $V_{qq}^{\star} \triangleq \frac{\partial^2 V^{\star}}{\partial q^2}$.

- (*ii*) $V^{\star}(0,q) = 0$ for all $q \in [0,1]$, and $V^{\star}(x,0) = V^{\star}(x,1) = 0$ for all $x \in \mathbb{X}$.
- (iii) $V^{\star}(x,q) = V^{\star}(-x,q)$ for all $x \in \mathbb{X}$ and $q \in [0,1]$, and $V^{\star}(x,q)$ is increasing in x on [0,M].
- (iv) $\frac{V_x^{\star}(x,q)}{q}$ is increasing in x on X for each $q \in (0,1)$, and decreasing in q on (0,1) for each $x \in X$.
- (v) $\frac{x}{V_{qq}^{*}(x,q)}$ is decreasing in x on [0, M].

Then, $V(x,q) = V^{\star}(x,q)$ for all $x \in \mathbb{X}$ and $q \in [0,1]$.

In Theorem 4, conditions (i) and (ii) specify the HJB equation and the boundary conditions that the value function has to satisfy. In fact, the value function V can be uniquely determined by these two conditions. However, the other conditions, (iii)–(v), are also necessary to show that this value is indeed achievable within our definition of the admissible policies, $\Pi(x)$ and $\Gamma(q)$. More specifically, condition (iii) asserts the symmetry and the monotonicity of the value function with respect to position size x, and conditions (iv) and (v) assert certain behaviors of the optimal policies that are implied from the HJB equation (e.g., the optimal liquidation strategy should trade more aggressively when liquidating a larger quantity). While these properties of the value function are natural given the problem structure, they serve as regularity conditions in our proof to resolve technical issues arising in the convergence analysis.

Observe that the optimization terms in the HJB equation (10) are separated and each of them is a trivial quadratic optimization problem. By solving these optimizations explicitly, the HJB equation can be translated into the following partial differential equation:

$$V_x^2(x,q) \times V_{qq}(x,q) = -\sigma^2 \eta \times x^2 \times q.$$

It turns out that this differential equation with the boundary condition (ii) admits a separable

solution. The value function V(x,q) can be represented as a product of a function of x and a function of q, as identified in the following theorem.

Theorem 5 (Value function). Consider a function $V^* : \mathbb{X} \times [0,1] \to \mathbb{R}$ defined as

$$V^{\star}(x,q) \triangleq (3/4)^{\frac{2}{3}} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times |x|^{\frac{4}{3}} \times \varphi(q), \tag{11}$$

where $\varphi : [0,1] \to \mathbb{R}_+$ is the solution in $\mathcal{C}((0,1))$ to the following differential equation:

$$\varphi^2(q) \times \varphi''(q) = -q, \ \forall q \in (0,1), \quad and \quad \varphi(0) = \varphi(1) = 0.$$
(12)

Then, V^* satisfies the conditions of Theorem 4, and hence $V(x,q) = V^*(x,q)$ for all $x \in \mathbb{X}$ and $q \in [0,1]$.



Figure 2: (Left) Value function V(x,q) that represents the minimal S-CVaR loss (i.e., S-CVaR_q[C_{∞}]) at a quantile level q when liquidating x units of an asset given that $\sigma = \eta = 1$. The closed-form expression is provided in (11). (Right) Function $\varphi(q)$ that is derived in Proposition 5. The curve $\{(q,\varphi(q))\}_{q\in[0,1]}$ is represented in parametric form, in which the left part of the curve is represented as $\{(q_L(\theta),\varphi_L(\theta))\}_{\theta\in[0,\infty]}$, and the right part of the curve is represented as $\{(q_R(\theta),\varphi_R(\theta))\}_{\theta\in[0,\overline{\alpha}]}$.

The differential equation of form (12) is known as the *Emden-Fowler equation* (Polyanin and Zaitsev, 2003, 2.3.27), and its solution can be expressed in parametric form as follows:

Proposition 5 (Parametric representation of $\varphi(q)$). The function $\varphi(q)$ can be represented in a parametric form that admits a closed-form expression. Define

$$Z_L(\theta) \triangleq -\frac{2}{\pi} K_{1/3}(\theta), \quad Z_R(\theta) \triangleq \sqrt{3} J_{1/3}(\theta) - Y_{1/3}(\theta)$$

where J and Y are the first and second kinds of Bessel functions, and K is the second kind of

modified Bessel function. Further define

$$\bar{\theta} \triangleq \inf\{\theta > 0 : Z_R(\theta) = 0\} \approx 2.3834, \quad a \triangleq \frac{1}{\bar{\theta}^{\frac{4}{3}} \left(Z'_R(\bar{\theta})\right)^2} \approx 0.2910, \quad b \triangleq a \left(9/2\right)^{\frac{1}{3}} \approx 0.1763.$$

Then, the curve $\{(q,\varphi(q))\}_{q\in[0,1]}$ is parameterized as

$$\{(q,\varphi(q))\}_{q\in[0,1]} = \{(q_L(\theta),\varphi_L(\theta))\}_{\theta\in[0,\infty]} \bigcup \{(q_R(\theta),\varphi_R(\theta))\}_{\theta\in[0,\bar{\theta}]},$$

 $where^7$

$$q_L(\theta) \triangleq a\theta^{-\frac{2}{3}} \left[\left(\theta Z'_L(\theta) + \frac{1}{3} Z_L(\theta) \right)^2 - \theta^2 Z_L^2(\theta) \right], \quad \varphi_L(\theta) \triangleq b\theta^{\frac{2}{3}} Z_L^2(\theta), \tag{13}$$

and

$$q_R(\theta) \triangleq a\theta^{-\frac{2}{3}} \left[\left(\theta Z_R'(\theta) + \frac{1}{3} Z_R(\theta) \right)^2 + \theta^2 Z_R^2(\theta) \right], \quad \varphi_R(\theta) \triangleq b\theta^{\frac{2}{3}} Z_R^2(\theta).$$
(14)

4.2. Optimal Adaptive Liquidation Strategy

Let $f^*(x,q)$ and $g^*(x,q)$ be, respectively, the minimizer and the maximizer of the optimization terms in the HJB equation (10), i.e.,

$$f^{\star}(x,q) \triangleq \underset{v \in \mathbb{R}}{\operatorname{argmin}} \left\{ \frac{\eta}{2} q v^2 - V_x(x,q) v \right\} = \frac{V_x(x,q)}{\eta q} = (3/4)^{-\frac{1}{3}} \times \sigma^{\frac{2}{3}} \eta^{-\frac{2}{3}} \times x^{\frac{1}{3}} \times \frac{\varphi(q)}{q}, \quad (15)$$

$$g^{\star}(x,q) \triangleq \operatorname*{argmax}_{w \in \mathbb{R}} \left\{ \frac{1}{2} V_{qq}(x,q) \, w^2 - \sigma x w \right\} = \frac{\sigma x}{V_{qq}(x,q)} = -(3/4)^{-\frac{2}{3}} \times \sigma^{\frac{1}{3}} \eta^{-\frac{1}{3}} \times x^{-\frac{1}{3}} \times \frac{\varphi^2(q)}{q} (16)$$

where we define $x^{\frac{1}{3}} = -|x|^{\frac{1}{3}}$ for x < 0. The function $f^{\star}(x,q)$ specifies the trader's optimal liquidation rate when the current position size is x and the current quantile level is q, and similarly, the function $g^{\star}(x,q)$ specifies the adversary's optimal quantile diffusion rate in that situation. We can naturally postulate mutually coupled (X,Q)-Markov policies π^{\star} and γ^{\star} induced by these functions f^{\star} and g^{\star} , which are characterizing the equilibrium of the stochastic game.⁸ Under policies π^{\star} and γ^{\star} , the system is described by the following stochastic differential equations:

$$dX_t = -f^{\star}(X_t, Q_t) dt, \quad dQ_t = g^{\star}(X_t, Q_t) dW_t, \tag{17}$$

with $X_0 = x$ and $Q_0 = q$.

⁷The values of $Z_L(\theta)$ and $Z_R(\theta)$ are not defined when $\theta = 0$. However, the limit points (e.g., $\lim_{\theta \to 0} (q_L(\theta), \varphi_L(\theta))$), $\lim_{\theta \neq \infty} (q_L(\theta), \varphi_L(\theta))$) do exist, and our parametric representation includes those limit points.

⁸This does not mean that the policy π^* is optimal against any adversary's policy γ , nor an (X, Q)-Markov policy induced by f^* and coupled with γ is the best response against γ . It merely means that π^* is the best response against γ^* only, and vice versa. In order to obtain a best response against an arbitrary adversary's policy $\gamma \in \Gamma(q)$, we may need to characterize the best possible performance against γ , e.g., $V(x,q;\gamma) \triangleq \inf_{\pi \in \Pi(x)} J(\pi,\gamma;x,q)$, and derive and solve the HJB equation associated with it. Nevertheless, the policy π^* is the optimal liquidation strategy that minimizes the CVaR value of implementation shortfall.



Figure 3: Optimal trading rate function $f^*(x,q)$ (left) and optimal quantile diffusion rate function $g^*(x,q)$ (right) given that $\sigma = \eta = 1$. The closed-form expressions are provided in (15) and (16).

However, we cannot directly show that the policies π^* and γ^* satisfy the admissibility conditions introduced in (2) and (5), because the functions f^* and g^* exhibit extreme behaviors near the boundaries, as shown in Figure 3. For example, if $Q_t \searrow 0$, the liquidation rate π_t may increase unboundedly since $\lim_{q \searrow 0} f^*(x,q) = \infty$, and thus may violate the condition $\mathbb{E}\left[\left(\int_{t=0}^{\infty} \pi_t^2 dt\right)^2\right] < \infty$. If $Q_t \nearrow 1$, on the other hand, the liquidation rate π_t may vanish since $\lim_{q \nearrow 1} f^*(x,q) = 0$, and thus may violate the condition $\mathbb{E}\left[\int_{t=0}^{\infty} |X_t^{x,\pi}|^2 dt\right] < \infty$.

We instead prove that the optimal value function can be achieved "asymptotically" by a sequence of admissible policies that approximate π^* and γ^* .

Theorem 6 (Policy optimality). There exists a sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ with $f^{(n)}$: $\mathbb{X} \times [0, 1] \to \mathbb{R}$ and $g^{(n)} : \mathbb{X} \times [0, 1] \to \mathbb{R}$ such that

$$\lim_{n \to \infty} f^{(n)}(x,q) = f^{\star}(x,q), \quad \lim_{n \to \infty} g^{(n)}(x,q) = g^{\star}(x,q), \quad \forall (x,q) \in \mathbb{R} \setminus \{0\} \times (0,1],$$

and it satisfies the following properties for any $x \in \mathbb{X}$ and $q \in [0, 1]$:

(a) For any given $\gamma \in \Gamma(q)$, let $\pi^{(n),\gamma}$ be an (X,Q)-Markov trader's policy induced by $f^{(n)}$ and coupled with γ . Then, $\pi^{(n),\gamma}$ is admissible and

$$\limsup_{n \to \infty} J(\pi^{(n),\gamma},\gamma;x,q) \le V(x,q), \quad \forall \gamma \in \Gamma(q).$$

(b) For any given $\pi \in \Pi(x)$, let $\gamma^{(n),\pi}$ be an (X,Q)-Markov adversary's policy induced by $g^{(n)}$ and coupled with π . Then, $\gamma^{(n),\pi}$ is admissible and

$$\liminf_{n \to \infty} J(\pi, \gamma^{(n), \pi}; x, q) \ge V(x, q), \quad \forall \pi \in \Pi(x).$$

(c) Let $(\pi^{(n)}, \gamma^{(n)})$ be a mutually coupled (X, Q)-Markov policy pair induced by $(f^{(n)}, g^{(n)})$. Then, $\pi^{(n)}$ and $\gamma^{(n)}$ are admissible and

$$\lim_{n \to \infty} J(\pi^{(n)}, \gamma^{(n)}; x, q) = V(x, q).$$

Theorem 6 shows that we can construct a sequence of functions $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ that converges to (f^*, g^*) pointwise except at the boundaries, and further induces (X, Q)-Markov policies that are admissible and asymptotically optimal. More precisely, against any adversary's policy γ , the sequence of functions $(f^{(n)})_{n \in \mathbb{N}}$ induces a sequence of admissible policies $(\pi^{(n)}, \gamma)_{n \in \mathbb{N}}$ for the trader, and in the limit, the trader achieves an outcome that is no worse than the equilibrium outcome. And vice versa, against any trader's policy π , the sequence of functions $(g^{(n)})_{n \in \mathbb{N}}$ induces a sequence of admissible policies $(\gamma^{(n),\pi})_{n \in \mathbb{N}}$ for the adversary, and in the limit, the adversary achieves an outcome that is no worse than the equilibrium outcome. As a combination, the sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ induces a sequence of mutually admissible policy pairs $(\pi^{(n)}, \gamma^{(n)})_{n \in \mathbb{N}}$ that yields the equilibrium outcome asymptotically.

The construction of such a sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ is straightforward. We consider a vanishing subset of the augmented state space that contains the boundaries, i.e., $\{(x,q)||x| \leq \frac{1}{n}, q \leq \frac{1}{n} \text{ or } q \geq 1 - \frac{1}{n}\}$, and obtain $f^{(n)}$ and $g^{(n)}$ by suppressing the extreme behaviors of f^* and g^* arising in this subset. Roughly speaking, the liquidation strategy induced by $(f^{(n)}, g^{(n)})$ mimics the optimal strategy until it clears almost all positions (i.e., $X_t \leq \frac{1}{n}$) or it becomes almost risk-neutral (i.e., $Q_t \geq 1 - \frac{1}{n}$) or extremely risk-averse (i.e., $Q_t \leq \frac{1}{n}$), and then liquidates according to a deterministic schedule thereafter. We can show that the gap between the outcome of this approximated strategy and the theoretical equilibrium outcome is diminishing as n goes to infinity. We refer the readers to Appendix D for the details.

Aggressiveness-in-the-money. Despite that the admissibility of the optimal liquidation policy π^* is not guaranteed, we can still characterize its behaviors by inspecting the stochastic differential equations (17). Without loss of generality, let us consider the task of liquidation (i.e., $x \ge 0$).

First, we observe that the optimal policy liquidates only (i.e., $f^*(x,q) \ge 0$), until it completes the execution⁹ (i.e., $f^*(0,q) = 0$). Note that we have not imposed any constraint on the trading direction. This formally shows that winding back the position during the liquidation process will never be helpful in reducing the CVaR loss.

Second, when the trader becomes more risk-averse (i.e., $q \searrow 0$), the optimal policy trades more aggressively (i.e., $f^*(x,q) \nearrow \infty$). The opposite also holds true. This is because, by liquidating the position more quickly, it can reduce the risk exposure to the changes in price more effectively. Even though it will be more costly in terms of market impact, it can make sure that the transactions

⁹We are not sure if the completion time is almost surely finite even though the position will be vanishing eventually (i.e., $X_t \searrow 0$). In particular, when $Q_t \approx 1$, the optimal policy trades very slowly ($\pi_t \approx 0$) and the position process may never hit zero. We believe that the completion time is finite with a probability of at least 1 - q.

will be made at a certain level, which is more beneficial to a risk-averse trader than a risk-neutral trader. We can also understand this behavior based on the alternative interpretation of the problem discussed in Remark 2: when the risk-averse liquidation problem is cast as a risk-neutral execution problem that involves an adverse price drift, being more risk-averse is equivalent to facing a more adverse price drift, which encourages the risk-neutral trader to trade more aggressively.

Most interestingly, we observe that when the price moves in a favorable direction toward in-themoney (i.e., $dW_t > 0$), the policy becomes more risk-averse (i.e., $dQ_t < 0$ since $g^*(x,q) < 0$) and hence it trades more aggressively. This formally characterizes aggressiveness-in-the-money, which has been observed by Lorenz and Almgren (2007, 2011); Gatheral and Schied (2011); Forsyth et al. (2012). Intuitively, if the trader has made some "free" money due to the price change, he would have an additional incentive to complete the liquidation early so as to secure his current profit, and thus he would be willing to pay an additional deterministic cost for aggressive execution.

This behavior can also be understood in the context of a more general risk-sensitive optimal control problem. Recall that the optimal adversary's martingale Q_t^{q,γ^*} represents the likelihood that the current sample path leads to one of the worst q-fraction of outcomes. When something favorable happens, it becomes less likely that the sample path is in the worst q^{th} quantile, and therefore Q_t decreases. This means that the trader will need to pay attention to a smaller fraction of adverse scenarios; i.e., he will become more risk-averse.

Threshold behavior. While we do not have a formal characterization here, we observe that the optimal strategy exhibits some threshold behavior, particularly near the end of the liquidation. We observe that the policy trades aggressively when the cumulative cost C_t is below some threshold, and it trades passively when the cumulative cost C_t is above the threshold. Near the end of the liquidation, such a threshold corresponds to $\operatorname{VaR}_q[C_\infty]$ (the value-at-risk, i.e., the q^{th} quantile of the loss distribution), and the liquidation rate sharply changes around $\operatorname{VaR}_q[C_\infty]$. This behavior is related to an alternative representation of CVaR (Rockafellar and Uryasev, 2002): $\operatorname{CVaR}_q[C_\infty] = \max_{c \in \mathbb{R}} \{c + \frac{1}{q} \mathbb{E} [(C_\infty - c)^+]\}$, where the maximizer c^* is in fact $\operatorname{VaR}_q[C_\infty]$, and $\operatorname{CVaR}_q[C_\infty]$ only concerns the cases where $C_\infty > c^*$.

To better illustrate, suppose that the trader is currently left with a small amount of position to liquidate. If $C_t < c^*$, the trader is willing to pay a large transaction cost (up to $c^* - C_t$) to complete the execution as soon as possible, thereby making sure that the total loss C_{∞} won't exceed the threshold c^* . If $C_t > c^*$, the trader may believe that the total loss C_{∞} will inevitably exceed the threshold c^* , and then tries to minimize the expected future cost by slowing down the liquidation. One can make a connection with aggressiveness-in-the-money, since having $C_t < c^*$ implies that the price has moved in a favorable direction.

We believe that the threshold value is a function of remaining position size such that it increases as the position size decreases and converges to $\operatorname{VaR}_q[C_\infty]$ as the position vanishes. Moreover, the change in the aggressiveness around the threshold also depends on the remaining position size. To formalize this behavior, one may adopt an alternative formulation of the problem with an extra state variable representing the cost realized so far (i.e., a Markov policy defined on the augmented state space (X_t, C_t)), which is in fact the approach suggested by Bäuerle and Ott (2011) for a general class of control problems with a CVaR objective. This might be a topic of future research.

5. Cost Analysis: Adaptive vs. Deterministic Strategy

In this section, we provide a comparison between the optimal adaptive strategy derived in §4 and the (optimized) deterministic schedules under which the liquidation is executed according to a deterministic schedule committed at the beginning of the liquidation process.

