OPTIMAL ACCELERATED SHARE REPURCHASE

S. JAIMUNGAL, D. KINZEBULATOV, AND D.H. RUBISOV

ABSTRACT. An accelerated share repurchase (ASR) allows a firm to repurchase a significant portion of its shares immediately, while shifting the burden of reducing the impact and uncertainty in the trade to an intermediary. The intermediary must then purchase the shares from the market over several days, weeks, or as much as several months. Some contracts allow the intermediary to specify when the repurchase ends, at which point the corporation and the intermediary exchange the difference between the arrival price and the TWAP over the trading period plus a spread. Hence, the intermediary effectively has an American option embedded within an optimal execution problem. As a result, the firm receives a discounted spread relative to the no early exercise case. In this work, we address the intermediary's optimal execution and exit strategy taking into account the impact that trading has on the market. We demonstrate that it is optimal to exercise when the TWAP exceeds $\zeta(t)S_t$ where S_t is the fundamental price of the asset and $\zeta(t)$ is deterministic. Moreover, we develop a dimensional reduction of the stochastic control and stopping problem and implement an efficient numerical scheme to compute the optimal trading and exit strategies.

SJ and DK would like to thank NSERC for partially funding this work. SJ would also like to thank GRI for partially funding this work. The authors would like to thank Álvaro Cartea, Matt Davison, Ryan Donnelly, and Jason Ricci as well as participants at the Recent Advances in High Frequency and Algorithmic Trading conference at UCL for useful comments on various stages of this work. The opinions in this article are solely those of the authors and do not necessarily express those of BMO Capital Markets.

S. Jaimungal, Department of Statistical Sciences, University of Toronto.

D. Kinzebulatov, The Fields Institute for Mathematical Sciences.

D.H. Rubisov, BMO Capital Markets.

1. Introduction

Accelerated share repurchases (ASRs) allow a firm to repurchase a significant portion of its shares immediately, while shifting the burden of reducing the impact and uncertainty in the trade to an intermediary, typically an investment bank (see Bargeron et al. (2011) for an empirical analysis of ASR adoption by firms). Hence, for the intermediary it is a form of agency trading. Naturally, the immediate repurchase of a signification portion of a corporation's shares would cost the firm more than the prevailing share price – as other traders would demand a premium to sell their shares. This is sometimes called immediate price impact and can also be viewed from the perspective of orders walking through the limit order book (LOB) – the collection of time and price prioritized quotes. In addition to immediate price impact, once other traders notice a sequence of buy orders, this may induce an upward drift on the asset's midprice – often called a permanent price impact – although prices typically revert back to their un-adjusted prices after the sequence of trades. There is a very strong resemblance to the now classical optimal liquidation problem, first studied by Bertsimas and Lo (1998) (in discrete time), Almgren and Chriss (1999) and Almgren and Chriss (2001) (as continuous time limit of a discrete model), whereby a large order is broken into smaller orders throughout the day to balance price impact with the risk of uncertain prices in the future. There have been many subsequent studies of the optimal liquidation problem in the literature ranging from ever more complicated models of price impact (see e.g., Obizhaeva and Wang (2012), Gatheral (2010), Predoiu et al. (2011), Gatheral et al. (2012) and Graewe et al. (2013)) to those who instead allow the agent to trade using limit orders to avoid price impact, see e.g., Bayraktar and Ludkovski (2012) and Guéant and Lehalle (2013). Optimal liquidation problems can also be viewed as one side of the classical market making problem as in Ho and Stoll (1981), Avellaneda and Stoikov (2008), Cartea et al. (2011), Cartea and Jaimungal (2012), Guéant et al. (2012) and Cartea and Jaimungal (2013),.

In this article, we investigate for the first time the optimal trading strategies associated with an ASR containing an embedded early exercise option. Specifically, the firm pays up front the arrival price for the asset, and the intermediary provides them with those shares immediately. The intermediary then proceeds to reacquire shares from the market by trading over the next few days, weeks or as much as months. At any point in time, up to an agreed upon maturity date, the intermediary may call upon the firm to pay the difference between the arrival price and the time weighted average price (TWAP) of the asset[1], up to that point in time, plus a spread. As such, the intermediary has an embedded American option on the average price. The firm in exchange receives a reduced spread relative to the case when the

[1]In many real world contracts, the TWAP is often replaced by the volume weighted average price VWAP. It is possible to generalize our approach to this case, however, in the interest of maintaining simplicity in the approach for opt to work with TWAP.

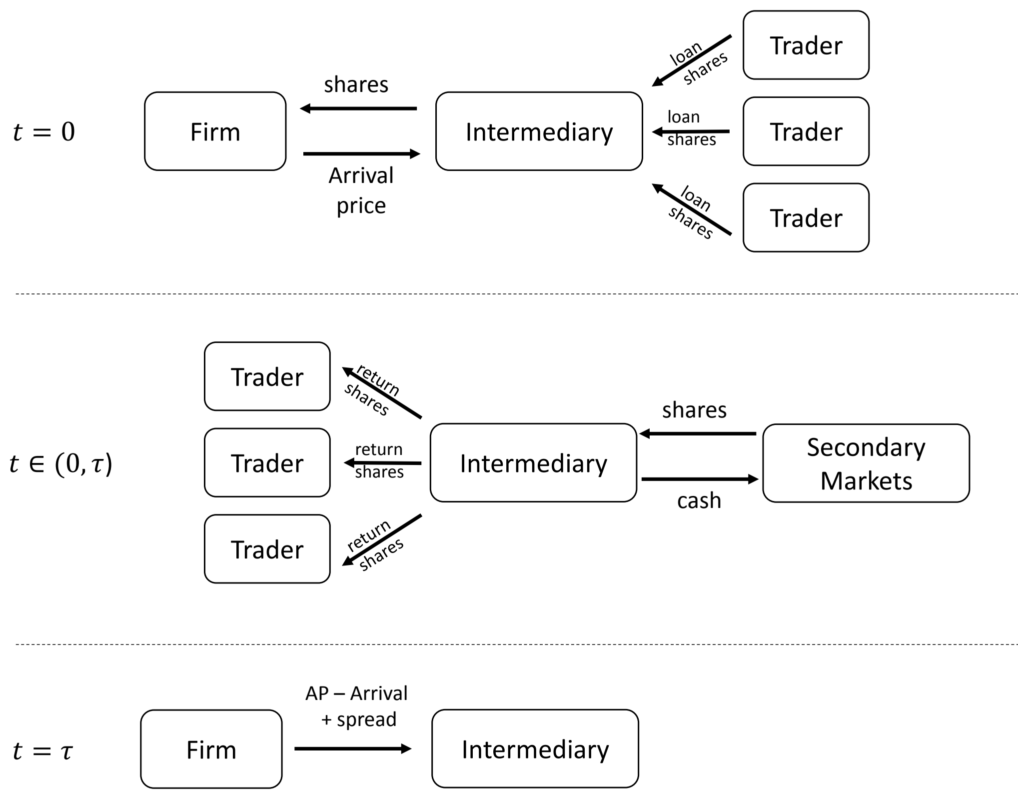

FIGURE 1. Schematic of the ASR transactions showing cash-flows and how shares are transferred. Note that the cash-flow in panel 3 may in fact stream from the intermediary to the firm.

intermediary has no early exercise clause, and may in fact receive a spread rather than pay one. Figure 1 summarizes the manner in which shares and cash are transferred between the intermediary and the firm. Note that the intermediary may initially borrow the shares from traders within the same institution (the typical scenario) or from external institutions. The shares are eventually replaced by purchasing them from the secondary markets. One may be tempted to view this embedded option simply as an Asian styled option on an asset, however, that approach would lead to undervaluation. The reason being that since a large number of shares are to be repurchased from the market, it is important to account for the impact that trading has on the asset's price. Moreover, the intermediary needs to determine the optimal trading strategy which realizes the best net expected payoff – including the cost of acquiring the shares and the TWAP spread over arrival price. Since often the embedded option and the optimal acquisition are managed by different desks, we will also investigate a simplified version of the problem which decouples the embedded optionality from the optimal acquisition problem.

We are able to derive dynamic programming equations (DPEs) for the valuation function H and demonstrate that these equations can be dimensionally reduced. In particular, the value functions are shown to all be of the form $S\,h(t, \frac{A}{S}, q)$ where S is the state variable corresponding to the asset's price, A is the state variable corresponding to TWAP up to time t, and q is the state variable corresponding to the remaining inventory to acquire. Moreover, we show that under some simplifying assumptions that the optimal stopping boundary corresponding to when the intermediary should end trading and receive the option payoff is characterised by the ratio of the TWAP to the fundamental price as follows

$$\tau = T \wedge \inf\left\{t : \tfrac{A_t}{S_t} \geq \zeta^*(t, q_t)\right\}$$

for some function $\zeta^* : {\mathbb{R}_+}^2 \mapsto \mathbb{R}_+$. Finally, we analyze the existence and uniqueness of the DPEs, provide bounds on the trading speed, and demonstrate that the trading speed is convex in A/S, and may monotonically increase or decrease in time and in volatility, depending on the urgency of the intermediary.

2. Market model

To this end, we assume that the security's *execution price* $\hat{S} = (\hat{S}_t)_{0\leq t\leq T}$ (per unit) is affected by the agent's (the intermediary) trading at time $0 < t < T$ via a *temporary market impact* so that

$$\hat{S}_t = S_t + a(S_t)\,\nu_t, \qquad \text{(execution price)} \tag{2.1}$$

where $S = (S_t)_{0\leq t\leq T}$ is the *fundamental price* of the security, $\nu = (\nu_t)_{0\leq t\leq T}$ is the rate at which the agent acquires the units of security at time t, and $a : \mathbb{R}_+ \mapsto \mathbb{R}_+$, continuous, determines the impact's strength. We further assume that the fundamental price is affected by the agent's trading through a *permanent market impact*. Specifically, S is a perturbed geometric Brownian motion, and satisfies the SDE

$$\frac{dS_t}{S_t} = b(S_t)\,\nu_t\,dt + \sigma\,dW_t, \qquad \text{(fundamental price)} \tag{2.2}$$

where $b : \mathbb{R}_+ \mapsto \mathbb{R}_+$, continuous, controls the permanent impact strength, $W = (W_t)_{0\leq t\leq T}$ is a standard Brownian motion, and $\sigma > 0$ is volatility. In the remainder, we work on the completed filtered probability space $(\Omega, \mathbb{P}, \mathcal{F} = \{\mathcal{F}_t\}_{0\leq t\leq T})$, where $\mathcal{F}$ is the natural filtration generated by W, and all stochastic processes are adapted to this filtration. Note that the Brownian motion W can be viewed as stemming from the actions that all other trading in the asset has on the assets's price. An important point is that typically $b(S) \ll a(S)$, $\forall S \in \mathbb{R}_+$, so that the permanent impact is much smaller than the temporary impact. Due to this order of magnitude difference in permanent versus temporary impact, later on we assume permanent impact is zero since the resulting analysis becomes much simpler and the main features of the solution still remain.

