

Hedge and Speculate: Replicating Option Payoffs with Limit and Market Orders*

Álvaro Cartea[†], Luhui Gan[‡], and Sebastian Jaimungal[‡]

Abstract. We consider an agent who takes a short position in a contingent claim and employs limit orders (LOs) and market orders (MOs) to trade in the underlying asset to maximize expected utility of terminal wealth. The agent solves a combined optimal stopping and control problem where trading has frictions: MOs (executed by the agent and other traders) have permanent price impact and pay exchange fees, and LOs earn the spread (relative to the midprice of the asset) and pay no exchange fees. We show how the agent replicates the payoff of the claim and also speculates in the asset to maximize expected utility of terminal wealth. In the strategy, MOs are used to keep the inventory on target, to replicate the payoff, and LOs are employed to build the inventory at favorable prices and boost expected terminal wealth by executing roundtrip trades that earn the spread. We calibrate the model to the E-mini contract that tracks the S&P 500 index, provide numerical examples of the performance of the strategy, and prove that our scheme converges to the viscosity solution of the dynamic programming equation.

Key words. optimal hedging, market impact, impulse control, algorithmic trading, high-frequency trading

AMS subject classifications. 91G80, 49L20, 60G55, 60G99

DOI. 10.1137/18M1192706

1. Introduction. In the classical approach to pricing and hedging contingent claims, a widespread approach is to use delta hedging. The objective of this hedge-portfolio is to replicate the payoff of the claim, which typically requires dynamically rebalancing the portfolio, and only under ideal conditions is the claim perfectly replicated.

The theory of dynamic replication of payoffs relies on a number of assumptions including the absence of price impact stemming from taking positions in the underlying and other trading costs. Traditional models assume that the buying and selling activity of the hedger does not affect the dynamics of the asset. This assumption contradicts the empirical finding that trades may have an impact on the prices the hedger receives, i.e., a temporary impact on execution prices, and may have a permanent impact on the price of the asset.

Furthermore, assuming the hedger can take positions in the underlying asset at no cost other than the price of the asset is also incorrect in the vast majority of cases. A clear example where this assumption fails is in order-driven equity markets. In these markets, trading with aggressive market orders (MOs) is expensive because they incur exchange fees and the assets normally trade with a spread about their midprice. That is, the price paid for one share is

*Received by the editors June 7, 2018; accepted for publication (in revised form) June 14, 2019; published electronically September 17, 2019.

<https://doi.org/10.1137/18M1192706>

Funding: The work of the third author was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), funding references RGPIN-2018-05705 and RGPAS-2018-522715.

[†]Mathematics, University of Oxford, Oxford, OX2 9HB, UK (Alvaro.Cartea@maths.ox.ac.uk).

[‡]Department of Statistical Sciences, University of Toronto, Toronto M5T 1P5, ON, Canada (ganlh6@gmail.com, sebastian.jaimungal@utoronto.ca).

more than the cash received from selling one share when MOs are used.

Earlier models of pricing and replication of claims have been designed without considering how assets are traded in modern electronic markets. In this paper we show how an agent who shorts a European-style claim replicates its terminal payoff while taking speculative positions in the underlying asset to maximize expected utility of terminal wealth.

We consider an agent who solves a combined optimal stopping and control problem where the underlying asset is traded in an order-driven market. The agent employs a combination of limit orders (LOs) and MOs during the life of the strategy; see Guéant, Lehalle, and Fernandez-Tapia (2012) and Cartea and Jaimungal (2015a). LOs are executed at better prices than the quoted midprice: buy LOs are executed at the midprice minus a premium, and sell LOs are executed at the midprice plus a premium. The premium earned by the LOs is the depth at which the agent posts sell and buy LOs. Thus, LOs are desirable in the agent's strategy, but their execution is uncertain. If the agent requires certainty in execution, she employs MOs which are more costly because they pay exchange fees and are filled at prices worse than the quoted midprice of the asset; see Cartea and Jaimungal (2015a).

We show how the agent aims to replicate the terminal payoff of the contingent claim and simultaneously speculates in the asset to maximize expected utility of terminal wealth. In the optimal strategy, MOs are used to keep the inventory on target to replicate the payoff, and LOs are employed to (i) build the inventory at favorable prices and without price impact, and (ii) boost expected terminal wealth by executing roundtrip trades that earn the spread.

Our model incorporates a number of important features such as permanent price impact and adverse selection costs. Orderflow from MOs typically affects the midprice of the underlying asset; see, for instance, Cartea and Jaimungal (2016). When MOs arrive in the exchange, the midprice of the asset jumps by a random amount in the direction of the trade. Under the optimal strategy, the agent accounts for this price impact when she posts LOs to protect herself from adverse selection costs. These costs arise when a market participant sends a MO that is filled by the agent's LO, and immediately after the midprice jumps in the same direction as the MO. We find that the volume of shares the agent is willing to post in the limit order book (LOB) depends on the exposure to adverse selection costs.

Our work is closest to that of Guéant and Pu (2015), who consider the pricing and hedging of European call options in illiquid markets (see also Guéant (2016)). Guéant and Pu assume a Bachelier model for the midprice dynamics of the underlying asset of the options. In their model the hedger only executes MOs, and does so by choosing a continuous trading rate. Illiquidity in the market precludes frictionless hedging of the options. Thus, the authors assume that MOs executed by the hedger have (linear in the speed of trading) permanent impact on the price of the underlying and incur execution costs (also a function of the speed of trading MOs).

We depart from Guéant and Pu in a number of aspects. The main difference is that in our model the hedger employs both LOs and MOs in the trading strategy. LOs are continuously updated in the market, and MOs are modeled as an impulse control. In practice, hedgers draw on both types of order to replicate options at the most favorable prices, and more so in illiquid markets where MOs affect prices, may walk the LOB, and pay exchange fees. On the other hand, LOs are filled at better prices than the midprice of the underlying—LOs earn at least half the bid-ask spread and pay no fees. In our model, MOs from all market participants (not

only the hedger's) have a permanent impact on the price of the underlying asset. Conditional on a MO arriving in the exchange, the underlying price jumps in the direction of the trade. Finally, in our model the hedger not only targets an amount of shares or cash to settle the options, but also executes speculative trades to maximize expected utility of wealth.

In recent years, advances in optimal execution have inspired a new stream of literature. Rogers and Singh (2010) formulated a stochastic control problem where an agent minimizes the integrated squared error between her inventory process and the option delta. A similar formulation can be found in Naujokat and Westray (2011) and Bank, Soner, and Voß (2017), where the former allows the agent to post passive orders and the latter allows for non-Markovian controls. Almgren and Li (2016) consider intraday hedging of an option for an agent whose trades have price impact. The authors formulate a general model and provide an explicit solution when the option's gamma is approximately constant and the agent's trades have permanent and temporary price impact. The optimal hedge trades smoothly towards the classical Black–Scholes delta, where the trading intensity is proportional to the instantaneous mishedge and is inversely proportional to illiquidity.

Moreover, in the extant literature, several attempts have been made to incorporate transaction costs when trading assets. Leland (1985) proposes a discrete-time hedging strategy as an alternative to the Black–Scholes model when the transaction costs are proportional to the volumes traded in a given time interval. Soner, Shreve, and Cvitanić (1995) show how to hedge a European call option in a Black–Scholes model with proportional transaction costs; see also Mohamed (1994) and Marco and Antonio (1994). Barles and Soner (1998) use indifference pricing to derive pricing formulae for European call options under proportional transaction costs and also assume that transaction costs are proportional to the volumes traded. Other examples include Boyle and Vorst (1992), Davis, Panas, and Zariphopoulou (1993), and Cvitanić and Karatzas (1996).

Another stream of literature examines the effects of market impact on option pricing and volatility. Cetin, Jarrow, and Protter (2004) extended the Black–Scholes framework by adding a stochastic supply curve for the underlying asset. Their model was later extended by Roch (2011) to incorporate permanent price impact. A similar set-up was also proposed by Bank and Baum (2004). Lions and Lasry (2007) show how the trading activity of a large investor affects the volatility of the asset.

The remainder of the paper is organized as follows. Section 2 describes the model components and the optimal control problem solved by the agent, and in section 3 we derive its associated dynamic programming equation (DPE). Section 4 presents a numerical example of the strategy's performance when the agent sells a call option written on the S&P E-mini Futures. Section 5 contains proofs for the convergence of the numerical scheme we adopt. Section 6 offers conclusions, and we collect other proofs in Appendix A.

2. Model. Fix a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq \hat{T}})$, where \mathbb{F} is the natural filtration generated by the collection of observable processes we describe below. At time $t = 0$ the agent sells an option written on a traded asset (the case where the agent buys an option is similar). The option expires at time \hat{T} and can only be exercised by the buyer at expiry, i.e., it is a European-style claim, and we assume that the option is physically settled. The agent is risk-averse and trades the underlying asset of the option over the time

window $[0, T]$ to maximize her expected utility at time T , where $T \leq \hat{T}$.

The asset is traded in an order-driven exchange that matches orders from participants who provide liquidity using LOs with those who take liquidity by executing MOs. The agent employs four types of orders to trade the underlying asset: market buy order (MBO), market sell order (MSO), limit buy order (LBO), and limit sell order (LSO).

Below we propose the dynamics of the midprice of the traded asset and describe the dynamics of trading activity, including that of LOs and MOs.

Bid-ask spread and tick size. We denote by $S = (S_t)_{t \in [0, T]}$ the midprice of the underlying asset. Midprices take on values on a grid determined by the tick size in the exchange. Here we denote the tick size by $\sigma > 0$ and assume that the asset trades with a bid-ask spread fixed at σ .

Agent's limit orders. The agent's LOs are always posted at the best bid or best ask, i.e., LSOs are posted at the price level $S_t + \sigma/2$, and LBOs at the price level $S_t - \sigma/2$. The volume of the LBOs posted by the agent is denoted by the process $\ell^+ = (\ell_t^+)_{t \in [0, T]}$, and the volume of the LSOs is denoted by the process $\ell^- = (\ell_t^-)_{t \in [0, T]}$.

Agent's market orders. We denote by $\tau^{0+} = (\tau_k^{0+})_{k \geq 1}$ the times at which the agent sends MBOs, and we denote their volume by $m^+ = (m_k^+)_{k \geq 1}$. Similarly, $\tau^{0-} = (\tau_k^{0-})_{k \geq 1}$ are the times at which the agent sends MSOs, and we denote their volume by $m^- = (m_k^-)_{k \geq 1}$. Here, $\tau^{0\pm}$ are sequences of increasing stopping times and m_k^+ (resp., m_k^-) are $\mathcal{F}_{\tau_k^{0+}}$ -measurable (resp., $\mathcal{F}_{\tau_k^{0-}}$ -measurable) strictly positive integer-valued random variables. In subsection 2.1 we specify the values that m_k^\pm takes on.

