Optimal Trade Execution Problem Based on Hamilton-Jacobi-Bellman Equations

Yan Zhang*, Peng Li
School of Mathematics and Statistics, North China University of Water Resources and Electric Power, Zhengzhou, Henan, China
**Corresponding Author*

Abstract: The optimal trade execution problem in financial markets is investigated within a risk-return trade-off perspective evaluated at the initial time. A linear-quadratic stochastic optimal control formulation is obtained by introducing the Lagrange multiplier method, from which the corresponding nonlinear Hamilton-Jacobi-Bellman equation is derived via the dynamic programming principle. To solve this equation, a semi-Lagrangian numerical scheme is employed. Its implementation offers considerable flexibility with respect to the particular specification of the price impact function. Under a suitable comparison principle, the theoretical convergence of this numerical scheme to the viscosity solution of the HJB equation follows. Numerical experiments are performed to recover the corresponding efficient trading frontiers and to analyze the optimal execution strategies. In particular, a sensitivity analysis is performed with respect to the market volatility parameter, for which four distinct values are considered, namely σ = 0.5, 1.0, 1.5, and 2.0. As the volatility increases, the efficient frontier shifts to the right and downward, indicating that risk rises for a given level of expected return, or alternatively, that expected return falls for a given level of risk. When the volatility tends toward zero, the efficient frontier shrinks to a single risk-free point, which helps confirm the limiting consistency of the proposed numerical algorithm. In addition, under some parameter settings, distinct trading strategies may generate nearly identical risk-return frontiers, a finding that sheds new light on the selection of practical trading strategies.

Keywords: Optimal Trade Execution; Mean-variance; HJB Equation; Semi-Lagrangian Method; Efficient Frontier

1. Introduction

Large-scale asset trading in financial markets is often accompanied by substantial market impact. For institutional investors, executing a trade rapidly over a short time interval may itself adversely move the asset price, which consequently erodes the realized trading profits. On the other hand, while spreading the execution over a longer horizon through multiple smaller orders can reduce price impact, it at the same time exposes the trader to greater price fluctuation risk over the extended trading period. Achieving a proper balance among trading speed, market impact, and price risk thus is central to the optimal trade execution problem [1,2].

The advancement of algorithmic trading technology has brought mathematically based automated trading strategies into widespread use in real-world financial markets. [3-7] Such strategies are typically guided by well-specified optimization objectives and employ systematic methodologies for controlling and decision-making across the trading process. [8,9] This paper considers a stochastic price model that includes price impact effects, in which the asset price evolves with random fluctuations while being simultaneously influenced by trading activity. Within this framework, the trade execution problem admits a formulation as a stochastic optimal control problem, whose objective is to liquidate the asset within a given horizon while achieving a trade-off between the expected return and the associated risk [10-12]. The mean-variance criterion is adopted here to formulate this problem, a choice that has gained broad acceptance in practical applications due to its clarity and ease of interpretation [13].

Notably, alternative modeling approaches based on utility functions have also been developed in the existing literature. However, optimal strategies derived from different objective functions are intrinsically distinct in nature [14]. In addition, within the mean-variance framework, whether risk and return are assessed consistently from the initial time directly impacts the optimality of the resulting strategy [15]. To address this issue, the present paper concentrates on the pre-commitment strategy, which performs a global optimization at the initial time, and examines the corresponding optimal trade execution problem on this basis.

In stochastic optimal control theory, the Hamilton-Jacobi-Bellman (HJB) equation characterizes the evolution of the value function and constitutes a key link between optimal control problems and their corresponding optimal strategies. For the optimal trade execution problem, the dynamic programming principle enables the investor's decision-making process to be reformulated as a value function problem governed by the associated HJB equation. Analysis and solution of this equation then provide a systematic means of characterizing the structural properties of the optimal trading strategy. The remainder of this paper is concerned with the specific formulation of this equation and its numerical solution methods.