Another important ingredient in the analysis is the TWAP process $A = (A_t)_{0\le t\le T}$ which is directly computable from the asset's midprice

$$A_t = \frac{1}{t}\int_0^t S_u du\,. \qquad \text{(TWAP)} \tag{2.3}$$

This average is computed from the asset's fundamental price rather than the execution price $\hat{S}$ for two reasons: (i) the resulting strategies do not differ significantly between the two approaches, and (ii) the resulting analysis simplifies considerable when the fundamental price is used. In fact, on the time scales that trades are executed, orders are first broken into lots of say 100 and (for relatively liquid assets) rarely walk the through the book, and if they do, they walk through at most one or two levels. The temporary impact is therefore the result of several small orders being traded consecutively and not a walking of the book. Hence, using the midprice will be a good approximation to using all executed trade prices to compute TWAP.

The last two ingredients are the agent's inventory $Q = (Q_t)_{0\le t\le T}$,

$$Q_t = \int_0^t \nu_u\, du \qquad \text{(inventory)} \tag{2.4}$$

and the agent's remaining inventory to acquire $q = (q_t)_{0\le t\le T}$, i.e.,

$$q_t = \mathfrak{n} - Q_t, \qquad \text{(inventory remaining)} \tag{2.5}$$

where $\mathfrak{n} > 0$ is the total number of units of security that the agent has to acquire.

3. The Optimal Acquisition Problem

In this section we describe the details of the optimal control and stopping problem that underlie the ASR with TWAP optionality. As describe in the introduction, the agent (i.e., the intermediary) provides the firm with $\mathfrak{n}$ shares at the contract initiation. These shares are borrowed from traders and must be returned to them prior to the end of the contract. The agent is allowed to pick the time τ at which the contract terminates (up to a maximum time of T). At time τ the firm pays the agent the difference between the arrival price (the midprice of the asset at the contract start) and the TWAP up to τ. The agent returns the shares to their original owners by time $\tau + \varepsilon$, where $\varepsilon > 0$ is fixed a priori, and is small compared to T. For example, T may be several weeks, while ε may be a few days. After the exercise of the contract at time τ, the agent's P&L is no longer affected by changes in A_t (TWAP), hence, from τ to $\tau + \varepsilon$ she uses an Almgren-Chriss strategy in to acquire *all remaining* q_τ shares. This strategy balances the market impact with the need to have $q_{\tau+\varepsilon} = 0$. Furthermore, since ε is small compared to T, the agent may assume that on $[\tau, \tau + \varepsilon]$ the permanent market impact is negligible. As shown in Appendix

B, the value function associated with acquiring these remaining shares is $\ell_\tau = \ell(q_{\tau-}, S_\tau)$ where

$$\ell(q,S) = q\,(S + \alpha(S)\,q),$$

and $\alpha(S)$ is provided in Appendix B.

In all, at time τ the agent's profit and loss (PnL) can be written

$$\varphi^\nu_\tau = \underbrace{\mathfrak{n}\,(A_\tau - S_0)}_{\text{TWAP spread}} - \left(\underbrace{\int_0^\tau \nu_u\,\hat{S}^\nu_u\,du}_{\text{cost of trading to }\tau} + \underbrace{q_{\tau^-}\,(S^\nu_\tau + \alpha(S^\nu_\tau)\,q_{\tau^-})}_{\text{cost of acquiring remaining shares}} \right). \tag{3.1}$$

The agent's goal is to maximize the expected value of this payoff subject to a penalty on inventories different from the target. To this end, the agent's value function is defined as

$$H = \sup_{\tau\in\mathcal{T},\,\nu\in\mathcal{A}} \mathbb{E}_{0,S,S,\mathfrak{n}}\left[\varphi^\nu_\tau - \int_0^\tau \phi(S^\nu_u)\,q_u^2\,du\right], \tag{3.2}$$

where the notation $\mathbb{E}_{t,S,A,q}$ represents expectation conditional on $S_t = S$, $A_t = A$, $q_{t^-} = q$, $\mathcal{T}$ is the collection of $\mathcal{F}$-stopping times bounded by T and $\mathcal{A}$ is the set of admissible strategies consisting of non-negative, bounded, $\mathcal{F}$-predictable processes. The second term under the expectation represents the penalization on inventories different from the target inventory $\mathfrak{n}$, it may also be viewed as representing the impatience of the agent in which large ϕ induces the agent to trade faster. Cartea et al. (2013) also demonstrate that penalties of this kind can be viewed as stemming from an agent's aversion to model uncertainty, i.e. ambiguity aversion in the sense of Knightian uncertainty, on the asset's mid price process. In particular, they demonstrate, through a robust optimization approach with relative entropy penalization, that an ambiguity averse agent trades as an agent who is certain of their model but penalizes inventories different from their target.

In Section 8, we demonstrate that if $a(S) = a_0 S$, $\phi(S) = \phi_0 S$ for some a_0, $\phi_0 > 0$ and b is constant, then the optimal stopping time τ^*, has the representation (cf. Proposition 5.1 below)

$$\tau^* = \inf\{t\,:\,A_t \geq \zeta^*(t,q_t)S_t\},$$

where $\zeta^* : [0,T]\times[0,\mathfrak{n}] \to [1,\infty)$ is a continuous function, see Figure 2. This linearity assumption stems from the fundamental price of the asset being modeled as a GBM. If the penalty functions are not linear in S, then ζ^* will in general depend on (t,S,q). It is possible to restate all of our results in this case, however, we opt to reduce the problem to this more tractable form.

To solve the control problem, as usual, we introduce the more general problem where the agent starts acquisition at an arbitrary time $t \in [0,T]$ with arbitrary asset midprice, TWAP and inventory remaining. Without loss of generality, we can restrict the class of optimal controls ν_t to Markov controls $\nu_t = \nu(t, S_t, A_t, q_t)$, where a slight abuse of notation is used whereby in the r.h.s. ν is a deterministic function

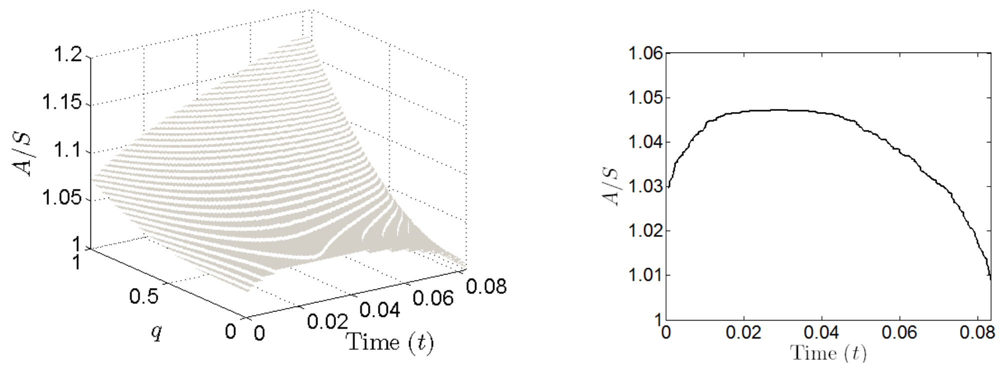

FIGURE 2. (Left panel) Graph of function $\zeta^*(t,q)$ (Right panel) Graph of function $\zeta(t)$.

of four variables. That is, the optimal trading rate ν_t is completely determined by current time t, current fundamental price S_t, current average A_t, and the remaining shares to acquire q_t. For a given admissible strategy $\nu \in \mathcal{A}$, we define the performance criteria as

$$H^{(\nu,\tau)}(t,S,A,q) = \mathbb{E}_{t,S,A,q}\left[\varphi_\tau^\nu - \int_t^\tau \phi(S_u^\nu)\, q_u^2\, du\right],$$

and seek the *value function*

$$\textbf{(P1)} \qquad H(t,S,A,q) = \sup_{\tau\in\mathcal{T},\, \nu\in\mathcal{A}} H^{(\nu,\tau)}(t,S,A,q)\,,$$

the optimal acquisition strategy ν^* and the optimal stopping time τ^* which attain the supremum (if they are within the set of admissible strategies and stopping times). The Dynamic Programming Principle (DPP), implies that the optimization problem (**P1**) can be reduced to a quasi-variational inequality for a 3-dimensional semi-linear hyperbolic PDE, called the *dynamic programming equation*, see Section 4. For now, suppose we have found the value function, then it is not difficult to see that the solution is characterised by the continuation , stopping, and exercise regions

$$\text{(3.3a)} \qquad C = \{\, (t,S,A,q)\;:\; H(t,S,A,q) > \mathfrak{n}(A - S_0) - \ell(q,S)\,\} \qquad \text{continuation region,}$$

$$\text{(3.3b)} \qquad D = \{\, (t,S,A,q)\;:\; H(t,S,A,q) \le \mathfrak{n}(A - S_0) - \ell(q,S)\,\} \qquad \text{stopping region,}$$

$$\text{(3.3c)} \qquad E = \{\, (t,S,A,q)\;:\; H(t,S,A,q) = \mathfrak{n}(A - S_0) - \ell(q,S)\,\} \qquad \text{exercise boundary.}$$

With these definitions, $\tau^*(t,S,A,q)$ is the first time the processes (t, S_t, A_t, q_{t-}) enter the stopping region D, or if the process begins within the continuation region C, then it is the first time that the processes touches the exercise boundary E.

A simplified problem – supplied exit strategy. In many cases, one desk in the intermediary firm is responsible for valuing the embedded option while another is responsible for the optimal acquisition problem. In such cases, the desk which values the embedded option will find an optimal exit strategy and provide this strategy to the desk which is acquiring the asset. In this sense, the combined control and stopping problems becomes sequential and the problem simplifies considerably. Such a treatment of the problem is clearly sub-optimal, however, it may be of interest to intermediaries who have so-called "Chinese-walls" between the various functions of the firm and is quite commonplace. Moreover, we find that the optimal strategies produce quite similar profit and losses, and hence this simplification which provides large efficiency gains in the computations is of practically interest.