We define the processes $M^{0+} = (M_t^{0+})_{t \in [0, T]}$ and $M^{0-} = (M_t^{0-})_{t \in [0, T]}$ via

$$M_t^{0+} = \sum_{k=1}^{\infty} \mathbb{1}_{\tau_k^{0+} \leq t} \quad \text{and} \quad M_t^{0-} = \sum_{k=1}^{\infty} \mathbb{1}_{\tau_k^{0-} \leq t}$$

as the total number of MBOs and MSOs that the agent has executed up to time t , respectively. We also define the processes $\mathbb{M}^{0+} = (\mathbb{M}_t^{0+})_{t \in [0, T]}$ and $\mathbb{M}^{0-} = (\mathbb{M}_t^{0-})_{t \in [0, T]}$ via

$$\mathbb{M}_t^{0+} = \sum_{k=1}^{\infty} m_k^+ \mathbb{1}_{\tau_k^{0+} \leq t} \quad \text{and} \quad \mathbb{M}_t^{0-} = \sum_{k=1}^{\infty} m_k^- \mathbb{1}_{\tau_k^{0-} \leq t}$$

as the accumulated volume of MBOs and MSOs the agent has executed up to time t , respectively.

Other market participants' market orders. We assume that there is enough liquidity at the best prices to fill incoming MOs. Thus, all MSOs are executed at the price $S_t - \Upsilon$ and all MBOs are executed at $S_t + \Upsilon$, where Υ consists of the half-spread $\sigma/2$ (i.e., half a tick) plus a liquidity taking fee.

MBOs and MSOs sent by other market participants arrive in the exchange at the arrival times of a homogeneous Poisson process. Buy orders arrive with intensity λ^+ and sell orders with intensity λ^- . The counting processes for buy and sell MOs are denoted by $M^+ = (M_t^+)_{t \in [0, T]}$ and $M^- = (M_t^-)_{t \in [0, T]}$, respectively, with associated event times denoted by $\tau^+ = (\tau_k^+)_{k \geq 1}$ and $\tau^- = (\tau_k^-)_{k \geq 1}$, respectively. We assume the agent's short position in the

contingent claim is relatively small compared to the trading volume of the underlying asset. Hence it is sufficient for her to post LOs of relatively small size, and we can assume that her LOs are always executed in full.

Conditional on the arrival of the k th MO, the agent's LO is filled according to a Bernoulli random variable ζ_k with probability of success ρ . The random variables ζ_k are independent of all other random variables and independent and identically distributed (i.i.d.). We denote by $L^+ = (L_t^+)_{t \in [0, T]}$ and $L^- = (L_t^-)_{t \in [0, T]}$ the counting processes for the agent's filled LBOs and LSOs, respectively. Our assumption on ζ_k represents uncertainty in LO execution. For example, under the time-priority rule, each level of the LOB is executed with a first-in-first-out rule. Thus, the fill probability of a LO depends on its price, position in the queue, and size of the MO. This randomness is captured by ζ_k .

Midprice dynamics. Changes in the midprice of the asset reflect the arrival of information that affects the value of the asset. We assume that buying and selling pressure from MOs has an impact on the midprice, and other information and news also affect the midprice of the asset; see, for example, [Cartea and Jaimungal \(2015b\)](#).

Changes in the midprice are given by

$$(1) \quad dS_t = \sigma (dP_t^+ - dP_t^-),$$

where $P_t^+ = (P_t^+)_{t \in [0, T]}$ is the upward pressure on the midprice exerted by market participants as a result of MBOs and reshuffling of LSOs. Similarly, $P_t^- = (P_t^-)_{t \in [0, T]}$ represents the downward pressure on the midprice as a result of MSOs and changes in LOs. For more general price impact functions, see [Almgren \(2003\)](#), [Bouchard, Loeper, and Zou \(2017\)](#), and [Cetin, Soner, and Touzi \(2010\)](#).

We assume that P_t^\pm have the following form:

$$(2) \quad P_t^+ = Z_t^+ + \sum_{k=1}^{M_t^+} \xi_k^+ + \sum_{k=1}^{M_t^{0+}} \xi_k^{0+},$$

$$(3) \quad P_t^- = Z_t^- + \sum_{k=1}^{M_t^-} \xi_k^- + \sum_{k=1}^{M_t^{0-}} \xi_k^{0-}.$$

We explain the various objects in the three terms on the right-hand side of (2).

The first term represents exogenous changes in the midprice due to reshuffling of LSOs as a result of information and news that are impounded in the midprice of the asset. Specifically, $Z^\pm = (Z_t^\pm)_{t \in [0, T]}$ are independent Poisson processes with intensity θ , and are independent of all other processes and random variables.

The second term, $\sum_{k=1}^{M_t^+} \xi_k^+$, is a compound Poisson process that represents the price impact of MBOs sent by all market participants (excluding the agent). Moreover, $\xi^+ = (\xi_k^+)_{k \geq 1}$ is a sequence of i.i.d. Bernoulli random variables with constant success probability α . Recall that M_t^+ is a homogeneous Poisson process with arrival rate λ^+ ; thus, every time a MBO arrives, the midprice jumps by ξ^+ .

The third term, $\sum_{k=1}^{M_t^{0+}} \xi_k^{0+}$, represents the price impact of the agent's MBOs. Upon the agent sending a buy MO, the midprice jumps up by the random amount ξ^{0+} . These

midprice innovations are characterized by $\xi^{0+} = (\xi_k^{0+})_{k \geq 1}$, which is a sequence of $\mathcal{F}_{\tau_k^{0+}}$ -measurable Bernoulli random variables. Each ξ_k^{0+} has probability of success $\beta^+(m_k^+)$, where $\beta^+ : \mathbb{Z}_+ \mapsto [0, 1]$ is a deterministic function. We could assume that the price impact of other market participants' MOs is a function of the volume of the order.

The three terms on the right-hand side of (3) have a similar interpretation. The first term represents changes in the midprice due to reshuffling of LOs in the sell side of the book. The second term represents the price impact of market participants' MSOs, where $\xi^- = (\xi_k^\pm)_{k \geq 1}$ is a sequence of i.i.d. Bernoulli random variables with constant probability of success α . And the third term is the price impact of the agent's MSOs, where $(\xi_k^{0-})_{k \geq 1}$ is a sequence of $\mathcal{F}_{\tau_k^{0-}}$ -measurable Bernoulli random variables. Each ξ_k^{0-} has probability of success $\beta^-(m_k^-)$.

2.1. The agent's optimization problem. The agent's objective is to maximize expected utility of wealth at time T . Recall that the agent sells a contingent claim expiring at $\hat{T} \geq T$, trades the underlying asset over the time window $[0, T]$, and the claim is physically settled. We denote by $Q = (Q_t)_{t \in [0, T]}$ the agent's inventory, which takes on integer values, and we restrict the strategies so that $\underline{q} \leq Q_t \leq \bar{q}$, where \bar{q} and \underline{q} are the maximum and the minimum inventory level the agent is willing to hold. As a result of the agent's LO and MO activity, the inventory Q_t process satisfies the stochastic differential equation (SDE)

$$(4) \quad dQ_t = \ell_{t-}^- dL_t^- - \ell_{t-}^+ dL_t^+ + d\mathbb{M}_t^{0+} - d\mathbb{M}_t^{0-},$$

and the volumes of the agent's LOs and MOs are such that

$$(5a) \quad \ell_t^+ \in \{0, 1, \dots, \bar{q} - Q_t\},$$

$$(5b) \quad \ell_t^- \in \{0, 1, \dots, Q_t - \underline{q}\},$$

$$(5c) \quad m_k^+ \in \{0, 1, \dots, \bar{q} - Q_{\tau_k^+}\},$$

$$(5d) \quad m_k^- \in \{0, 1, \dots, Q_{\tau_k^-} - \underline{q}\}$$

for $0 \leq t \leq T$ and $k \geq 1$, so the strategy obeys the inventory constraints.

The agent's cash process is represented by $X = (X_t)_{t \in [0, T]}$ and satisfies the SDE

$$(6) \quad dX_t = -\left(S_{t-} - \frac{\sigma}{2}\right) \ell_{t-}^- dL_t^- + \left(S_{t-} + \frac{\sigma}{2}\right) \ell_{t-}^+ dL_t^+ + (S_{t-} - \Upsilon) d\mathbb{M}_t^{0-} - (S_{t-} + \Upsilon) d\mathbb{M}_t^{0+}.$$

The first two terms in the right-hand side of (6) represent changes in the agent's cash position due to filled LOs. The last two terms represent changes in cash due to the agent's executed MOs. Note that the cash process does not take into account the value of the agent's inventory.

We denote by $G : (\sigma, \mathbb{Z}) \mapsto \mathbb{R}$ the value function of the contingent claim. At time $t = 0$ the agent receives cash for the contingent claim, and at expiry the value of the claim is $G(S_T)$. We assume the agent's initial wealth X_0 consists of the premium she obtains from selling the contingent claim. Below, in subsection 3.1 we show how the agent employs indifference pricing to calculate the premium charged for the claim. In the remainder of the analysis, for simplicity we assume the expiry of the contingent claim and the agent's horizon are the same, i.e., $\hat{T} = T$.

The agent's preferences are given by an exponential utility function of wealth, and her value function is

$$(7) \quad H(t, x, s, q) = \sup_{(\ell^\pm, \tau^{0\pm}, m^\pm) \in \mathcal{A}} \mathbb{E}_{t,x,s,q} \left[-e^{-\gamma (X_T + S_T Q_T - G(S_T) - C(Q_T, S_T))} \right],$$

where $\gamma > 0$ is the agent's risk-aversion parameter and $\mathbb{E}_{t,x,s,q}[\cdot]$ denotes the \mathbb{P} expectation conditional on the initial condition $X_{t-} = x$, $Q_{t-} = q$, and $S_{t-} = s$. Here, \mathcal{A} denotes the set of admissible strategies in which $\ell^\pm = (\ell_t^\pm)_{0 \leq t \leq T}$ are \mathcal{F} -adapted processes, $\tau^{0\pm} = (\tau_k^{0\pm})_{k \geq 1}$ are sequences of \mathcal{F} -stopping times, and $m^\pm = (m_k^\pm)_{k \geq 1}$ are sequences of $\mathcal{F}_{\tau_k^{0\pm}}$ -measurable random variables. Moreover, ℓ^\pm and m^\pm satisfy (5) pointwise. The function $C : \mathbb{Z} \times (\sigma \mathbb{Z}) \mapsto \mathbb{R}$ represents other costs the agent incurs at the terminal date to account for any additional terminal transaction costs. For example, if the option is physically settled, the agent may have to adjust her terminal inventory position to settle the option or unwind inventory that is not required (see the example in (14a) which corresponds to the case of physically settling a call option).