5.1. Optimized Deterministic Schedules

First observe that any deterministic schedule will yield a normally distributed implementation shortfall; i.e., given a deterministic schedule π , the total implementation shortfall $C_{\infty}^{x,\pi} = \int_0^\infty \frac{\eta}{2} \pi_t^2 dt - \int_0^\infty \sigma X_t dW_t$ follows a normal distribution with mean and variance given by

$$\mathbb{E}[C_{\infty}^{x,\pi}] = \int_0^\infty \frac{\eta}{2} \pi_t^2 dt, \qquad \operatorname{Var}[C_{\infty}^{x,\pi}] = \int_0^\infty \sigma^2 X_t^2 dt.$$

To understand the performance of deterministic schedules, therefore, it suffices consider the CVaR value of a normal distribution.

Given a normal random variable $C \sim \mathcal{N}(\mu, \sigma^2)$, it is easy to verify that

$$CVaR_q[C] = \mu + \sigma \frac{\kappa(q)}{q}, \qquad (18)$$

where

$$\kappa(q) \triangleq \phi(\Phi^{-1}(1-q)),\tag{19}$$

and $\phi(\cdot)$ and $\Phi(\cdot)$ are the p.d.f. and the c.d.f. of a standard normal distribution, respectively. Thus,

$$\operatorname{CVaR}_{q}[C_{\infty}^{x,\pi}] = \int_{0}^{\infty} \frac{\eta}{2} \pi_{t}^{2} dt + \frac{\kappa(q)}{q} \sqrt{\int_{0}^{\infty} \sigma^{2} X_{t}^{2} dt},$$

for any deterministic schedule π .

We first focus on the set of all deterministic schedules and find the optimal one that minimizes the CVaR cost. The next proposition shows that such an optimal schedule has the form of an exponential schedule under which the trader's position decays exponentially over time (i.e., the liquidation rate is proportional to the current position size).

Proposition 6 (Optimized deterministic schedule). Given an initial position $x \in \mathbb{R}$ and a target quantile $q \in (0, 1)$, the optimal deterministic schedule is given by an exponential schedule $X_t =$

 $X_0 e^{-t/\tau^*}$ where the optimal time constant τ^* is given by

$$\tau^{\star} = \left(\frac{\eta x q}{\sqrt{2}\sigma\kappa(q)}\right)^{\frac{2}{3}}.$$

Let EXP be this optimized exponential schedule. Then, its performance $V^{\text{EXP}}(x,q)$ is given by

$$V^{EXP}(x,q) \triangleq \text{S-CVaR}_{q} \left[C_{\infty}^{x,EXP} \right] = \frac{3}{2^{\frac{5}{3}}} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times |x|^{\frac{4}{3}} \times q^{\frac{1}{3}} (\kappa(q))^{\frac{2}{3}}.$$
 (20)

While the proof is provided in Appendix A, we remark that the optimality of the exponential schedule can be directly inferred from the result of Almgren and Chriss (2000): it was shown that, given a finite time-horizon of length T, a mean-variance optimization results in a trajectory $X_t = \frac{\sinh(\kappa(T-t))}{\sinh(\kappa T)}X_0$ for some constant $\kappa > 0$. As $T \nearrow \infty$, we can observe that the optimized trajectory converges to an exponential schedule $X_t = e^{-\kappa t}X_0$.

We next examine the volume-weighted average price (VWAP) schedules, under which the trader liquidates the asset at a constant rate until completion so that the trader's position decreases linearly over time.¹⁰ The next proposition identifies the optimized VWAP schedule (see Appendix A for the proof).

Proposition 7 (Optimized VWAP schedule). Given an initial position $x \in \mathbb{R}$ and a target quantile $q \in (0,1)$, the best VWAP schedule is given by $X_t = X_0 \left(1 - \frac{t}{T^*}\right)^+$, where the optimal execution period T^* is as follows:

$$T^{\star} = \left(\frac{\sqrt{3}\eta xq}{\sigma\kappa(q)}\right)^{\frac{2}{3}}.$$

Let VWAP be this optimized VWAP schedule. Then, its performance $V^{VWAP}(x,q)$ is given by

$$V^{VWAP}(x,q) \triangleq \text{S-CVaR}_{q} \left[C_{\infty}^{x,VWAP} \right] = \frac{3^{\frac{2}{3}}}{2} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times |x|^{\frac{4}{3}} \times q^{\frac{1}{3}} (\kappa(q))^{\frac{2}{3}}.$$
 (21)

5.2. Cost Analysis

We now compare three liquidation strategies: the optimal adaptive strategy (OPT) derived in §4, the optimized deterministic strategy (EXP), and the optimized VWAP strategy (VWAP). We have derived closed-form expressions (11), (20), and (21) that represent their S-CVaR performance V, V^{EXP} , and V^{VWAP} , respectively. Given that $V(x,q) \leq V^{\text{EXP}}(x,q) \leq V^{\text{VWAP}}(x,q)$ by their definitions,

 $^{^{10}}$ A VWAP schedule typically refers to a liquidation schedule that is proportional to the average market volume profile, aiming to make its average transaction price close to the market volume-weighted average price. As we assume a time-stationary market in this paper, the VWAP schedule is equivalent to a constant liquidation rate schedule (also known as a TWAP schedule).

we particularly consider the following ratios that are useful for pairwise comparison:

$$\Upsilon_{\text{OPT}}^{\text{EXP}} \triangleq \frac{V^{\text{EXP}}(x,q)}{V(x,q)} - 1 = \frac{(3/2)^{\frac{1}{3}}q^{\frac{1}{3}}(\kappa(q))^{\frac{2}{3}}}{\varphi(q)} - 1,$$

$$\Upsilon_{\text{OPT}}^{\text{VWAP}} \triangleq \frac{V^{\text{VWAP}}(x,q)}{V(x,q)} - 1 = \frac{2^{\frac{1}{3}}q^{\frac{1}{3}}(\kappa(q))^{\frac{2}{3}}}{\varphi(q)} - 1,$$

$$\Upsilon_{\text{EXP}}^{\text{VWAP}} \triangleq \frac{V^{\text{VWAP}}(x,q)}{V^{\text{EXP}}(x,q)} - 1 = (4/3)^{\frac{1}{3}} - 1 \approx 10\%.$$
(22)

Note that these ratios do not change even if we compare CVaR performance instead of S-CVaR performance since S-CVaR is merely a scaled version of CVaR (see Remark 1(i)).



Figure 4: The pairwise comparison among three liquidation strategies – the optimal adaptive strategy, the optimal deterministic strategy (EXP), and the optimized VWAP strategy – measured with the ratios $\Upsilon_{\text{OPT}}^{\text{EXP}}$, $\Upsilon_{\text{OPT}}^{\text{WAP}}$, and $\Upsilon_{\text{EXP}}^{\text{WAP}}$ defined in (22). These ratios represent the relative percentage reductions in CVaR cost, and depend on the target quantile q only. Each curve shows the value of each ratio when q ranges from zero to one.

These ratios can be expressed in closed form, and we observe the following. First, all the ratios depend only on the quantile q, but not on the other problem parameters such as the quantity to liquidate x, the price volatility σ , and the market impact factor η . Figure 4 plots these ratios as functions of q. Second, the optimal deterministic schedule always outperforms to the best VWAP schedule by $\approx 10.0\%$, irrespective of the value of q. Finally, in a moderate range of q, from 0.2 to 0.8, the adaptive strategy outperforms the optimal deterministic strategy by 5% to 15%, and outperforms the optimized VWAP strategy by 15% to 25%. This gap increases as q increases (i.e., becomes more risk-neutral). More specifically, it explodes as q approaches one, and vanishes as q approaches zero.¹¹ Note that most traders using optimal execution algorithms are large investors trading a small portion of their overall portfolio over a short time horizon. Thus, from the perspective of optimal decision making, their utility functions are nearly linear, hence the

¹¹The absolute performance (i.e., the CVaR cost) of all three policies converges to zero as $q \nearrow 1$, and diverges to infinity as $q \searrow 0$.

nearly risk neutral regime ($q \approx 1$), where the relative benefits of dynamic trading are greatest, is also the most practically relevant regime. See also §6.2 for a more illustrative comparison between the adaptive strategy and the deterministic strategies.

6. Numerical Simulations

Throughout this section, we denote the optimal adaptive strategy by OPT, the optimized deterministic schedule by EXP, and the optimized VWAP schedule by VWAP.

6.1. Illustration of Optimal Adaptive Strategy

We consider a situation where the trader wants to liquidate x = 1 unit of an asset given the volatility $\sigma = 5.06$ basis points per minute (1% per day), the market impact factor $\eta = 1.56 \times 10^3$ (liquidating one unit over a day at a constant rate incurs 2 basis points loss in average), and the target quantile q varying from 0.10 to 0.99.

We simulate the optimal policy OPT as follows. We first discretize the time horizon into subintervals of equal length $\Delta t = 10^{-4}$, and generate a sample path of the standard Brownian motion $W_0, W_{\Delta t}, W_{2\Delta t}, \ldots$ Starting from $X_0 = x$ and $Q_0 = q$, at each time $t = 0, \Delta t, 2\Delta t, \ldots$, we compute the liquidation rate $\pi_t = f^*(X_t, Q_t)$ and the quantile diffusion rate $\gamma_t = g^*(X_t, Q_t)$ and then update the position size $X_{t+\Delta t} = X_t - \pi_t \Delta t$ and the quantile $Q_{t+\Delta t} = Q_t + \gamma_t \Delta W_t$ accordingly, where $\Delta W_t \triangleq W_{t+\Delta t} - W_t$. The expressions for f^* and g^* are given in (15) and (16), and the value of $\varphi(q)$ can be computed using linear interpolation based on its parametric representation derived in Proposition 5. In order to prevent numerical instability, we keep the value of quantile process Q_t between ϵ and $1 - \epsilon$ via truncation (we take $\epsilon = 10^{-5}$); i.e., if $Q_t < \epsilon$ or $Q_t > 1 - \epsilon$, it is set to ϵ or $1 - \epsilon$, respectively. This procedure is repeated until the remaining position size X_t becomes smaller than 10^{-2} .

Figure 5 illustrates the sample paths of the price process σW_t , the position process X_t , and the quantile process Q_t under OPT for different values of target quantile $q \in \{0.1, 0.2, \ldots, 0.99\}$ in the following two scenarios: when the price moves in an adverse direction (left), and when the price moves in a favorable direction (right). From these results, we confirm the behaviors of the optimal strategy characterized in §4.2. In every case, the position monotonically decreases over time; i.e., the optimal policy keeps trading in one direction. Also observe that the policy liquidates the position more aggressively as we take a smaller value for q (i.e., as the policy becomes more risk-averse). In a comparison between two scenarios (left vs. right), we observe "aggressiveness-in-the-money"; i.e., the policy trades more aggressively when the price moves in a favorable direction (right). This behavior can also be observed within each sample path: during the execution process, the quantile process Q_t decreases when the price moves upward and the policy trades more aggressively. In addition, the quantile process Q_t converges to either zero or one, indicating whether the realized price process is among the worst q-fraction of the scenarios. While not reported here, we observe



that Q_t converges to one in the q-fraction of simulations and converges to zero in the other (1-q)-fraction of simulations (recall that Q_t is a martingale starting at q).

Figure 5: Illustration of the optimal adaptive liquidation processes with different values for target quantile $q \in \{0.1, 0.2, 0.5, 0.8, 0.9\}$ in two scenarios: when the price moves in an adverse direction (left), and when the price moves in a favorable direction (right). The plots in the top row show the realized price process over time, the plots in the middle row show the position processes X_t , and the plots in the bottom row show the quantile processes Q_t associated with these strategies.

6.2. Comparison with Deterministic Strategies

We provide the detailed simulation results of the optimal adaptive strategy (OPT) in a comparison with those of the deterministic strategies (EXP, VWAP) introduced in §5.1.

Figure 6 illustrates the position process trajectories under these three strategies with target quantile q = 0.5 in the two scenarios as in Figure 5. One can immediately observe that the deterministic strategies are not adaptive to the price changes. The optimal adaptive strategy OPT liquidates at a similar rate to the exponential schedule EXP during the initial periods, but it deviates as soon as it adjusts its aggressiveness adaptively to the price changes. In particular, it slows down when the price moves in the adverse direction (i.e., the quantile process Q_t moves toward one). Figure 7 shows the average position trajectories under these three strategies given the target quantile $q \in \{0.1, 0.5, 0.8\}$, aggregated across 100,000 runs of simulations. Similarly to the above, we observe that OPT trades more aggressively than EXP when the target quantile is small and less aggressively when the target quantile is large.



Figure 6: Illustration of liquidation processes under the optimal adaptive strategy (OPT, red solid lines), the optimized deterministic schedule (EXP, blue dashed lines), and the optimized VWAP schedule (VWAP, green dashed lines) with the target quantile q = 0.5, and in two scenarios: when the price moves in an adverse direction (left), and when the price moves in a favorable direction (right).



Figure 7: Average liquidation processes under the optimal adaptive strategy (OPT, red), the optimized deterministic schedule (EXP, blue), and the optimized VWAP schedule (VWAP, green) given the target quantile q = 0.1 (left), q = 0.5 (middle), and q = 0.8 (right). Each curve reports the average time that the position process under each strategy falls behind a certain level, i.e., $\{(\mathbb{E}[\min_t \{X_t^{\pi} \leq y\}], y)\}_{y \in \{0.9, 0.8, \dots, 0.05\}}$.

Figure 8 shows the implementation shortfall distributions (i.e., the histograms of C_{∞}) resulting from OPT (top) and EXP (bottom) with different values of the target quantile $q \in \{0.1, 0.5, 0.8\}$. These histograms are obtained from 100,000 simulation trials, where all the strategies see the same price process realization per simulation. The resulting distributions are visually very different: EXP yields a normal distribution whereas OPT yields a distribution that has a sharp peak at the q^{th} quantile. Such a sharp peak can be explained by the threshold behavior of the optimal adaptive strategy, discussed at the end of §4.2. Figure 9 visualizes these distributions for a wider range of target quantiles $q \in \{0.1, 0.2, \dots, 0.99\}$. We observe that the implementation shortfall distribution induced by OPT is more concentrated than the ones induced by EXP and VWAP when the target quantile q is small, and it is the opposite when the target quantile q is large (roughly speaking, it has a longer right tail, visually similar to an exponential distribution).

Table 1 reports in detail the summary statistics of those implementation shortfall distributions. We first confirm that the simulation results are consistent with our theoretic predictions, i.e., the measured CVaR values match with the values calculated from the expressions (11), (20), and (21). We also see that the optimal adaptive strategy does not make an improvement over the deterministic strategies in terms of the average nor the variance as it specifically targets to minimize the CVaR value at a given target quantile. But, it significantly improves the median value particularly in the risk-neutral regime (when $q \approx 1$). See also a discussion on Figure 10 below.



Figure 8: The distributions of the implementation shortfall (i.e., the histogram of C_{∞}) incurred by the optimal adaptive liquidation strategy (top) and the optimized deterministic schedule (bottom) given the target quantile $q \in \{0.1, 0.5, 0.8\}$ (left, middle, and right, respectively). In each plot, the solid vertical line represents the q-quantile (i.e., the value-at-risk $\operatorname{VaR}_q[C_{\infty}]$), and the dashed vertical line represents the tail average beyond the q-quantile (i.e., the conditional value-at-risk $\operatorname{CVaR}_q[C_{\infty}]$). The highlighted area represents the worst q-fraction of the outcomes whose average corresponds to $\operatorname{CVaR}_q[C_{\infty}]$. These are obtained from 100,000 runs of simulations.



Figure 9: Alternative visualization of the implementation shortfall distributions induced by the optimal adaptive strategy (OPT, red), the optimized deterministic schedule (EXP, blue), and the optimized VWAP schedule (VWAP, green) given the target quantile $q \in \{0.1, 0.2, \ldots, 0.99\}$. Each distribution is represented with a colored bar with seven segments whose edges are $0.1, 0.2, \ldots, 0.99$. Each distribution is brightest segment on the top represents the range of the implementation shortfall between 0.1-quantile and 0.2-quantile, and the darkest segment represents the range of implementation shortfall between 0.4-quantile and 0.6-quantile. The diamond-shaped dots report the CVaR values of these distributions.

Target quantile	Policy	CVaR (theo.)	VaR	Average	Median	Std. dev.	50%time	95%time
q = 0.10	OPT	45.37 (44.96)	36.42	28.61	29.11	10.34	14	71
	EXP	47.24 (47.02)	38.75	15.79	15.77	17.86	17	75
	VWAP	52.00 (51.75)	42.61	17.37	17.30	19.65	23	43
q = 0.20	OPT	38.55(38.26)	27.40	23.95	23.62	12.93	18	145
	EXP	40.62(40.44)	29.72	13.60	13.57	19.26	20	87
	VWAP	44.69(44.51)	32.74	14.96	14.88	21.18	26	50
q = 0.30	OPT	33.62(33.40)	20.28	20.33	18.52	16.29	22	241
	EXP	$35.80 \ (35.66)$	22.76	12.02	11.99	20.51	23	98
	VWAP	39.41 (39.25)	24.98	13.21	13.15	22.56	30	57
q = 0.40	OPT	29.44(29.28)	13.67	17.20	13.02	20.23	27	363
	EXP	31.72(31.58)	16.22	10.66	10.66	21.80	26	111
	VWAP	34.91(34.75)	17.76	11.72	11.64	23.98	34	64
q = 0.50	OPT	25.65(25.48)	6.95	14.38	6.95	24.86	34	518
	EXP	27.96(27.80)	9.39	9.41	9.39	23.24	29	126
	VWAP	30.74(30.60)	10.30	10.34	10.30	25.56	38	73
q = 0.60	OPT	21.95(21.79)	-0.42	11.72	0.40	30.46	44	722
	EXP	24.27(24.10)	1.79	8.19	8.20	24.97	34	145
	VWAP	26.67(26.52)	1.94	8.98	9.00	27.46	44	84
q = 0.70	OPT	18.17(18.01)	-9.23	9.17	-6.21	37.56	60	998
	EXP	20.44(20.27)	-7.46	6.92	6.95	27.24	40	173
	VWAP	22.48(22.31)	-8.21	7.59	7.62	29.97	52	100
q = 0.80	OPT	14.07(13.92)	-21.12	6.62	-13.14	47.48	90	1404
	EXP	$16.24\ (16.05)$	-20.25	5.53	5.51	30.62	51	218
	VWAP	17.89(17.66)	-22.28	6.09	6.07	33.71	66	126
q = 0.90	OPT	9.16(9.03)	-41.76	3.92	-21.65	64.62	169	2125
	EXP	11.08(10.87)	-43.93	3.83	3.78	37.22	75	323
	VWAP	12.22(11.96)	-48.20	4.23	4.20	41.02	98	186
q = 0.95	OPT	5.92(5.85)	-64.35	2.35	-28.15	82.52	294	2874
	\mathbf{EXP}	7.59(7.35)	-71.53	2.68	2.59	45.23	110	477
	VWAP	8.34 (8.09)	-78.99	2.93	2.93	49.82	145	275
q = 0.99	OPT	1.94(2.09)	-132.60	0.59	-37.62	129.72	873	4740
	EXP	3.16(2.90)	-167.00	1.20	1.12	71.96	279	1207
	VWAP	3.47(3.19)	-183.10	1.33	1.49	79.16	366	696

Table 1: Summary statistics of implementation shortfall C_{∞} incurred by the optimal adaptive strategy (OPT), the optimized deterministic schedule (EXP), and the optimized VWAP schedule (VWAP), given $x = 1, \sigma = 5.06, \eta = 1.56 \times 10^3$, and the target quantile $q \in \{0.1, 0.2, \ldots, 0.99\}$. For each combination of a policy and a target quantile, it reports the CVaR value, the VaR value, the average, the median, and the standard deviation of implementation shortfall C_{∞} , represented in basis points. It additionally reports the average time to complete 50% of the execution and the average time to complete 95% of the execution, represented in minutes. These statistics are measured from 100,000 runs of simulations. The numbers in parentheses in the third column report the theoretically predicted CVaR values, computed with the expressions (11), (20), and (21).