In this case, assuming that the functions a, b, α, ϕ are linear/constant, it is possible to reduce the 3-dimensional free boundary problem to a 1-dimensional fixed boundary problem for a semi-linear (backward) *parabolic* PDE. We now describe precisely the approximate version of the optimization problem (**P1**). The intermediary first solves for the optimal stopping time $\widetilde{\tau}^*$ that maximizes only the expected payoff $\mathfrak{n}\, A_t$, and then uses this stopping time to find the optimal acquisition strategy $\widetilde{\nu}_t^*$ that maximizes the expected *final* payoff. More precisely, let $\widetilde{\tau}^*$ be a stopping time such that[2]

$$\mathbb{E}[A_{\widetilde{\tau}^*}] = \sup_{\tau\in\mathcal{T}} \mathbb{E}[A_\tau] \tag{3.4}$$

where A_τ is computed under the assumption that $dS_t = \sigma S_t dW_t$, i.e. we ignore the permanent price impact of the agent's trading on the dynamic of S_t ($b = 0$). The problem of finding $\widetilde{\tau}^*$ coincides with the problem of finding the optimal exercise strategy for the American-style option (with zero interest rates) paying A_τ at the time of exercise τ, which we address in Section B. This is a somewhat standard problem and it is well known that that $\widetilde{\tau}^*$ is characterized by an exercise boundary in the asset-price as a function of time:

$$\widetilde{\tau}^* = \inf\{t : A_t \geq \zeta(t)\, S_t\}$$

for a continuous function $\zeta : [0, T] \to [1, \infty)$ (cf. Figure 2). Next, having found the stopping time $\widetilde{\tau}^*$, the intermediary seeks the optimal rate of acquisition ν_t^* which maximizes the performance criteria

$$\widetilde{H}^{(\nu)}(t, S, A, q) = \mathbb{E}_{t,S,A,q}\left[\varphi^{\nu}_{\widetilde{\tau}^*} - \int_t^{\widetilde{\tau}^*} \phi(S_u^\nu)\, q_u^2\, du\right],$$

and the intermediaries value function is

$$\widetilde{H}(t, S, A, q) = \sup_{\nu\in\mathcal{A}} \widetilde{H}^{(\nu)}(t, S, A, q)\,. \tag{\textbf{P2}}$$

In Section 8 we will compare the solution of the simplified problem to the full problem.

---

[2]We assume that supremum is attained within $\mathcal{T}$.

4. Dynamic programming equations

We now proceed to solve the optimization problems. First, since the arrival price of the asset S_0 plays no part in the optimization, without loss of generalization, we set it to 0. Next, a straight forward application of the DPP, shows that a trial solution to the optimization problem (**P1**) should satisfy the quasi-variational inequality (where we have already substituted the optimal trading speed in feedback control form)

$$\begin{cases} \partial_t H + \frac{1}{2}\sigma^2 S^2 \partial_{SS} H + \frac{S-A}{t}\partial_A H - \phi(S) q^2 \\ \qquad\qquad + \frac{1}{4a(S)}\left(b(S) S\, \partial_S H - \partial_q H - S\right)^2 & \leq 0\,, \\ \qquad\qquad\qquad\qquad H(t,S,A,q) & \geq \mathfrak{n}\, A - \ell(q,S)\,, \end{cases} \tag{4.1a}$$

for all $0 < t < T$, $S > 0$, $A > 0$, $0 < q < \mathfrak{n}$, where only one of the inequalities in (4.1a) is strict at all times, subject to terminal condition

$$H(T,S,A,q) = \mathfrak{n}\, A - \ell(q,S) \tag{4.1b}$$

and boundary conditions, along the hyperplane $q = 0$,

$$H(t,S,A,0) = \mathfrak{n}\, S\, w\big(t, \tfrac{A}{S}\big)\,, \tag{4.1c}$$

and where $S\, w(\cdot,\cdot)$ is the value of an American option paying $\mathfrak{n}\, A_\tau$ at the time of exercise – the explicit form of the function w is provided in Proposition B.1. This boundary condition represents the fact that once the intermediary has acquired all inventory, they cease trading, but the stopping clause of the ASR might not be exercised at this time, hence the intermediary still holds an American option and this value is what the intermediary receives along the boundary $q = 0$. Moreover, the optimal acquisition strategy ν^* and the optimal stopping time τ^* for the optimization problem (**P1**) are given in feedback control form as follows:

$$\nu^*(t,S,A,q) = \frac{1}{2\,a(S)}\left(b(S) S\, \partial_S H - \partial_q H - S\right) \quad (0 < q \leq \mathfrak{n})\,, \quad \nu^*(t,S,A,0) = 0\,,$$

$$\tau^*(t,S,A,q) = \inf\{u \in [t,T] : H(u,S_u,A_u,q_u) \geq \mathfrak{n} A_u \mid S_t = S, A_t = A, q_t = q\}\,.$$

To solve the TBV problem (4.1a)–(4.1c) numerically, we make a change of variables $(S,A) \mapsto (S,z)$, $z = \frac{A}{S}$, which is equivalent to using the asset's midprice as a numeraire, and equip the resulting TBV problem with the additional boundary conditions in Section 6(4). In our numerical experiments, we have found that the optimal solution is typically to exercise the stopping clause prior to T and prior to acquiring all $\mathfrak{n}$ shares. Once the stopping clause is exercised, the agent then continues to acquire the remaining shares (as in Almgren-Chriss) up to a maximum amount of time ε later. Various features of the solution will be explored in section 8.

Next, for the simplified problem (**P2**), through the DPP, the trial solution for the value function $\widetilde{H}$ solves the fixed-boundary-value problem (where again we have already substituted the optimal trading strategy in feedback control form)

$$
\begin{cases}
\partial_t \widetilde{H} + \frac{1}{2}\sigma^2 S^2\, \partial_{SS}\widetilde{H} + \frac{S-A}{t}\,\partial_A \widetilde{H} - \phi(S)\, q^2 \\
\qquad + \frac{1}{4a(S)} \left(b(S) S\, \partial_S \widetilde{H} - \partial_q \widetilde{H} - S \right)^2 & = \quad 0, \\
\widetilde{H}(T,S,A,q) & = \quad \mathfrak{n}\, A - \ell(q,S), \\
\widetilde{H}(t,S,\zeta(t)\, S, q) & = \quad \mathfrak{n}\, \zeta(t)\, S - \ell(q,S), \\
\widetilde{H}(t,S,A,0) & = \quad \mathfrak{n}\, A.
\end{cases}
\tag{4.2}
$$

The terminal and boundary conditions reflect the fact that the agent acquires the remaining shares at the stopping time characterized by the boundary $A^*(t) = \zeta(t)\, S(t)$. In order to solve TBV problem (4.1a)–(4.1c) numerically, we impose additional boundary conditions on $\widetilde{H}$ in Section 6(5). Moreover, the optimal acquisition strategy $\widetilde{\nu}^*$ in feedback control form is given by

$$
\widetilde{\nu}^*(t,S,A,q) = \frac{1}{2\,a(S)} \left(b(S) S\, \partial_S \widetilde{H} - \partial_q \widetilde{H} - S \right).
$$

5. Dimensional reduction

The DPEs for (**P1**) and (**P2**) have $3+1$ dimensions, however, under certain simplifying assumptions, the dimension of the DPEs can be reduced considerably and lead to more tractable equations. Here, we provide those assumptions and the resulting dimensional reductions.

Proposition 5.1 (Dimensional reduction of **P1**). *Assume that $a(S) = a_0 S$, $\phi = \phi_0 S$, and $b(S) = b_0$. Then $\alpha(S) = \alpha_0\, S$, where α_0 is given in* (A.4)*, and*

$$
H^*(t,S,A,q) = S\, h(t,z,q), \qquad z = A/S,
$$

where h solves

$$
\begin{cases}
\partial_t h + \frac{1}{2}\sigma^2 z^2 \partial_{zz} h + \frac{1-z}{t}\partial_z h - \phi_0 q^2 \\
\qquad + \frac{1}{4a_0}\left(b_0 (h - z\partial_z h) - \partial_q h - 1\right)^2 & \leq 0 \\
h(t,z,q) & \geq z - q(1 + \alpha_0 q)
\end{cases}
\tag{5.1a}
$$

for all $0 < t < T$, $0 < q < \mathfrak{n}$, and only one of the inequalities in (5.1a) is strict at all times, subject to terminal and boundary conditions

$$
h(T,z,q) = \mathfrak{n}\, z - q(1 + \alpha_0\, q), \ \ \textit{and } h(t,z,0) = \mathfrak{n}\, w(t,z),
\tag{5.1b}
$$

where w is the function of Proposition B.1. The optimal acquisition strategy ν^* is given by

$$\nu^*(t,S,A,q) = \tfrac{1}{2a_0}\left(b_0\left(h(t,A/S,q) - \tfrac{A}{S}(\partial_z h)(t,A/S,q)\right) - \partial_q h(t,A/S,q) - 1\right) \quad \textit{for } q>0,$$

and $\nu^*(t,S,A,0)=0$.

To solve the QVI (5.1) numerically, we impose additional boundary conditions on h in Section 6.2 at the boundaries of the domain that the numerical solution is sought.

Interestingly, the ansatz of Proposition 5.1 is also applicable to the optimization problem (**P2**).

Proposition 5.2 (Dimensional reduction of (**P2**))**.** *Suppose that the assumptions of Proposition 5.1 are satisfied. Then $\alpha(S)=\alpha_0 S$, where α_0 is given in* (A.4)*, and*

$$\widetilde{H}^*(t,S,A,q) = S\,\widetilde{h}(t,z,q), \quad z = A/S,$$

where $\widetilde{h}$ solves

$$\text{(5.2)} \qquad \begin{cases} \partial_t\widetilde{h} + \frac{1}{2}\sigma^2 z^2\partial_{zz}\widetilde{h} + \frac{1-z}{t}\partial_z\widetilde{h} - \phi_0 q^2 & \\ \qquad + \frac{1}{4a_0}\left(b_0(\widetilde{h} - z\partial_z\widetilde{h}) - \partial_q\widetilde{h} - 1\right)^2 &= 0, \\ \widetilde{h}(T,z,q) &= \mathfrak{n}\,z - q(1+\alpha_0 q), \\ \widetilde{h}(T,z,0) &= \mathfrak{n}\,z, \\ \widetilde{h}(T,\zeta(t),0) &= \mathfrak{n}\,\zeta(t) - q(1+\alpha_0 q), \end{cases}$$

for all $0<t<T$, $0<q<\mathfrak{n}$. The optimal acquisition strategy ν^* is given by the formula

$$\nu^*(t,S,A,q) = \tfrac{1}{2a_0}\left(b_0\left(h(t,A/S,q) - \tfrac{A}{S}(\partial_z h)(t,A/S,q)\right) - \partial_q h(t,A/S,q) - 1\right), \quad q>0$$

and $\nu^*(t,S,A,0)=0$.