Table 1 summarizes the parameters of the model. It also shows parameter values that we employ below in section 4 when we illustrate the performance of the model.

Table 1

Model parameters. The last column shows parameter values that we use below when we illustrate features and performance of the strategy.

Parameter	Description	Value
θ	rate of exogenous midprice changes	107.7/hour
λ^\pm	rate of arrival of other participants' MOs (+: buy, -: sell)	4492.4/hour
ρ	probability of agent's LO being filled	0.0171
α	probability of price impact of other participants' MOs	0.094
T	time horizon of optimization problem to maturity	1 hour
\hat{T}	expiry of the option	1 hour
σ	bid-ask spread (tick size)	0.25
Υ	half-spread	0.125
	(assuming no MO fee)	
$\beta^\pm(m)$	agent's price impact	$0.01 \times m$
γ	agent's risk-aversion parameter	0.1
q_0	initial inventory	0
\bar{q}	maximum inventory	10
\underline{q}	minimum inventory	0

3. The dynamic programming equations. To solve the combined optimal stopping and control problem, we employ the dynamic programming principle. Standard results imply that the value function (7) is the unique viscosity solution to the quasi-variational inequality (QVI)

$$(8) \quad \begin{aligned} & \max \left\{ (\partial_t + \mathcal{L})H(t, x, s, q) \right. \\ & \quad + \max_{\ell^+ \in \{0, \dots, \bar{q}-q\}} \left\{ \lambda^- \mathbb{E} \left[H \left(t, x - \ell^+ \left(s - \frac{\sigma}{2} \right) \zeta, s - \sigma \xi^-, q + \ell^+ \zeta \right) - H(t, x, s, q) \right] \right\} \\ & \quad + \max_{\ell^- \in \{0, \dots, q-\underline{q}\}} \left\{ \lambda^+ \mathbb{E} \left[H \left(t, x + \ell^- \left(s + \frac{\sigma}{2} \right) \zeta, s + \sigma \xi^+, q - \ell^- \zeta \right) - H(t, x, s, q) \right] \right\} \Big\}; \end{aligned}$$

$$\begin{aligned} & \max_{m^+ \in \{1, \dots, \bar{q}-q\}} \left\{ \mathbb{E} \left[H(t, x - m^+ (s + \Upsilon), s + \sigma \xi^{0+}, q + m^+) - H(t, x, s, q) \right] \right\}; \\ & \max_{m^- \in \{1, \dots, q-\underline{q}\}} \left\{ \mathbb{E} \left[H(t, x + m^- (s - \Upsilon), s - \sigma \xi^{0-}, q - m^-) - H(t, x, s, q) \right] \right\} = 0, \end{aligned}$$

subject to the terminal condition

$$H(t, x, s, q) = -e^{-\gamma(x+sq-G(s)-C(s,q))}.$$

The infinitesimal generator \mathcal{L} in (8) acts on a smooth function $g(t, x, s, q)$ as follows:

$$\mathcal{L}g(t, x, s, q) = \theta (g(t, x, s + \sigma, q) + g(t, x, s - \sigma, q) - 2g(t, x, s, q)).$$

The expectation operators in the second and third lines of the QVI are with respect to the random variables ζ, ξ , and the expectation operator in the last two lines of the QVI is with respect to the random variables $\xi^{0\pm}$. Recall that $\zeta, \xi, \xi^{0\pm}$ are independent Bernoulli random variables with probability of success ρ, α , and $\beta^\pm(m^\pm)$, respectively.

The terms in the QVI have the following interpretation. The term on the right-hand side of the first line of (8) represents the changes in the value function due to time and exogenous changes in the midprice that result from reshuffling of LOs.

The second line represents the changes in the value function due to other market participants' incoming MSOs, which are filled by the agent's LBOs with probability ρ , and which generate price impact on the midprice (represented by ξ).

Similarly, the third line represents the changes in the value function due to other agents' MBOs. Finally, the fourth and fifth lines account for changes in the value function due to the agent's MOs. Recall that Υ consists of the half-spread $\sigma/2$ plus a liquidity taking fee.

To simplify (8), we adopt the ansatz

$$(9) \quad H(t, x, s, q) = -e^{-\gamma(x+sq+h(t,s,q))}$$

and derive the QVI satisfied by $h(t, s, q)$:

$$\begin{aligned} (10) \quad 0 = \min & \left\{ -\gamma \partial_t h(t, s, q) + \theta \left(e^{-\gamma(q\sigma+h(t,s+\sigma,q)-h(t,s,q))} + e^{-\gamma(-q\sigma+h(t,s-\sigma,q)-h(t,s,q))} - 2 \right) \right. \\ & + \lambda^- \min_{\ell^+ \in \{0,1,\dots,\bar{q}-q\}} \left\{ \mathbb{E} \left[e^{-\gamma(\ell^+ \frac{\sigma}{2} \zeta - (q+\ell^+ \zeta) \sigma \xi + h(t,s-\sigma \xi, q+\ell^+ \zeta) - h(t,s,q))} - 1 \right] \right\} \\ & + \lambda^+ \min_{\ell^- \in \{0,1,\dots,q-\underline{q}\}} \left\{ \mathbb{E} \left[e^{-\gamma(\ell^- \frac{\sigma}{2} \zeta + (q-\ell^- \zeta) \sigma \xi + h(t,s+\sigma \xi, q-\ell^- \zeta) - h(t,s,q))} - 1 \right] \right\}; \\ & \min_{m^+ \in \{0,1,\dots,\bar{q}-q\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + h(t,s+\sigma \xi^{0+}, q+m^+) - h(t,s,q))} - 1 \right] \right\}; \\ & \left. \min_{m^- \in \{0,1,\dots,q-\underline{q}\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^- \Upsilon + h(t,s-\sigma \xi^{0-}, q-m^-) - h(t,s,q))} - 1 \right] \right\} \right\}, \end{aligned}$$

subject to the terminal condition $h(T, s, q) = -G(s) - C(s, q)$.

We can extend our approach by assuming other (concave) utility functions. The advantage of employing an exponential utility function is that wealth factors out using the ansatz (9), and hence the dimension of the QVI is reduced.

3.1. Indifference price of contingent claim. The agent determines the premium for the options she sells using indifference pricing (see, e.g., [Carmona \(2008\)](#)). We define the *indifference price* of the options as the amount $I \geq 0$ that makes the agent indifferent between the following two scenarios:

1. The agent does not sell contingent claims. She employs all the aforementioned order types (MBO, MSO, LBO, and LSO) to maximize her expected utility at time T , subject to the same inventory constraint ($q \leq Q_t \leq \bar{q}$).
2. At time $t = 0$ the agent receives the cash amount I and commits to delivering the payoffs at time T , again employing all order types to maximize her expected utility.

In scenario 1, the agent solves the optimal control problem:

$$(11) \quad \tilde{H}(t, x, s, q) = \sup_{(\ell^\pm, \tau^{0\pm}, m^\pm) \in \mathcal{A}} \mathbb{E}_{t, x, s, q} \left[-e^{-\gamma(X_T + S_T Q_T - \tilde{C}(Q_T))} \right],$$

where $\tilde{C}(Q_T)$ represents other (transaction) costs that arise from unwinding terminal inventory.

Classical results (e.g., [Pham \(2009\)](#)) suggest that \tilde{H} satisfies (8), subject to the terminal condition

$$\tilde{H}(T, x, s, q) = -e^{-\gamma(x + s q - \tilde{C}(q))}.$$

By adopting the ansatz $\tilde{H}(t, x, s, q) = -e^{-\gamma(x + s q + \tilde{h}(t, s, q))}$, we can show that $\tilde{h}(t, s, q)$ satisfies (10) subject to the terminal condition $\tilde{h}(T, s, q) = -\tilde{C}(q)$.

On the other hand, in scenario 2, the agent solves the problem where the indifference price $I(t, x, s, q)$ satisfies

$$H(t, x + I(t, x, s, q), s, q) = \tilde{H}(t, x, s, q).$$

It is straightforward to see that I can be expressed in terms of $h(t, s, q)$ and $\tilde{h}(t, s, q)$:

$$I(t, s, q) = \tilde{h}(t, s, q) - h(t, s, q),$$

where we suppress x in $I(t, x, s, q)$ because there is no dependency on wealth.

3.1.1. The value of limit orders. The agent maximizes expected utility of terminal wealth by (i) trading in the underlying asset to ensure that the payoff of the options is settled, and (ii) taking speculative positions in the underlying asset to benefit from roundtrip trades.¹ A key driver of revenues in the agent's strategy is the use of LOs, in both (i) and (ii), because they are priced at the midprice plus or minus half-spread when selling or buying the asset, respectively.

Analogous to finding the value of the option using indifference pricing, we may also use indifference pricing to determine the value of employing LOs in the agent's strategy. To this end, consider the following problem:

$$(12) \quad \hat{H}(t, x, s, q) = \sup_{(\tau^{0\pm}, m^\pm) \in \hat{\mathcal{A}}} \mathbb{E}_{t, x, s, q} \left[-e^{-\gamma(X_T + S_T Q_T - G(S_T) - C(S_T, Q_T))} \right],$$

¹The agent behaves like a market maker; see [Avellaneda and Stoikov \(2008\)](#), [Cartea, Jaimungal, and Penalva \(2015\)](#), and [Guéant \(2017\)](#).

where $\hat{\mathcal{A}}$ denotes the set of strategies $\{(\ell^\pm, \tau^{0\pm}, m^\pm) \in \mathcal{A} : (\ell_t^\pm)_{t \in [0, T]} = 0\}$, i.e., the set of admissible strategies where the agent uses MOs only. By adopting the ansatz $\hat{H}(t, x, s, q) = -e^{-\gamma(x + sq + \hat{h}(t, s, q))}$, we can show that $\hat{h}(t, s, q)$ satisfies a modification of (10) in which the optimization over ℓ^\pm is removed, subject to the terminal condition $\hat{h}(T, s, q) = -G(s) - C(s, q)$. The agent's indifference price of LOs is defined as the function $I^{lo}(t, x, s, q)$ such that

$$H(t, x, s, q) = \hat{H}(t, x + I^{lo}(t, x, s, q), s, q).$$

Thus, the value of the LOs is given by

$$(13) \quad I^{lo}(t, s, q) = h(t, s, q) - \hat{h}(t, s, q),$$

where $h(t, s, q)$ solves the QVI (10), and provides the agent's value function when she can employ LOs and MOs in the strategy. We suppress x in $I^{lo}(t, x, s, q)$ because there is no dependency on wealth.