Figure 10 demonstrates some additional advantages of the adaptive strategy other than minimizing the CVaR value. Here, the target quantile q is considered as a control parameter to the strategies on the behalf of a trader who may not be particularly interested in minimizing the CVaR value, and we compare three families of strategies (OPT, EXP, and VWAP with the target quantile ranging from 0.1 to 0.99) in terms of the mean, the median, and the tail probability of the resulting implementation shortfall distributions. It is shown in the left plot that, by implementing the adaptive strategy OPT with a suitably chosen target quantile, it can achieve a smaller median value than any of deterministic strategies that yield the same average loss value. Similarly, it is shown in the right plot that a smaller tail probability can be achieved by OPT for a given target threshold level: for example, when wanting to avoid the event that the implementation shortfall exceeds 25 basis points, the trader can implement OPT with the target quantile q = 0.4 so that such event takes place with probability 12.9%, whereas the probability is 25.3% under the best deterministic schedule (EXP with q = 0.6).



Figure 10: Plots that characterize three families of implementation shortfall distributions produced by the optimal adaptive strategy (OPT, red), the optimized deterministic schedule (EXP, blue), and the optimized VWAP schedule (VWAP, green). The left plot shows the average and the median loss values that can be obtained by a strategy (either OPT, EXP, or VWAP) when q varies, i.e., $\{(\mathbb{E}[C_{\infty}^{\pi_q}], \operatorname{Median}[C_{\infty}^{\pi_q}])\}_{q\in\{0,1,\dots,0.99\}}$ where π_q denotes the strategy with a target quantile q. The right plot reports the pointwise minimum of the complementary cumulative distribution functions induced by a strategy with the different values of target quantiles, i.e., $\min_{q\in\{0,1,\dots,0.99\}} \mathbb{P}[C_{\infty}^{\pi_q} > c]$, the lowest tail probability that the implementation shortfall exceeds a given threshold level c when the target quantile q for the strategy can be arbitrarily chosen.

References

Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. 2000.

- Philippe Artzner, Freddy Delbaen, Jean-Mark Eber, and David Heath. Coherent measures of risk. Mathematical Finance, 9(3):203-228, 1999.
- Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, David Heath, and Hyejin Ku. Coherent multiperiod risk adjusted values and Bellman's principle. Annals of Operations Research, 152:5–22, 2007.

- Julio Backhoff-Veraguas and Ludovic Tangpi. On the dynamic representation of some time-inconsistent risk measures in a Brownian filtration. *Mathematics and Financial Economics*, 14:433–460, 2020.
- Viorel Barbu and Teodor Precupanu. Convexity and Optimization in Banach Spaces. Springer, 4th edition, 2012.
- Nicole Bäuerle and Jonathan Ott. Markov decision processes with average-value-at-risk criteria. Mathematical Methods of Operational Research, 74(4):361–379, 2011.
- Margaret P. Chapman, Jonathan P. Lacotte, Kevin M. Smith, Insoon Yang, Yuxi Han, Marco Pavone, and Claire J. Tomlin. Risk-sensitive safety specifications for stochastic system using conditional value-at-risk. 2018.
- Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-sensitive and robust decision-making: a CVaR optimization approach. Advances in Neural Information Processing Systems 28 (NIPS 2015), 2015.
- Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. Journal of Machine Learning Research, 18:1–51, 2018.
- R. Courant and D. Hilbert. *Methods of Mathematical Physics*, volume I. New York and London (Interscience Publishers), 1953.
- Rick Durrett. Probability: Theory and Examples. Cambridge University Press, 4th edition, 2010.
- Peter A. Forsyth. A Hamilton–Jacobi–Bellman approach to optimal trade execution. Applied Numerical Mathematics, 61:241–265, 2011.
- Peter A. Forsyth, Shannon Kennedy, Shu Tong Tse, and Heath Windcliff. Optimal trade execution: A mean quadratic variation approach. *Journal of Economic Dynamics & Control*, 36:1971–1991, 2012.
- Jim Gatheral and Alexander Schied. Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. *International Journal of Theoretical and Applied Finance*, 14(3):353–368, 2011.
- Paul Glasserman and Xingbo Xu. Robust portfolio control with stochastic factor dynamics. Operations Research, 61(4):874–893, 2013.
- Yonghui Huang and Xianping Guo. Minimum average value-at-risk for finite horizon semi-Markov decision processes in continuous time. SIAM Journal on Optimization, 26(1):1–28, 2016.
- Robert Kissell and Roberto Malamut. Understanding the profit and loss distribution of trading algorithms. In B. R. Bruce, editor, *Algorithmic Trading*, pages 41–49. Institutional Investor, 2005.
- Xiaocheng Li, Huaiyang Zhong, and Margaret L. Brandeau. Quantile Markov decision processes. 2020.
- Qihang Lin, Xi Chen, and Javier Peña. A trade execution model under a composite dynamic coherent risk measure. Operations Research Letters, 43:52–58, 2015.
- Julian Lorenz and Robert Almgren. Adaptive arrival price. In B. R. Bruce, editor, *Algorithmic Trading III*, pages 59–66. Institutional Investor, 2007.
- Julian Lorenz and Robert Almgren. Mean-variance optimal adaptive execution. Applied Mathematical Finance, 18(5):395–422, 2011.
- Christopher W. Miller and Insoon Yang. Optimal control of conditional value-at-risk in continuous time. SIAM Journal on Control and Optimization, 55(2):856–884, 2017.
- Georg Ch. Pflug and Alois Pichler. Time-inconsistent multistage stochastic programs: Martingale bounds. European Journal of Operational Research, 249:155–163, 2016.
- Andrei D. Polyanin and Valentin F. Zaitsev. Handbook of Exact Solutions for Ordinary Differential Equations. Chapman & Hall/CRC Press, 2003.

Philip Protter. Stochastic Integration and Differential Equations. Springer, 2015.

- R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. Journal of Banking & Finance, 26:1443–1471, 2002.
- Alexander Schied and Torsten Schöneborn. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, 13(2):181–204, 2009.
- Alexander Shapiro. On a time consistency concept in risk averse multistage stochastic programming. Operations Research Letters, 37:143–147, 2009.

Maurice Sion. On general minimax theorems. Pacific Journal of Mathematics, 8(1):171–176, 1958.

Organization of appendix. The appendix is organized as follows. In Appendix A, we identify the optimal deterministic strategy and its performance for §5.1. The CVaR performance of the optimal deterministic strategy is utilized as an upper bound on the CVaR performance of the optimal adaptive strategy. In Appendix B, we provide the basic characterizations of S-CVaR measure introduced in §2. In Appendix C, we provide the preliminary characterizations of the value function, and by applying Sion's minimax theorem, we prove Theorem 2 and Theorem 3 stated in §3. The main challenge here is to verify the conditions of Sion's minimax theorem. In Appendix D, we provide proofs for §4. We first state and prove Theorem 7 from which Theorem 4 and Theorem 6 follow almost immediately. Proposition 4, Proposition 5 and Theorem 5 are proven separately.

A. Optimal Deterministic Schedules

Lemma 2. For any a, b > 0,

$$\min_{x \in \mathbb{R}_+} \left\{ \frac{a}{x} + b\sqrt{x} \right\} = \frac{a}{x} + b\sqrt{x} \Big|_{x = \left(\frac{2a}{b}\right)^{\frac{2}{3}}} = \frac{3a^{\frac{1}{3}}b^{\frac{2}{3}}}{2^{\frac{2}{3}}}.$$

Proof. Let $f(x) \triangleq \frac{a}{x} + b\sqrt{x}$. Since $f'(x) = -\frac{a}{x^2} + \frac{b}{2\sqrt{x}}$, the equation f'(x) = 0 has a unique solution at $x = \left(\frac{2a}{b}\right)^{\frac{2}{3}}$.

Proof of Proposition 6. We prove the optimality of exponential schedules and identify the optimal decaying rate.

First we consider a mean-variance optimization problem:

$$\underset{\pi \in \mathcal{D}(x)}{\text{minimize}} \quad \mathbb{E}[C_{\infty}^{x,\pi}] + \lambda \text{Var}[C_{\infty}^{x,\pi}],$$

where $\mathcal{D}(x)$ is the set of all deterministic policies and $\lambda \in (0, \infty)$ is a penalty for variance term. Applying (18), this is equivalent to an optimization over the deterministic trajectories of $(X_t)_{t\geq 0}$:

$$\inf_{X:X_0=x} \int_{t=0}^{\infty} \left(\frac{\eta}{2} \dot{X}_t^2 + \lambda \sigma^2 X_t^2\right) dt$$

where $\dot{X}_t \triangleq dX_t/dt$. By applying standard calculus of variations arguments (Courant and Hilbert, 1953), we deduce that the optimal schedule X^* has to satisfy the Euler-Lagrange equation $2\lambda\sigma^2 X_t^* - \eta \ddot{X}_t^* = 0$, at each time t, with boundary conditions $X_0^* = x$ and $\lim_{t\to\infty} X_t^* = 0$. The solution is uniquely given by an exponential schedule

$$X_t^{\star} = x \exp\left(-t/\rho_{\lambda}\right),$$

with the decay rate $\rho_{\lambda} \triangleq \sqrt{\frac{\eta}{2\lambda\sigma^2}} \in (0,\infty)$, and such a schedule yields

$$\mathbb{E}[C_{\infty}^{x,\pi^{\star}}] = \frac{\eta x^2}{4\rho_{\lambda}}, \quad \operatorname{Var}[C_{\infty}^{x,\pi^{\star}}] = \frac{\sigma^2 x^2 \rho_{\lambda}}{2}.$$
(23)

In other words, the efficient frontier of the range of mean and variance achievable by deterministic schedules, $\{(\mathbb{E}[C_{\infty}^{x,\pi}], \operatorname{Var}[C_{\infty}^{x,\pi}])\}_{\pi \in \mathcal{D}(x)}$, is given by $\{(\frac{\eta x^2}{4\rho}, \frac{\sigma^2 x^2 \rho}{2})\}_{\rho \in (0,\infty)}$ and is attained by exponential schedules.

Let us now consider the achievable range of mean and standard deviation, $\{(\mathbb{E}[C_{\infty}^{x,\pi}], \sqrt{\operatorname{Var}[C_{\infty}^{x,\pi}]})\}_{\pi \in \mathcal{D}(x)}$. Observe that its efficient frontier is still characterized by $\{(\frac{\eta x^2}{4\rho}, \sqrt{\frac{\sigma^2 x^2 \rho}{2}})\}_{\rho \in (0,\infty)}$. Therefore, the optimal solution of the following mean-standard deviation optimization problem

$$\underset{\pi \in \mathcal{D}(x)}{\text{minimize}} \quad \mathbb{E}[C_{\infty}^{x,\pi}] + \theta \sqrt{\text{Var}[C_{\infty}^{x,\pi}]},$$

is also given by an exponential schedule, for any given $\theta \in (0, \infty)$.

Finally, observe that for any deterministic schedule $\pi \in \mathcal{D}(x)$ the resulting cost $C_{\infty}^{x,\pi}$ is normally distributed, thus $\operatorname{CVaR}_q[C_{\infty}^{x,\pi}]$ is minimized by an exponential schedule. Applying (18), for any $\pi \in \mathcal{D}(x)$ and $q \in (0, 1)$, we have

$$\operatorname{CVaR}_{q}\left[C_{\infty}^{x,\pi}\right] = \mathbb{E}\left[C_{\infty}^{x,\pi}\right] + \frac{\kappa(q)}{q} \sqrt{\operatorname{Var}\left[C_{\infty}^{x,\pi}\right]},$$

Using (23), the optimal time constant can be determined as

$$\tau^{\star} \triangleq \underset{\rho}{\operatorname{argmin}} \left\{ \frac{\eta x^2}{4\rho} + \frac{\kappa(q)}{q} \sqrt{\frac{\sigma^2 x^2 \rho}{2}} \right\} = \left(\frac{\eta x q}{\sqrt{2} \sigma \kappa(q)} \right)^{\frac{2}{3}}$$

This concludes the proof.

Proof of Proposition 7. With some calculation, it can be easily shown that a VWAP schedule $X_t = x \left(1 - \frac{t}{T}\right)^+$ yields

$$\mathbb{E}[C_{\infty}^{x,\pi}] = \frac{\eta x^2}{2T}, \quad \operatorname{Var}[C_{\infty}^{x,\pi}] = \frac{\sigma^2 x^2 T}{3}.$$

Therefore, the optimal execution horizon T^* is given by

$$T^{\star} \triangleq \underset{T}{\operatorname{argmin}} \left\{ \frac{\eta x^2}{2T} + \frac{\kappa(q)}{q} \sqrt{\frac{\sigma^2 x^2 T}{3}} \right\} = \left(\frac{\sqrt{3}\eta x q}{\sigma \kappa(q)} \right)^{\frac{2}{3}}.$$

We state the following lemma that identifies the boundary values of S-CVaR_q[$C_{\infty}^{x,\text{EXP}}$], which is useful to characterize the optimal value function.

Lemma 3. The function $\kappa(q)$ given in (19) satisfies

$$\sup_{q \in [0,1]} \kappa(q) < \infty, \quad \lim_{n \to \infty} \sqrt{n} \kappa(1/n) = \lim_{n \to \epsilon} \sqrt{n} \kappa(1-1/n) = 0.$$

Proof. Recall that $\kappa(q) \triangleq \phi\left(\Phi^{-1}(1-q)\right)$. Since $\phi(z) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2} \leq \frac{1}{\sqrt{2\pi}}$ for any $z \in \mathbb{R}$, we have $\sup_{q \in [0,1]} \kappa(q) \leq \frac{1}{\sqrt{2\pi}} < \infty$. Also note that $\kappa(q) = \kappa(1-q)$ since $\Phi^{-1}(q) = -\Phi^{-1}(1-q)$ and $\phi(z) = \phi(-z)$. Therefore, it suffices to show that $\lim_{n \to \infty} \sqrt{n\kappa(1/n)} = 0$.

We have the following tail bounds of standard normal distribution (Durrett, 2010, Theorem 1.2.3): for any z > 0,

$$(z^{-1} - z^{-3})\phi(z) \le 1 - \Phi(z) \le z^{-1}\phi(z).$$

Define $z_n \triangleq \Phi^{-1}(1-1/n) > 0$, and then we have

$$\frac{\phi(z_n)}{2z_n} \le \frac{1}{n} \le \frac{\phi(z_n)}{z_n},$$

for large enough n (such that $z_n^{-3} \leq \frac{1}{2} z_n^{-1}$) since $\lim_{n\to\infty} z_n = \infty$. Observe that $1/n \leq \phi(z_n)/z_n \leq \phi(z_n) = \exp(-z_n^2/2)/\sqrt{2\pi} \leq \exp(-z_n^2/2)$ and thus $z_n \leq \sqrt{2\log n}$. We further deduce that, since $\kappa(1/n) = \phi(\Phi^{-1}(1-1/n)) = \phi(z_n)$,

$$\sqrt{n}\kappa(1/n) = \sqrt{n}\phi(z_n) = 2\sqrt{n}z_n \times \frac{\phi(z_n)}{2z_n} \le 2\sqrt{n}z_n \times \frac{1}{n} \le \frac{2z_n}{\sqrt{n}} \le 2\sqrt{\frac{2\log n}{n}}.$$

Therefore, $\lim_{n\to\infty}\sqrt{n}\kappa(1/n)=0.$

B. Preliminary Characterizations of S-CVaR

We begin with a technical lemma:

Lemma 4. The risk envelope $\mathcal{Q}(q)$ is a non-empty, convex, and weak-* compact subset of \mathcal{L}^{∞} .

Proof. It is non-empty because $Q(\omega) = q$ is always feasible.

Consider $Q_1, Q_2 \in \mathcal{Q}(q)$ and $Q_{\lambda} \triangleq \lambda Q_1 + (1 - \lambda)Q_2$ for some $\lambda \in [0, 1]$. Since $Q_1(\omega), Q_2(\omega) \in [0, 1]$, we have $Q_{\lambda}(\omega) \in [0, 1]$, and by the linearity of expectation, $\mathbb{E}[Q_{\lambda}] = \lambda \mathbb{E}[Q_1] + (1 - \lambda)\mathbb{E}[Q_2] = q$. Therefore, $Q_{\lambda} \in \mathcal{Q}(q)$ and thus $\mathcal{Q}(q)$ is convex.

Finally, note that $\mathcal{Q}(q)$ is a closed subset of the unit ball in $\mathcal{L}^{\infty}(\Omega, \mathcal{F}, \mathbb{P})$. Given that \mathcal{L}^{∞} is the dual space of \mathcal{L}^1 , it is weak-* compact by Banach-Alaoglu theorem.

Proof of Proposition 1. *Proof of claim (i) and (v).* Claim (i) immediately follows from our definition of CVaR. Claim (v) follows from the following identity (Shapiro, 2009, Theorem 6.2):

$$CVaR_{q}[C] = \mathbb{E}\left[C \left| C \ge F_{C}^{-1}(1-q)\right] = \frac{\mathbb{E}\left[C \mathbb{I}_{\{C \ge F_{C}^{-1}(1-q)\}}\right]}{\mathbb{P}[C \ge F_{C}^{-1}(1-q)]} = \frac{\mathbb{E}\left[C \mathbb{I}_{\{C \ge F_{C}^{-1}(1-q)\}}\right]}{q}.$$

Proof of claim (ii). When q = 0, the risk envelope $\mathcal{Q}(q)$ has a single element $Q(\omega) = 0$, and hence, $\sup_{Q \in \mathcal{Q}(0)} \mathbb{E}[CQ] = 0$. When q = 1, the risk envelope $\mathcal{Q}(q)$ also has a single element $Q(\omega) = 1$, and hence, $\sup_{Q \in \mathcal{Q}(0)} \mathbb{E}[CQ] = \mathbb{E}C$.

Proof of claim (iii). For any $Q \in \mathcal{Q}(q)$, we have $|\mathbb{E}[CQ]| \leq \mathbb{E}[|CQ|] \leq \mathbb{E}|C|$ since $|Q| \leq 1$ almost surely, and therefore, $|S-CVaR_q[C]| \leq \mathbb{E}|C|$. Furthermore, $\mathcal{Q}(q)$ contains $Q(\omega) = q$, and therefore, $CVaR_q[C] \geq \mathbb{E}[qC] = q\mathbb{E}C$.