In fact, for problem (**P2**) we can go one step further, and reduce (5.2) to a TBV problem for a $(1+1)$-dimensional PDE, subject to an additional simplifying assumption. If we ignore the effect of agent's trading when computing the term $\mathfrak{n}\,A_{\widetilde{\tau}^*}$ in the optimization problem (**P2**) – i.e., the derivatives desk only models the fundamental price, then the problem reduces further.

Proposition 5.3 (Further dimensional reduction of (**P2**))**.** *Assume that $a(S)=a_0S$, $\phi=\phi_0S$, and $b(S)=0$, i.e. the permanent market impact is negligible. Then $\alpha(S)=\alpha_0 S$, where α_0 is given in* (A.4)*, and the value function*

$$\text{(5.3)} \qquad \widetilde{H}^*(t,S,A,q) = \mathfrak{n}\,S\,\omega(t,S,A) - S\,(q + q^2 g(t,z)), \quad z = A/S,$$

where

$$
(5.4)\qquad \begin{cases} \partial_t g + \frac{1}{2}\sigma^2 z^2 \partial_{zz} g + \frac{1-z}{t}\partial_z g + \phi_0 - \frac{1}{a_0} g^2 & = & 0, \\ g(T,z) & = & \alpha_0, \\ g(t,\zeta(t)) & = & \alpha_0, \end{cases}
$$

with $0 < t < T$, $z < \zeta(t)$. *Moreover, the optimal trading speed simplifies to*

$$
(5.5)\qquad \widetilde{\nu}^*(t,S,A,q) = \frac{1}{a_0}\, q\, g\,(t, A/S)\,.
$$

We note that in addition to being a $1+1$ PDE, equation (5.4), unlike (5.1) and (5.2), is a (backward) *parabolic* semi-linear PDE, which leads to a considerably simpler numerical procedure for finding the optimal acquisition strategy.

6. Solutions over bounded domain. Existence and uniqueness.

In this section, we address the numerical solution of the DPEs and focus only the dimensional reduced cases.

We begin with investigating the dimensionally reduced TBV problem (5.1) for problem (**P1**) arising in Proposition 5.1. We restrict the problem to a bounded domain, large enough to include the, as yet, unknown boundary. In this case, the region is given by

$$
0 < t < T, \quad 0 < q < \mathfrak{n}, \quad \underline{z} < z < \overline{z},
$$

and we impose additional the boundary conditions

$$
(6.1a)\qquad \partial_z h(t,\underline{z},q) = 0, \quad \partial_{qq} h(t,z,\mathfrak{n}) = 0, \quad \text{and}
$$

$$
(6.1b)\qquad h(t,\zeta^*(t,q),q) = \mathfrak{n}\, z - q(1+\alpha_0 q).
$$

where $\zeta^*(t,q) < \overline{z}$ is the minimal value of z such that $h(t,z,q) = \mathfrak{n}\, z - q(1+\alpha_0\, q)$, i.e., the optimal exercise boundary (cf. Figure 2). The first boundary condition is independent of z for small values of z. The rationale for this additional boundary condition is that when $z = \underline{z}$, where $\underline{z}$ is sufficiently small, the trading strategy is far away from the early exercise barrier $z = \zeta(t)$, hence the intermediary behaves as if there is no early exercise option and g becomes independent of z. The second boundary condition imposes that h is linear in q when the maximal inventory is remaining. The third boundary condition in (6.1) is the continuous pasting condition. The numerical solution of the QVI (5.1), (6.1) is solved using an explicit-implicit scheme where the non-linear terms are treated explicitly.

Next, we focus on the TBV problem (5.2) of the further dimensional reduction of (**P2**) provided in Proposition 5.2. In this case, we select the bounded domain to be

$$0 < t < T, \quad 0 < q < \mathfrak{n}, \quad \underline{z} < z < \zeta(t), \tag{6.2}$$

and require the analogous boundary conditions

$$\partial_z \widetilde{h}(t, \underline{z}, q) = 0, \quad \partial_{qq} \widetilde{h}(t, z, \mathfrak{n}) = 0.$$

Standard results imply that TBV problem (5.2) considered over bounded domain (6.2), complemented with the above boundary conditions, has a unique continuous viscosity solution Crandall et al. (1992).

Finally, we investigate the dimensionally reduced TBV problem (5.4) for problem (**P2**) (arising in Proposition 5.3). For this case, we will solve the problem on the bounded domain $0 < t < T, \quad \underline{z} < z < \zeta(t)$, where the lower bound $0 \leq \underline{z} < 1$, and require that along the lower boundary

$$\partial_z g(t, \underline{z}) = 0, \tag{6.3}$$

The first boundary condition is motivated by the same considerations as in the previous domain restriction – far from the boundary, the intermediary behaves as if there is no early exercise clause and the solution is independent of z. Indeed when $\underline{z} = 0$, on substituting (6.3) into the dimensionally reduced DPE (5.4) and writing

$$g(t, 0) = \omega(t), \tag{6.4}$$

then $\omega(t)$ satisfies the ODE $\omega' + \phi_0 - \frac{1}{a_0}\omega^2 = 0$ with $\omega(T) = \alpha_0$. This ODE can be solved explicitly to find

$$\omega(t) = \begin{cases} \sqrt{\phi_0 a_0}\, \dfrac{\xi\, e^{2\gamma(T-t)} + 1}{\xi\, e^{2\gamma(T-t)} - 1}, & \phi > 0, \\ \left(\dfrac{1}{\alpha_0} + \dfrac{T-t}{a_0} \right)^{-1}, & \phi = 0, \end{cases} \tag{6.5}$$

with $\xi = (\alpha_0 - \sqrt{\phi_0 a_0})/(\alpha_0 + \sqrt{\phi_0 a_0})$ and $\gamma = \sqrt{\phi_0 / a_0}$. Moreover, $g(t, z) := \omega(t)$ for all $0 < t < T$, $0 < z < \infty$ is the solution of problem (5.4) when there is no barrier (i.e., $\zeta(t) = +\infty$ for all t). We will see later on that ω in fact provides a bound on the optimal acquisition rate (cf Theorem 7.1).

We assume that the boundary function $\zeta(t)$ is smooth, or if not, we can approximate arbitrarily close with a smooth approximation. Then, standard results (see, e.g. (Friedman, 1964, 7.4)) imply that TBV problem (5.4) with $0 < \underline{z} \leq z \leq \overline{z}$, complemented with boundary condition $g(t, \underline{z}) = \omega(t)$, has a unique strong solution for $0 < t \leq T$, $\underline{z} \leq z \leq \zeta(t)$ that is C^2-smooth (in z) in the interior of the domain. For the numerical implementation, the TBV problem (5.4), equipped with the above additional boundary conditions, uses an explicit scheme for the nonlinear terms and an implicit-explicit scheme for the linear terms.

Remark 6.1. The full QVI (4.1) can be solved numerically as well. For this purpose, we construct its discrete approximations by first allowing early exercise only at the fixed set of times $0 = t_0 < t_1 < \cdots < t_n = T$. The approximate value function is then given by

$$H^{(n)}(t,S,A,q) := H_i^{(n)}(t,S,A,q), \quad \text{for } t \in (t_{i-1}, t_i]$$

where within each interval, the agent solves the optimal control problem

$$\begin{aligned}(6.6)\quad H_i^{(n)}(t,S,A,q) = \sup_{\{\nu_{u,i}\}_{u\in[t,t_i)}\in\mathcal{A}_i} \mathbb{E}_{t,S,A,q}\Bigg[&-\int_t^{t_i}\nu_{u,i}\big(a(S_u)\nu_{u,i}+S_u\big)du - \int_t^{t_i}\phi(S_u)\,q_u^2 du \\ &+ \max\big\{\mathfrak{n}\,A_{t_i} - \ell(q_{t_i},S_{t_i})\,,\, H_{i+1}^{(n)}(t_i,S_{t_i},A_{t_i},q_{t_i})\big\}\Bigg],\end{aligned}$$

where $\mathcal{A}_i$ is the admissible set of trading strategies (on the interval $[t_{i-1},t_i)$) in which ν is non-negative and uniformly bounded from above. Let $\nu_{t,i}^{(n)*}$ be the solution of problem (6.6). We define

$$\nu_t^{(n)*} := \nu_{t,i}^{(n)*} \qquad \text{for } t_{i-1} < t \leq t_i,$$

i.e., the approximate optimal control consists of pasting together the optimal controls on the sub-intervals, and the optimal exercise time is defined as

$$\tau^{(n)*} := \min\{t_i : \mathfrak{n}A_{t_i} - \ell(q_{t_i},S_{t_i}) \geq H_{i+1}(t_i,S_{t_i},A_{t_i},q_{t_i})\}.$$

As the norm of the partition, $\max_i\{t_i - t_{i-1}\}$, tends to zero, $H^{(n)}$, $\nu^{(n)*}$, $\tau^{(n)*}$ converge to, respectively, the value function H, the optimal acquisition strategy ν^* and the optimal stopping time τ^* for problem (4.1).

It remains to describe how to compute the sequence of approximations $H_i^{(n)}$. For this, we employ the DPP within the interval $[t_{i-1},t_i)$ to obtain the sequence of DPEs

$$(6.7)\quad \begin{cases} \partial_t H_i^{(n)} + \frac{1}{2}\sigma^2 S^2\,\partial_{SS}H_i^{(n)} + \frac{S-A}{t}\,\partial_A H_i^{(n)} - \phi(S)\,q^2 & \\ \qquad +\frac{1}{4a(S)}\left(b(S)\,S\,\partial_S H_i^{(n)} - \partial_q H_i^{(n)} - S\right)^2 & = \; 0 \\ \qquad\qquad H_i^{(n)}(t_i,S,A,q) & = \; \max\{\mathfrak{n}\,A - \ell(q,S)\,,\, H_{i+1}^{(n)}(t_i,S,A,q)\}, \end{cases}$$

for all $t_{i-1} \leq t \leq t_i$, $S \geq 0$, $A \geq 0$, $0 \leq q \leq \mathfrak{n}$, subject to the boundary condition (4.1c). To solve this TBV problem numerically, it is convenient to make a change of variables $(S,A) \mapsto (S,z)$, $z = A/S$ which reduces the PDE to $1+1$-dimensions. This TBV is then solved numerically by implementing an explicit-implicit finite-difference scheme, whereby the non-linear terms are treated explicitly, and the linear terms are treated with a Crank-Nicholson scheme.

For the collection of parameters used in the numerical experiments (cf. Section 8) the implicit-explicit finite difference scheme applied to the linear part of (5.1a), with the nonlinear part treated explicitly, appears to converge, however, the question of whether such a finite difference scheme for a nonlinear PDE does indeed converge to the viscosity solution is not trivial, and requires careful study, cf. Forsyth and Vetzal (2012). We do not focus on analyzing this issue here and instead assume the scheme does converge to the viscosity solution and explore the resulting financial intuition gained from the simulations.