The rest of this paper proceeds as follows. Section 2 establishes the stochastic control model for optimal trade execution, covering the price impact mechanism, the state evolution equations, and the transformation of the mean-variance optimization framework. Section 3 derives the corresponding Hamilton-Jacobi-Bellman equation and presents the construction of the efficient frontier via similarity scaling. Section 4 elaborates on the semi-Lagrangian discretization scheme, the solution of the local optimization problem, and the boundary treatment, along with a proof of the convergence of the numerical scheme. Building on this foundation, Section 5 analyzes the effect of volatility on the efficient frontier through numerical experiments and verifies the effectiveness of the algorithm. Section 6 concludes the paper and discusses directions for future research.

2. Optimal Control Problem Framework

This chapter constitutes the core modeling component of the paper. It takes a realistic optimal trade execution problem and, through a set of well-defined assumptions and carefully crafted mathematical transformations, reduces it to a standard stochastic optimal control framework that is amenable to dynamic programming and numerical solution.

2.1 Overview and Modeling Choices

Consider a continuous-time financial market consisting of a single risky asset, whose price process is denoted by $S(t)$, and a cash account $B(t)$ that grows at the risk-free interest rate r. Let the investor's holdings of the risky asset at time t be denoted by $\alpha(t)$. The portfolio value is then expressed as $\Pi(t) = B + \alpha S$.

It is assumed that trading takes place over a finite time horizon [0, T]. At time zero, the investor holds a given position α and is assumed to hold no cash, i.e., $B = 0$, $\alpha = \alpha_I$, and $S = S_0$. By the terminal time T, the entire position must be liquidated, so that $\alpha_T = 0$. The final cash proceeds obtained upon liquidation are denoted by B_L. The scenario $\alpha_I > 0$ corresponds to a sell execution problem, whereas $\alpha_I < 0$ corresponds to a buy execution problem. Through a unified modeling approach, both cases can be analyzed within the same framework.

2.2 Trading Mechanism and State Evolution

Let the trading rate be denoted by $v = \frac{d\alpha}{dt}$, which serves as the control variable characterizing the intensity of trade execution. The price of the risky asset experiences both random fluctuations and the impact of trading activity. Its dynamics are modeled as $\mathrm{d}S(t) = (\eta + kpv(t))S(t)\mathrm{d}t + \sigma S(t)\mathrm{d}Z(t)$ where η is the price drift, σ represents the volatility, and K_p captures the permanent price impact arising from trading.

By design, this specification ensures that the price impact is symmetric with respect to the trade direction, which rules out the possibility of earning systematic gains through reverse trading. The evolution of the cash account is governed by both interest accrual and trading costs, with its dynamics given by $\frac{\mathrm{d}B(t)}{\mathrm{d}t} = rB(t) - v(t)S(t)f(v(t))$ where the function $f(v)$ describes the temporary price impact and nonlinear transaction costs, assuming the general form $f(v) = (1 + \kappa_s \operatorname{sgn}(v))\exp(\kappa_t |v|^{\beta} \operatorname{sgn}(v))$. This functional form reflects the market characteristic that, under high-frequency or large-scale trading, the cost per unit increases with the trading rate. If any portion of the position remains unliquidated prior to the terminal time, it must be disposed of within an extremely short interval at the end.

Let the final liquidation rate be defined as $v_T = -\frac{\alpha(T^-)}{\Delta t_T}$. The final cash proceeds is then given by $B_L = B(T^-) - v_T \frac{\Delta t}{T} S(T^-) f(v_T)$ Since Δt_T is extremely small, failing to complete the transaction in time results in considerable costs, which are thus effectively discouraged within the optimal strategy.

2.3 Risk–Return Optimization Framework

For a given trading strategy $v(.)$, the resulting terminal cash proceeds B_L constitute a random variable. The mean-variance criterion is employed here to evaluate the trading outcome, with its expectation and risk characterized respectively by the mean $E_0[B_L]$ and the variance $Var_0[B_L]$.

Given a target expected return fixed at d, the optimal trading problem reduces to finding the admissible strategy that minimizes risk, namely $\min_{v(\cdot)} \mathrm{Var}_0[B_L]$ s.t. $\text{ã}_0[B_L] = d$ The solution to this problem identifies the minimal risk attainable for a given return level. The convexity of this problem enables us to introduce a Lagrange multiplier, recasting it as an unconstrained optimization problem. Equivalently, the optimal strategy follows from minimizing the following objective function: $E_0\left[(B_L - 2\gamma)^2\right]$ .The parameter γ determines the trade-off between return and risk. For different choices of γ, the resulting optimal strategies form a set of Pareto-optimal solutions in the expected return-risk space. The collection of the corresponding points defines the efficient frontier for the trade execution problem.