Proof of claim (iv). Consider $q_1, q_2 \in [0, 1]$ and $q_{\lambda} \triangleq \lambda q_1 + (1 - \lambda)q_2$ for some $\lambda \in [0, 1]$. Let $Q_1^* \in \operatorname{argmax}_{Q \in \mathcal{Q}(q_1)} \mathbb{E}[CQ]$ and $Q_2^* \in \operatorname{argmax}_{Q \in \mathcal{Q}(q_2)} \mathbb{E}[CQ]$ (the maximum is attained since $\mathcal{Q}(q)$ is weak-* compact). Let $Q_{\lambda} \triangleq \lambda Q_1^* + (1 - \lambda)Q_2^*$. Observe that $Q_{\lambda} \in [0, 1]$ a.s. and $\mathbb{E}[Q_{\lambda}] = \lambda \mathbb{E}[Q_1^*] + (1 - \lambda)\mathbb{E}[Q_2^*] = \lambda q_1 + (1 - \lambda)q_2 = q_{\lambda}$. Therefore, $Q_{\lambda} \in \mathcal{Q}(q_{\lambda})$. Trivially, $\mathbb{E}[CQ_{\lambda}] = \mathbb{E}[\lambda C Q_1^* + (1 - \lambda)CQ_2^*] = \lambda \mathbb{E}[CQ_1^*] + (1 - \lambda)\mathbb{E}[CQ_2^*]$, and thus

$$\begin{aligned} \mathbf{S}\text{-}\mathbf{C}\mathbf{VaR}_{q_{\lambda}}[C] &= \sup_{Q \in \mathcal{Q}(q_{\lambda})} \mathbb{E}[CQ] \geq \mathbb{E}[CQ_{\lambda}] = \lambda \mathbb{E}[CQ_{1}^{*}] + (1-\lambda)\mathbb{E}[CQ_{2}^{*}] \\ &= \lambda \mathbf{S}\text{-}\mathbf{C}\mathbf{VaR}_{q_{1}}[C] + (1-\lambda) \mathbf{S}\text{-}\mathbf{C}\mathbf{VaR}_{q_{2}}[C]. \end{aligned}$$

Continuity follows from concavity since the function is bounded on its domain from (iii).

C. Proofs for §3

From now on, we characterize S-CVaR_q[·] as a mapping from \mathcal{L}^2 to \mathbb{R} . This can be done without loss generality since we have $C_{\infty}^{x,\pi} \in \mathcal{L}^2$ for any feasible policy $\pi \in \Pi(x)$. This is to utilize the fact that \mathcal{L}^2 is reflexive so that its weak-* topology coincides with its weak topology.

Lemma 5. The risk envelope Q(q) is a weakly compact subset of \mathcal{L}^2 .

Proof. As stated in the proof of Lemma 4, $\mathcal{Q}(q)$ is a closed subset of the unit ball in \mathcal{L}^{∞} , which is a subset of \mathcal{L}^2 . By Banach–Alaoglu theorem, it is weak-* compact in \mathcal{L}^2 and hence weakly compact since \mathcal{L}^2 is reflexive.

Proof of Proposition 2. Proof of claim (i). Note that $(Q_t^{q,\gamma})_{t\geq 0}$ is a continuous local martingale since it is a stochastic integral of a progressively measurable process with respect to Brownian motion (Protter, 2015, Theorem 33 in Chap. III). Since $Q_t^{q,\gamma} \in [0,1]$ for any $t \in [0,\infty)$ by the definition of $\Gamma(q)$, it is a bounded local martingale, which is indeed a martingale (Protter, 2015, Thm. 51 in Chap. I).

Proof of claim (ii). The limit $\lim_{t\to\infty} Q_t^{q,\gamma}$ exists due to martingale convergence theorem (Protter, 2015, Theorem 10 in Chap. I).

Proof of claim (iii). Recall that $(Q_t^{q,\gamma})_{t\geq 0}$ is a martingale taking values in [0,1]. Define a stopping time $\tau \triangleq \inf_{t\geq 0} \{Q_t^{q,\gamma} = 0\}$. Then, for any $\tau' \geq \tau$, we have $\mathbb{E}[Q_{\tau'}^{q,\gamma}|\mathcal{F}_{\tau}] = Q_{\tau}^{q,\gamma} = 0$ and therefore $Q_{\tau'}^{q,\gamma} = 0$ almost surely (otherwise, we would have $\mathbb{E}[Q_{\tau'}^{q,\gamma}|\mathcal{F}_{\tau}] > 0$). In words, once $(Q_t^{q,\gamma})_{t\geq 0}$ hits zero, it never escapes thereafter, and the same argument holds for the other boundary.

C.1. Proof of Theorem 2

Within this subsection, we characterize the trader's policy with its position process $(X_t)_{t\geq 0}$ rather than dealing with the liquidation rate process $(\pi_t)_{t\geq 0}$: the set of admissible policies is represented as

$$\mathcal{X}(x) \triangleq \left\{ X : \mathbb{T} \times \Omega \to \mathbb{R} \left| X \in \mathcal{P}, X_0 = x, \mathbb{E} \left[\left(\int_{t=0}^{\infty} \dot{X}_t^2 dt \right)^2 \right] < \infty, \mathbb{E} \left[\int_{t=0}^{\infty} X_t^2 dt \right] < \infty, \sup_{t \ge 0} |X_t| \le M \right\}$$

where we require that each sample path $X \in \mathcal{X}(x)$ is differentiable so that we can define $\dot{X}_t \triangleq dX_t/dt$. Then, $\Pi(x) = \{\pi = \dot{X} | X \in \mathcal{X}(x)\}$. Accordingly, we represent the loss process as

$$C_t^X \triangleq \int_{s=0}^t \frac{1}{2}\eta |\dot{X}_s|^2 ds - \int_{s=0}^t \sigma X_s dW_s,$$

so that we have $C_t^X = C_t^{x,\pi}$ if $X \in \mathcal{X}(x)$ and $\pi = \dot{X}$.

We aim to prove Theorem 2 by utilizing Sion's minimax theorem:

Lemma 6 (Sion's minimax theorem (Sion, 1958)). Let \mathcal{X} be a convex subset of a linear topological

space and \mathcal{Y} a compact convex subset of a linear topological space. If f is a real-valued function on $\mathcal{X} \times \mathcal{Y}$ with $f(x, \cdot)$ upper semicontinuous and quasi-concave on Y, for each $x \in \mathcal{X}$, and $f(\cdot, y)$ lower semicontinuous and quasi-convex on X, for each $y \in \mathcal{Y}$, then,

$$\inf_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) = \max_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} f(x, y),$$

where an optimal solution $y \in \mathcal{Y}$ must exist for each maximization.

Lemma 7. $\mathcal{X}(x)$ is a convex subset of a linear space endowed with a norm $\|\cdot\|_{\mathcal{X}}$:

$$\|X\|_{\mathcal{X}} \triangleq \sqrt{\mathbb{E}\left[\int_{t=0}^{\infty} \dot{X}_t^2 dt\right]} + \sqrt{\mathbb{E}\left[\int_{t=0}^{\infty} X_t^2 dt\right]}.$$

Proof. It can be easily verified that $\|\cdot\|_{\mathcal{X}}$ is a valid norm (as a norm of a Sobolev space). Also note that $\|X\|_{\mathcal{X}} < \infty$ for any $X \in \mathcal{X}(x)$. Now consider $X^{(1)}, X^{(2)} \in \mathcal{X}(x)$ and $X^{(\lambda)} \triangleq \lambda X^{(1)} + (1-\lambda)X^{(2)}$ for some $\lambda \in [0,1]$. Observe that (i) $X_0^{(\lambda)} = x$, (ii) $\sqrt{\int_{t=0}^{\infty} (\dot{X}_t^{(\lambda)})^2 dt} \leq \lambda \sqrt{\int_{t=0}^{\infty} (\dot{X}_t^{(1)})^2 dt} + (1-\lambda)\sqrt{\int_{t=0}^{\infty} (\dot{X}_t^{(2)})^2 dt} \in \mathcal{L}^4$, and thus $\mathbb{E}\left[\left(\int_{t=0}^{\infty} (\dot{X}_t^{(\lambda)})^2 dt\right)^2\right] < \infty$, (iii) similarly, $\mathbb{E}\left[\int_{t=0}^{\infty} (X_t^{(\lambda)})^2 dt\right] < \infty$, and (iv) $\sup_{t\geq 0} |X_t^{(\lambda)}| \leq \lambda \sup_{t\geq 0} |X_t^{(1)}| + (1-\lambda) \sup_{t\geq 0} |X_t^{(2)}| \leq M$. Therefore, $X^{(\lambda)} \in \mathcal{X}(x)$, and hence $\mathcal{X}(x)$ is a convex set.

Proof of Theorem 2. By Lemma 1, it is equivalent to show that

$$\inf_{X \in \mathcal{X}(x)} \max_{Q \in \mathcal{Q}(q)} \mathbb{E}\left[C_{\infty}^{X}Q\right] = \max_{Q \in \mathcal{Q}(q)} \inf_{X \in \mathcal{X}(x)} \mathbb{E}\left[C_{\infty}^{X}Q\right].$$

We aim to prove this by utilizing Sion's minimax theorem (Lemma 6), so it suffices to check the conditions of Lemma 6. We have shown that $\mathcal{X}(x)$ is a convex subset of a linear space endowed with a norm $\|\cdot\|_{\mathcal{X}}$ (Lemma 7), and $\mathcal{Q}(q)$ is a compact convex subset of $\mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P})$ endowed with the weak topology (Lemma 5). In the rest of the proof, we verify the followings:

- (i) The mapping $X \mapsto \mathbb{E}[C_{\infty}^{X}Q]$ is convex and continuous in norm $\|\cdot\|_{\mathcal{X}}$ for any given $Q \in \mathcal{Q}(q)$.
- (ii) The mapping $Q \mapsto \mathbb{E}[C_{\infty}^{X}Q]$ is concave and weakly continuous on $\mathcal{Q}(q)$ for any given $X \in \mathcal{X}(x)$.

Proof of claim (i). The convexity immediately follows from the fact that the mapping $X \mapsto C_{\infty}^X$ is quadratic and convex. We now focus on the continuity.

Fix $X \in \mathcal{X}(x)$ and consider a sequence $(X^{(n)} \in \mathcal{X}(x))_{n \in \mathbb{N}}$ such that $\lim_{n \to \infty} ||X - X^{(n)}||_{\mathcal{X}} = 0$. That is,

$$\lim_{n \to \infty} \mathbb{E}\left[\int_{t=0}^{\infty} \left(\dot{X}_t - \dot{X}_t^{(n)}\right)^2 dt\right] = 0, \quad \lim_{n \to \infty} \mathbb{E}\left[\int_{t=0}^{\infty} \left(X_t - X_t^{(n)}\right)^2 dt\right] = 0.$$
(24)

Let

$$A \triangleq \sqrt{\int_{t=0}^{\infty} \dot{X}_t^2 dt}, \quad A_n \triangleq \sqrt{\int_{t=0}^{\infty} (\dot{X}_t^{(n)})^2 dt}, \quad \Delta_n \triangleq \sqrt{\int_{t=0}^{\infty} \left(\dot{X}_t - \dot{X}_t^{(n)} \right)^2 dt}$$

We then have $A \in \mathcal{L}^2$, $A_n \in \mathcal{L}^2$, and $|A - A_n| \leq \Delta_n$ (triangle inequality) where $\Delta_n \xrightarrow{\mathcal{L}^2} 0$ due to the above condition (24). Further observe that

$$\begin{aligned} \left| C_{\infty}^{X}Q - C_{\infty}^{X^{(n)}}Q \right| &\leq \left| C_{\infty}^{X} - C_{\infty}^{X^{(n)}} \right| \\ &\leq \frac{\eta}{2} \left| \int_{t=0}^{\infty} \dot{X}_{t}^{2} dt - \int_{t=0}^{\infty} (\dot{X}_{t}^{(n)})^{2} dt \right| + \sigma \left| \int_{t=0}^{\infty} X_{t} dW_{t} - \int_{t=0}^{\infty} X_{t}^{(n)} dW_{t} \right| \\ &= \frac{\eta}{2} \left| A^{2} - A_{n}^{2} \right| + \sigma \left| \int_{t=0}^{\infty} \left(X_{t} - X_{t}^{(n)} \right) dW_{t} \right|, \end{aligned}$$
(25)

where the first inequality holds since $|Q| \leq 1$. For the first term of (25), we have

$$\mathbb{E}\left[\left|A_{n}^{2}-A^{2}\right|\right] = \mathbb{E}\left[\left|(A_{n}-A)^{2}+2A(A_{n}-A)\right|\right]$$

$$\leq \mathbb{E}\left[(A_{n}-A)^{2}\right]+2\mathbb{E}\left[|A|\cdot|A_{n}-A|\right]$$

$$\leq \mathbb{E}\left[(A_{n}-A)^{2}\right]+2\sqrt{\mathbb{E}\left[A^{2}\right]}\sqrt{\mathbb{E}\left[|A_{n}-A|^{2}\right]}$$

$$= \mathbb{E}\left[\Delta_{n}^{2}\right]+2\sqrt{\mathbb{E}\left[A^{2}\right]}\sqrt{\mathbb{E}\left[\Delta_{n}^{2}\right]},$$
(26)

where for (26) we apply the Cauchy-Schwartz inequality. Therefore, $\mathbb{E}\left[|A_n^2 - A^2|\right] \to 0$ since $\Delta_n \stackrel{\mathcal{L}^2}{\to} 0$, as $n \to \infty$. For the second term of (25), using the Itô isometry,

$$\mathbb{E}\left[\left(\int_{t=0}^{\infty} \left(X_t - X_t^{(n)}\right) dW_t\right)^2\right] = \mathbb{E}\left[\int_{t=0}^{\infty} \left(X_t - X_t^{(n)}\right)^2 dt\right]$$

which vanishes as $n \to \infty$ due to condition (24). Therefore,

$$\lim_{n \to \infty} \mathbb{E}\left[\left| C_{\infty}^{X} Q - C_{\infty}^{X^{(n)}} Q \right| \right] = 0,$$

which implies that $\lim_{n\to\infty} \mathbb{E}\left[C_{\infty}^{X^{(n)}}Q\right] = \mathbb{E}\left[C_{\infty}^{X}Q\right]$. This concludes the proof.

Proof of claim (ii). The concavity immediately follows from the fact that the mapping $Q \mapsto C_{\infty}^{X}Q$ is linear. Fix $Q \in \mathcal{Q}(q)$, consider a sequence $(Q_n \in \mathcal{Q}(q))_{n \in \mathbb{N}}$ converging to Q in \mathcal{L}^2 -weak topology: i.e., $\lim_{n\to\infty} \mathbb{E}[ZQ_n] = \mathbb{E}[ZQ]$ for any $Z \in \mathcal{L}^2$. Since $C_{\infty}^X \in \mathcal{L}^2$ for any $X \in \mathcal{X}(x)$, the continuity of the mapping $Q \mapsto \mathbb{E}[C_{\infty}^X Q]$ immediately follows.

C.2. Preliminary Characterizations of Value Function

We begin with a preliminary characterization of the value function (*) that will be used in later proofs in this section.

Proposition 8. The value function $V : \mathbb{X} \times [0,1] \to \mathbb{R}$ satisfies the followings:

(i)
$$0 \le V(x,q) \le \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times |x|^{\frac{4}{3}} \times q^{\frac{1}{3}} (\kappa(q))^{\frac{2}{3}}.$$

- (ii) V(x,q) is convex in x on X for each $q \in [0,1]$.
- (iii) V(x,q) is concave in q on [0,1] for each $x \in \mathbb{X}$.
- (iv) V(x,0) = V(x,1) = V(0,q) = 0 for each $x \in \mathbb{X}$ and $q \in [0,1]$.

Proof. Proof of claim (i). Note that $\mathbb{E}[C_{\infty}^{x,\pi}] \geq 0$ for any $\pi \in \Pi(x)$. By Proposition 1(iii), we have S-CVaR_q $[C_{\infty}^{x,\pi}] \geq 0$ for any $\pi \in \Pi(x)$, and hence $V(x,q) \geq 0$. The upper bound follows from Proposition 6 given that $\frac{3}{2^{5/3}} < 1$.

Proof of claim (ii). Fix $q \in [0, 1]$ and define $\varphi_{\gamma} : \mathbb{R} \to \mathbb{R}$ as follows:

$$\varphi_{\gamma}(x) \triangleq \inf_{\pi \in \Pi(x)} J(\pi, \gamma; x, q),$$

so that we have $V(x,q) = \max_{\gamma \in \Gamma(q)} \varphi_{\gamma}(x)$. Since the pointwise maximum of a set of convex functions is convex, it suffices to show that $\varphi_{\gamma}(\cdot)$ is convex for each $\gamma \in \Gamma(q)$.

Fix $\gamma \in \Gamma(q)$ and consider any $x_1, x_2 \in \mathbb{X}$. For any $\epsilon > 0$, by definition of the infimum, there exist $\pi^{1,\epsilon} \in \Pi(x_1)$ and $\pi^{2,\epsilon} \in \Pi(x_2)$ such that

$$J(\pi^{1,\epsilon},\gamma;x_1,q) \le \varphi_{\gamma}(x_1) + \epsilon, \quad J(\pi^{2,\epsilon},\gamma;x_2,q) \le \varphi_{\gamma}(x_2) + \epsilon.$$

Given some $\lambda \in [0,1]$, let $x_{\lambda} \triangleq \lambda x_1 + (1-\lambda)x_2$ and $\pi^{\lambda,\epsilon} \triangleq \lambda \pi^{1,\epsilon} + (1-\lambda)\pi^{2,\epsilon}$. Note that

$$X_t^{x_{\lambda},\pi^{\lambda,\epsilon}} = x_{\lambda} - \int_{s=0}^t \pi_s^{\lambda,\epsilon} ds = \lambda X_t^{x_1,\pi^{1,\epsilon}} + (1-\lambda) X_t^{x_2,\pi^{2,\epsilon}}.$$

Similarly to the proof of Lemma 7, it can be shown that $\mathbb{E}\left[\left(\int_{t=0}^{\infty} |\pi_t^{\lambda,\epsilon}|^2 dt\right)^2\right] < \infty$ and $\mathbb{E}\left[\int_{t=0}^{\infty} |X_t^{\lambda,\pi^{\lambda,\epsilon}}|^2 dt\right] < \infty$, and therefore $\pi^{\lambda,\epsilon} \in \Pi(x_{\lambda})$. Consequently,

$$C_{\infty}^{x_{\lambda},\pi^{\lambda,\epsilon}} = \int_{s=0}^{t} \frac{1}{2} \eta \left| \pi_{s}^{\lambda,\epsilon} \right|^{2} ds - \int_{s=0}^{t} \sigma X_{s}^{x_{\lambda},\pi^{\lambda,\epsilon}} dW_{s} \le \lambda C_{\infty}^{x_{1},\pi^{1,\epsilon}} + (1-\lambda) C_{\infty}^{x_{2},\pi^{2,\epsilon}},$$

and therefore,

$$J(\pi^{\lambda,\epsilon},\gamma;x_{\lambda},q) = \mathbb{E}\left[C_{\infty}^{x_{\lambda},\pi^{\lambda,\epsilon}}Q_{\infty}^{q,\gamma}\right] \leq \lambda \mathbb{E}\left[C_{\infty}^{x_{1},\pi^{1,\epsilon}}Q_{\infty}^{q,\gamma}\right] + (1-\lambda)\mathbb{E}\left[C_{\infty}^{x_{2},\pi^{2,\epsilon}}Q_{\infty}^{q,\gamma}\right]$$
$$= \lambda J(\pi^{1,\epsilon},\gamma;x_{1},q) + (1-\lambda)J(\pi^{2,\epsilon},\gamma;x_{2},q)$$
$$\leq \lambda \varphi_{\gamma}(x_{1}) + (1-\lambda)\varphi_{\gamma}(x_{2}) + \epsilon.$$

As a result, we have

$$\varphi_{\gamma}(\lambda x_1 + (1 - \lambda)x_2) \le J(\pi^{\lambda, \epsilon}, \gamma; x_{\lambda}, q) \le \lambda \varphi_{\gamma}(x_1) + (1 - \lambda)\varphi_{\gamma}(x_2) + \epsilon.$$

Since $x_1, x_2, \lambda, \epsilon$ were arbitrarily chosen, $\varphi_{\gamma}(\cdot)$ is convex on \mathbb{R} .