7. Estimates on the optimal acquisition strategy

In this section, we investigate various bounds on the value function as well as the behaviour of the optimal strategy with respect to key model parameters.

Under the assumptions of Proposition 5.3, the value function can be bounded by the analog of the classical Almgren-Chriss strategy and, moreover, it exhibits monotonicity in parameters (subject to a relative ordering) as stated in the theorem below.

Theorem 7.1 (Bounds on the optimal strategy)**.** *Suppose that the assumptions of Proposition 5.3 are satisfied, and* $\underline{z} = 0$.

(*i*) *If* $\phi_0 < \frac{\alpha_0^2}{a_0}$, *then* $\widetilde{\nu}^*(t, S, A, q)$ *is monotonically increasing in* $\frac{A}{S}$ *and*

$$\frac{\omega(t)}{a_0}\, q \le \widetilde{\nu}^*(t, A, S, q) \le \frac{\alpha_0}{a_0}\, q\,, \tag{7.1a}$$

where ω *is given by* (6.5).

If $\phi_0 > \frac{\alpha_0^2}{a_0}$, *then* $\widetilde{\nu}^*(t, S, A, q)$ *is monotonically decreasing in* $\frac{A}{S}$ *and*

$$\frac{\alpha_0}{a_0}\, q \le \widetilde{\nu}^*(t, A, S, q) \le \frac{\omega(t)}{a_0}\, q\,. \tag{7.1b}$$

If $\phi_0 = \frac{\alpha_0^2}{a_0}$, *then*

$$\widetilde{\nu}^*(t, A, S, q) = \frac{\alpha_0}{a_0} q\,. \tag{7.1c}$$

(*ii*) *Let* $\widetilde{\nu}^*_{(i)}$, $i = 1, 2$, *be the optimal acquisition strategies, cf.* (5.5), *corresponding to the pair of collection of parameters* $\{a_0^{(i)},\, \alpha_0^{(i)},\, \phi_0^{(i)}\}$, $i = 1, 2$. *If*

$$\frac{\phi_0^{(1)}}{a_0^{(1)}} \le \frac{\phi_0^{(2)}}{a_0^{(2)}}\,, \quad \text{and} \quad \frac{\alpha_0^{(1)}}{a_0^{(1)}} \le \frac{\alpha_0^{(2)}}{a_0^{(2)}}\,, \tag{7.2}$$

then

$$\widetilde{\nu}^*_{(1)}(t, S, A, q) \le \widetilde{\nu}^*_{(2)}(t, S, A, q) \tag{7.3}$$

for all (t, S, A, q). If one of the inequalities in (7.2) is strict, then the inequality in (7.3) is strict as well.

Under the assumptions of Proposition 5.3, we can reduce the optimal strategy to the solution of a simple non-linear ODE.

Theorem 7.2 (The perpetual case)**.** *Suppose that the assumptions of Proposition 5.3 are satisfied, $\underline{z} = 0$, $\zeta(t) \equiv z_* > 1$. Then*

$$\widetilde{\nu}^*(t, S, A, q) \xrightarrow[T\to\infty]{} \frac{1}{a_0} q\,\gamma(A/S) \quad \textit{for } t < +\infty$$

uniformly with respect to $\frac{A}{S} \in [0, z_*]$, $q \in [0, \mathfrak{n}]$, and γ satisfies

$$\tfrac{1}{2}\sigma^2 z^2\ \gamma''(z) + \phi_0 - \tfrac{1}{a_0}\gamma^2(z) = 0, \qquad \gamma(0) = \sqrt{\phi_0 a_0}, \quad \textit{and} \qquad \gamma(z_*) = \alpha_0. \tag{7.4}$$

In particular, it follows that

(1) *if $\phi_0 = 0$, then $\widetilde{\nu}^*$ tends to a convex function of $z = A/S$.*
(2) *if $\phi_0 \geq \frac{\alpha_0^2}{a_0}$, then $\widetilde{\nu}^*$ tends to a concave function of $z = A/S$.*

Once again under the assumptions of Proposition 5.3, the optimal strategy exhibits monotonicity in the ratio of average price to fundamental price. Moreover, as volatility increases, the strategy exhibits monotonicity in time, fundamental price, average price and inventory remaining.

Theorem 7.3. *Suppose that the assumptions of Proposition 5.3 are satisfied, and $\underline{z} = 0$.*

If $\phi_0 < \frac{\alpha_0^2}{a_0}$, then for all $0 < t \leq T$ such that

$$\frac{2}{t} + \frac{2}{a_0}\omega(t) \geq \sigma^2$$

the optimal acquisition strategy $\widetilde{\nu}^*(t, S, A, q)$ is convex in $\frac{A}{S}$, and is monotonically increasing for all t as above, S, A and q, as volatility $\sigma > 0$ increases.

If $\phi_0 > \frac{\alpha_0^2}{a_0}$, then for all $0 < t \leq T$ such that

$$\frac{2}{t} + \frac{2}{a_0}\alpha \geq \sigma^2$$

the optimal acquisition strategy $\widetilde{\nu}^*(t, S, A, q)$ is concave in $\frac{A}{S}$, and is monotonically decreasing for all t as above, S, A and q, as volatility $\sigma > 0$ increases.

If $b_0 > 0$ and, hence, the assumptions of Proposition 5.3 are not satisfied, then both assertions of Theorem 7.3 fail. For instance, if $b_0 \gg a_0$, then the agent has to slow down their trading as the ratio $\frac{A}{S}$ approaches the barrier $\zeta(t)$, so as not to drive the asset price further up, and avoid paying a penalty.

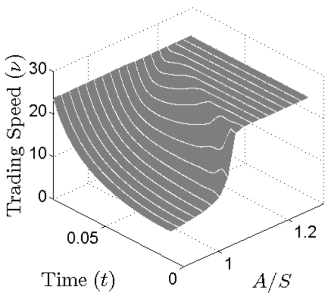

(a) (**P1**) trading speed

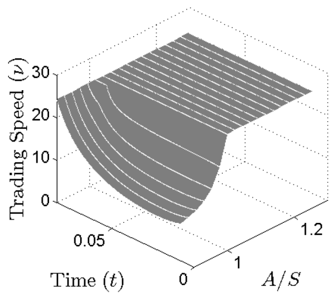

(b) (**P2**) trading speed

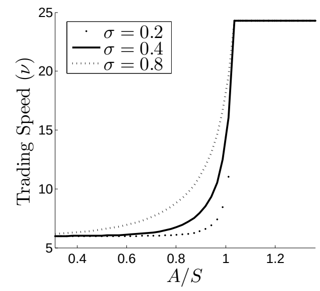

(c) (**P2**) trading speed

FIGURE 3. The optimal rate of acquisition when $q = 0.5$ (half the inventory remaining). Panels (a) and (b) for problems (**P1**) and (**P2**), respectively, as a function of time and A/S. Panel (c) for (**P2**) when $t = T/5$ for different volatilities.

8. Numerical experiments

In this section we illustrate several features of the optimal acquisition strategies ν^* and $\widetilde{\nu}^*$ for the full optimization problem (**P1**) and the simplified optimization problem (**P2**) (where the exit strategy is provided by the embedded American option), respectively.

To compute the optimal strategies for (**P1**), we apply the ansatz in Proposition 5.1 and numerically solve the TBV problem (5.1) as outlined in section 6. To compute the optimal strategies for (**P2**), we apply the ansatz in Propostion 5.3 and solve the TBV problem (5.4). Throughout this section we assume that volatility $\sigma = 0.4$, target inventory $\mathfrak{n} = 1$, running penalty $\phi(S) = \phi_0 S$ where $\phi = 0.1$, temporary market impact $a(S) = a_0 S$ where $a_0 = 5 \cdot 10^{-3}$, penalty $\alpha(S) = \alpha_0 S$, where α_0 is given by formula (A.4) with $\phi_1 = 1.5$ and $\phi_0 = 0.15$. As well, we set the permanent market impact $b(S) = 0$, since the optimal strategies exhibits similar qualitative behaviour with or without the permanent impact and in sufficiently liquid markets this impact is typically small.

To begin, we first focus on (**P2**) since the results are easier to explain and interpret. Panel (b) of Figure 3 shows the optimal execution strategy $\widetilde{\nu}^*$ for fixed $q = 0.5$. The strategy has two interesting features (i) as maturity approaches the rate increases and approaches a finite limit, and (ii) as the relative price A/S approaches the pre-specified exercise boundary $\widetilde{E}_t = \{(S, A, q) : \frac{A}{S} = \zeta(t)\}$, the rate increases and approaches the same finite limit. The rate of trading increases in both cases simply because for a fixed number of shares remaining, the agent must acquire shares faster as the strategy nears its exit time (either at maturity or at the boundary $\widetilde{E}_t$) to avoid the terminal penalty cost. While the finite limiting rate of trading are those implied by Theorem 7.1(i).

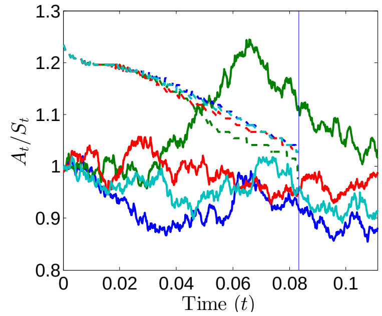

(a) Price ratio

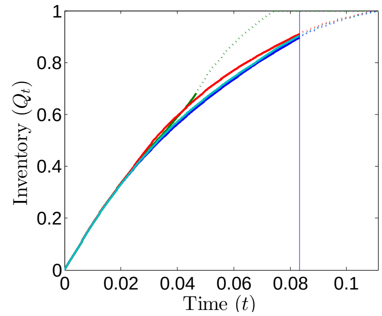

(b) Inventory

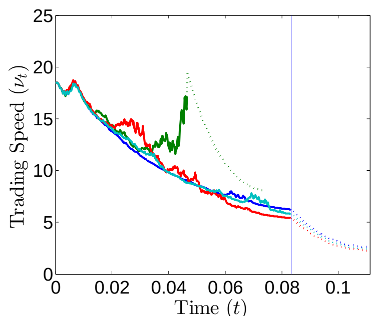

(c) Trading Speed

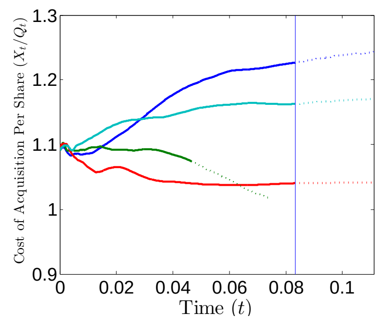

(d) Cost / Share

FIGURE 4. Sample paths of fundamental price S_t, time-weighted average price A_t, optimal acquisition rate ν_t^* and inventory $Q_t = \mathfrak{n} - q_t$ for problem (**P1**). The dashed lines in the last figure show the evolution of the exercise barrier $A_t/\zeta^*(t, q_t)$ versus the evolution of the fundamental price F_t.