3. Numerical Solution of Optimal Execution Strategies and Construction of the Efficient Frontier

Proceeding from the modeling and transformation established in Chapter 2, this section derives the core partial differential equation that underlies the solution of the optimal control problem, and discusses how these equations can be used to compute the entire efficient frontier in an efficient manner. This constitutes the bridge connecting the theoretical model with the numerical algorithm.

3.1 HJB Formulation of the Optimal Control Problem

In the model presented in Section 2, the state of the system at any time is completely specified by the triple (S, B, α) . To quantify the risk associated with the trading process, a value function is introduced as $V(S,B,\alpha,\tau) = \text{ã}[B_L^2], \quad \tau = T - t$ where B_L represents the terminal liquidation value defined in Section 2.

Given the assumptions concerning the asset price dynamics and the account evolution equations, the dynamic programming principle dictates that the value function V satisfies the following Hamilton-Jacobi-Bellman equation:

$$\frac{\partial V}{\partial \tau} = LV + rBV_B + \min_{v\in Z}\left[-vSf(v)V_B + vV_\alpha + g(v)SV_S\right]$$

where the differential operator L is defined as $\mathrm{L}\,V = \frac{1}{2}\sigma^2 S^2 V_{SS} + \eta S V_S$ with terminal condition $V(S,B,\alpha,0) = B_L^2$ Solving the above HJB equation then gives the corresponding optimal trading strategy $v^*(S,B,\alpha,\tau)$ in the state space.

3.2 Computation of Expected Liquidation Value under the Optimal Strategy

Once the optimal strategy v has been determined, we proceed to define the expected liquidation value function $U(S,B,\alpha,\tau) = \text{ã}^{v^*}[B_L]$.

With the same assumptions regarding the state dynamics, the function U satisfies the following linear partial differential equation

$$\frac{\partial U}{\partial \tau} = \mathrm{L}\,U + rBU_B - v^* S f(v^*) U_B + v^* U_\alpha + g(v^*) S U_S$$

with the terminal condition $U(S,B,\alpha,0) = B_L$

Since the optimal control is already known in this equation, its numerical solution is relatively straightforward to implement, which makes it possible to obtain the expected return under the optimal strategy efficiently. To identify a point on the efficient frontier, we require not only the variance information but also the expected return associated with that strategy. Notably, once the optimal control function has been obtained from the nonlinear HJB equation, computing the expected return under this fixed strategy amounts to solving a linear PDE. The computational cost of solving the linear PDE is substantially lower than that of solving the nonlinear HJB equation.

3.3 Construction of the Efficient Frontier

In the preceding section, we have obtained both the variance function V and the expected return function U under the optimal trading strategy v^*, by solving the variance minimization problem and the expected return equation, respectively. The present section is devoted to constructing the efficient frontier in the risk-return sense from these two functions.

Consider the initial state fixed as $S = S_0$, $\alpha = \alpha_I$ and evaluate the value of the value function at backward time $\tau = T$. Define $V_0(B) = V(S_0, B, \alpha_I, T)$, $U_0(B) = U(S_0, B, \alpha_0, T)$.

By the definitions of the value function and the expected return function, we have $V_0(B) = \text{ã}^{v^*}[B_L^2]$, $U_0(B) = \text{ã}^{v^*}[B_L]$. According to the Lagrange multiplier method introduced in Section 2, the relation between the risk aversion parameter γ and the initial account variable B is given by $\gamma = -2e^{rT}B$.

Thus, solving the system of equations for different initial values of B amounts to solving the optimal execution problem under different values of the risk preference parameter γ.