Proof of claim (iii). Fix $x \in \mathbb{X}$ and define $\varphi_x : [0,1] \mapsto \mathbb{R}$ as follows:

$$\varphi_{\pi}(q) \triangleq \sup_{\gamma \in \Gamma(q)} J(\pi, \gamma; x, q),$$

so that we have $V(x,q) = \inf_{\pi \in \Pi(x)} \varphi_{\pi}(q)$. Since the point-wise infimum of a set of concave functions is concave, it suffices to show that $\varphi_{\pi}(\cdot)$ is concave on [0,1] for each $\pi \in \Pi(x)$.

Fix $\pi \in \Pi(x)$ and consider any $q_1, q_2 \in [0, 1]$. By Lemma 1 and Lemma 5, there exist $\gamma^1 \in \Gamma(q_1)$ and $\gamma^2 \in \Gamma(q_2)$ such that

$$J(\pi, \gamma^1; x, q_1) = \varphi_{\pi}(q_1), \quad J(\pi, \gamma^2; x, q_2) = \varphi_{\pi}(q_2).$$

Given some $\lambda \in [0,1]$, let $q_{\lambda} \triangleq \lambda q_1 + (1-\lambda)q_2$ and $\gamma^{\lambda} \triangleq \lambda \gamma^1 + (1-\lambda)\gamma^2$ (i.e., $\gamma_t^{\lambda}(\omega) = \lambda \gamma_t^1(\omega) + (1-\lambda)\gamma_t^2(\omega)$ for $\forall t, \omega$). Note that

$$Q_t^{q_{\lambda},\gamma^{\lambda}} = q_{\lambda} + \int_{s=0}^t \gamma_s^{\lambda} dW_s = \lambda Q_t^{q_1,\gamma^1} + (1-\lambda)Q_t^{q_2,\gamma^2},$$

and hence $Q_t^{q_{\lambda},\gamma^{\lambda}} \in [0,1]$ almost surely for any t, i.e., $\gamma^{\lambda} \in \Gamma(q_{\lambda})$. Therefore,

$$J(\pi, \gamma^{\lambda}; x, q_{\lambda}) = \mathbb{E}\left[C_{\infty}^{x, \pi}Q_{\infty}^{q_{\lambda}, \gamma^{\lambda}}\right] = \lambda \mathbb{E}\left[C_{\infty}^{x, \pi}Q_{\infty}^{q_{1}, \gamma^{1}}\right] + (1 - \lambda)\mathbb{E}\left[C_{\infty}^{x, \pi}Q_{\infty}^{q_{2}, \gamma^{2}}\right]$$
$$= \lambda J(\pi, \gamma^{1}; x, q_{1}) + (1 - \lambda)J(\pi, \gamma^{2}; x, q_{2})$$
$$= \lambda \varphi_{\pi}(q_{1}) + (1 - \lambda)\varphi_{\pi}(q_{2}).$$

As a result, we have

$$\varphi_{\pi}(\lambda q_1 + (1 - \lambda)q_2) \ge J(\pi, \gamma^{\lambda}; x, q_{\lambda}) = \lambda \varphi_{\pi}(q_1) + (1 - \lambda)\varphi_{\pi}(q_2).$$

Since q_1, q_2, λ were arbitrarily chosen, $\varphi_{\pi}(\cdot)$ is concave on [0, 1].

Proof of claim (iv). The claim immediately follows from claim (i).

C.3. Proof of CVaR Dynamic Programming Principle

We first state a proposition that is useful for proving upper-semicontinuity of a mapping with respect to the weak topology.

Lemma 8 (Barbu and Precupanu (2012), Proposition 2.10. Restated for upper-semicontinuity). Let \mathcal{Y} be a locally convex space. A proper concave function $f : \mathcal{Y} \to [-\infty, \infty)$ is upper-semicontinuous on \mathcal{Y} if and only if it is upper-semicontinuous with respect to the weak topology on \mathcal{Y} .

We next prove a minimax theorem that includes the value function and a stopping time, which extends Theorem 2.

Proposition 9. For any stopping time $\tau : \Omega \to \mathbb{T}$ with $\mathbb{E}[\tau] < \infty$, we have

$$\inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{x,\pi}_{\tau} Q^{q,\gamma}_{\tau} + V(X^{x,\pi}_{\tau}, Q^{q,\gamma}_{\tau}) \right] = \sup_{\gamma \in \Gamma(q)} \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{x,\pi}_{\tau} Q^{q,\gamma}_{\tau} + V(X^{x,\pi}_{\tau}, Q^{q,\gamma}_{\tau}) \right].$$

Moreover, an optimal solution $\gamma \in \Gamma(q)$ must exist for each maximization, i.e.,

$$\inf_{\pi \in \Pi(x)} \max_{\gamma \in \Gamma(q)} \mathbb{E}\left[C^{x,\pi}_{\tau} Q^{q,\gamma}_{\tau} + V(X^{x,\pi}_{\tau}, Q^{q,\gamma}_{\tau})\right] = \max_{\gamma \in \Gamma(q)} \inf_{\pi \in \Pi(x)} \mathbb{E}\left[C^{x,\pi}_{\tau} Q^{q,\gamma}_{\tau} + V(X^{x,\pi}_{\tau}, Q^{q,\gamma}_{\tau})\right].$$

Proof. Let us first define

$$\mathcal{Q}_{\tau}(q) \triangleq \{ \mathbb{E}(Q|\mathcal{F}_{\tau}) : Q \in \mathcal{Q}(q) \}$$

As analogous to Lemma 1, it can be easily shown that $Q_{\tau}(q) = \{Q^{q,\gamma}_{\tau} | \gamma \in \Gamma(q)\}$. With the notation introduced in the proof of Theorem 2 (§C.1), it suffices to show that

$$\inf_{X \in \mathcal{X}(x)} \max_{Q \in \mathcal{Q}_{\tau}(q)} \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau}, Q)\right] = \max_{Q \in \mathcal{Q}_{\tau}(q)} \inf_{X \in \mathcal{X}(x)} \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau}, Q)\right],$$

for which we need to verify the followings:

- (i) The mapping $X \mapsto \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau}, Q)\right]$ is convex and continuous on $\mathcal{X}(x)$ with respect to the norm $\|\cdot\|_{\mathcal{X}}$ for any given $Q \in \mathcal{Q}_{\tau}(q)$.
- (ii) $\mathcal{Q}_{\tau}(q)$ is a non-empty, convex, and weakly compact subset of $\mathcal{L}^2(\Omega, \mathcal{F}_{\tau}, \mathbb{P})$.
- (iii) The mapping $Q \mapsto \mathbb{E}[V(X_{\tau}, Q)]$ is continuous on $\mathcal{Q}_{\tau}(q)$ endowed with \mathcal{L}^2 -norm¹² for any given $X \in \mathcal{X}(x)$.
- (iv) The mapping $Q \mapsto \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau}, Q)\right]$ is concave and upper-semicontinuous on $\mathcal{Q}_{\tau}(q)$ endowed with \mathcal{L}^{2} -weak topology for any given $X \in \mathcal{X}(x)$.

Together with the convexity of $\mathcal{X}(x)$ (Lemma 7), we obtain the desired claim by utilizing Sion's minimax theorem (Lemma 6).

¹²Note that what we ultimately want to show is that the mapping $Q \mapsto \mathbb{E}[V(X_{\tau}, Q)]$ is upper-semicontinuous on $\mathcal{Q}_{\tau}(q)$ endowed with \mathcal{L}^2 -weak topology. We first show its continuity with respect to \mathcal{L}^2 -norm topology and then utilize Lemma 8 to achieve this goal.

Proof of claim (i). The convexity of the mapping $X \mapsto \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau},Q)\right]$ immediately follows from the arguments in the proof of Theorem 2, and the convexity of $V(\cdot,Q)$, which was shown in Proposition 8(ii). Now fix $X \in \mathcal{X}(x)$ and consider a sequence $(X^{(n)} \in \mathcal{X}(x))_{n \in \mathbb{N}}$ such that $\|X - X^{(n)}\|_{\mathcal{X}} \to 0$ as $n \to \infty$. Following the arguments in the proof of Theorem 2, we can show that $C_{\tau}^{X^{(n)}}Q \xrightarrow{\mathcal{L}^{1}} C_{\tau}^{X}Q$ as $n \to \infty$. We also have that

$$\mathbb{E}\left[\left|X_{\tau}^{(n)} - X_{\tau}\right|\right] = \mathbb{E}\left[\left|\int_{t=0}^{\tau} (\dot{X}_{t}^{(n)} - \dot{X}_{t})dt\right|\right]$$

$$\leq \mathbb{E}\left[\int_{t=0}^{\tau} \left|\dot{X}_{t}^{(n)} - \dot{X}_{t}\right|dt\right]$$

$$\leq \left(\mathbb{E}\left[\int_{t=0}^{\tau} \left|\dot{X}_{t}^{(n)} - \dot{X}_{t}\right|^{2}dt\right]\right)^{\frac{1}{2}} \times \left(\mathbb{E}\left[\int_{t=0}^{\tau} 1^{2}dt\right]\right)^{\frac{1}{2}}$$

$$\leq \|X^{(n)} - X\|_{\mathcal{X}} \times \sqrt{\mathbb{E}[\tau]},$$

where the second inequality holds due to Hölder's inequality. Since $\mathbb{E}[\tau] < \infty$ and $||X^{(n)} - X||_{\mathcal{X}} \searrow 0$ as $n \to \infty$, we deduce that $X_{\tau}^{(n)} \stackrel{\mathcal{L}^1}{\to} X_{\tau}$. By the continuous mapping theorem, we further obtain that $V(X_{\tau}^{(n)}, Q) \stackrel{p}{\to} V(X_{\tau}, Q)$ as $n \to \infty$. Given that $\sup_{t \ge 0} |X_t^{(n)}| \le M$ for all $n \in \mathbb{N}$, the sequence $(|V(X_{\tau}^{(n)}, Q)|)_{n \in \mathbb{N}}$ is uniformly bounded by $\sup_{q' \in [0,1]} \{\sigma^{\frac{2}{3}}\eta^{\frac{1}{3}} \times M^{\frac{4}{3}} \times q'^{\frac{1}{3}}\kappa(q')^{\frac{2}{3}}\} < \infty$, and hence uniformly integrable. Therefore, $V(X_{\tau}^{(n)}, Q) \stackrel{\mathcal{L}^1}{\to} V(X_{\tau}, Q)$. Combining these results, we obtain $\lim_{n\to\infty} \mathbb{E}\left[C_{\tau}^{X^{(n)}}Q + V(X_{\tau}^{(n)}, Q)\right] = \mathbb{E}\left[C_{\tau}^XQ + V(X_{\tau}, Q)\right]$, which concludes the proof.

Proof of claim (ii). The proof is identical to that of Lemma 4 and 5 except that $\mathcal{L}^{\infty}(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{L}^{2}(\Omega, \mathcal{F}, \mathbb{P})$ are replaced with $\mathcal{L}^{\infty}(\Omega, \mathcal{F}_{\tau}, \mathbb{P})$ and $\mathcal{L}^{2}(\Omega, \mathcal{F}_{\tau}, \mathbb{P})$, respectively.

Proof of claim (iii). Fix $Q \in \mathcal{Q}_{\tau}(q)$ and consider a sequence $(Q^{(n)} \in \mathcal{Q}_{\tau}(q))_{n \in \mathbb{N}}$ such that $Q^{(n)} \stackrel{\mathcal{L}^2}{\to} Q$ as $n \to \infty$. By the continuous mapping theorem, we have $V(X_{\tau}, Q^{(n)}) \stackrel{p}{\to} V(X_{\tau}, Q)$. Since the sequence $(|V(X_{\tau}, Q^{(n)})|)_{n \in \mathbb{N}}$ is uniformly integrable (uniformly bounded by $\sup_{q' \in [0,1]} \{\sigma^{\frac{2}{3}}\eta^{\frac{1}{3}} \times |M|^{\frac{4}{3}} \times q'^{\frac{1}{3}}\kappa(q')^{\frac{2}{3}}\} < \infty$), we further have $V(X_{\tau}, Q^{(n)}) \stackrel{\mathcal{L}^1}{\to} V(X_{\tau}, Q)$, which concludes the proof.

Proof of claim (iv). The concavity of the mapping $Q \mapsto \mathbb{E}\left[C_{\tau}^{X}Q + V(X_{\tau},Q)\right]$ immediately follows from the concavity of $V(X_{\tau},\cdot)$, which was shown in Proposition 8(iii). On the other hand, by combining the result of claim (iii) with Lemma 8, we deduce that the mapping $Q \mapsto \mathbb{E}\left[V(X_{\tau},Q)\right]$ is upper-semicontinuous on $\mathcal{Q}_{\tau}(q)$ endowed with \mathcal{L}^{2} -weak topology. Therefore, it suffices to show that the mapping $Q \mapsto \mathbb{E}\left[C_{\tau}^{X}Q\right]$ is upper-semicontinuous with respect to \mathcal{L}^{2} -weak topology.

Fix $Q \in \mathcal{Q}_{\tau}(q)$ and consider a sequence $(Q^{(n)} \in \mathcal{Q}_{\tau}(q))_{n \in \mathbb{N}}$ such that $\mathbb{E}[ZQ^{(n)}] \to \mathbb{E}[ZQ]$ as $n \to \infty$ for any $Z \in \mathcal{L}^2$. Since $C_{\tau}^X \in \mathcal{L}^2$, we have $\lim_{n\to\infty} \mathbb{E}[C_{\tau}^X Q^{(n)}] = \mathbb{E}[C_{\tau}^X Q]$, which implies the continuity of the mapping $Q \mapsto \mathbb{E}\left[C_{\tau}^X Q\right]$.

Before proving Theorem 3, we introduce additional notation to describe the Markovian structure of problem. We denote by $(X_s^{t,s,\pi})_{s\geq t}$ the trader's position process under control π that begins from

the value x at time t, i.e.,

$$X_s^{t,x,\pi} \triangleq x - \int_{u=t}^s \pi_u du.$$

We define the adversary's martingale process $(Q_s^{t,q,\gamma})_{s\geq t}$ and the loss process $(C_s^{t,x,\pi})_{s\geq t}$ analogously,

$$Q_s^{t,q,\gamma} \triangleq q + \int_{u=t}^s \gamma_u dW_u, \quad C_s^{t,x,\pi} \triangleq \int_{u=t}^s \frac{1}{2} \eta \pi_u^2 du - \int_{u=t}^s \sigma X_u^{t,x,\pi} dW_u$$

With this notation, we can describe the aforementioned processes in a recursive way,

$$X_{s}^{0,x,\pi} = X_{s}^{t,X_{t}^{0,x,\pi},\pi}, \quad Q_{s}^{0,q,\gamma} = Q_{s}^{t,Q_{t}^{0,q,\gamma},\gamma}, \quad C_{s}^{0,q,\gamma} = C_{t}^{0,x,\pi} + C_{s}^{t,X_{t}^{0,x,\pi},\pi}.$$

The policy spaces are defined analogously as well,

$$\Pi_t(x) \triangleq \{ \pi \in \Pi(x) \mid \pi_s = 0, \forall s < t \}, \quad \Gamma_t(q) \triangleq \{ \gamma \in \Gamma(q) \mid \gamma_s = 0, \forall s < t \}.$$

We prove Theorem 3 by utilizing Proposition 9 and Proposition 3.

Proof of Theorem 3. Define

$$U(x,q) \triangleq \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau}^{0,x,\pi} Q_{\tau}^{0,q,\gamma} + V \left(X_{\tau}^{0,x,\pi}, Q_{\tau}^{0,q,\gamma} \right) \right].$$

By Lemma 1 and Lemma 5, we can equivalently write

$$U(x,q) = \inf_{\pi \in \Pi(x)} \max_{\gamma \in \Gamma(q)} \mathbb{E}\left[C_{\tau}^{0,x,\pi} Q_{\tau}^{0,q,\gamma} + V\left(X_{\tau}^{0,x,\pi}, Q_{\tau}^{0,q,\gamma}\right)\right].$$

We aim to prove V(x,q) = U(x,q).

Proof of " $V(x,q) \leq U(x,q)$ ". Fix $\epsilon > 0$. By definition of U(x,q), there exists $\pi^{\circ} \in \Pi(x)$ such that

$$U(x,q) + \epsilon \ge \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau}^{0,x,\pi^{\circ}} Q_{\tau}^{0,q,\gamma} + V \left(X_{\tau}^{0,x,\pi^{\circ}}, Q_{\tau}^{0,q,\gamma} \right) \right].$$

On the other hand, we have that for any $t \ge 0$, $\hat{x} \in \mathbb{X}$, and $\hat{q} \in [0, 1]$,

$$V(\hat{x}, \hat{q}) = \inf_{\hat{\pi} \in \Pi_t(\hat{x})} \max_{\hat{\gamma} \in \Gamma_t(\hat{q})} \mathbb{E} \left[C_{\infty}^{t, \hat{x}, \hat{\pi}} Q_{\infty}^{t, \hat{q}, \hat{\gamma}} \right],$$

by time homogeneity of the problem. Consequently, for each $\omega \in \Omega$ and $\gamma \in \Gamma(q)$, there exists $\hat{\pi}^{\omega,\gamma} \in \Pi_{\tau(\omega)}(X^{0,x,\pi^{\circ}}_{\tau}(\omega))$ such that

$$\begin{split} V\left(X^{0,x,\pi^{\circ}}_{\tau}(\omega),Q^{0,q,\gamma}_{\tau}(\omega)\right) + \epsilon &\geq \max_{\hat{\gamma}\in\Gamma_{\tau(\omega)}(Q^{0,q,\gamma}_{\tau}(\omega))} \mathbb{E}\left[C^{\tau,X^{0,x,\pi^{\circ}}_{\infty},\hat{\pi}^{\omega,\gamma}}_{\infty}Q^{\tau,Q^{0,q,\gamma}_{\tau},\hat{\gamma}}_{\infty}\middle|\,\mathcal{F}_{\tau}\right](\omega) \\ &\geq \mathbb{E}\left[C^{\tau,X^{0,x,\pi^{\circ}}_{\infty},\hat{\pi}^{\omega,\gamma}}_{\infty}Q^{\tau,Q^{0,q,\gamma}_{\tau},\gamma}_{\infty}\middle|\,\mathcal{F}_{\tau}\right](\omega), \end{split}$$

where the second inequality follows from the fact that the given γ may not be optimal for the period $t \geq \tau$.