4. Numerical example. In this section we employ numerical methods to illustrate the performance and features of the agent's strategy. At time $t = 0$ the agent sells European calls written on the S&P 500 E-mini Futures and trades in the underlying to maximize expected utility of terminal wealth.

Data. We employ high-frequency data from the Chicago Mercantile Exchange (CME) to estimate the model parameters. We focus on the S&P E-mini contract that matures on July 20, 2014 and use messages sent to the exchange to build the LOB on March 24, 2014. Our data set contains all messages (FIX format) that traders see to track the liquidity makers and takers in the exchange.² Finally, we employ data between 10:00 and 15:30 Eastern time to exclude the excessive trading activity which normally occurs around market open and close.

The E-mini contract is worth $\$50 \times (\text{S\&P 500 index})$ and trades with tick size $\sigma = \$12.5$. In what follows, we report the price of the contract as the original price divided by 50 to reflect the value of the index.

Parameter estimation.

- *Arrival rate of MOs.* To estimate the rate of MO arrivals, we divide the total number of MOs (buy and sell) in the data set by the total time of 5.5 hours. Next, we count the total number of price changes not due to an incoming MO, divide it by the total time 5.5 hours, and obtain an estimate of the rate of exogenous price changes, which in our model is represented by the parameter θ . The estimates for the rate of MO arrivals and the rate of exogenous price changes are for both sides of the LOB.
- *Permanent price impact.* For an estimate of the probability of midprice changes after MO arrivals from other traders we use the proportion of MOs that results in a midprice change (i.e., consume at least the first level of the LOB) to estimate the parameter α .
- *Fill rate of LOs.* Finally, to estimate the probability of the agent's LO being filled by an incoming MO, i.e., the parameter ρ , we assume that the position of the agent's LO in the queue is uniformly distributed in the best bid or ask level when a MO

²FIX is the protocol that CME employs to communicate changes in the LOB. See <https://www.cmegroup.com/market-data/files/fix-fast-market-data-message-specification.pdf>.

arrives. Therefore, the probability that the agent's LO is filled by an incoming MO is $p_i = M_i/L_{\tau_i}$, where M_i is the size of the i th MO and L_{τ_i} is the total volume posted in the first level of the LOB when the i th MO arrives. We report the median of p_i as our estimate of ρ .

The first four rows in Table 1 report parameter estimates for θ , λ , ρ , α .

Agent's optimal trading. We assume that the agent sells N call options, and denote by S the price of the underlying S&P E-mini Futures. All options expire at time T and have the same strike price K . The options are physically settled; i.e., if the option expires in-the-money, the agent delivers N shares and collects the cash amount $N \times K$.

The options' payoffs and other costs incurred by the agent at the terminal date T are given by

$$(14a) \quad G(S_T) = N(S_T - K)_+,$$

$$(14b) \quad C(S_T, Q_T) = \mathcal{C}(Q_T, N) \mathbb{1}_{S_T \geq K} + \mathcal{C}(Q_T, 0) \mathbb{1}_{S_T < K},$$

where the operation $(x)_+$ represents $\max(x, 0)$, $\mathcal{C} : \mathbb{Z}^2 \mapsto \mathbb{R}$, and $\mathcal{C}(q_1, q_2)$ denotes other costs and fees the agent incurs when changing her inventory position from q_1 to q_2 . For example, if the option expires in-the-money, and the accumulated inventory is $Q_T < N$, the agent must purchase $N - Q_T$ shares to physically settle the options. The cost of these extra shares consists of the sum of (i) $(N - Q_T) S_T$, which is the market price of shares purchased at the terminal date, plus (ii) the cost Υ , which includes the half-spread and liquidity taking fees, in addition to any other fees that may arise from executing the trade—these costs are included in (14b).

In the results of the simulations reported below we assume that the agent is short $N = 10$ call options. At the time of selling the options the underlying asset is trading at $S_0 = 1,856.50$ and we assume that the options are written at-the-money; thus $K = S_0$. Transaction fees and the penalty imposed by the agent stemming from adjusting the inventory position at time T are given by the function

$$(15) \quad \mathcal{C}(q_1, q_2) = \Upsilon \varphi |q_1 - q_2|,$$

where $\varphi \geq 1$ is a terminal penalty parameter. Recall that Υ consists of the half-spread $\sigma/2$ plus a liquidity taking fee. Thus, $\varphi = 1$ represents the costs of unwinding the terminal inventory position in the exchange. When $\varphi > 1$, \mathcal{C} contains additional inventory penalties which do not affect the agent's wealth, but do affect the agent's trading strategy. For example, a high value of the penalty parameter φ forces the strategy to reach the terminal date T with the correct number of shares to settle the option. In the simulations we assume $\varphi = 1.05$, so the strategy reaches the terminal date with the required level of inventory, i.e., $Q_T = N$ if the options end up in-the-money, or $Q_T = 0$ if the options expire out-of-the-money.

Table 1 presents the remaining parameters required to perform the simulations. In particular, it shows the following: expiry of the options, tick size of exchange, half-spread, price impact function of the agent's MOs, the agent's risk-aversion parameter, initial inventory, and inventory constraints.

4.1. Features of the optimal strategy. In this section we solve the QVI system (10) numerically to illustrate a number of features of the agent's strategy. Below, in subsection 4.2




Figure 1. Optimal strategy when $q = 0$. Colors represent size of orders.

we perform simulations to explore the performance of the agent's strategy, and in section 5 we provide details of our numerical scheme and proof of its convergence.

Here we discuss how the agent employs MOs and LOs depending on the level of inventory, midprice, and remaining time to expiry of the contingent claims to maximize her expected utility of terminal wealth.

Figure 1 shows heat-maps of the agent's optimal strategy when the inventory is fixed at $q = 0$. On the left panel, the colored areas show the regions in which the agent executes MOs of various volumes, and the empty (i.e., white) area shows the region in which she does not execute any MOs. Similarly, the right panel shows the optimal posting of buy LOs across time and moneyness of the options.

In particular, the left panel of Figure 1 shows that as the option moves deeper in-the-money, the agent executes larger volumes of MBOs in anticipation of the option expiring in-the-money. The left panel also shows that everything else equal, when the option is in-the-money, as the option expiry approaches, the agent executes larger volumes of MBOs.

On the other hand, the right panel of Figure 1 shows that when the option is out-of-the-money and the stock price is trading not "too far" from the strike, the agent posts LBOs with volume of 5 shares. As the option moves further out-of-the-money, the agent posts smaller LBOs. Moreover, at a fixed asset price, as maturity approaches, the posted volume decreases.

Figure 2 shows the optimal strategy when the inventory is fixed at $q = 5$; recall that the agent is short $N = 10$ call options. When the option is close to at-the-money, the agent does not execute MOs, but rather posts LOs on both sides of the LOB to maximize revenues by earning the spread from speculative roundtrip trades. As time approaches maturity, if the option is in-the-money, the agent increases the volume of the LBOs and eventually executes a MBO, but still posts sell LOs to maximize revenues from speculative trades. Similarly, if the option is out-of-the-money, the agent increases the volume of the LSOs and, nearer expiry, the strategy eventually executes a MSO, but still posts buy LOs to maximize revenues from speculative roundtrip trades. Moreover, there are regions where the agent posts only LBOs (when the option is moderately in-the-money) or only LSOs (when the option is moderately out-of-the-money), but if the option moves significantly in- or out-of-the money, she executes MBOs or MSOs.

The value the agent attaches to the use of LOs depends on the agent's level of risk-aversion.




Figure 2. Optimal strategy when $q = 5$. Colors represent size of orders, where a positive (negative) value means a buy (sell) order.




Figure 3. Innermost boundaries of the region where agent immediately executes a MO. At each point in the left (right) panel, the color represents the minimum inventory level such that agent immediately executes a MBO (MSO) when the midprice lies below (above) the boundary for that inventory level.

LOs can increase both the (i) expected revenues that stem from speculative roundtrip trades, and (ii) volatility of the agent's terminal cash position. Thus, as the level of risk-aversion increases, the strategy employs more MOs and relies less on LOs to boost revenues with speculative roundtrip trades, i.e., trades off expected revenue against volatility of revenues. We return to this point in the discussion of Figure 4 and at the end of this section.

There are other notable features here. Figure 2(a) contains two inner boundaries separating the plane into three different regions: the upper colored region, the middle white region, and the lower colored region. Inside the middle region, the agent does not execute MOs, and as soon as a boundary is reached, the agent immediately executes a MBO (upper region) or a MSO (lower region) and the inventory jumps to a different level. As a result, any point (t, S) in the interior of the colored regions is not attainable unless the agent starts with this particular inventory level at time t and $S_t = S$. Figure 3 shows the strategy's innermost boundaries for different levels of inventory. The left (right) panel shows the boundaries for a given inventory level at which the agent would execute a MBO (MSO) when the asset price lies above (below) that boundary.

Next we focus on a sample path to illustrate how the strategy behaves dynamically. Figure 4(a) shows a sample path of the midprice, with the strike price depicted by a horizontal line. For this particular realization of the midprice, the option stays in-the-money early on, then switches moneyness a few times before maturity and eventually expires out-of-the-money.

Panel 4(b) shows the strategy's acquired inventory, where solid circles denote the execution of MOs. For this particular price path, most of the changes in the strategy's inventory are due to filled LOs. Moreover, early on, the strategy sends a MO (indicated by the red dot) taking the inventory to $q = 2$, and near expiry, when the midprice is dropping, and in anticipation of the options expiring out-of-the-money, the strategy executes a series of sell MOs (indicated by the green dots), so at expiry $q_T = 0$.

To glean additional insight into what the strategy does, panel 4(c) shows the inventory path which the strategy targets using LOs. The red and blue lines represent the levels at which the inventory would jump if the posted LOs were filled, i.e., $Q_t + \ell_t^+$ if the LBO is filled and $Q_t - \ell_t^-$ if the LSO is filled. For this simulation, the agent posts both LBOs and LSOs simultaneously most of the time. In doing so, the agent is aiming to profit from the bid-ask spread while simultaneously hedging the option payoff.