Using the variable transformation relations introduced in Section 3, the expected value and the variance can be recovered as statistics of the original liquidation value B_L :

$$\text{ã}[B_L^2] \quad = V_0(B) + \gamma U_0(B) - \frac{\gamma^2}{4},$$

$$\text{ã}[B_L] \quad = U_0(B) + \frac{\gamma}{2}.$$

with γ determined by the above mapping. By evaluating the corresponding $(\text{ã}[B_L], \mathrm{Var}[B_L])$ over a discrete set of B values, one obtains a variance-minimization frontier in the expected return-risk plane. This point set is then screened according to the criterion of decreasing expected return and increasing variance, from which the efficient frontier is constructed. Every point on this frontier corresponds to a class of optimal execution strategies that are non-dominated at the given level of risk.

3.4 Similarity Scaling and Dimensionality Reduction

Under the model specification employed in this paper, the permanent impact function and the temporary impact function take linear and power-law forms, respectively, while the asset price follows a geometric Brownian motion. Under these assumptions, the value function and the expected return function satisfy the following similarity scaling relations with respect to the price and account variables:

$$V(\xi S, \xi B, \alpha, \tau) \quad = \xi^2 V(S, B, \alpha, \tau),$$

$$U(\xi S, \xi B, \alpha, \tau) \quad = \xi U(S, B, \alpha, \tau).$$

where $\xi > 0$ represents an arbitrary scaling factor.

This observation leads us to introduce a reference account level B*, such that the function values at any state can be expressed as

$$V(S, B, \alpha, \tau) \quad = \left(\frac{B^*}{B}\right)^2 V\left(\frac{B^*}{B}S, B^*, \alpha, \tau\right),$$

$$U(S, B, \alpha, \tau) \quad = \frac{B^*}{B} U\left(\frac{B^*}{B}S, B^*, \alpha, \tau\right).$$

The above result indicates that it is not required to repeatedly solve the control problem over the full three-dimensional state space (S, B, α) during numerical computation. For practical implementation, one only needs to select a finite number of representative reference account values B^* (For instance, one positive and one negative), from which the solutions at all other states can be recovered via the scaling relations.

This similarity-based dimensionality reduction approach achieves a significant reduction of computational complexity. It facilitates the efficient construction of the complete risk-return efficient frontier while maintaining numerical accuracy, thereby providing a computationally feasible foundation for the subsequent numerical experiments and parameter analysis.

4. Discretization

This chapter describes in detail the numerical discretization schemes for the HJB equation and the linear expectation equation introduced in Section 3. A semi-Lagrangian formulation is adopted to handle the convection (control) terms, while an implicit finite difference method is used for the diffusion terms. Additionally, a discrete version of the similarity scaling is introduced to achieve dimensionality reduction.

4.1 Semi-Lagrangian Discretization Scheme

With the change of variables $\tau = T - \mathrm{t}$ for backward time, the HJB equation takes the form $V_\tau = \mathcal{L}V + rBV_B + \min_{v\in\mathbb{Z}}\left[-vSf(v)V_B + vV_\alpha + g(v)SV_S\right]$, where $\mathcal{L}V = \frac{1}{2}\sigma^2 S^2 V_{SS} + \eta SV_S$

To track the value function along the characteristic curves, one introduces the Lagrangian derivative $\frac{DV}{D\tau} = \mathcal{L}V$ with the characteristics determined by the following system of ordinary differential equations: $\frac{dS}{d\tau} = -g(v)S,\ \frac{dB}{d\tau} = -(rB - vSf(v)),\ \frac{d\alpha}{d\tau} = -v.$

Let the time step be denoted by $\Delta\tau$. At a grid point on the τ^{n+1} layer, the point traced back to the τ^n layer is obtained using analytical or approximate formulas: $S_i^n = S_i\exp\left(g(v)\Delta\tau\right),\quad \alpha_k^n = \alpha_k + v\Delta\tau,$ $B_j^n = B_j e^{r\Delta\tau} - vS_i f(v)\cdot\frac{e^{r\Delta\tau} - e^{g(v)\Delta\tau}}{r - g(v)},$ Taking the limit as $\Delta\tau \to 0$, we have $\frac{e^{r\Delta\tau} - e^{g(v)\Delta\tau}}{r - g(v)} \to \Delta\tau e^{r\Delta\tau}.$ The spatial discretization operator is then defined as a finite difference approximation of the diffusion term.