When the exercise boundary $\zeta(t) = +\infty$, so that there is no upper barrier, the agent's optimal strategy $\widetilde{\nu}^*$ is independent of the asset's volatility. More precisely, it is given by the explicit formula $\widetilde{\nu}_t^* = \frac{\omega(t)}{a_0} q_t$ where $\omega(t)$ is provided in (6.5) . Contrastingly, when the boundary $\zeta(t) < \infty$, the agent's behaviour does indeed depend on volatility. Panel (c) of Figure 3 shows this dependence at $t = \frac{T}{2}$ and $q = 0.5$. As the diagram shows, the more volatile the market, the faster the agent trades, however, far from the barrier, volatility has little effect.

Next, panel (a) of Figure 3 shows the optimal acquisition strategy ν^* for problem (**P1**) at $q = 0.5$ fixed. Its qualitative and numerical behaviour is very similar to that of $\widetilde{\nu}^*$ in panel (b). A priori, it is not obvious whether the reduced problem would be a good approximation for the complete problem, yet, these numerical experiments demonstrate that this is generally the case (at least with the set of model parameters we explored). This is somewhat surprising because the complete problem contains an exercise boundary which depends on both A/S and q, while the reduced problem depends only on A/S (compare left and right panels of Figure 2).

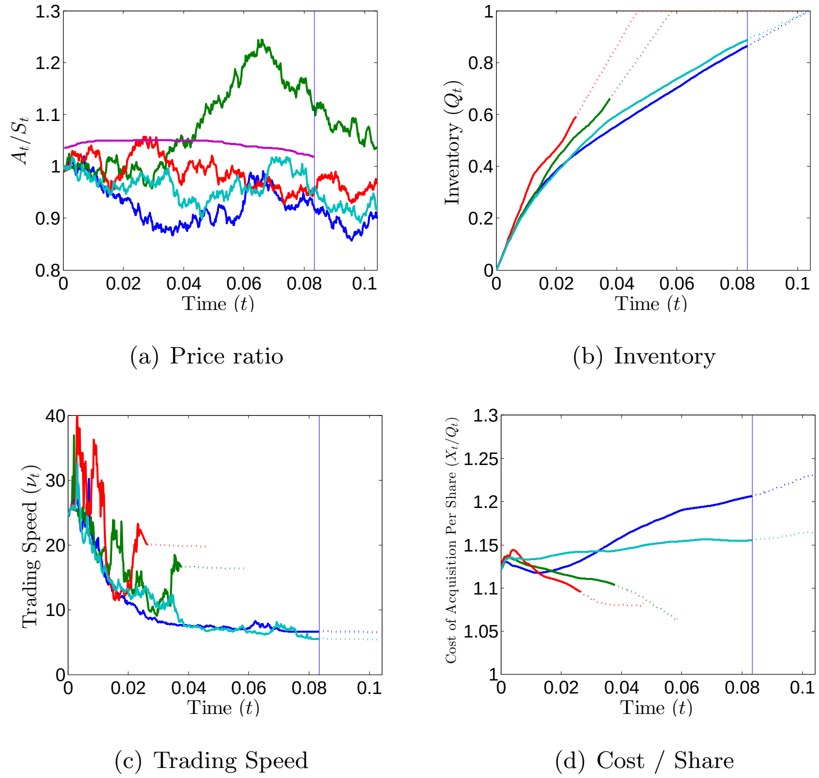

(a) Price ratio

(b) Inventory

(c) Trading Speed

(d) Cost / Share

FIGURE 5. Sample paths of fundamental price S_t, time-weighted average price A_t, optimal acquisition rate $\widetilde{\nu}_t^*$ and inventory $Q_t = \mathfrak{n} - q_t$ for problem (**P2**). The dashed lines in the last figure show the evolution of the exercise barrier $A_t/\zeta(t)$ versus the evolution of the fundamental price F_t.

We next explore the dynamical features of the strategies. When the agent runs a strategy, the inventory and fundamental price will all evolve, hence the static views in Figures 3 do not tell the entire story. To gain additional insight into the dynamic behaviour, in Figures 4 and 5 we plot four sample paths of the fundamental price S_t, ratio of the time-weighted average price A_t and S_t, the rate of acquisition ν_t, inventory $Q_t = \mathfrak{n} - q_t$ and cost per share X_t/Q_t, for (**P1**) and (**P2**), respectively. In the simulations, when a path hits the barrier or terminal time T is reached, the agent terminates the contract and switches to an Almgren-Chriss strategy to acquire the remaining inventory. The dotted lines in the figures indicate the corresponding fragment of the path. The blue and cyan paths are mostly far from the barrier and after some initial noise, due to the starting point being somewhat close to the barrier, the agent follows an essentially deterministic strategy similar to Almgren-Chris. In Figure 5 the red and green paths hit the pre-specified exercise boundary $\widetilde{E}_t$ early on and have spikes in their acquisition rates near the time in which the barrier is breached. These paths also attains less inventory prior to hitting the barrier. While in Figure 4, the green path hits the optimal exercise boundary later on and the red path is not exercised

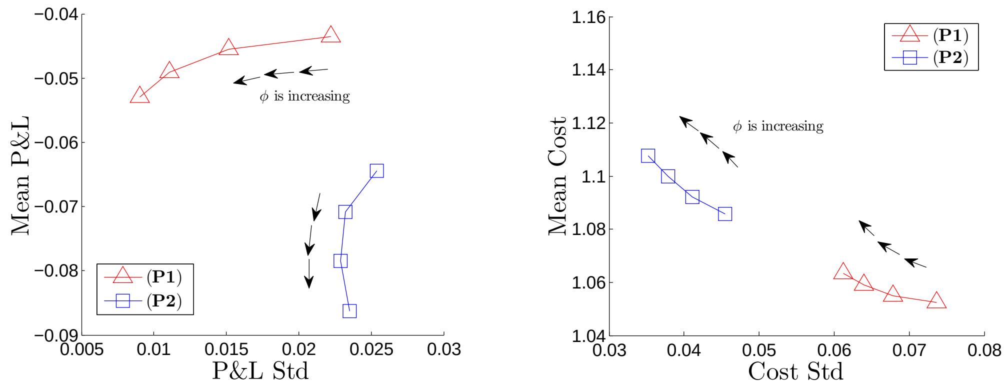

FIGURE 6. The profit-and-loss and acquisition cost risk-return profiles as the running penalty ϕ_0 varies ($\phi_0 = 0.01, 0.5, 1, 1.5$ and $\phi_1 = 10\,\phi_0$).

early at all. Thus, as expected, although the trading speeds in the two problems (**P1**) and (**P2**) have similar behaviour at each fixed inventory (see Figure 3), their dynamic behaviour can differ.

Finally, to gain a sense of the performance of the two strategies, in Figure 6 we show risk-reward profiles for the profit-and-loss and cost of acquisition using $10,000$ simulated paths. As the diagram shows the mean cost for the full strategy (**P1**) is lower than the reduced strategy (**P2**), and once the option payoff is taken into account, the mean PnL for the full strategy is higher. Note also that risk (measured by standard deviation) of the PnL in the full strategy is considerably lower than the risk in the reduced problem. Finally, the mean PnL provides the intermediary with a fair spread at which to charge the firm. For example, at the lower running penalty ϕ, the intermediary could charge a spread of about 4.5cents when using the full strategy to acquire the asset, while the intermediary would charge a spread of about 6.4cents when using the reduced strategy. With larger penalties, the risk of the full strategy reduces considerably (from a standard deviation of 2.15 cents down to 0.75 cents), while the costs increase incrementally (by about 1 cent). For the reduced strategy, the costs increases more significantly (by about 2.5 cents) while the risk changes little.

9. Conclusions

Here we developed a framework for determining the optimal exit and trading strategy for early exercise accelerated share repurchases where the intermediary is allowed to stop the process at any time prior to maturity. At this stopping time, the intermediary and the firm exchange TWAP up to that point in time. The early exercise clause adds an interesting layer of optionality into the classical optimal acquisition problem of Almgren-Chris. In such early exercise problems, one cannot hope to have analytical closed form solutions (indeed even for the simple American option pricing problem this is still an open problem), however, we were able to characterise the solution as a QVI and can establish bounds for the optimal trading strategy. Moreover, simple numerical schemes lead to optimal trading strategies and exit boundaries that exhibit interesting, but intuitively appealing, behaviours. We also demonstrated that a reduced problem, where the intermediary supplies an exit strategy dictated by the optimal exit policy of the embedded American Asian option (in which case the optimal exit strategy is independent of inventory), has similar behaving optimal strategies to the full problem. However, the dynamical results for the the two approaches differ and agents can squeeze extra value out of solving the combined optimal control and stopping problem.

One feature which we did not account for in this study is over-night risk. Such risks, however, can be incorporated by allowing the midprice process to jump at predictable times (end of day/beginning of day) by a random amount. This would result in an extra pasting conditions in the time dimension which can be dealt with in a straightforward manner. Another feature that we have not accounted for is the feedback effect of our own trading on the market dynamics. More specifically, we had set the permanent impact to zero. This simplifies the problem somewhat and allowed us to carry out a dimensional reduction, if the permanent impact were not zero, such a dimensional reduction would not be feasible. Nonetheless, a QVI can still be developed and the resulting system treated numerically, albeit it would be one dimension higher.

Appendix A. Almgren-Chriss acquisition strategy

As discussed in the setup of the problem in Section 2, once the intermediary has exercised the ASR option they are no longer exposed to TWAP risk. The agent, however, has a time limit of ϵ by which to acquire the remaining shares. Therefore, over the period $[\tau, \tau+\varepsilon]$ the agent adopts an Almgren-Chris like strategy with zero permanent impact. We further take the temporary market impact $a(S) = a_0\, S$ for some $a_0 > 0$. To enforce complete acquistion by $\tau + \varepsilon$, i.e., $q_{\tau+\varepsilon} = 0$, we adopt a terminal penalty approach and study the performance criteria

$$\text{(A.1)}\qquad H^{\nu}(t,S,q) = \mathbb{E}_{t,S,q}\left[\int_t^{\tau+\varepsilon} \widehat{S}_u\, \nu_u\, du + \, q_{\tau+\varepsilon}\left(S_{\tau+\varepsilon} + \phi_2\, S_{\tau+\varepsilon}\, q_{\tau+\varepsilon}\right) + \phi_1 \int_t^{\tau+\varepsilon} q_u^2 du\right].$$

The first term is the cost of acquiring shares (recall that $\widehat{S}_t = S_t + a(S_t)\,\nu_t$ is the execution price), the second term is the cost of acquiring all remaining shares at the terminal time. More specifically, the term proportional to ϕ_2 acts as a penalty, and we will take the limit in which this penalty tends to infinity, thus forcing the strategy to acquire all shares by $\tau+\varepsilon$. The third term corresponds to a running penalty which penalizes holding inventory through time and ϕ_1 acts as an urgency parameter. We seek the *optimal acquisition strategy* ν_t^* which attains the minimum, and the value function

$$H(t,S,q) = \inf_{\nu\in\mathcal{A}} H^{(\nu)}(t,S,q)$$

for all $\tau \le t \le \tau+\varepsilon$, $S > 0$, $1 \le q \le \mathfrak{n}$. Here, $\mathcal{A}$ is the set of admissible strategies consisting of non-negative, bounded, $\mathcal{F}$-predictable processes.