For illustrative purposes, panel 4(c) also shows the inventory path that would result if the agent employs the classical delta strategy to hedge the options using MOs. Without market impact and adverse selection, the high-frequency limit of our model results in a Brownian motion midprice dynamics. Hence, the delta-hedge strategy we show here is that resulting from the Bachelier model. We observe that the inventory position targeted by LOs is always above or below that targeted by a delta hedge because the strategy seeks to profit from the spread that accrues each time a roundtrip is completed.

To calculate the delta-hedge positions we assume that the midprice follows the arithmetic Brownian motion

$$(16) \quad S_t = \sigma^B W_t,$$

where $W = (W_t)_{t \in [0, T]}$ is a standard Brownian motion, and we choose the volatility parameter σ^B so that it matches the volatility of the midprice in model (1) in the absence of the agent's MOs. Furthermore, we assume that the Brownian motion is independent of all other random variables in the model. Thus, one can show that

$$\begin{aligned} (\sigma^B)^2 &= \mathbb{V} \left[Z_t^+ + \sum_{k=1}^{M_t^+} \xi_k^+ - Z_t^- - \sum_{k=1}^{M_t^-} \xi_k^- \right] / t \\ &= \sigma^2 (2\theta + 2\lambda\alpha), \end{aligned}$$

where $\mathbb{V}[\cdot]$ denotes the variance operator.

Interestingly, in Figure 4, the LOs in the strategy target inventory levels which contain the delta-hedge position, but this is not always the case. Figure 5 is similar to Figure 4, but the value of the risk-aversion parameter is $\gamma = 1$ and panel (c) shows the Bachelier delta is not always enclosed by the targeted optimal level of inventory. For example, at around $t = 0.65$, the Bachelier delta requires an inventory position larger than that sought by the agent who targets an inventory level with LOs.

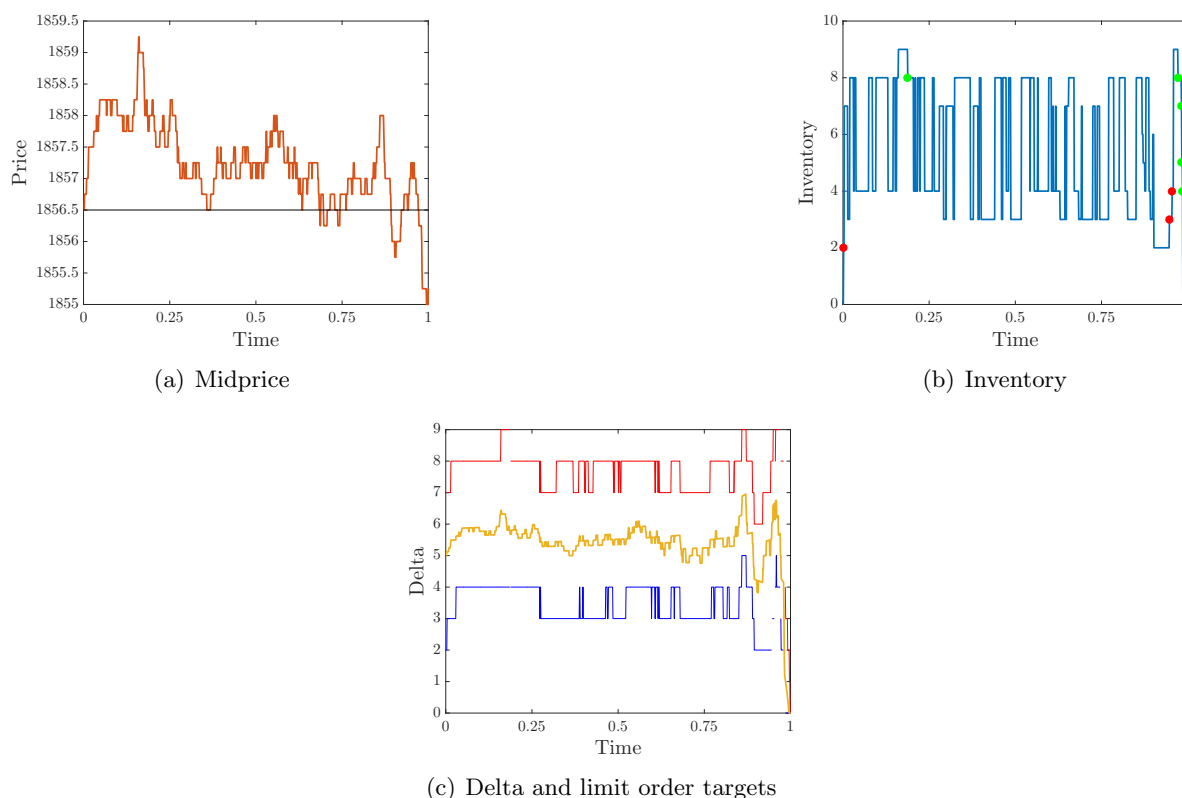


Figure 4. Sample path of price (a) and inventory (b). The red and blue lines in (c) show the inventory levels that the strategy targets using LOs. The posted buy volume is such that if the buy LO is filled, the inventory will jump to the quantity shown in the red line. Similarly, if the posted sell volume is such that if the sell LO is filled, the inventory will jump to the inventory shown in the blue line. The orange line is the option delta calculated from a Bachelier model.

Standard deviation of the optimal inventory. We employ the expression for the dynamics of the inventory (4) to derive the drift of the agent's inventory and to derive the standard deviation of the inventory when it is not optimal to execute MOs. In the region where the agent does not execute MOs, the drift of the inventory is given by

$$(17) \quad \mu(t, s, q) = \lim_{\delta \rightarrow 0} \frac{\mathbb{E}_{t, x, s, q} [Q_{t+\delta}] - q}{\delta} = \rho (\lambda^+ \ell^-(t, s, q) - \lambda^- \ell^+(t, s, q))$$

for $t \in [0, T]$, $q \in \{q, \dots, \bar{q}\}$, and $s \in \sigma \mathbb{Z}$. Recall that ℓ^\pm represents the number of shares the agent posts on the buy and sell sides of the book. We obtain values of $\mu(t, s, q)$ for any $q \in \mathbb{R}$ by linear interpolation, and we define $q^*(t, s)$ as the unique point satisfying $\mu(t, s, q^*(t, s)) = 0$, so that $q^*(t, s)$ represents the optimal inventory to hold at time t given $S_t = s$.

Moreover, we define the variance of inventory at the optimal inventory level to hold as

$$v(t, s) = \lim_{\delta \rightarrow 0} \frac{\mathbb{V}_{t, x, s, q^*(t, s)} [Q_{t+\delta} - q^*(t, s)]}{\delta}$$

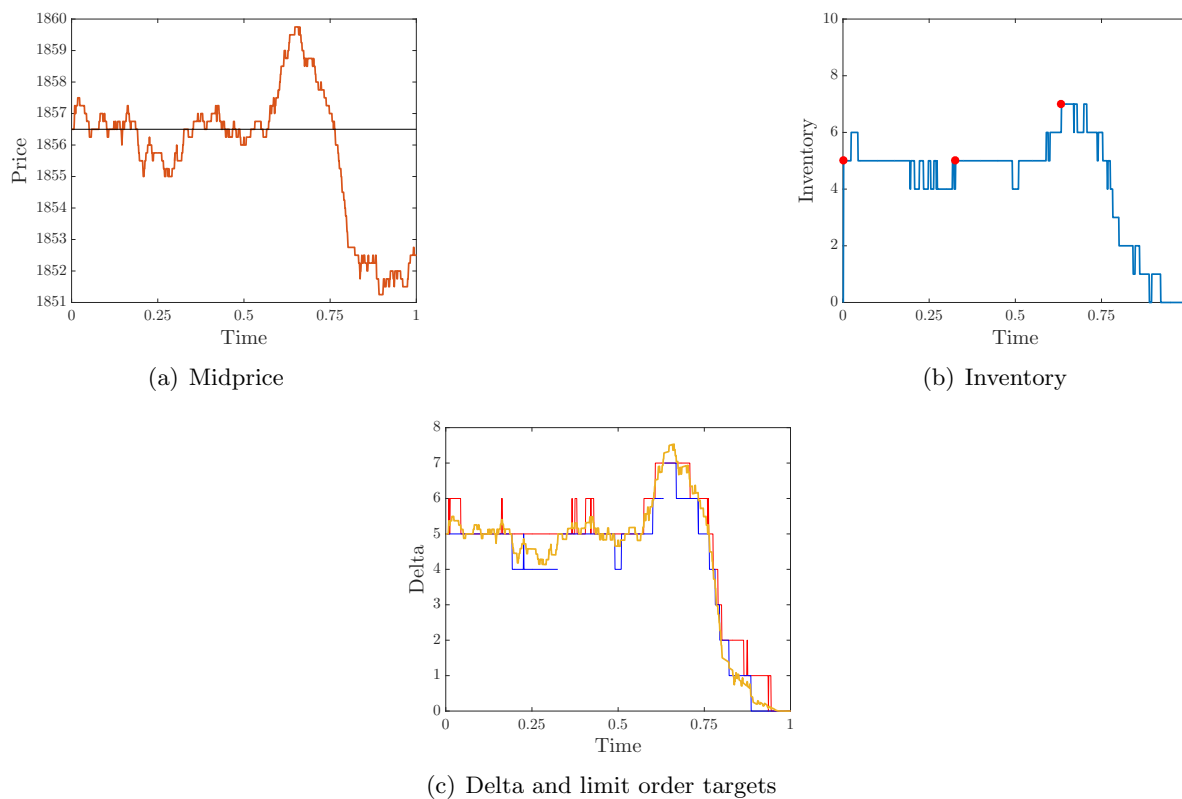


Figure 5. Example of an instance where the optimal inventory diverges from the Bachelier delta, $\gamma = 1$. At time $t = 0.65$, the Bachelier delta is outside the band.

$$(18) \quad = \rho \left[\lambda^+ \left(\ell^-(t, s, q^*(t, s)) \right)^2 + \lambda^- \left(\ell^+(t, s, q^*(t, s)) \right)^2 \right].$$

We define the standard deviation of the optimal inventory as $\sqrt{v(t, s)}$. The quantity $\sqrt{v(t, s)}$ measures the variability of the agent's inventory when she holds the optimal inventory level.

Panels (a) and (b) of Figure 6 show the optimal level of inventory and the contour plot of the standard deviation of the optimal inventory $v(t, s)$. The standard deviation is largest when the option is at-the-money.