A positive coefficient discretization is employed here to ensure that the corresponding discrete matrix is an M-matrix. The implicit time discretization of the HJB equation then takes the form $V_{i,j,k}^{n+1} = \min_{v\in\mathbb{Z}}\left\{\hat{V}_{i,j,k}^n + \Delta\tau(\mathcal{L}_h V^{n+1})_{i,j,k}\right\}$, where $V_{i,j,k}^n$ represents the value obtained through linear interpolation from the backtracked point. Upon rearrangement, one obtains: $\left(I - \Delta\tau\mathcal{L}_h\right)V^{n+1} = \min_v \hat{V}^n.$

4.2 Discrete Similarity Scaling and Dimensionality Reduction

The homogeneity property established in Section 3.4 makes it possible to reduce the original three-dimensional problem to two dimensions. On the discrete grid, one only needs to store two B-slices. For any query point, we have $V(S,B,\alpha,\tau) = \left(\frac{B}{B^*}\right)^2 V\left(\frac{B^*S}{B}, B^*, \alpha, \tau\right),\quad (B > 0)$ with a similar relation for U, though with a homogeneity power of one. Thus, during time stepping, the values at all nodes are computed only at the chosen reference B levels.

4.3 Discrete Solution of the Local Optimization Problem

At each grid point $\left(S_i, B_j, \alpha_k, \tau_{n+1}\right)$, a one-dimensional optimization problem of the form $\min_{v\in\mathbb{Z}} \hat{V}_{i,j,k}^n(v)$ must be solved. Since the value function can possess multiple local minima — especially in flat regions — gradient-based descent methods are unsuitable for this task. A control discretization approach is employed: the continuous control set $[v_{\min}, v_{\max}]$ is discretized into a discrete grid $\hat{Z}$, and then a linear search over (i.e., a complete search over the discrete set) is performed to compute the corresponding backtracking interpolant $\hat{V}^n$ and select the minimum value. As the grid is refined, the discrete control set should also be refined accordingly to guarantee consistency.

4.4 Boundary Treatment

At the boundaries of the holding domain, the control set is restricted depending on the trade direction. For instance, in a sell execution problem, only controls with $v \le 0$ are permitted at $\alpha_{\max}$, while only $v \ge 0$ are allowed at $\alpha_{\min}$. These natural boundary conditions are implemented directly by modifying the admissible control set Z, with no need for additional equations.

At $S = 0$, the diffusive term drops out and the equation becomes purely convective, which is solved directly. At the far-field boundary $S \to \infty$, an asymptotic convenience condition is adopted by assuming $V_{SS}, V_S \to 0$, which simplifies the equation to $V_\tau = rBV_B + \min_v[-vSf(v)V_B + vV_\alpha].$

In practice, $S_{\max}$ is chosen large enough so that the truncation error is kept small.

4.5 Convergence of the Numerical Scheme

If the HJB equation satisfies a strong comparison principle and the discrete scheme is stable, monotone, and consistent, then the scheme converges to the viscosity solution of the original problem. A detailed proof of this convergence can be found in the references cited therein.

4.6 Numerical Results

Market volatility is one of the key factors affecting the optimal trade execution strategy. To explore its impact on the mean-variance efficient frontier, we keep all other parameters consistent with the short-term liquidation example documented in Forsyth (2011), using the same parameter settings as in that study (specifically, a trading horizon of $T = 1/250$, a permanent impact function $K_P = 0$, a temporary impact coefficient $K_L = 2 \times 10^{-6}$, an impact exponent $\beta = 1$, an initial price $S_0 = 100$, and an initial holding $\alpha_I = 1$), while taking volatility values of $\sigma = 0.5, 1.0, 1.5, 2.0$. The semi-Lagrangian numerical scheme described above is subsequently applied to solve the corresponding HJB equations, from which the efficient frontier is constructed for each volatility scenario. Figure 1 presents the efficient frontiers obtained under different volatility levels.