For each $\gamma \in \Gamma(q)$, consider a trader's policy $\pi^{\gamma} : \mathbb{T} \times \Omega \to \mathbb{R}$ constructed as follows:

$$\pi_t^{\gamma}(\omega) \triangleq \begin{cases} \pi_t^{\circ}(\omega) & \text{if } t < \tau(\omega), \\ \hat{\pi}_t^{\omega,\gamma}(\omega) & \text{if } t \ge \tau(\omega), \end{cases}$$

i.e., it implements an ϵ -optimal solution to the Bellman equation before τ , and then implements another ϵ -optimal solution for the remaining horizon. We observe that $\pi^{\gamma} \in \Pi(x)$ since $\pi_t^{\circ}(\omega) \in$ $\Pi(x)$ and $\hat{\pi}_t^{\omega,\gamma} \in \Pi(x)$. Combining these results, we obtain

$$\begin{split} U(x,q) &\geq \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{0,x,\pi^{\circ}}_{\tau} Q^{0,q,\gamma}_{\tau} + V \left(X^{0,x,\pi^{\circ}}_{\tau}, Q^{0,q,\gamma}_{\tau} \right) \right] - \epsilon \\ &\geq \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{0,x,\pi^{\circ}}_{\tau} Q^{0,q,\gamma}_{\tau} + \mathbb{E} \left[C^{\tau,X^{0,x,\pi^{\circ}}_{\infty},\hat{\pi}^{\omega,\gamma}}_{\infty} Q^{\tau,Q^{0,q,\gamma}_{\infty},\gamma}_{\infty} \right| \mathcal{F}_{\tau} \right] \right] - 2\epsilon \\ &= \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{0,x,\pi^{\gamma}}_{\tau} Q^{0,q,\gamma}_{\tau} + C^{\tau,X^{0,x,\pi^{\gamma}}_{\tau},\pi^{\gamma}}_{\infty} Q^{\tau,Q^{0,q,\gamma}_{\tau},\gamma}_{\infty} \right] - 2\epsilon \\ &= \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{0,x,\pi^{\gamma}}_{\infty} Q^{0,q,\gamma}_{\infty} \right] - 2\epsilon \\ &\geq \max_{\gamma \in \Gamma(q)} \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,x,\pi}_{\infty} Q^{0,q,\gamma}_{\infty} \right] - 2\epsilon \\ &= V(x,q) - 2\epsilon. \end{split}$$

Since the choice of ϵ was arbitrary, we deduce that $U(x,q) \ge V(x,q)$.

Proof of " $V(x,q) \ge U(x,q)$ ". By minimax equality result for U(x,q) (Proposition 9), there exists $\gamma^{\circ} \in \Gamma(q)$ such that

$$U(x,q) = \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,x,\pi}_{\tau} Q^{0,q,\gamma^{\circ}}_{\tau} + V \left(X^{0,x,\pi}_{\tau}, Q^{0,q,\gamma^{\circ}}_{\tau} \right) \right].$$

By minimax equality result for V(x,q) (Theorem 2) and by time homogeneity of the problem, we further have that for any $t \ge 0$, $\hat{x} \in \mathbb{X}$, and $\hat{q} \in [0,1]$,

$$V(\hat{x}, \hat{q}) = \max_{\hat{\gamma} \in \Gamma_t(\hat{q})} \inf_{\hat{\pi} \in \Pi_t(\hat{x})} \mathbb{E} \left[C_{\infty}^{t, \hat{x}, \hat{\pi}} Q_{\infty}^{t, \hat{q}, \hat{\gamma}} \right].$$

Therefore, for each $\omega \in \Omega$ and $\pi \in \Pi(x)$, there exists $\hat{\gamma}^{\omega,\pi} \in \Gamma_{\tau(\omega)}(Q^{0,q,\gamma^{\circ}}_{\tau}(\omega))$ such that

$$V\left(X_{\tau}^{0,x,\pi}(\omega), Q_{\tau}^{0,q,\gamma^{\circ}}(\omega)\right) = \inf_{\hat{\pi} \in \Pi_{\tau(\omega)}(X_{\tau}^{0,x,\pi}(\omega))} \mathbb{E}\left[C_{\infty}^{\tau,X_{\tau}^{0,x,\pi},\hat{\pi}}Q_{\infty}^{\tau,Q_{\tau}^{0,q,\gamma^{\circ}},\hat{\gamma}^{\omega,\pi}} \middle| \mathcal{F}_{\tau}\right](\omega)$$
$$\leq \mathbb{E}\left[C_{\infty}^{\tau,X_{\tau}^{0,x,\pi},\pi}Q_{\infty}^{\tau,Q_{\tau}^{0,q,\gamma^{\circ}},\hat{\gamma}^{\omega,\pi}} \middle| \mathcal{F}_{\tau}\right](\omega).$$

For each $\pi \in \Pi(x)$, consider an adversary's policy $\gamma^{\pi} : \mathbb{T} \times \Omega \to \mathbb{R}$ constructed as

$$\gamma_t^{\pi}(\omega) \triangleq \begin{cases} \gamma_t^{\circ}(\omega) & \text{if } t < \tau(\omega), \\ \hat{\gamma}_t^{\omega,\pi}(\omega) & \text{if } t \ge \tau(\omega). \end{cases}$$

We observe that $\gamma^{\pi} \in \Gamma(q)$. Combining these results, we obtain

$$\begin{split} U(x,q) &= \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,x,\pi}_{\tau} Q^{0,q,\gamma^{\circ}}_{\tau} + V \left(X^{0,x,\pi}_{\tau}, Q^{0,q,\gamma^{\circ}}_{\tau} \right) \right] \\ &\leq \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,x,\pi}_{\tau} Q^{0,q,\gamma^{\circ}}_{\tau} + \mathbb{E} \left[C^{\tau,X^{0,x,\pi}_{\tau},\pi}_{\infty} Q^{\tau,Q^{0,q,\gamma^{\circ}}_{\tau}}_{\infty}, \hat{\gamma}^{\omega,\pi}_{\infty} \middle| \mathcal{F}_{\tau} \right] \right] \\ &= \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,x,\pi}_{\tau} Q^{0,q,\gamma^{\pi}}_{\tau} + C^{\tau,X^{0,x,\pi}_{\tau},\pi}_{\infty} Q^{\tau,Q^{0,q,\gamma^{\pi}}_{\infty}}_{\infty}, \gamma^{\pi}_{\infty} \right] \\ &= \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{0,X^{0,x,\pi}_{0},\pi}_{\infty} Q^{0,q,\gamma^{\pi}}_{\infty} \right] \\ &\leq \inf_{\pi \in \Pi(x)} \max_{\gamma \in \Gamma(q)} \mathbb{E} \left[C^{0,X^{0,x,\pi},\pi}_{\infty} Q^{0,q,\gamma}_{\infty} \right] = V(x,q). \end{split}$$

This concludes the proof.

D. Proofs for §4

D.1. Proofs of Theorem 4 and Theorem 6

In this section, we prove Theorem 4 and Theorem 6 together in Theorem 7 below. We do this by constructing the sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ from a function V^* satisfying the conditions of Theorem 4 and by characterizing the outcome of the game when the trader or the adversary employs a policy induced by $f^{(n)}$ or $g^{(n)}$. Before making a formal statement of Theorem 7, we clarify what we aim to show, address some technical difficulty, and illustrate how to detour such a difficulty.

Consider a function $V^*(x,q)$ that satisfies the conditions of Theorem 4, and define

$$f^{\star}(x,q) \triangleq \frac{V_x^{\star}(x,q)}{\eta q}, \quad g^{\star}(x,q) \triangleq \frac{\sigma x}{V_{qq}^{\star}(x,q)},$$

which solve the trader's and the adversary's optimization problems in the HJB equation:

$$f^{\star}(x,q) = \operatorname*{argmin}_{v \in \mathbb{R}} \left\{ \frac{\eta}{2} q v^2 - V_x^{\star}(x,q) v \right\}, \quad g^{\star}(x,q) = \operatorname*{argmax}_{w \in \mathbb{R}} \left\{ \frac{1}{2} V_{qq}^{\star}(x,q) w^2 - \sigma x w \right\},$$

for any $x \in \mathbb{X}$ and $q \in (0, 1)$. We would argue that they specify the trader's optimal liquidation rate, $\pi_t^{\star} = f^{\star}(X_t, Q_t)$, and the adversary's optimal quantile diffusion rate, $\gamma_t^{\star} = g^{\star}(X_t, Q_t)$, and this combination of policies, $(\pi^{\star}, \gamma^{\star})$, achieves the equilibrium of this game at which the outcome equals $V^{\star}(x, q)$.

First note that we aim to characterize the equilibrium of this game. We will not argue that the trader's policy induced by f^* is optimal against any adversary's policy γ . Instead, we will argue is that such a trader's policy yields an outcome that is no worse than $V^*(x,q)$ against any adversary's policy γ , i.e., $J(\pi^{*,\gamma},\gamma) \leq V^*(x,q)$ for any adversary's policy γ , where $\pi^{*,\gamma}$ is the trader's policy induced by f^* and coupled with γ . Similarly, we will also argue that the adversary's policy induced by g^* yields an outcome that is no better than $V^*(x,q)$ against any trader's policy π , i.e., $J(\pi, \gamma^{*,\pi}) \geq V^*(x,q)$ for any trader's policy π where $\gamma^{*,\pi}$ is the adversary's policy induced by g^* and coupled with π . By combining two results, we would be able to argue that a mutually coupled (X,Q)-Markov policy pair (π^*, γ^*) induced by (f^*, g^*) is the best response to each other and yields an outcome that equals $V^*(x,q)$ as a saddle point.

To complete the above arguments, however, we need to show that the policies π^* and γ^* are admissible (i.e., they satisfy the feasibility conditions that we have introduced in $\Pi(x)$ and $\Gamma(q)$). Unfortunately, their feasibility is uncertain due to the pathological behavior of f^* and g^* near the boundaries when $x \approx 0$, $q \approx 0$, or $q \approx 1$: for instance, we may have $\mathbb{E}[\int_{t=0}^{\infty} \pi_t^2 dt] = \infty$ since $\lim_{q \searrow 0} f^*(x,q) = \infty$; $\mathbb{E}[\int_{t=0}^{\infty} |X_t|^2 dt] = \infty$ since $\lim_{q \nearrow 1} f^*(x,q) = 0$; or $\mathbb{E}[\int_{t=0}^{\infty} \gamma_t^2 dt] = \infty$ since $\lim_{x \searrow 0} g^*(x,q) = -\infty$, which contradicts to the fact that Q_t is a bounded martingale.

To avoid this feasibility issue, we construct a sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ that

approximates (f^*, g^*) while inducing feasible policies. More specifically, we define a sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$:

$$f^{(n)}(x,q) \triangleq \begin{cases} f^{\star}(x,q) & \text{if } (x,q) \in \mathcal{A}_n, \\ x/n & \text{if } (x,q) \notin \mathcal{A}_n, \end{cases} g^{(n)}(x,q) \triangleq \begin{cases} g^{\star}(x,q) & \text{if } (x,q) \in \mathcal{A}_n, \\ 0 & \text{if } (x,q) \notin \mathcal{A}_n, \end{cases}$$

where

$$\mathcal{A}_n \triangleq \left\{ (x,q) \in \mathbb{X} \times [0,1] \, \middle| \, |x| > \frac{1}{n}, \, \frac{1}{n} < q < 1 - \frac{1}{n} \right\}.$$

It can be easily seen that $f^{(n)}$ and $g^{(n)}$ converge to f^* and g^* at every interior point $(x,q) \in \mathbb{R}_+ \times (0,1)$ as $n \to \infty$. Moreover, these functions suppress the extreme behavior of f^* and g^* near the boundaries, i.e., outside the set \mathcal{A}_n . The policies induced by $f^{(n)}$ and $g^{(n)}$ can be shown to be feasible. This is formally established in Theorem 7 (i)–(iii). Finally, we can complete the equilibrium argument made above by showing that this sequence of function pairs $(f^{(n)}, g^{(n)})_{n \in \mathbb{N}}$ achieves the equilibrium asymptotically (Theorem 7 (vi)–(viii)). In the midst of the proof, we confirm that the equilibrium outcome, V(x,q), coincides the solution of HJB equation, $V^*(x,q)$ (Theorem 7 (iv) and (v)), by showing that the gap between them can be tightened arbitrarily small as $n \to \infty$.

Theorem 7. Fix $x \in (0, M]$ and $q \in (0, 1)$, and consider functions V^* , f^* , g^* , $f^{(n)}$, and $g^{(n)}$ introduced above. For any $n > \max\{\frac{1}{x}, \frac{1}{q}, \frac{1}{1-q}\}$, we have:

- (i) For any adversary's policy $\gamma \in \Gamma(q)$, a (X,Q)-Markov trader's policy $\pi^{(n),\gamma}$ induced by $f^{(n)}$ and coupled with γ is admissible, i.e., $\pi^{(n),\gamma} \in \Pi(x)$.
- (ii) For any trader's policy $\pi \in \Pi(x)$, a (X, Q)-Markov adversary's policy $\gamma^{(n),\pi}$ induced by $g^{(n)}$ and coupled with π is admissible, i.e., $\gamma^{(n),\pi} \in \Gamma(q)$.
- (iii) Mutually coupled (X, Q)-Markov policy pair $(\pi^{(n)}, \gamma^{(n)})$ induced by $(f^{(n)}, g^{(n)})$ is admissible, i.e., $\pi^{(n)} \in \Pi(x)$ and $\gamma^{(n)} \in \Gamma(q)$.
- (iv) $V(x,q) \leq V^{\star}(x,q)$.
- (v) $V(x,q) \ge V^{\star}(x,q).$
- (vi) For any $\gamma \in \Gamma(q)$, $(\pi^{(n),\gamma})_{n \in \mathbb{N}}$ defined in (i) satisfies $\limsup_{n \to \infty} J(\pi^{(n),\gamma},\gamma;x,q) \leq V(x,q)$.

(vii) For any $\pi \in \Pi(x)$, $(\gamma^{(n),\pi})_{n \in \mathbb{N}}$ defined in (ii) satisfies $\liminf_{n \to \infty} J(\pi, \gamma^{(n),\pi}; x, q) \ge V(x, q)$.

(viii) $(\pi^{(n)}, \gamma^{(n)})_{n \in \mathbb{N}}$ defined in (iii) satisfies $\lim_{n \to \infty} J(\pi^{(n)}, \gamma^{(n)}; x, q) = V(x, q).$

Proof. Throughout the proof, we define a sequence of hitting times $(\tau_n)_{n \in \mathbb{N}}$ as follows:

$$\tau_n \triangleq \inf_{t \ge 0} \{ (X_t, Q_t) \notin \mathcal{A}_n \}$$

Since \mathcal{A}_n^c is a closed set and $(X_t)_{t\geq 0}$ and $(Q_t)_{t\geq 0}$ are continuous, τ_n is a stopping time for each

n. Also note that if the trader's policy is induced by $f^{(n)}$ or the adversary's policy is induced by $g^{(n)}$, then the process (X_t, Q_t) never returns to \mathcal{A}_n once it leaves \mathcal{A}_n , since $|X_t|$ is monotonically decreasing or Q_t remains unchanged after τ_n .

Proof of claim (i). Under the trader's policy $\pi^{(n),\gamma}$, the position process $(X_t)_{t\geq 0}$ is described by

$$X_t = \begin{cases} x - \int_{s=0}^t f^{\star}(X_s, Q_s) \, ds, & \forall t \le \tau_n, \\ X_{\tau_n} e^{-(t - \tau_n)/n}, & \forall t \ge \tau_n. \end{cases}$$
(27)

Observe that $X_t \in [0, x]$ for any $t \ge 0$, and in particular, X_t is monotonically decreasing over time. Given that $(Q_t)_{t\ge 0}$ has a continuous sample path and n is large enough such that $q \in \left[\frac{1}{n}, 1 - \frac{1}{n}\right]$, we also have $Q_t \in \left[\frac{1}{n}, 1 - \frac{1}{n}\right]$ for any $t \in [0, \tau_n]$. Then, the sample path of $(X_t)_{t\ge 0}$ for each ω is uniquely determined (i.e., the associated SDE has a unique strong solution): the uniqueness of sample path on $[0, \tau_n)$ follows from the fact that $f^*(\cdot, \cdot)$ is Lipschitz continuous on $\left(\frac{1}{n}, x\right] \times \left(\frac{1}{n}, 1 - \frac{1}{n}\right) \subset \mathcal{A}_n$ due to condition (iv), and the uniqueness on $[\tau_n, \infty)$ immediately follows from the fact that $X_t = X_{\tau_n} e^{-(t-\tau_n)/n}$ for $t \ge \tau_n$.

We next show that τ_n is almost surely bounded. For any $t < \tau_n$ (so that $(X_t, Q_t) \in \mathcal{A}_n$), by condition (iv), we have

$$f^{(n)}(X_t, Q_t) \ge f^{\star}(X_t, 1 - 1/n) \ge f^{\star}(1/n, 1 - 1/n) := \alpha_n > 0.$$

Then, for any $t < \tau_n$, we have $dX_t/dt = -f^{(n)}(X_t, Q_t) \leq -\alpha_n$, and hence $X_{\tau_n} \leq x - \alpha_n \tau_n$. Together with the fact that $X_{\tau_n} \geq 0$, we deduce that

$$\tau_n \le \frac{x}{\alpha_n},\tag{28}$$

on any sample path.

We further have that, by condition (iv), for any $t < \tau_n$,

$$f^{(n)}(X_t, Q_t) \le f^*(X_t, 1/n) \le f^*(x, 1/n) := \beta_n < \infty$$

Then,

$$\int_{t=0}^{\tau_n} \left(\pi_t^{(n),\gamma}\right)^2 dt = \int_{t=0}^{\tau_n} \left(f^{\star}(X_t, Q_t)\right)^2 dt \le \tau_n \beta_n^2 \le \frac{x\beta_n^2}{\alpha_n}, \quad \int_{t=0}^{\tau_n} X_t^2 dt \le \tau_n x^2 \le \frac{x^3}{\alpha_n}$$

on any sample path. Further, since $\pi_t^{(n),\gamma} = X_t/n$ for any $t \ge \tau_n$,

$$\int_{t=\tau_n}^{\infty} \left(\pi_t^{(n),\gamma}\right)^2 dt = \frac{1}{n^2} \int_{t=\tau_n}^{\infty} X_t^2 dt = \frac{1}{n^2} \int_{t=\tau_n}^{\infty} X_{\tau_n}^2 e^{-2(t-\tau_n)/n} dt \le \frac{x^2}{2n}, \quad \int_{t=\tau_n}^{\infty} X_t^2 dt \le \frac{nx^2}{2}.$$

As a result,

$$\mathbb{E}\left[\left(\int_{t=0}^{\infty} \left(\pi_t^{(n),\gamma}\right)^2 dt\right)^2\right] \le \left(\frac{x\beta_n^2}{\alpha_n} + \frac{x^2}{2n}\right)^2 < \infty, \quad \mathbb{E}\left[\int_{t=0}^{\infty} X_t^2 dt\right] \le \frac{x^3}{\alpha_n} + \frac{nx^2}{2} < \infty.$$

Also note that $\sup_{t\geq 0} |X_t| \leq x \leq M$. Therefore, $\pi^{(n),\gamma}$ is admissible, i.e., $\pi^{(n),\gamma} \in \Pi(x)$.