Figure 7 shows how the agent values the use of LOs in the trading strategy. The vertical axis shows the indifference price of LOs, i.e., $I^{lo}(0, s, q)$ computed using (13), where we assume the inventory is $q = 0$ and $s = \$1856.5$, and the horizontal axis represents different values of the risk-aversion parameter γ . Observe that as the agent becomes more risk-averse, the value she attaches to the use of LOs diminishes. When the agent's degree of risk-aversion is extremely high, the optimal trading strategy relies much less on LOs and mainly employs MOs to guarantee that at the terminal date the agent has the correct amount of the underlying asset to settle the option. The use of LOs increases the gains from speculative roundtrip trades and incurs no exchange fees; however, LOs introduce additional tracking error and make the agent's terminal wealth more volatile.

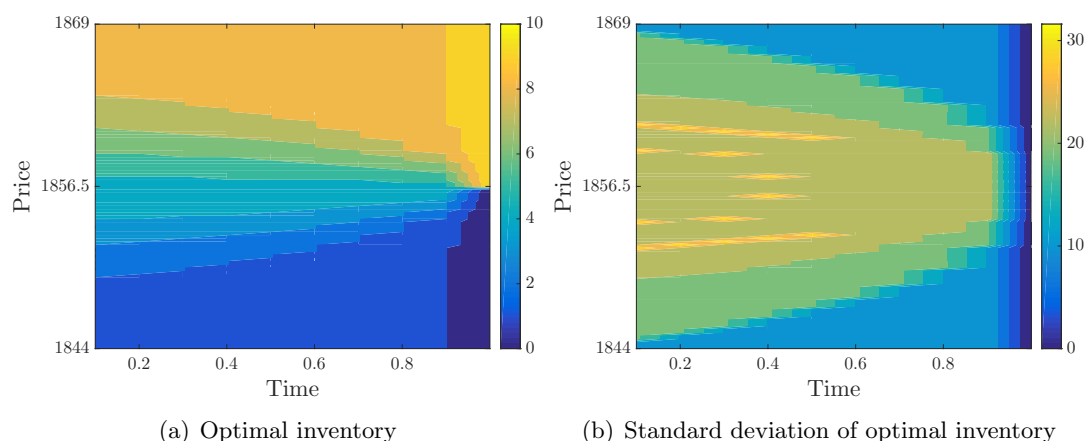


Figure 6. Sample path of price and inventory.




Figure 7. Indifference price for being able to post LOs ($s = 1856.5$ and $q = 0$).

Finally, we discuss how the optimal strategy accounts for adverse selection costs. Recall that the agent's and other market participants' MOs have permanent impact on the midprice of the underlying asset. This price impact is captured by the random variables ξ^\pm , $\xi^{0\pm}$. As this impact becomes larger (on average), the agent's strategy relies less on LOs. That is, everything else being equal, the agent will execute more MOs, and LOs are adjusted in two ways: (i) posting fewer LOs, and (ii) adjusting the volume of LOs downwards. In this way the agent strikes the optimal balance between how often and how many shares she posts at the best bid and best ask of the LOB, and the number of MOs she executes. (In the interest of space we do not include figures to show the effect of adverse selection on the agent's LOs and MOs.)

4.2. Simulations: Financial performance of the strategy. To illustrate the financial performance of the strategy we simulate 10^5 sample paths of the midprice. We show how the strategy maximizes the agent's expected utility of wealth while replicating the payoff of the N call options. Figure 8 shows the value of the payoff of the 10 call options (solid line) and the strategy's terminal wealth (blue circles), assuming low risk-aversion ($\gamma = 0.1$) in the left panel, and high risk-aversion ($\gamma = 0.3$) in the right panel. In both cases, for all simulations shown




Figure 8. Terminal value of the agent's portfolio.




Figure 9. Terminal value of the agent's portfolio when she adopts the Bachelier strategy.

here, the agent super-replicates the call option. Part of what drives the “super-replication” is due to the indifference price the agent charges the client. This price accounts for jump-risk, adverse selection costs, bid-ask spreads, and exchange fees, all of which increase the indifference price. On the other hand, the potential profits from roundtrip trades using LOs lower the price. Nonetheless, we see the overall effect is that the strategy lies above the payoff in all scenarios.

Clearly, as the agent becomes more risk-averse, the strategy becomes more conservative, so it relies more on MOs and less on LOs. By executing more MOs the agent increases the certainty with which she achieves the desired level of inventory, but pays higher costs due to crossing the spread, in addition to MO fees, and forgoes opportunities to earn the spread with LOs. On the other hand, as the agent become less risk-averse, the strategy relies more on LOs and less on MOs. In this case the agent earns the spread when LOs are filled but bears more volatility in the wealth because there is no guarantee that the LOs will be filled. These effects can be appreciated in Figure 8, and are also seen in Figure 10, where we show the histogram of total number of executed MOs and matched LOs for the agent's strategy. Clearly, when $\gamma = 0.3$, the agent uses more MOs and fewer LOs. Thus, she pays higher transaction costs on average and has a lower terminal portfolio value. These costs are, however, transferred to the client through the indifference price.

Figure 9 shows the cost of replicating the options in the Bachelier model. Here, the premium charged by the agent for the options is calculated using the Bachelier model, and the delta-hedge strategy is implemented as follows. At every point in time the agent calculates and rounds the Bachelier delta to the nearest integer. If the strategy's inventory is different from this integer, the agent executes a MO so the inventory remains on target. For every simulation shown in the figure we see the strategy is a sub-replication one because MOs cross the spread and pay a fee.

Figure 11 shows the agent's mean terminal wealth as a function of the standard deviation of terminal wealth when γ varies from 0.1 to 0.3. As γ increases, both the mean and the standard deviation of terminal wealth decrease. As the agent becomes more risk-averse the strategy employs more MOs and fewer LOs. This mix of market and limit orders reduces the variance of terminal wealth, but as a tradeoff also reduces the mean of terminal wealth




Figure 10. Histogram of number of market/limit orders.




Figure 11. Risk-reward plot for the agent's terminal wealth.

because the strategy does not reap the benefits from employing LOs.

5. Convergence results. In this section, we describe the numerical algorithm we employ to solve for the value function and obtain the optimal trading strategy. We restrict the asset price to take values on a finite grid $\bar{\mathbb{S}} := \{k\sigma : k \in \mathbb{Z}, k_{\min} \leq k \leq k_{\max}\}$ for some $k_{\min} \leq k_{\max} \in \mathbb{Z}$. In this way, the system of QVIs shown in (10) becomes a coupled system of ordinary differential equations indexed by (s, q) on $\bar{\mathbb{S}} \times \mathbb{Q}$, where $\mathbb{Q} = \{q, \dots, \bar{q}\}$. We write $\mathbb{S} := \bar{\mathbb{S}} / \{k_{\min}\sigma, k_{\max}\sigma\}$ (i.e., excluding the end points) and denote by n (resp., m) the number of elements in the set \mathbb{S} (resp., \mathbb{Q}). We seek a bounded viscosity solution $u : [0, T] \rightarrow \mathbb{R}^{n \times m}$ to the following system of QVIs:

$$(19) \quad \begin{cases} F_{s,q}(t, u, \partial_t u_{s,q}) = 0, & 0 \leq t < T, \\ u_{s,q}(T) = -G(s) - C(s, q), \end{cases}$$

for all $(s, q) \in \bar{\mathbb{S}} \times \mathbb{Q}$. The mapping F is defined as

$$(20) \quad F_{s,q}(t, u, p) = \min\{L_{s,q}(t, u, p), M_{s,q}^+(t, u), M_{s,q}^-(t, u)\},$$

where

$$(21a) \quad L_{s,q}(t, u, p) = -\gamma p + \kappa u_{s,q} + \theta J_{s,q}(t, u) + \lambda^- K_{s,q}^+(t, u) + \lambda^+ K_{s,q}^-(t, u),$$

$$(21b) \quad M_{s,q}^+(t, u) = \min_{m^+ \in \{0, 1, \dots, \bar{q}-q\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + u_{s+1} \mathbb{1}_{s \neq k_{\max}\sigma} \sigma \xi^{0+, q+m^+} - (1+\kappa)u_{s,q})} - 1 \right] \right\},$$

$$(21c) \quad M_{s,q}^-(t, u) = \min_{m^- \in \{0, 1, \dots, q-\underline{q}\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^- \Upsilon + u_{s-1} \mathbb{1}_{s \neq k_{\min}\sigma} \sigma \xi^{0-, q-m^-} - (1+\kappa)u_{s,q})} - 1 \right] \right\},$$

$$(21d) \quad J_{s,q}(t, u) = \begin{cases} e^{-\gamma(q\sigma + u_{s+\sigma, q} - u_{s,q})} + e^{-\gamma(-q\sigma + u_{s-\sigma, q} - u_{s,q})} - 2, & s \in \mathbb{S}, \\ 2e^{-\gamma(q\sigma + u_{s+\sigma, q} - u_{s,q})} - 2, & s = k_{\min}\sigma, \\ 2e^{-\gamma(-q\sigma + u_{s-\sigma, q} - u_{s,q})} - 2, & s = k_{\max}\sigma, \end{cases}$$

$$(21e) \quad K_{s,q}^+(t, u) = \min_{\ell^+ \in \{0, 1, \dots, \bar{q} - q\}} \left\{ \mathbb{E} \left[e^{-\gamma \left(\ell^+ \frac{\sigma}{2} \zeta - (q + \ell^+ \zeta) \sigma \xi + u_{s-\sigma \xi, q+\ell^+ \zeta} - u_{s,q} \right)} - 1 \right] \right\},$$

$$(21f) \quad K_{s,q}^-(t, u) = \min_{\ell^- \in \{0, 1, \dots, q - \underline{q}\}} \left\{ \mathbb{E} \left[e^{-\gamma \left(\ell^- \frac{\sigma}{2} \zeta + (q - \ell^- \zeta) \sigma \xi + u_{s+\sigma \xi, q-\ell^- \zeta} - u_{s,q} \right)} - 1 \right] \right\},$$

where $\kappa \downarrow 0$ is a robustness parameter. In all the above equations, we define $\min\{\emptyset\} = +\infty$.

Our first result is a comparison principle for the QVI (19).