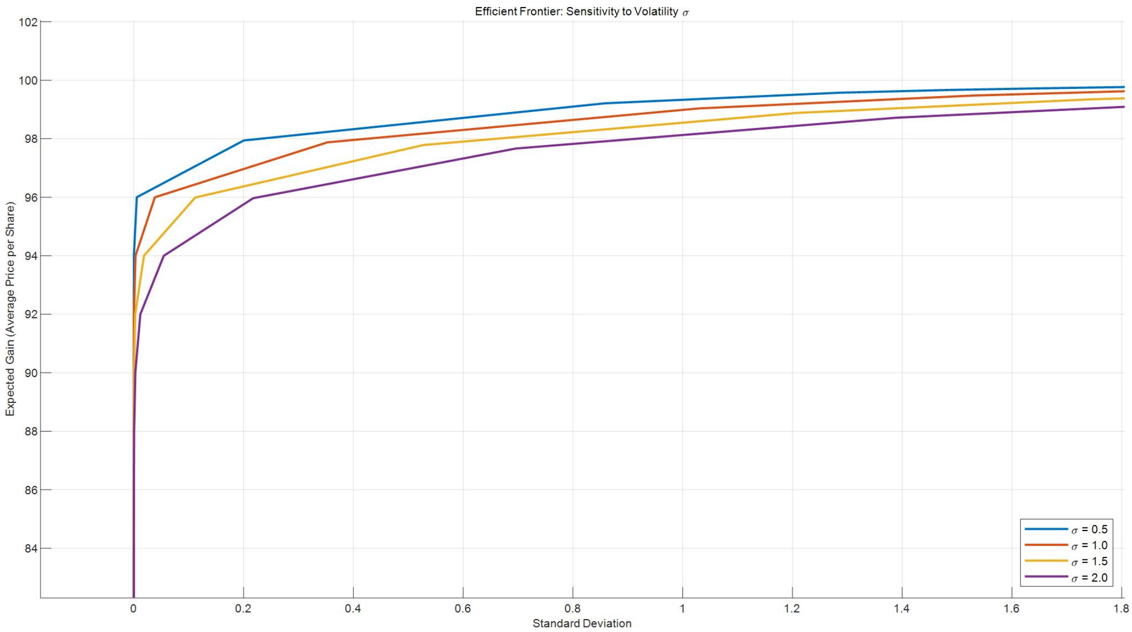

Figure 1. Efficient Frontiers under Different Volatility Levels

The horizontal axis represents the standard deviation (used here as the risk metric), while the vertical axis gives the expected per-share proceeds, $\tilde{a}[B_L]$ defined as the expected liquidation value divided by the initial number of shares. In the baseline case with σ = 1.0, the efficient frontier occupies the upper-left portion of the plot; for instance, an expected return of 99.29 corresponds to a standard deviation of approximately 0.75, which is qualitatively consistent with the results reported in Forsyth (2011). As the volatility increases to $\sigma = 1.5$ and $\sigma = 2.0$, the entire frontier moves to the right and downward: at a given level of expected return, the standard deviation rises considerably, and conversely, to achieve a given level of risk, the attainable expected return drops noticeably. When the volatility is reduced to $\sigma = 0.5$, the frontier moves to the upper-left, indicating that lower market uncertainty allows the trader to obtain higher expected returns with less risk.

This phenomenon follows from the structure of the HJB equation. The volatility parameter σ appears through the diffusion term $\frac{1}{2}\sigma^2 S^2 V_{SS}$, which influences the value function V. An increase in σ magnifies the randomness of the price process, producing a larger variance in the final proceeds under the same trading strategy. Within the mean-variance optimization framework, this is equivalent to requiring the trader to reduce positions more aggressively in order to protect against price risk, which in turn leads the trader to incur greater temporary impact costs. The net effect is reflected as a degradation of the efficient frontier. Importantly, when $\sigma = 0$, the price process becomes purely deterministic. If the permanent impact function $\kappa_P = 0$, the trader can liquidate the asset at any rate without incurring any cost. In this limiting case, the efficient frontier degenerates to a single point with expected return of 100 and standard deviation of zero. Our numerical scheme can approximate this limiting behavior as σ approaches zero, which further confirms the validity of the algorithm.