Proof of claim (ii). Under the adversary's policy $\gamma^{(n),\pi}$, the martingale process $(Q_t)_{t\geq 0}$ is described by

$$Q_t = \begin{cases} q + \int_{s=0}^t g^{\star} \left(X_s, Q_s \right) dW_s, & \forall t \le \tau_n, \\ Q_{\tau_n}, & \forall t \ge \tau_n. \end{cases}$$
(29)

From condition (v), we can show that $g^{\star}(\cdot, \cdot)$ is Lipschitz continuous and bounded on $\left(\frac{1}{n}, \infty\right) \times \left(\frac{1}{n}, 1 - \frac{1}{n}\right) \subset \mathcal{A}_n$. Therefore, for each ω , the sample path of $(Q_t)_{t\geq 0}$ is unique and continuous (given that n is large enough so that $q \in \left[\frac{1}{n}, 1 - \frac{1}{n}\right]$, and $(X_t)_{t\geq 0}$ is continuous), and hence $Q_t \in \left[\frac{1}{n}, 1 - \frac{1}{n}\right] \subset [0, 1]$ almost surely for any $t \in [0, \infty)$, and hence $\gamma^{(n), \pi}$ is admissible.

For later use, we show that $\mathbb{E}[\tau_n] < \infty$ under the adversary's policy $\gamma^{(n),\pi}$. Observe that, for any $t < \tau_n$,

$$\left|g^{(n)}(X_t, Q_t)\right| = \left|g^{\star}(X_t, Q_t)\right| \ge \left|g^{\star}(M, Q_t)\right| \ge \min_{q \in [1/n, 1-1/n]} \left\{g^{\star}(M, q)\right\} := \xi_n.$$

In other words, the diffusion rate of the quantile process $(Q_t)_{t\geq 0}$ is lower bounded by ξ_n (up to time τ_n). Therefore, its quadratic variation process $(\langle Q \rangle_t)_{t\geq 0}$ satisfies

$$\langle Q \rangle_t = \int_{s=0}^{t \wedge \tau_n} \left(g^{\star}(X_s, Q_s) \right)^2 ds \ge \xi_n^2(t \wedge \tau_n),$$

for any $t \in \mathbb{R}$. On the other hand, by Burkholder-Davis-Gundy inequality, there exists c such that $c\mathbb{E}[\langle Q \rangle_{t\wedge\tau_n}] \leq \mathbb{E}[(\sup_{s \leq t\wedge\tau_n} Q_s)^2]$ for any t, and thus $\xi_n^2 \mathbb{E}[t\wedge\tau_n] \leq \mathbb{E}[\langle Q \rangle_{t\wedge\tau_n}] \leq \mathbb{E}[(\sup_{s \leq t\wedge\tau_n} Q_s)^2]/c \leq 1/c$. By monotone convergence theorem, we deduce that $\mathbb{E}[\tau_n] = \lim_{t\to\infty} \mathbb{E}[t\wedge\tau_n] \leq 1/(c\xi_n^2) < \infty$.

Proof of claim (iii). Under the mutually coupled policy pair $(\pi^{(n)}, \gamma^{(n)})$, the process pair $(X_t, Q_t)_{t\geq 0}$ satisfies (27) and (29) simultaneously. By the same argument above (claim (ii)), we can show that $\gamma^{(n)}$ is admissible, and also that $\pi^{(n)}$ is admissible by claim (i).

Proof of claim (iv). Fix $\gamma \in \Gamma(q)$ and $n \in \mathbb{N}$. For simplicity, we let π denote $\pi^{(n),\gamma} \in \Pi(x)$. For any

 $t < \tau_n$, since $f^{(n)}(X_t, Q_t) = f^*(X_t, Q_t)$, by condition (i), we have

$$\begin{split} &\left(\frac{\eta}{2}Q_{t}\pi_{t}^{2}-V_{x}^{\star}\left(X_{t},Q_{t}\right)\pi_{t}\right)+\left(\frac{1}{2}V_{qq}^{\star}\left(X_{t},Q_{t}\right)\gamma_{t}^{2}-\sigma X_{t}\gamma_{t}\right)\\ &=\left(\frac{\eta}{2}Q_{t}\left(f^{\star}\left(X_{t},Q_{t}\right)\right)^{2}-V_{x}^{\star}\left(X_{t},Q_{t}\right)f^{\star}\left(X_{t},Q_{t}\right)\right)+\left(\frac{1}{2}V_{qq}^{\star}\left(X_{t},Q_{t}\right)\gamma_{t}^{2}-\sigma X_{t}\gamma_{t}\right)\\ &=\min_{v\in\mathbb{R}}\left\{\frac{\eta}{2}Q_{t}v^{2}-V_{x}^{\star}\left(X_{t},Q_{t}\right)v\right\}+\left(\frac{1}{2}V_{qq}^{\star}\left(X_{t},Q_{t}\right)\gamma_{t}^{2}-\sigma X_{t}\gamma_{t}\right)\\ &\leq\min_{v\in\mathbb{R}}\left\{\frac{\eta}{2}Q_{t}v^{2}-V_{x}^{\star}\left(X_{t},Q_{t}\right)v\right\}+\max_{w\in\mathbb{R}}\left\{\frac{1}{2}V_{qq}^{\star}\left(X_{t},Q_{t}\right)w^{2}-\sigma X_{t}w\right\}\\ &=0. \end{split}$$

By plugging this result into Proposition 4, we obtain

$$\mathbb{E}\left[C_{\tau_{n}}Q_{\tau_{n}}\right] \tag{30}$$

$$\leq \mathbb{E}\left[C_{\tau_{n}}Q_{\tau_{n}} + V^{\star}\left(X_{\tau_{n}}, Q_{\tau_{n}}\right)\right]$$

$$= V^{\star}\left(x, q\right) + \mathbb{E}\left[\int_{t=0}^{\tau_{n}} \left\{ \left(\frac{\eta}{2}Q_{t}\pi_{t}^{2} - V_{x}^{\star}\left(X_{t}, Q_{t}\right)\pi_{t}\right) + \left(\frac{1}{2}V_{qq}^{\star}\left(X_{t}, Q_{t}\right)\gamma_{t}^{2} - \sigma X_{t}\gamma_{t}\right) \right\} dt \right]$$

$$\leq V^{\star}\left(x, q\right),$$

where the first inequality follows from the non-negativity of $V^*(\cdot, \cdot)$. On the other hand, by the definition of τ_n and the continuity of sample paths, we have either $X_{\tau_n} = \frac{1}{n}$ or $Q_{\tau_n} \in \{\frac{1}{n}, 1 - \frac{1}{n}\}$, and therefore,

$$\mathbb{E}\left[V\left(X_{\tau_n}, Q_{\tau_n}\right)\right] \le \mathbb{E}\left[\max\left\{V(1/n, Q_{\tau_n}), V(X_{\tau_n}, 1/n), V(X_{\tau_n}, 1-1/n)\right\}\right]$$
$$\le \sup_{q' \in [0,1]} \left\{V(1/n, q')\right\} + V(x, 1/n) + V(x, 1-1/n) := A_n(x), \tag{31}$$

where the second inequality follows from the fact that $X_{\tau} \leq x$ under policy $\pi^{(n),\gamma}$ and $V(\cdot,q)$ is increasing on $[0,\infty)$. It can be shown easily that $\lim_{n\to\infty} A_n(x) = 0$ for any x due to Proposition 8.

Since γ was chosen to be arbitrary, from (30), we deduce that

$$\sup_{\gamma \in \Gamma(q)} \mathbb{E}\left[C^{x,\pi^{(n),\gamma}}_{\tau_n} Q^{q,\gamma}_{\tau_n}\right] \le V^*(x,q).$$
(32)

Utilizing Theorem 3 (we have shown that $\mathbb{E}[\tau_n] < \infty$ in (28)) and Proposition 9 with (31) and (32),

we further obtain

$$V(x,q) \stackrel{\text{Thm 3}}{=} \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau_n}^{x,\pi} Q_{\tau_n}^{q,\gamma} + V \left(X_{\tau_n}^{x,\pi}, Q_{\tau_n}^{q,\gamma} \right) \right]$$

$$\stackrel{\text{Prop 9}}{=} \sup_{\gamma \in \Gamma(q)} \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C_{\tau_n}^{x,\pi} Q_{\tau_n}^{q,\gamma} + V \left(X_{\tau_n}^{x,\pi}, Q_{\tau_n}^{q,\gamma} \right) \right]$$

$$\leq \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau_n}^{x,\pi^{(n),\gamma}} Q_{\tau_n}^{q,\gamma} + V \left(X_{\tau_n}^{x,\pi^{(n),\gamma}}, Q_{\tau_n}^{q,\gamma} \right) \right]$$

$$\stackrel{(31)}{\leq} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau_n}^{x,\pi^{(n),\gamma}} Q_{\tau_n}^{q,\gamma} \right] + A_n(x)$$

$$\stackrel{(32)}{\leq} V^{\star}(x,q) + A_n(x).$$

Since $\lim_{n\to\infty} A_n(x) = 0$, we obtain $V(x,q) \leq V^*(x,q)$, which concludes the proof.

Proof of claim (v). The proof is almost symmetric to that of claim (iv). Fix $\pi \in \Pi(x)$ and $n \in \mathbb{N}$. For simplicity, we let γ denote $\gamma^{(n),\pi} \in \Gamma(q)$. For any $t < \tau_n$, since $g^{(n)}(X_t, Q_t) = g^*(X_t, Q_t)$, by condition (i), we have

$$\begin{pmatrix} \frac{\eta}{2}Q_t\pi_t^2 - V_x^{\star}\left(X_t, Q_t\right)\pi_t \end{pmatrix} + \begin{pmatrix} \frac{1}{2}V_{qq}^{\star}\left(X_t, Q_t\right)\gamma_t^2 - \sigma X_t\gamma_t \end{pmatrix}$$

$$= \begin{pmatrix} \frac{\eta}{2}Q_t\pi_t^2 - V_x^{\star}\left(X_t, Q_t\right)\pi_t \end{pmatrix} + \max_{w\in\mathbb{R}} \left\{ \frac{1}{2}V_{qq}^{\star}\left(X_t, Q_t\right)w^2 - \sigma X_tw \right\}$$

$$\ge \min_{v\in\mathbb{R}} \left\{ \frac{\eta}{2}Q_tv^2 - V_x^{\star}\left(X_t, Q_t\right)v \right\} + \max_{w\in\mathbb{R}} \left\{ \frac{1}{2}V_{qq}^{\star}\left(X_t, Q_t\right)w^2 - \sigma X_tw \right\}$$

$$= 0.$$

Consequently, by Proposition 4,

$$\mathbb{E}\left[C_{\tau_n}Q_{\tau_n} + V^{\star}(X_{\tau_n}, Q_{\tau_n})\right]$$

$$= V^{\star}(x, q) + \mathbb{E}\left[\int_{t=0}^{\tau_n} \left\{ \left(\frac{\eta}{2}Q_t \pi_t^2 - V_x^{\star}(X_t, Q_t) \pi_t\right) + \left(\frac{1}{2}V_{qq}^{\star}(X_t, Q_t) \gamma_t^2 - \sigma X_t \gamma_t\right) \right\} dt \right]$$

$$\geq V^{\star}(x, q).$$

$$(33)$$

By the definition of τ_n and the continuity of sample paths, we have either $X_{\tau_n} = \frac{1}{n}$ or $Q_{\tau_n} \in \{\frac{1}{n}, 1 - \frac{1}{n}\}$, and hence,

$$\mathbb{E}\left[V^{\star}(X_{\tau_n}, Q_{\tau_n})\right] \leq \mathbb{E}\left[\max\left\{V^{\star}(1/n, Q_{\tau_n}), V^{\star}(X_{\tau_n}, 1/n), V^{\star}(X_{\tau_n}, 1-1/n)\right\}\right]$$
(34)
$$\leq \sup_{q' \in [0,1]} \{V^{\star}(1/n, q')\} + V^{\star}(M, 1/n) + V^{\star}(M, 1-1/n) := B_n,$$

where the last inequality follows from the feasibility of π (i.e., $|X_{\tau_n}| \leq M$) and the monotonicity of $V^*(\cdot, q)$ (condition (iii)). By condition (ii), we have $\lim_{n\to 0} B_n = 0$. Since π was chosen arbitrary, we further deduce that

$$V^{\star}(x,q) \stackrel{(33)}{\leq} \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{x,\pi}_{\tau_n} Q^{q,\gamma^{(n),\pi}}_{\tau_n} + V^{\star} \left(X^{x,\pi}_{\tau_n}, Q^{q,\gamma^{(n),\pi}}_{\tau_n} \right) \right] \stackrel{(34)}{\leq} \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C^{x,\pi}_{\tau_n} Q^{q,\gamma^{(n),\pi}}_{\tau_n} \right] + B_n.$$
(35)

By utilizing Theorem 3 (we have shown that $\mathbb{E}[\tau_n] < \infty$ at the end of the proof of claim (ii)) and the above results, we obtain

$$V(x,q) \stackrel{\text{Thm 3}}{=} \inf_{\pi \in \Pi(x)} \sup_{\gamma \in \Gamma(q)} \mathbb{E} \left[C_{\tau_n}^{x,\pi} Q_{\tau_n}^{q,\gamma} + V \left(X_{\tau_n}^{x,\pi}, Q_{\tau_n}^{q,\gamma} \right) \right]$$

$$\geq \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C_{\tau_n}^{x,\pi} Q_{\tau_n}^{q,\gamma^{(n),\pi}} + V \left(X_{\tau_n}^{x,\pi}, Q_{\tau_n}^{q,\gamma^{(n),\pi}} \right) \right]$$

$$\geq \inf_{\pi \in \Pi(x)} \mathbb{E} \left[C_{\tau_n}^{x,\pi} Q_{\tau_n}^{q,\gamma^{(n),\pi}} \right]$$

$$\stackrel{(35)}{\geq} V^{\star}(x,q) - B_n,$$

where the first inequality follows from the definition of supremum, and the second inequality follows from the non-negativity of $V(\cdot, \cdot)$. By taking $n \nearrow \infty$, we obtain the desired result.

Proof of claim (vi). To simplify notation, let $\pi := \pi^{(n),\gamma}$ and $\tau := \tau_n$ temporarily. By Proposition 3, we have

$$J(\pi,\gamma;x,q) = \mathbb{E}\left[C^{0,x,\pi}_{\tau}Q^{0,q,\gamma}_{\tau} + C^{\tau,X^{0,x,\pi},\pi}_{\infty}Q^{\tau,Q^{0,q,\gamma},\gamma}_{\infty}\right]$$

(Recall that $C_{\infty}^{\tau,X_{\tau}^{0,x,\pi},\pi} = C_{\infty}^{0,x,\pi} - C_{\tau}^{0,x,\pi}$.) In the proof of claim (iv), in (30), we have shown that

$$\mathbb{E}\left[C^{0,x,\pi}_{\tau}Q^{0,q,\gamma}_{\tau}\right] \le V^{\star}(x,q) = V(x,q),$$

where the last equality follows from the claims (iv) and (v). Therefore,

$$\begin{split} J(\pi,\gamma;x,q) &\leq V(x,q) + \mathbb{E}\left[C_{\infty}^{\tau,X_{\tau}^{0,x,\pi},\pi}Q_{\infty}^{\tau,Q_{\tau}^{0,q,\gamma},\gamma}\right] \\ &\leq V(x,q) + \mathbb{E}\left[\mathrm{S}\text{-}\mathrm{CVaR}_{Q_{\tau}^{0,q,\gamma}}\left[C_{\infty}^{\tau,X_{\tau}^{0,x,\pi},\pi}\right]\right]. \end{split}$$

To obtain the desired result, it suffices to show that $\lim_{n\to\infty} \mathbb{E}\left[\text{S-CVaR}_{Q^{0,q,\gamma}_{\tau}}\left[C^{\tau,X^{0,x,\pi}_{\tau},\pi}_{\infty}\right]\right] = 0.$

Note that after time τ_n , the policy $\pi^{(n),\gamma}$ trades according to the exponential schedule (i.e., $X_t = X_{\tau_n} e^{-(t-\tau_n)/n}$), and thus $C_{\infty}^{\tau_n, X_{\tau_n}, \pi}$ is normally distributed conditional on X_{τ_n} : more specifically, as derived in (23), we have

$$\mathbb{E}\left[\left.C_{\infty}^{\tau_n, X_{\tau_n}, \pi}\right| X_{\tau_n}\right] = \frac{\eta X_{\tau_n}^2}{4n}, \quad \operatorname{Var}\left[\left.C_{\infty}^{\tau_n, X_{\tau_n}, \pi}\right| X_{\tau_n}\right] = \frac{n\sigma^2 X_{\tau_n}^2}{2}.$$

Consequently,

S-CVaR_{$$Q_{\tau_n}$$} $\left[C_{\infty}^{\tau_n, X_{\tau_n}, \pi}\right] = \frac{\eta X_{\tau_n}^2 Q_{\tau_n}}{4n} + \kappa(Q_{\tau_n}) \times \sqrt{\frac{n\sigma^2 X_{\tau_n}^2}{2}}.$

By the definition of τ_n , we have either $X_{\tau_n} = \frac{1}{n}$ or $Q_{\tau_n} \in \left\{\frac{1}{n}, 1 - \frac{1}{n}\right\}$, and thus

$$S-CVaR_{Q_{\tau_n}}\left[C_{\infty}^{\tau_n,X_{\tau_n},\pi}\right] \leq \max\left\{\frac{\eta Q_{\tau_n}}{4n^3} + \kappa(Q_{\tau_n}) \times \sqrt{\frac{\sigma^2}{2n}}, \frac{\eta X_{\tau_n}^2 Q_{\tau}}{4n} + \max\left\{\kappa(1/n),\kappa(1-1/n)\right\} \times \sqrt{\frac{n\sigma^2 X_{\tau_n}^2}{2}}\right\}$$
$$\leq \max\left\{\frac{\eta}{4n^3} + \sup_{q'}\left\{\kappa(q')\right\} \times \sqrt{\frac{\sigma^2}{2n}}, \frac{\eta x^2}{4n} + \max\left\{\kappa(1/n),\kappa(1-1/n)\right\} \times \sqrt{\frac{n\sigma^2 x^2}{2}}\right\}.$$

$$(36)$$

Since $\sup_{q'} \kappa(q') < \infty$, $\lim_{n\to\infty} \sqrt{n}\kappa(1/n) = 0$, and $\lim_{n\to\infty} \sqrt{n}\kappa(1-1/n) = 0$ (Lemma 3), the right side of (36) vanishes as $n \to \infty$. Therefore, we have $\lim_{n\to\infty} \mathbb{E}\left[\text{S-CVaR}_{Q_{\tau_n}}\left[C_{\infty}^{\tau_n,X_{\tau_n},\pi}\right]\right] = 0$ and this concludes the proof.