Proposition A.1. *The limiting value function*

$$\text{(A.2)}\qquad \lim_{\phi_2\to\infty} H^{(\nu^*)}(t,S,q) = q\,(S + \alpha_1(t)\, S\, q)$$

where

$$\text{(A.3)}\qquad \alpha_1(t) = \sqrt{\phi_1 a_0}\,\frac{e^{2\gamma\,(\tau+\varepsilon-t)}+1}{e^{2\gamma\,(\tau+\varepsilon-t)}-1}$$

where $\gamma = \sqrt{\phi_1/a_0}$.

Proof. Through the DPP, the value function H is the unique viscosity solution of the HJB equation over the domain $\tau \le t \le \tau+\varepsilon$, $S>0$, $1\le q\le \mathfrak{n}$:

$$\begin{cases} \partial_t H + \frac{1}{2} S^2\sigma^2 \partial_{SS} H + \phi_1 S\, q^2 & \\ \qquad\qquad + \min_{\nu}\{-\nu\,\partial_q H + \nu(a_0 S\ \nu + S)\} & = \ 0, \\ \qquad\qquad\qquad\qquad H(\tau+\varepsilon, S, q) & = \ q\,(F + \phi_2 S\, q), \end{cases}$$

The optimal strategy in feedback control form is obtained through the first order condition and is given by

$$\nu^* = \frac{S - \partial_q H}{a_0\, S}\,.$$

On substituting ν^* back into the HJB equation and writing, $H(t,S,q) = q\big(S + q\,S\,h(t)\big)$, one finds that $h(t)$ satisfies the ODE:

$$\partial_t h - \frac{1}{a_0} h^2 + \phi_1 = 0, \quad h(\tau+\epsilon) = \phi_2\,.$$

The solution to ODE is

$$h(t) = \sqrt{\phi_1 a_0}\,\frac{\xi\, e^{2\gamma(\tau+\varepsilon-t)} + 1}{\xi\, e^{2\gamma(\tau+\varepsilon-t)} - 1}$$

with $\xi = (\phi_2 - \sqrt{\phi_1 a_0})/(\phi_2 + \sqrt{\phi_1 a_0})$. By noting that $\xi \xrightarrow{\phi_2\to\infty} 1$ we obtain (A.2) and (A.3). □

At the time of option exercise, τ, the value function becomes independent of τ and we write

$$\ell(S,q) := H(\tau,S,q) = q\,S + \alpha_0\,S\,q^2\,,$$

where

$$\alpha_0 = \sqrt{\phi_1 a_0}\,\frac{e^{2\gamma\varepsilon} + 1}{e^{2\gamma\varepsilon} - 1}\,. \tag{A.4}$$

The value $\ell(S_\tau, q_{\tau-})$ is the cost we use in (3.1) to pose the agent's optimization problem prior to τ.

Appendix B. Derivative Valuation

Here, we consider the problem of finding the no-arbitrage price of the embedded early exercise option that pays the TWAP A_τ at the time of exercise τ. Recall that $A_t := \frac{1}{t}\int_0^t S_u\,du$ and here we assume there is no impact from trading, so that the asset price process S_t satisfies $dS_t = \sigma\,S_t\,dW_t$, where W_t is a standard Brownian motion, and $\sigma > 0$ is volatility.

Let $p(t,S,A)$ denote the no-arbitrage price of the option at time t, conditioned by $S_t = S,\ A_t = A$. Then, as usual, the (S,A)-plane splits into three regions:

$$C_t = \{(S,A) : p(t,S,A) > A\} \quad (\text{“continuation set”}),$$

$$D_t = \{(S,A) : p(t,S,A) \geq A\} \quad (\text{“stopping set”}),$$

We denote by $E_t \subset \{(S,A) : S, A \geq 0\}$ the boundary between sets C_t and D_t(“exercise boundary”).

Proposition B.1. *Let $p(t,S,A)$ be the no-arbitrage price at time t of the TWAP option, conditioned by $S_t = S,\ A_t = A$. Then*

$$p(t,S,A) = S\,w(t,z), \quad z = A/S, \tag{B.1}$$

where the function w satisfies the linear complementarity equations

$$w(t,z) \geq z\,, \tag{B.2}$$

$$\partial_t w + \frac{1}{2}\sigma^2 z^2 \partial_{zz} w + \frac{1-z}{t}\partial_z w \leq 0\,, \tag{B.3}$$

for all $0 < t < T$, $z > 0$, $w(T,z) = z$, with equality either in (B.2) or (B.3) at all times.

Proof. Indeed, it follows from standard results in optimal stopping of diffusion processes (see e.g., Peskir and Shiryaev (2006)) that p must satisfy

$$p(t,S,A) \geq A\,,$$

$$\partial_t p + \frac{1}{2}\sigma^2 S^2 \partial_{SS} p + \frac{A-S}{t}\partial_A p \leq 0\,,$$

for all $0 < t < T$, $z > 0$, $p(T,S,A) = A$, with equality either in the first or in the second inequality at all times. It is now easy to verify that ansatz (B.1) yields (B.2). $\square$

Corollary B.2. *For fixed t the exercise boundary E_t is a cone in variables (S,A):*

$$\text{if } (S,A) \in E_t \quad \Rightarrow \quad (cS,cA) \in E_t \text{ for all } c > 0.$$

Proof. Let w be the solution of linear complementarity equations (B.2), (B.3). We have

$$C_t = \{(t,S,A) : \text{there is strict inequality in (B.2) for } z = A/S\},$$

$$D_t = \{(t,S,A) : \text{there is strict inequality in (B.3) for } z = A/S\}.$$

Since inequalities (B.2), (B.3) are invariant with respect to transformation $(A,S) \mapsto (cA,cS)$, $c > 0$, so are domains C_t, D_t, and the boundary between them E_t. $\square$

Appendix C. Proofs

C.1. **Proof of Proposition 5.1.** We have

$$\partial_t H(t,S,A,q) = S\partial_t h(t,z,q), \quad \partial_S H(t,S,A,q) = h(t,z,q) - (A/S)\partial_z h(t,z,q),$$

$$\partial_{SS} H(t,S,A,q) = (A^2/S^3)\partial_{zz} h(t,z,q), \quad \partial_A H(t,S,A,q) = \partial_z h(t,z,q),$$

$$\partial_q H(t,S,A,q) = S\partial_q h(t,z,q), \quad \partial_{qq} H(t,S,A,q) = S\partial_{qq} h(t,z,q).$$

It remains to substitute these expressions into TBV problem (5.1a), (5.1b). $\square$

C.2. **Proof of Proposition 5.2.** Obviously, the identities stated in the proof of Proposition 5.2 still hold for $\tilde{H}$, $\tilde{h}$ (in place of H, h). It suffices to substitute these identities into TBV problem (5.2). $\square$

C.3. **Proof of Proposition 5.3.** Using the ansatz of Proposition 5.2, we first obtain the same TBV problem as (5.2) but with the $\mathfrak{n}z$ term absent. Hence, the ansatz of the present proposition reduces to

$$\widetilde{h}(t,z,q) = -q(1+g(t,z)q).$$

We have

$$\partial_t \widetilde{h} = -q^2 \partial_t g, \quad \partial_z \widetilde{h} = -q^2 \partial_z g, \quad \partial_{zz} \widetilde{h} = -q^2 \partial_{zz} g, \quad \partial_q \widetilde{h} = -1 - 2qg.$$

Substituting these expressions in the above TBV problem, we obtain TBV problem (5.4), as needed. □

C.4. **Proof of Theorem 7.1.** (i) By Proposition 5.3 the optimal acquisition strategy $\widetilde{\nu}^*$ is given by formula (5.5), where the solution g of the TBV problem (5.4) satisfies the boundary condition $g(t,0) = \omega(t)$, cf. (6.4). Equivalently, $\partial_z g(t,0) = 0$, see Section 6 for details.

Let $f(t,z) := g(t,z) - \sqrt{\phi_0 a_0}$. Then (5.4) becomes

$$\text{(C.1)} \qquad \begin{cases} \partial_t f + \frac{1}{2}\sigma^2 z^2 \partial_{zz} f + \frac{1-z}{t}\partial_z f - \frac{1}{a_0}(f + 2\sqrt{\phi_0 a_0}) f &= 0, \\ f(T,z) &= \alpha_0 - \sqrt{\phi_0 a_0}, \\ f(t,\zeta(t)) &= \alpha_0 - \sqrt{\phi_0 a_0}, \end{cases}$$

and

$$\text{(C.2)} \qquad f(t,0) = \omega(t) - \sqrt{\phi_0 a_0}.$$

By (Friedman, 1964, Theorem 6, Sect 4.4) there exists a (unique) classical solution f to (C.1), (C.2). Denote $c(t,z) := \frac{1}{a_0}(f + 2\sqrt{\phi_0 a_0})$; function c is bounded and continuous in $0 \le z \le \zeta(t)$, $0 \le t \le T$. We may assume that function c is a priori specified, so (C.1), (C.1) becomes a TBV problem for *linear* parabolic equation

$$\partial_t f + \frac{1}{2}\sigma^2 z^2 \partial_{zz} f + \frac{1-z}{t}\partial_z f - c(t,z) f = 0.$$

By our assumption $\alpha_0 - \sqrt{\phi_0 a_0} > 0$, $\omega(t) - \sqrt{\phi_0 a_0} > 0$, i.e. the terminal and boundary conditions for f are determined by positive functions. Therefore, by (Friedman, 1964, Lemma 5, Sect 2.4) $f(t,z) \ge 0$ for all $0 \le z \le \zeta(t)$, $0 \le t \le T$. Hence, by the weak maximum principle (cf. (Friedman, 1964, Theorem 6, Sect 2.1)) f can attain its maximum only at $t = T$ or at $z = 0$, $z = \zeta(t)$ for $0 \le t \le T$. Namely, the (global) maximum of f is equal to $\alpha_0 - \sqrt{\phi_0 a_0}$, and therefore the (global) maximum of g is α_0.