Proposition 1. *The system (19) admits a comparison principle, i.e., if u (resp., v) is an upper (resp., lower) semicontinuous subsolution (resp., supersolution), then $u \leq v$.*

Proof. For the proof see Appendix A. ■

We introduce the finite difference scheme $F^\epsilon = (F_{s,q}^\epsilon) : [0, T] \times \mathbb{R}^{m \times n} \times C([0, T]) \rightarrow \mathbb{R}^{m \times n}$ by the formula

$$F_{s,q}^\epsilon(t, u, \phi(\cdot)) = \min\{L_{s,q}^\epsilon(t, u, \phi(\cdot)), M_{s,q}^+(t, u), M_{s,q}^-(t, u)\},$$

where

$$L_{s,q}^\epsilon(t, r, \phi(\cdot)) = -\gamma \frac{\phi(t + \epsilon) - r_{s,q}}{\epsilon} + \theta J_{s,q}(t, r) + \lambda^- K_{s,q}^+(t, r) + \lambda^+ K_{s,q}^-(t, r).$$

Let

$$\mathbb{T}_\epsilon := \{t_j := \epsilon j : 0 \leq j \leq N\},$$

where $\epsilon > 0$ is chosen so that $N := \frac{T}{\epsilon} \in \mathbb{Z}_+$. We show that the solution $u^\epsilon = (u_{s,q}^\epsilon)$, $u_{s,q}^\epsilon : \mathbb{T}_\epsilon \rightarrow \mathbb{R}$, of the discrete problem

$$(22) \quad \begin{cases} F_{s,q}^\epsilon(t_j, u_{s,q}^\epsilon(t_j), u_{s,q}^\epsilon) &= 0, & 0 \leq j \leq N-1, \\ u_{s,q}^\epsilon(T) &= -G(s) - C(s, q) \end{cases}$$

can be used to approximate the solution of (19).

Proposition 2. *The following are true:*

- (C1) *For all bounded functions $u = (u_{s,q})$, $v = (v_{s,q})$ with $u_{s,q}, v_{s,q} \in C[0, T]$ and $u_{s,q} \leq v_{s,q}$ on $[0, T]$, suppose $r, r' \in \mathbb{R}^{m \times n}$ satisfies $r_{s^*, q^*} - r'_{s^*, q^*} = \max_{s,q} \{r_{s,q} - r'_{s,q}\} = \delta \geq 0$ and $M_{s^*, q^*}(t, r) \geq 0$; then we have*

$$F_{s^*, q^*}^\epsilon(t, r, u_{s^*, q^*} + \delta) - F_{s^*, q^*}^\epsilon(t, r', v_{s^*, q^*}) \geq \min\{(M_{s^*, q^*}(t, r') + 1)\gamma, 1\} \kappa \delta.$$

- (C2) *For any $\epsilon > 0$ and any bounded function $\phi \in C([0, T])$, fix $t \in [0, T]$ and $r = (r_{s,q}) \in \mathbb{R}^{m \times n}$; then the functions*

$$r \mapsto F_{s,q}^\epsilon(t, r, \phi)$$

are uniformly continuous for r in a bounded set, uniformly in $t \in [0, T]$.

- (C3) *For every $\psi = (\psi_{s,q})$, $\psi_{s,q} \in C([0, T]) \cap C^1((0, T))$, $t \in [0, T]$, and $\delta > 0$, there exists $\epsilon > 0$ such that*

$$|F_{s,q}^\epsilon(t, \psi(t), \psi_{s,q}) - F_{s,q}(t, \psi, \partial_t \psi_{s,q})| \leq \delta, \quad \underline{q} \leq q \leq \bar{q}.$$

(C4) *There exists $\epsilon_0 > 0$ such that for any $\epsilon \in (0, \epsilon_0)$, there exists a bounded solution $u^\epsilon = (u_{s,q}^\epsilon, u_{s,q}^\epsilon : \mathbb{T}^\epsilon \rightarrow \mathbb{R}$ of the discrete problem (22).*

Proof. For the proof see Appendix A. ■

With the above results, we define the candidate super- and subsolution of (19) as

$$(23) \quad \underline{u}_{s,q}(t) = \liminf_{\epsilon \rightarrow 0, t' \rightarrow t} u_{s,q}^\epsilon(t'),$$

$$(24) \quad \bar{u}_{s,q}(t) = \limsup_{\epsilon \rightarrow 0, t' \rightarrow t} u_{s,q}^\epsilon(t')$$

for all $t \in [0, T]$, $q \leq q \leq \bar{q}$.

Proposition 2 allows us to apply the arguments of Proposition 3.3 of [Briani, Camilli, and Zidani \(2012\)](#) to show that $\underline{u} = (\underline{u}_{s,q})$ and $\bar{u} = (\bar{u}_{s,q})$ are supersolution and subsolution of (19). By Proposition 1, we have $\underline{u} \geq \bar{u}$. The opposite direction is true by definition and we have that $u = \underline{u} = \bar{u}$ is the viscosity solution of (19).

6. Conclusions. We show how an agent maximizes expected utility of wealth when she takes a position in a contingent claim and employs limit and market orders to trade in the underlying of the claim. The agent solves a combined optimal stopping and stochastic control problem, and we characterize the solution in terms of a Hamilton–Jacobi–Bellman quasi-variational inequality (QVI). We employ a numerical scheme to solve the QVI and prove a convergence result for the scheme.

In our model, market orders from all traders (including the agent) have price impact and pay other exchange related fees. The price impact is permanent and creates adverse selection costs when the agent’s limit orders are filled. That is, the midprice of the underlying asset jumps in the direction of the market order.

We discuss a particular case where the agent takes a short position in European options written on the E-mini that tracks the S&P500 index. MOs are expensive because they cross the spread and pay fees to the exchange, but guarantee execution, so the agent employs these to keep the inventory on target. The agent employs limit and market orders to hedge the exposure to the contingent claim, and also engages in speculative trades to maximize expected utility of wealth.

LOs, on the other hand, do not incur exchange fees, but there is no guarantee that they will be filled in time. We show that the agent’s strategy relies on LOs not only to achieve the desired inventory target, which depends on time to expiry and moneyness of the contingent claim, but are also employed in speculative trades to earn the spread from roundtrip trades.

The agent’s strategy accounts for adverse selection costs by (i) controlling how often she posts LOs on the bid and ask of the LOB, and (ii) the volume posted in the LO. If adverse selection costs increase (decrease), the strategy relies more (less) on MOs and less (more) on LOs.

Finally, we note that speculative trades may be curbed by regulation. Section 619 of the Dodd–Frank Act,³ known as the Volcker Rule, bans proprietary trading by commercial banks. Thus, the strategy derived in this paper cannot be employed by banks that either trade on

³See <https://www.govinfo.gov/content/pkg/PLAW-111publ203/html/PLAW-111publ203.htm>.

behalf of a client (i.e., hedge a client's position) or execute proprietary trades that are backed by deposits of their clients.

Appendix A. Proofs. The following proposition is later used in the proof of the comparison principle.

Proposition 3. *Let $u = (u_{s,q})$, $v = (v_{s,q}) \in \mathbb{R}^{n \times m}$. Suppose that $u_{s^*,q^*} - v_{s^*,q^*} = \max_{s,q} \{u_{s,q} - v_{s,q}\} \geq 0$ and $M_{s^*,q^*}^\pm(t, v) \geq 0$; then*

$$F_{s^*,q^*}(t, u, p) - F_{s^*,q^*}(t, v, p) \geq \min(1, \gamma) \kappa(u_{s^*,q^*} - v_{s^*,q^*}).$$

Proof of Proposition 3. It suffices to show the following:

$$(25) \quad L_{s^*,q^*}(t, u, p) - L_{s^*,q^*}(t, v, p) \geq \kappa(u_{s^*,q^*} - v_{s^*,q^*}),$$

$$(26) \quad M_{s^*,q^*}^+(t, u) - M_{s^*,q^*}^+(t, v) \geq \gamma \kappa(u_{s^*,q^*} - v_{s^*,q^*}),$$

$$(27) \quad M_{s^*,q^*}^-(t, u) - M_{s^*,q^*}^-(t, v) \geq \gamma \kappa(u_{s^*,q^*} - v_{s^*,q^*}).$$

To show (25), note that $u_{s^*,q^*} - v_{s^*,q^*} = \max_{s,q} \{u_{s,q} - v_{s,q}\}$ implies $u_{s,q} - u_{s^*,q^*} \leq v_{s,q} - v_{s^*,q^*}$ for all s and q . This further implies $J_{s^*,q^*}(t, u) \geq J_{s^*,q^*}(t, v)$, $K_{s^*,q^*}^+(t, u) \geq K_{s^*,q^*}^+(t, v)$, and $K_{s^*,q^*}^-(t, u) \geq K_{s^*,q^*}^-(t, v)$. The rest are straightforward calculations.

It remains to show (26). Showing (27) is similar. Let us write $(s, q) = (s^* + \sigma \xi^{0+}, q^* + m^+)$. Note that the left-hand side of (26) is greater than or equal to

$$\begin{aligned} & \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + u_{s,q} - (1+\kappa)u_{s^*,q^*})} - e^{-\gamma(-m^+ \Upsilon + v_{s,q} - (1+\kappa)v_{s^*,q^*})} \right] \right\} \\ & \geq \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + v_{s,q} - (1+\kappa)v_{s^*,q^*})} \left(e^{-\gamma(u_{s,q} - (1+\kappa)u_{s^*,q^*} - v_{s,q} + (1+\kappa)v_{s^*,q^*})} - 1 \right) \right] \right\} \\ & \geq \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + v_{s,q} - (1+\kappa)v_{s^*,q^*})} \left(e^{\gamma \kappa(u_{s^*,q^*} - v_{s^*,q^*})} - 1 \right) \right] \right\} \\ & \geq \gamma \kappa(u_{s^*,q^*} - v_{s^*,q^*}) \left(M_{s^*,q^*}^+(t, v) + 1 \right) \\ & \geq \gamma \kappa(u_{s^*,q^*} - v_{s^*,q^*}). \end{aligned}$$

Proof of Proposition 1. We prove the proposition by contradiction.

Suppose $u_{s^*,q^*}(\bar{t}) - v_{s,q}(\bar{t}) = \max_{s,q} \sup_{[0,T]} \{u_{s,q} - v_{s,q}\} = \eta > 0$. Without loss of generality, we may assume $0 < \bar{t} < T$. Consider the family of upper semicontinuous functions

$$\Phi_\epsilon(s, q, t, t') = u_{s,q}(t) - v_{s,q}(t') - \phi_\epsilon(t, t'),$$

where $\phi_\epsilon(t, t') = \frac{1}{\epsilon} |t - t'|^2$ and $\epsilon > 0$, and the bounded sequence $(q_\epsilon, s_\epsilon, t_\epsilon, t'_\epsilon)$ attains the maximum of Φ_ϵ . By standard arguments, we have

$$\eta_\epsilon = \max \Phi_\epsilon = \Phi_\epsilon(q_\epsilon, s_\epsilon, t_\epsilon, t'_\epsilon) \rightarrow \eta, \quad \frac{1}{\epsilon} |t_\epsilon - t'_\epsilon|^2 \rightarrow 0,$$

as $\epsilon \rightarrow 0$.