Proof of claim (vii). In the proof of claim (v), we have shown that

$$\mathbb{E}\left[C_{\tau_n}^{x,\pi}Q_{\tau_n}^{q,\gamma^{(n),\pi}} + V^{\star}\left(X_{\tau_n}^{x,\pi},Q_{\tau_n}^{q,\gamma^{(n),\pi}}\right)\right] \ge V^{\star}(x,q) = V(x,q),$$

where the last inequality follows from the claims (iv) and (v), and $\lim_{n\to\infty} \mathbb{E}\left[V^{\star}\left(X_{\tau_n}^{x,\pi}, Q_{\tau_n}^{q,\gamma^{(n),\pi}}\right)\right] = 0.$

Observe that we have $Q_{\infty} = Q_{\tau_n}$ since $g^{\star}(X_t, Q_t) = 0$ for all $t \ge \tau_n$ under $\gamma^{(n), \pi}$. Therefore,

$$\mathbb{E}\left[C_{\infty}^{x,\pi}Q_{\infty}^{q,\gamma^{(n),\pi}}\right] - \mathbb{E}\left[C_{\tau_{n}}^{x,\pi}Q_{\tau_{n}}^{q,\gamma^{(n),\pi}}\right] = \mathbb{E}\left[\left(C_{\infty} - C_{\tau_{n}}\right)Q_{\tau_{n}}\right]$$
$$= \mathbb{E}\left[\mathbb{E}\left(\int_{t=\tau_{n}}^{\infty}\frac{\eta}{2}\pi_{t}^{2}dt - \int_{t=\tau_{n}}^{\infty}\sigma X_{t}dW_{t}\middle|\mathcal{F}_{\tau_{n}}\right)Q_{\tau_{n}}\right]$$
$$= \mathbb{E}\left[\left(\int_{t=\tau_{n}}^{\infty}\frac{\eta}{2}\pi_{t}^{2}dt\right)Q_{\tau_{n}}^{q,\gamma^{(n)}}\right]$$
$$\geq 0.$$

Consequently,

$$J(\pi,\gamma^{(n),\pi};x,q) = \mathbb{E}\left[C_{\infty}^{x,\pi}Q_{\infty}^{q,\gamma^{(n),\pi}}\right] \ge \mathbb{E}\left[C_{\tau_n}^{x,\pi}Q_{\tau_n}^{q,\gamma^{(n),\pi}}\right] \ge V(x,q) - \mathbb{E}\left[V^{\star}\left(X_{\tau_n}^{x,\pi},Q_{\tau_n}^{q,\gamma^{(n),\pi}}\right)\right].$$

By taking $\liminf_{n\to\infty}$ on both sides, we obtain the desired result.

Proof of Theorem 4. We aim to show that $V(x,q) = V^*(x,q)$ for any $x \in \mathbb{R}$ and $q \in [0,1]$. From Theorem 7(iv) and (v), we deduce that $V(x,q) = V^*(x,q)$ for any x > 0 and $q \in (0,1)$. By symmetry, the same argument holds for any x < 0 and $q \in (0,1)$. From Proposition 8(iv) and condition (ii), we can also verify that their boundary values match, i.e., $V^*(x,q) = V(x,q) = 0$ if x = 0 or $q \in \{0,1\}$.

Proof of Theorem 6. For any $x \in (0, M]$ and $q \in (0, 1)$, the claims follow from Theorem 7(i), (ii),

(iii), (vi), (vii), and (viii). The claims also hold true for any $x \in [-M, 0)$ by symmetry.

We next examine the cases when x = 0, q = 0, or q = 1. Note that $V(0, \cdot) = V(\cdot, 0) = V(\cdot, 1) = 0$.

- Suppose x = 0. Under $\pi^{(n),\gamma}$, we have $X_t = 0$ for any $t \ge 0$ and thus $C_{\infty}^{x,\pi^{(n),\gamma}} = 0$ almost surely for any $\gamma \in \Gamma(q)$. Therefore, $J(\pi^{(n),\gamma},\gamma;x,q) = 0$, from which the claims (a) and (c) follow. Consequently, under $\gamma^{(n),\pi}$, we have $Q_t = q$ for any $t \ge 0$ since $g^{(n)}(0,Q_t) = 0$, and thus $\mathbb{E}[C_{\infty}^{x,\pi}Q_{\infty}^{q,\gamma^{(n),\pi}}] = q\mathbb{E}[C_{\infty}^{x,\pi}] \ge 0$ for any $\pi \in \Pi(x)$, from which the claim (b) follows.
- Suppose q = 0. The only admissible adversary's policy is the one satisfying that $Q_t = 0$ for all $t \ge 0$, and thus $J(\pi, \gamma; x, q) = 0$ for any $\pi \in \Pi(x)$, from which claims (a)–(c) follow.
- Suppose q = 1. Again, the only admissible adversary's policy is the one satisfying that $Q_t = 1$ for all $t \ge 0$. Then, $\pi^{(n),\gamma}$ is an exponential policy in which we have $X_t = x \exp(-t/n)$ and thus $J(\pi^{(n),\gamma},\gamma;x,q) = \mathbb{E}[C_{\infty}^{x,\pi^{(n),\gamma}}] = \frac{\eta x^2}{4n} \searrow 0$ as $n \to \infty$, from which claims (a) and (c) follow. Consequently, since $\mathbb{E}[C_{\infty}^{x,\pi}] \ge 0$ for any $\pi \in \Pi(x)$, the claim (b) immediately follows.

D.2. Other Proofs

Proof of Proposition 4. By applying Itô's formula, we obtain

$$\widehat{V}(X_t, Q_t) = \widehat{V}(x, q) - \int_{s=0}^t \widehat{V}_x(X_s, Q_s) \pi_s ds + \int_{s=0}^t \widehat{V}_q(X_s, Q_s) \gamma_s dW_s + \frac{1}{2} \int_{s=0}^t \widehat{V}_{qq}(X_s, Q_s) \gamma_s^2 ds.$$

Recall that $C_t \triangleq \int_{s=0}^t \frac{\eta}{2} \pi_s^2 ds - \int_{s=0}^t \sigma X_s dW_s$. By applying Itô's product rule, we further obtain

$$C_t Q_t = Q_0 C_0 + \int_{s=0}^t \left(\frac{\eta}{2} \pi_s^2 Q_s - \sigma X_s \gamma_s\right) ds + \int_{s=0}^t \left(-\sigma X_s Q_s + X_s \gamma_s\right) dW_s$$

Let $M_t^{(1)} \triangleq \int_{s=0}^t \widehat{V}_q(X_s, Q_s) \gamma_s dW_s$, $M_t^{(2)} \triangleq -\sigma \int_{s=0}^t X_s Q_s dW_s$, and $M_t^{(3)} \triangleq \int_{s=0}^t X_s \gamma_s dW_s$. We first verify that $(M_t^{(1)})_{t>0}$ is uniformly integrable:

$$\sup_{t} \mathbb{E}\left[|M_{t}^{(1)}|^{2}\right] \leq \mathbb{E}\left[\int_{s=0}^{\infty} \widehat{V}_{q}(X_{s}, Q_{s})^{2} \gamma_{s}^{2} ds\right] \leq \left(\sup_{x, q} \widehat{V}_{q}(x, q)\right)^{2} \mathbb{E}\left[\int_{s=0}^{\infty} \gamma_{s}^{2} ds\right] < \infty,$$

where $\sup_{x,q} \hat{V}_q(x,q) < \infty$ since we have assumed that the domain of \hat{V} is compact and \hat{V}_q is continuous, and $\mathbb{E}\left[\int_{s=0}^{\infty} \gamma_s^2 ds\right] < \infty$ since $(Q_t)_{t\geq 0}$ is a bounded martingale. Consequently,

$$\sup_{t} \mathbb{E}\left[|M_{t}^{(2)}|^{2}\right] \leq \mathbb{E}\left[\int_{s=0}^{\infty} X_{s}^{2} Q_{s}^{2} ds\right] \leq \mathbb{E}\left[\int_{s=0}^{\infty} X_{s}^{2} ds\right] < \infty,$$

where the last inequality directly follows from the definition of $\Pi(x)$. Similarly, we have $\sup_t \mathbb{E}\left[|M_t^{(3)}|^2\right] \leq \mathbb{E}\left[\int_{s=0}^{\infty} X_s^2 \gamma_s^2 ds\right] \leq M^2 \mathbb{E}\left[\int_{s=0}^{\infty} \gamma_s^2 ds\right] < \infty$. Therefore, $(M_t^{(1)})_{t\geq 0}$, $(M_t^{(2)})_{t\geq 0}$, and $(M_t^{(3)})_{t\geq 0}$ are all uniformly integrable, and hence, $\mathbb{E}[M_{\tau}^{(1)}] = \mathbb{E}[M_{\tau}^{(2)}] = \mathbb{E}[M_{\tau}^{(3)}] = 0$ for any stopping time τ .

Combining these results, we deduce that

$$\mathbb{E}\left[C_{\tau}Q_{\tau} + \widehat{V}(X_{\tau}, Q_{\tau})\right] = \widehat{V}(x, q) + \mathbb{E}\left[\int_{t=0}^{\tau} \left\{\frac{\eta}{2}\pi_t^2 Q_t - \widehat{V}_x(X_t, Q_t)\pi_t + \frac{1}{2}\widehat{V}_{qq}(X_t, Q_t)\gamma_t^2 - \sigma X_t\gamma_t\right\}dt\right].$$

We obtain the claim by rearranging terms.

Proof of Theorem 5. We prove that the function V^* , defined in (11), satisfies the conditions (i)–(v) in Theorem 4.

Verification of condition (i). Observe that $\varphi''(q) = -\frac{q}{\varphi^2(q)} < 0$ for any $q \in (0,1)$, and thus and $V_{qq}^{\star}(x,q) < 0$. The minimum/maximum in (10) are attainable, and we have

$$\min_{v \in \mathbb{R}} \left\{ \frac{\eta}{2} q v^2 - V_x^{\star}(x,q) v \right\} + \max_{w \in \mathbb{R}} \left\{ \frac{1}{2} V_{qq}^{\star}(x,q) w^2 - \sigma x w \right\} = -\frac{\left(V_x^{\star}(x,q)\right)^2}{2\eta q} - \frac{\sigma^2 x^2}{2V_{qq}^{\star}(x,q)} + \frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{\left(V_x^{\star}(x,q)\right)^2}{2\eta q} - \frac{\sigma^2 x^2}{2V_{qq}^{\star}(x,q)} + \frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{\left(V_x^{\star}(x,q)\right)^2}{2\eta q} - \frac{\sigma^2 x^2}{2V_{qq}^{\star}(x,q)} + \frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{\left(V_x^{\star}(x,q)\right)^2}{2\eta q} - \frac{\sigma^2 x^2}{2V_{qq}^{\star}(x,q)} + \frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{\left(V_x^{\star}(x,q)\right)^2}{2\eta q} - \frac{\sigma^2 x^2}{2V_{qq}^{\star}(x,q)} + \frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right) = -\frac{1}{2} \left(V_{qq}^{\star}(x,q) w^2 - \sigma x w\right)$$

Therefore, it suffices show that $V_x^{\star 2} \times V_{qq}^{\star} = -\sigma^2 \eta \times x^2 q$:

$$V_x^{\star 2} \times V_{qq}^{\star} = \left((3/4)^{\frac{2}{3}} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times \frac{4}{3} x^{\frac{1}{3}} \times \varphi(q) \right)^2 \times \left((3/4)^{\frac{2}{3}} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times |x|^{\frac{4}{3}} \times \varphi''(q) \right)$$
$$= \sigma^2 \eta \times x^2 \times \varphi^2(q) \times \varphi''(q) = -\sigma^2 \eta x^2 q.$$

Verification of conditions (ii)-(iii). These conditions can be easily verified by inspection.

Verification of condition (iv). Note that $\frac{V_x^*(x,q)}{q} = (3/4)^{-\frac{1}{3}} \times \sigma^{\frac{2}{3}} \eta^{\frac{1}{3}} \times x^{\frac{1}{3}} \times \frac{\varphi(q)}{q}$. Its continuity and monotonicity (with respect to x) can be verified immediately, and it suffices to show that $\frac{\varphi(q)}{q}$ is decreasing in q.

Observe that for any q_1, q_2 such that $0 < q_1 \le q_2 \le 1$, we have $\varphi(q_1) \ge \frac{q_1}{q_2}\varphi(q_2) + \left(1 - \frac{q_1}{q_2}\right)\varphi(0) = q_1 \times \frac{\varphi(q_2)}{q_2}$ due to its concavity, and therefore $\frac{\varphi(q_1)}{q_1} \ge \frac{\varphi(q_2)}{q_2}$. As a result, $\frac{\varphi(q)}{q}$ is monotonically decreasing in q on (0, 1).

Verification of condition (v). Note that $\frac{x}{V_{qq}^{\star}(x,q)} = (3/4)^{-\frac{2}{3}} \times \sigma^{-\frac{2}{3}} \eta^{-\frac{1}{3}} \times x^{-\frac{1}{3}} \times \frac{\varphi^2(q)}{q}$. The continuity and monotonicity (with respect to x) is trivial.

Proof of Proposition 5. Emden–Fowler equation deals with a differential equation with a form of $\frac{d^2\varphi}{dq^2} = Aq^n\varphi^m$, and the differential equation that we have corresponds to the case of A = -1, n = 1, and m = -1/2. In this case, the solutions can be written in parametric form (Polyanin and Zaitsev, 2003, 2.3.27): for constants a and b such that $A = -\frac{9}{2}(b/a)^3$,

$$q(\theta) = a\theta^{-\frac{2}{3}} \left[\left(\theta Z'(\theta) + \frac{1}{3} Z(\theta) \right)^2 - \theta^2 Z^2(\theta) \right], \ \varphi(\theta) = b\theta^{\frac{2}{3}} Z^2(\theta), \ Z(\theta) = C_1 I_{1/3}(\theta) + C_2 K_{1/3}(\theta),$$
(37)

$$q(\theta) = a\theta^{-\frac{2}{3}} \left[\left(\theta Z'(\theta) + \frac{1}{3} Z(\theta) \right)^2 + \theta^2 Z^2(\theta) \right], \ \varphi(\theta) = b\theta^{\frac{2}{3}} Z^2(\theta), \ Z(\theta) = C_1 J_{1/3}(\theta) + C_2 Y_{1/3}(\theta),$$
(38)

where $\theta \in \mathbb{R}_+$ is the parameter, and C_1 and C_2 are arbitrary constants.

Note that the expression (13) is obtained by taking $C_1 = -\frac{2}{\pi}$ and $C_2 = 0$ in (37), and the expression (14) is obtained by taking $C_1 = \sqrt{3}$ and $C_2 = -1$ in (38). Therefore it suffices to show that the curve $\{(q_L(\theta), \varphi_L(\theta))\}_{\theta \in (0,\infty]} \cup \{(q_R(\theta), \varphi_R(\theta))\}_{\theta \in (0,\overline{\theta}]}$ is a valid graph satisfying the boundary conditions, i.e.,

- (i) $\lim_{\theta \nearrow \infty} (q_L(\theta), \varphi_L(\theta)) = (0, 0).$
- (ii) $(q_R(\bar{\theta}), \varphi_R(\bar{\theta})) = (1, 0).$
- (iii) The left part and the right part meet at a point, i.e., $\lim_{\theta \searrow 0} (q_L(\theta), \varphi_L(\theta)) = \lim_{\theta \searrow 0} (q_R(\theta), \varphi_R(\theta)).$
- (iv) They have the same slope at the contact point, i.e., $\lim_{\theta \searrow 0} \frac{d\phi_L/d\theta}{dq_L/d\theta} = \lim_{\theta \searrow 0} \frac{d\phi_R/d\theta}{dq_R/d\theta}$

First, observe that $\lim_{\theta \nearrow \infty} Z_L(\theta) = -\frac{2}{\pi} \lim_{\theta \nearrow \infty} K_{1/3}(\theta) = 0$ and $\lim_{\theta \nearrow \infty} Z'_L(\theta) = \frac{1}{\pi} \lim_{\theta \nearrow \infty} (K_{-2/3}(\theta) + K_{4/3}(\theta)) = 0$. Therefore, $\lim_{\theta \nearrow \infty} q_L(\theta) = 0$ and $\lim_{\theta \nearrow \infty} \varphi_L(\theta) = 0$, which proves (i).

Next, observe that the value of $\bar{\theta}$ was chosen to satisfy $Z_R(\bar{\theta}) = 0$, and the value of a was chosen to satisfy $q_R(\bar{\theta}) = a \times \bar{\theta}^{\frac{4}{3}} \left(Z'_R(\bar{\theta}) \right)^2 = 1$, which proves (ii).

Finally, with some algebra, it can be shown that

$$\lim_{\theta \searrow 0} q_L(\theta) = \lim_{\theta \searrow 0} q_R(\theta) = \frac{2^{\frac{10}{3}}a}{3\Gamma^2(\frac{1}{3})}, \quad \lim_{\theta \searrow 0} \varphi_L(\theta) = \lim_{\theta \searrow 0} \varphi_R(\theta) = \frac{2^{\frac{8}{3}}b}{3\Gamma^2(\frac{2}{3})},$$

and

$$\lim_{N \to 0} \frac{d\phi_L(\theta)/d\theta}{dq_L(\theta)/d\theta} = \lim_{\theta \searrow 0} \frac{d\phi_L(\theta)/d\theta}{dq_L(\theta)/d\theta} = \frac{3 \times 2^{\frac{1}{3}} b\Gamma(\frac{2}{3})}{a\Gamma(\frac{1}{3})},$$

where Γ is the Gamma function. These results prove (iii) and (iv).

Some notes on the determination of constants. In the representation of the general solution, (37) and (38), there are seven constants that need to be identified: $a, b, \bar{\theta}, C_{L1}, C_{L2}, C_{R1}, and C_{R2}$. The upper limit $\bar{\theta}$ and the constant a are uniquely determined by $(C_{L1}, C_{L2}, C_{R1}, C_{R2})$ due to (ii), and the constant b is also uniquely determined by a due to the identity $(b/a)^3 = 9/2$. We can also observe that the curve $\{(q(\theta), \varphi(\theta))\}_{\theta \in \mathbb{R}_+}$ is invariant to a uniform scaling of $(C_{L1}, C_{L2}, C_{R1}, C_{R2})$, and thus we can set $C_{R4} = -1$ without loss of generality. We can further obtain a system of equations from the other conditions: $C_{L2} = 0$ from (i), $C_{L1}^2 \pi^2 = 4C_{R2}^2$ and $C_{L1}^2 \pi^2 = 3C_{R1}^2 + 2\sqrt{3}C_{R1}C_{R2} + 3C_{R2}^2$ from (ii), and $C_{R2} = -\sqrt{3}C_{R1}$ from (iv). These three equations uniquely determine the values of $C_{L1}, C_{L2}, and C_{R1}$.