Next, we prove that $z \mapsto g(t,z)$ attains its minimum at $z = 0$. Denote $u = \partial_z g$. By differentiating the PDE in (5.4) in variable z, we obtain that u must satisfy

$$\text{(C.3)} \qquad \partial_t u + \frac{1}{2}\sigma^2 z^2 \partial_{zz} u + \left(\sigma^2 z + \frac{1-z}{t}\right)\partial_z u - \left(\frac{1}{t} + \frac{2}{a_0} g\right) u = 0.$$

Since $\partial_z g(t,0) = 0$, we have $u(t,0) = 0$. Further, we have $u(t,\zeta(t)) \ge 0$ for every $0 \le t \le T$; assuming $u(t,\zeta(t)) = \partial_z(t,\zeta(t)) < 0$ for some t, we obtain a contradiction with the earlier result stating that

$g(t,\zeta(t))=\alpha$ is the global maximum of g over $0\le z\le\zeta(t)$, $0\le t\le T$. Since $\frac{1}{t}+\frac{2}{a_0}g(t,z)>0$ for all t and z (as $f(t,z)=g(t,z)-\sqrt{\phi_0 a_0}\ge 0$ by the previous result), we obtain $u(t,z)\ge 0$ by (Friedman, 1964, Lemma 5, Sect 2.4). Hence, $z\mapsto g(t,z)$ is monotonically increasing. This implies that $g(t,z)$ attains its minimum along the line $z=0$.

The proof for the case $\phi_0>\frac{\alpha_0^2}{a}$ is similar, with only some signs and inequalities needed to be reversed. If $\phi_0=\frac{\alpha_0^2}{a}$, then $g(t,z)\equiv\alpha_0$ is trivially the solution of (5.4). □

(ii) Let $g_{(i)}$, $i=1,2$, be the corresponding solutions of TBV problem (5.4) satisfying boundary condition (6.4). We define $h_{(i)}=\frac{g_{(i)}}{a_0^{(i)}}$, $i=1,2$. Then $h_{(i)}$ satisfies

$$\text{(C.4)}\qquad \begin{cases} \partial_t h_{(i)}+\frac{1}{2}\sigma^2 z^2\partial_{zz}h_{(i)}+\frac{1-z}{t}\partial_z h_{(i)}+\frac{\phi_0^{(i)}}{a_0^{(i)}}-h_{(i)}^2 &= 0,\\ h_{(i)}(T,z) &= \frac{\alpha_0^{(i)}}{a_0^{(i)}}\\ h_{(i)}(t,\zeta(t)) &= \frac{\alpha_0^{(i)}}{a_0^{(i)}}.\\ h_{(i)}(t,0) &= \frac{\omega(t)}{a_0^{(i)}}. \end{cases}$$

The required inequality now follows from the maximum principle for (C.4) (cf. the argument above) and formula (5.5). □

C.5. **Proof of Theorem 7.2.** Following the discussion in Section 6, we complement TBV of Proposition 5.3 with boundary condition $g(t,0)=\omega(t)$, where ω is given by formula (6.5). As $T\to+\infty$

$$g(t,0)=\omega(t)\to\sqrt{\phi_0 a_0}.$$

Also, $g(t,z^*)=\alpha_0$. Consider boundary value problem

$$\tfrac{1}{2}\sigma^2 z^2\ \gamma''(z)+\phi_0-\tfrac{1}{a_0}\gamma^2(z)=0,\qquad \gamma(0)=\sqrt{\phi_0 a_0},\quad\text{and}\qquad \gamma(z_*)=\alpha_0.$$

By (Friedman, 1964, Theorem 2, Sect 6.1) for any fixed t $g(t,\cdot)\to\gamma(\cdot)$ uniformly on $[0,z^*]$ as $T\to+\infty$.

If $\phi_0=0$, then $\gamma''(z)\ge 0$, and hence γ is convex. If $\phi_0\ge\frac{\alpha_0^2}{a_0}$, then $a_0\le\gamma(z)\le\sqrt{\phi_0 a_0}$ by the maximum principle, and so $\phi_0-\frac{1}{a_0}\gamma^2(z)\ge 0$; it follows that $\gamma''(z)\le 0$, i.e. γ is concave. □

C.6. **Proof of Theorem 7.3.** By Proposition 5.3 the optimal acquisition strategy $\widetilde{\nu}^*$ is given by formula (5.5), where function g additionally satisfies the boundary condition $g(t,0)=\omega(t)$, cf. (6.4).

Denote $v=\partial_{zz}g$. By differentiating the PDE in (5.4) in variable z twice, we obtain that v must satisfy

$$\text{(C.5)}\qquad \partial_t v+\frac{1}{2}\sigma^2 z^2\partial_{zz}v+\left(2\sigma^2 z+\frac{1-z}{t}\right)\partial_z v-\left(\frac{2}{t}+\frac{2}{a_0}g-\sigma^2\right)v=\frac{2}{a_0}\big(\partial_z g(t,z)\big)^2.$$

Since $g(t,z)$ is between $\omega(t)$ and α_0 for all t, z by Theorem 7.1(i), our assumptions imply that $\frac{2}{t}+\frac{2}{a_0}g-\sigma^2\ge 0$.

Let $\phi_0 < \frac{\alpha_0^2}{a_0}$. In the proof of Theorem 7.1(i) we have established that $z \mapsto \partial_z g(t,z)$ attains its global maximum α_0 at $z = \zeta(t)$ for each t. Hence, $\partial_{zz} g(t, \zeta(t))$ must be non-negative. Also, $\partial_{zz} g(t,0) \geq 0$ for all t; assuming otherwise, we obtain that g can take negative values, since $\partial_z g(t,0) = 0$ for all t. Therefore, $(\partial_{zz} g(t,z) =)\ v(t,z) \geq 0$ for all t, z by (Friedman, 1964, Lemma 5, Sect 2.4), so $z \mapsto g(t,z)$ is a convex function.

The proof for the case $\phi_0 > \frac{\alpha_0^2}{a_0}$ is analogous – with only some signs and inequalities needed to be reversed.

The assertions regarding increasing volatility $\sigma > 0$ are obtained by differentiating the PDE, terminal and boundary conditions in (5.4) in variable σ: let $r(t,z) = \partial_\sigma g(t,z)$, then

$$\begin{cases} \partial_t r + \frac{1}{2}\sigma^2 z^2 \partial_{zz} r + \frac{1-z}{t}\partial_z r - \frac{2}{a_0} g r + \sigma z^2 \partial_{zz} g &= 0, \\ r(T,z) &= 0, \\ r(t,\zeta(t)) &= 0, \\ r(t,0) &= 0 \end{cases}$$

If $\phi_0 < \frac{\alpha_0^2}{a_0}$, then by the last result $\partial_{zz} g(t,z) \geq 0$ for all t, z. Since $g \geq 0$, we obtain that $r(t,z) \geq 0$ for all t, z by (Friedman, 1964, Lemma 5, Sect 2.4). The case $\phi_0 > \frac{\alpha_0^2}{a_0}$ is treated similarly. $\square$

References

Almgren, R. and N. Chriss (1999). Value under liquidation. *Risk December.*

Almgren, R. and N. Chriss (2001). Optimal execution of portfolio transactions. *Risk 3*, 5–39.

Avellaneda, M. and S. Stoikov (2008, November). High-frequency trading in a limit order book. *Quantitative Finance 8*, 217–224.

Bargeron, L., M. Kulchania, and S. Thomas (2011). Accelerated share repurchases. *Journal of Financial Economics 101(1)*, 69–89.

Bayraktar, E. and M. Ludkovski (2012). Liquidation in limit order books with controlled intensity. *Mathematical Finance.*

Bertsimas, D. and A. W. Lo (1998). Optimal control of execution costs. *Journal of Financial Markets 1(1)*, 1–50.

Cartea, Á., R. Donnelly, and S. Jaimungal (2013). Robust market making. *Available at SSRN: http://ssrn.com/abstract=2310645.*

Cartea, Á. and S. Jaimungal (2012). Risk metrics and fine tuning of high frequency trading strategies. *Forthcoming in Mathematical Finance.*

Cartea, Á. and S. Jaimungal (2013). Modeling asset prices for algorithmic and high frequency trading. *Forthcoming in Applied Mathematical Finance.*

Cartea, Á., S. Jaimungal, and J. Ricci (2011). Buy low sell high: A high frequency trading perspective. *SSRN eLibrary, http://ssrn.com/abstract=1964781*.

Crandall, M., H. Ishii, and P.-L. Lions (1992). Users guide to viscosity solutions of second order partial differential equations. *Bulletin of the AMS 27*, 1–68.

Forsyth, P. and K. Vetzal (2012). *Handbook of Computational Finance*, Chapter Numerical methods for nonlinear PDEs in finance (Ch 18), pp. 503–528. Springer.

Friedman, A. (1964). *Partial Differential Equations of Parabolic Type*. Prentice-Hall.

Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance 10(7)*, 749–759.

Gatheral, J., A. Schied, and A. Slynko (2012). Transient linear price impact and Fredholm integral equations. *Mathematical Finance 22*(3), 445–474.

Graewe, P., U. Horst, and J. Qiu (2013). A non-Markovian liquidation problem and backward SPDEs with singular terminal conditions. *Working paper*.

Guéant, O. and C.-A. Lehalle (2013). General intensity shapes in optimal liquidation. *Mathematical Finance*, n/a–n/a.

Guéant, O., C.-A. Lehalle, and J. Fernandez-Tapia (2012). Dealing with the inventory risk: a solution to the market making problem. *Mathematics and Financial Economics*, 1–31.

Ho, T. and H. R. Stoll (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics 9*, 47–73.

Obizhaeva, A. A. and J. Wang (2012). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets In Press*.

Peskir, G. and A. Shiryaev (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser.

Predoiu, S., G. Shaikhet, and S. Shreve (2011). Optimal execution of a general one-sided limit-order book. *SIAM J. Financial Math 2*, 183–212.

E-mail address: `sebastian.jaimungal@utoronto.ca`

Department of Statistical Sciences, University of Toronto, 100 St. George St., Toronto, ON, M5S 3G3, Canada

The Fields Institute, 222 College Street, Toronto, ON, M5T 3J1, Canada
E-mail address: `dkinzebu@fields.utoronto.ca`

BMO Capital Markets, 1 First Canadian Place, 100 King Street West, Toronto, ON M5X 2A1 Canada
E-mail address: `dmitri.rubisov@bmo.com`