Note that $u_{s_\epsilon, q_\epsilon}(\cdot) - \phi_\epsilon(\cdot, t'_\epsilon)$ attains maximum at t_ϵ . By the fact that u is a subsolution, we have

$$F_{s_\epsilon, q_\epsilon}(t_\epsilon, u(t_\epsilon), \partial_t \phi_\epsilon) \leq 0.$$

Similarly we have

$$F_{s_\epsilon, q_\epsilon}(t'_\epsilon, v(t'_\epsilon), -\partial_{t'} \phi_\epsilon) \geq 0.$$

The above two inequalities yield

$$\begin{aligned} 0 &\geq F_{s_\epsilon, q_\epsilon}(t_\epsilon, u(t_\epsilon), \partial_t \phi_\epsilon) - F_{s_\epsilon, q_\epsilon}(t'_\epsilon, v(t'_\epsilon), -\partial_{t'} \phi_\epsilon) \\ &\geq F_{s_\epsilon, q_\epsilon}(t_\epsilon, u(t_\epsilon), \partial_t \phi_\epsilon) - F_{s_\epsilon, q_\epsilon}(t_\epsilon, v(t'_\epsilon), \partial_t \phi_\epsilon) + F_{s_\epsilon, q_\epsilon}(t_\epsilon, v(t'_\epsilon), \partial_t \phi_\epsilon) - F_{s_\epsilon, q_\epsilon}(t'_\epsilon, v(t'_\epsilon), -\partial_{t'} \phi_\epsilon) \\ &= \min(1, \gamma) \kappa(u_{s_\epsilon, q_\epsilon}(t_\epsilon) - v_{s_\epsilon, q_\epsilon}(t'_\epsilon)) + \underbrace{F_{s_\epsilon, q_\epsilon}(t_\epsilon, v(t'_\epsilon), \partial_t \phi_\epsilon) - F_{s_\epsilon, q_\epsilon}(t'_\epsilon, v(t'_\epsilon), -\partial_{t'} \phi_\epsilon)}_{(\mathfrak{A})}, \end{aligned}$$

where the second inequality follows from Proposition 3. It is clear from the definition that $F_{s_\epsilon, q_\epsilon}(\cdot, \cdot, \cdot)$ is a continuous mapping from \mathbb{R}^3 to \mathbb{R} . Therefore the term (\mathfrak{A}) vanishes when $\epsilon \rightarrow 0$. We have the required contradiction. ■

Proof of Proposition 2. The verification of conditions (C2), (C3), and (C4) is immediate. For (C1), it suffices to show the following:

$$(28) \quad L_{s^*, q^*}^\epsilon(t, r, u_{s^*, q^*} + \delta) - L_{s^*, q^*}^\epsilon(t, r', v_{s^*, q^*}) \geq \kappa \delta,$$

$$(29) \quad M_{s^*, q^*}^+(t, r) - M_{s^*, q^*}^+(t, r') \geq (M_{s^*, q^*}^+(t, r') + 1) \gamma \kappa \delta.$$

$$(30) \quad M_{s^*, q^*}^-(t, r) - M_{s^*, q^*}^-(t, r') \geq (M_{s^*, q^*}^-(t, r') + 1) \gamma \kappa \delta.$$

To prove (28), note that $r_{s^*, q^*} - r'_{s^*, q^*} = \max_{s, q} \{r_{s, q} - r'_{s, q}\}$ implies $r_{s, q} - r_{s^*, q^*} \leq r'_{s, q} - r'_{s^*, q^*}$ for all s, q , which further implies $e^{-\gamma(r_{s, q} - r_{s^*, q^*})} \geq e^{-\gamma(r'_{s, q} - r'_{s^*, q^*})}$ for all s, q . The rest are straightforward calculations.

It remains to prove (29). Proving (30) is similar. We write $(\mathfrak{s}, \mathfrak{q}) = (s^* + \sigma \xi^{0+}, q^* + m^+)$ and proceed as in the proof of Proposition 3:

$$\begin{aligned} &M_{s^*, q^*}(t, r) - M_{s^*, q^*}(t, r') \\ &\geq \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + r_{s, q} - (1+\kappa)r_{s^*, q^*})} - e^{-\gamma(-m^+ \Upsilon + r'_{s, q} - (1+\kappa)r'_{s^*, q^*})} \right] \right\} \\ &\geq \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + r'_{s, q} - (1+\kappa)r'_{s^*, q^*})} \left(e^{-\gamma(r_{s, q} - (1+\kappa)r_{s^*, q^*} - r'_{s, q} + (1+\kappa)r'_{s^*, q^*})} - 1 \right) \right] \right\} \\ &\geq \min_{m^+ \in \{0, 1, \dots, \bar{q} - q^*\}} \left\{ \mathbb{E} \left[e^{-\gamma(-m^+ \Upsilon + r'_{s, q} - (1+\kappa)r'_{s^*, q^*})} \left(e^{\gamma \kappa(r_{s^*, q^*} - r'_{s^*, q^*})} - 1 \right) \right] \right\} \\ &\geq \gamma \kappa(r_{s^*, q^*} - r'_{s^*, q^*}) \left(M_{s^*, q^*}^+(t, r') + 1 \right) \\ &\geq \gamma \kappa(r_{s^*, q^*} - r'_{s^*, q^*}) \\ &\geq \gamma \kappa \delta. \end{aligned}$$

REFERENCES

- R. ALMGREN (2003), *Optimal execution with nonlinear impact functions and trading-enhanced risk*, Appl. Math. Finance, 10, pp. 1–18.
- R. ALMGREN AND T. M. LI (2016), *Option hedging with smooth market impact*, Market Microstructure Liquidity, 2, 1650002.
- M. AVELLANEDA AND S. STOIKOV (2008), *High-frequency trading in a limit order book*, Quant. Finance, 8, pp. 217–224.
- P. BANK AND D. BAUM (2004), *Hedging and portfolio optimization in financial markets with a large trader*, Math. Finance, 14, pp. 1–18.
- P. BANK, H. M. SONER, AND M. VOß (2017), *Hedging with temporary price impact*, Math. Financ. Econ., 11, pp. 215–239.
- G. BARLES AND H. M. SONER (1998), *Option pricing with transaction costs and a nonlinear Black-Scholes equation*, Finance Stoch., 2, pp. 369–397.
- B. BOUCHARD, G. LOEPER, AND Y. ZOU (2017), *Hedging of covered options with linear market impact and gamma constraint*, SIAM J. Control Optim., 55, pp. 3319–3348, <https://doi.org/10.1137/15M1054109>.
- P. P. BOYLE AND T. VORST (1992), *Option replication in discrete time with transaction costs*, J. Finance, 47, pp. 271–293.
- A. BRIANI, F. CAMILLI, AND H. ZIDANI (2012), *Approximation schemes for monotone systems of nonlinear second order partial differential equations: Convergence result and error estimate*, Differ. Equ. Appl., 4, pp. 297–317.
- R. CARMONA (2008), *Indifference Pricing: Theory and Applications*, Princeton University Press.
- Á. CARTEA AND S. JAIMUNGAL (2015a), *Optimal execution with limit and market orders*, Quant. Finance, 15, pp. 1279–1291.
- Á. CARTEA AND S. JAIMUNGAL (2015b), *Risk metrics and fine tuning of high-frequency trading strategies*, Math. Finance, 25, pp. 576–611.
- Á. CARTEA AND S. JAIMUNGAL (2016), *Incorporating order-flow into optimal execution*, Math. Financ. Econ., 10, pp. 339–364.
- Á. CARTEA, S. JAIMUNGAL, AND J. PENALVA (2015), *Algorithmic and High-Frequency Trading*, Cambridge University Press.
- U. CETIN, R. A. JARROW, AND P. PROTTER (2004), *Liquidity risk and arbitrage pricing theory*, Finance Stoch., 8, pp. 311–341.
- U. CETIN, H. M. SONER, AND N. TOUZI (2010), *Option hedging for small investors under liquidity costs*, Finance Stoch., 14, pp. 317–341.
- J. CVITANIĆ AND I. KARATZAS (1996), *Hedging and portfolio optimization under transaction costs: A martingale approach*, Math. Finance, 6, pp. 133–165.
- M. H. A. DAVIS, V. G. PANAS, AND T. ZARIPHPOULOU (1993), *European option pricing with transaction costs*, SIAM J. Control Optim., 31, pp. 470–493, <https://doi.org/10.1137/0331022>.
- O. GUÉANT (2016), *The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making*, Vol. 33, CRC Press.
- O. GUÉANT (2017), *Optimal market making*, Appl. Math. Finance, 24, pp. 112–154.
- O. GUÉANT, C.-A. LEHALLE, AND J. FERNANDEZ-TAPIA (2012), *Optimal portfolio liquidation with limit orders*, SIAM J. Financial Math., 3, pp. 740–764, <https://doi.org/10.1137/110850475>.
- O. GUÉANT AND J. PU (2015), *Option pricing and hedging with execution costs and market impact*, Math. Finance, 27, pp. 803–831.
- H. E. LELAND (1985), *Option pricing and replication with transactions costs*, J. Finance, 40, pp. 1283–1301.
- P.-L. LIONS AND J.-M. LASRY (2007), *Large investor trading impacts on volatility*, in Paris-Princeton Lectures on Mathematical Finance 2004, Springer, pp. 173–190.
- A. MARCO AND P. ANTONIO (1994), *Dynamic hedging portfolios for derivative securities in the presence of large transaction costs*, Appl. Math. Finance, 1, pp. 165–194.
- B. MOHAMED (1994), *Simulations of transaction costs and optimal reheding*, Appl. Math. Finance, 1, pp. 49–62.
- F. NAUJOKAT AND N. WESTRAY (2011), *Curve following in illiquid markets*, Math. Financ. Econ., 4, pp. 299–335.

- H. PHAM (2009), *Continuous-Time Stochastic Control and Optimization with Financial Applications*, Stoch. Model. Appl. Probab. 61, Springer-Verlag.
- A. F. ROCH (2011), *Liquidity risk, price impacts and the replication problem*, Finance Stoch., 15, pp. 399–419.
- L. C. ROGERS AND S. SINGH (2010), *The cost of illiquidity and its effects on hedging*, Math. Finance, 20, pp. 597–615.
- H. M. SONER, S. E. SHREVE, AND J. CVITANIĆ (1995), *There is no nontrivial hedging portfolio for option pricing with transaction costs*, Ann. Appl. Probab., 5, pp. 327–355.