In summary, volatility plays a significant role in shaping the efficient frontier. In practical trading, accurate estimation of volatility matters greatly. Underestimating volatility may lead the trader into unexpectedly high levels of risk, whereas overestimating it may result in excessively prudent strategies that sacrifice too much expected return. It is therefore advisable, in practice, to incorporate volatility forecasting models for adaptive updating of the optimal strategy.

5. Conclusions

This paper presents a systematic treatment of the optimal trade execution problem under the pre-commitment strategy. A stochastic optimal control model based on the mean-variance criterion is established, which is then transformed into a linear-quadratic problem via the Lagrange multiplier method. The dynamic programming principle is applied to derive the nonlinear HJB equation, for which a semi-Lagrangian discretization scheme is developed. In addition, similarity scaling is employed to reduce the three-dimensional state space to two dimensions, yielding a substantial improvement in computational efficiency.

In numerical experiments, we first validate our algorithm against the short-term liquidation example from Forsyth (2011), using the parameter settings reported in Forsyth (2011). This comparison establishes the convergence of the scheme and verifies the fundamental shape of the resulting efficient frontier. We then examine the effect of market volatility σ on the efficient frontier, by computing the frontier for four distinct volatility levels: $\sigma = 0.5, 1.0, 1.5, 2.0$. It is observed that higher volatility shifts the frontier to the right and downward (meaning greater risk for a given expected return, or lower expected return for a given level of risk). This behavior is in line with the dominant role of the diffusion term in the HJB equation. Moreover, as σ tends to zero, the efficient frontier degenerates to a risk-free point with expected return of 100 and zero standard deviation, which further validates the numerical method. Several directions merit attention for future work, including time-consistent strategies, the incorporation of permanent impact, and the use of alternative risk measures for comparative analysis.

References

[1] Almgren, R., & Chriss, N. (2000). Optimal execution of portfolio transactions. Journal of Risk, 3(Winter), 5-39.

[2] Bertsimas, D., & Lo, A. (1998). Optimal control of execution costs. Journal of Financial Markets, 1, 1-50.

[3] Zhou, X. Y., & Li, D. (2000). Continuous time mean variance portfolio selection: A stochastic LQ framework. Applied Mathematics and Optimization, 42, 19-33.

[4] Na, A. S., & Wan, J. W. L. (2024). Residual U-net with Self-Attention to Solve Multi-Agent Time-Consistent Optimal Trade Execution. arXiv preprint arXiv:2312.09353.

[5] Barles, G., & Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear equations. Asymptotic Analysis, 4, 271-283.

[6] Barles, G. (1997). Convergence of numerical schemes for degenerate parabolic equations arising in finance. In Numerical Methods in Finance. Cambridge University Press.

[7] Palmari, G. (2024). Optimal execution with deterministically time varying liquidity: well posedness and price manipulation. arXiv:2410.04867.

[8] Bank, P., Cartea, Á., & Körber, L. (2025). Optimal Execution and Speculation with Trade Signals. Fields Institute Talk, May 2025.

[9] Huberman, G., & Stanzl, W. (2004). Price manipulation and quasi-arbitrage. Econometrica, 72, 1247-1275.

[10] Anonymous. (2024). Reducing Obizhaeva–Wang-type trade execution problems to LQ stochastic control problems. Finance and Stochastics, 28, 813–863.

[11] Schied, A., & Schoeneborn, T. (2009). Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. Finance and Stochastics, 13, 181-204.

[12] Vath, V. L., Mnif, M., & Pham, H. (2007). A model of optimal portfolio selection under liquidity risk and price impact. Finance and Stochastics, 11, 51-90.

[13] Lorenz, J., & Almgren, R. (2007). Adaptive arrival price. In Algorithmic Trading III: Precision, Control, Execution. Institutional

Investor Journals.
[14] Engle, R., & Ferstenberg, R. (2007). Execution risk. Journal of Trading, 2(2), 10-20.
[15] Forsyth, P. A., Kennedy, J. S., Tse, S. T., & Windcliff, H. (2009). Optimal trade execution: a mean quadratic variation approach. Quantitative Finance (submitted).
