

Optimal Liquidation under Partial Information with Price Impact

K. Colaneri and Z. Eksi and R. Frey and M. Szölgényi

Research Report No. 2018-17
April 2018

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

Optimal Liquidation under Partial Information with Price Impact

Katia Colaneri ^{*} Zehra Eksi [†] Rüdiger Frey [‡] Michaela Szölgényi [§]

Abstract

We study the optimal liquidation problem in a market model where the bid price follows a geometric pure jump process whose local characteristics are driven by an unobservable finite-state Markov chain and by the liquidation rate. This model is consistent with stylized facts of high frequency data such as the discrete nature of tick data and the clustering in the order flow. We include both temporary and permanent effects into our analysis. We use stochastic filtering to reduce the optimal liquidation problem to an equivalent optimization problem under complete information. This leads to a stochastic control problem for piecewise deterministic Markov processes (PDMPs). We carry out a detailed mathematical analysis of this problem. In particular, we derive the optimality equation for the value function, we characterize the value function as continuous viscosity solution of the associated dynamic programming equation, and we prove a novel comparison result. The paper concludes with numerical results illustrating the impact of partial information and price impact on the value function and on the optimal liquidation rate.

Keywords: Optimal liquidation, Stochastic filtering, Piecewise deterministic Markov process, Viscosity solutions and comparison principle.

1 Introduction

In financial markets, traders frequently face the task of selling a large amount of a given asset over a short time period. This has led to a large literature on optimal portfolio execution. The existing work can be divided into two classes: market impact models and order book models. In a market impact model one directly specifies the impact of a given trading strategy on the bid price of the asset. The fundamental price (the price if the trader is inactive) is usually modelled as a diffusion process such as Brownian motion. In an order book model, instead, one specifies the dynamics of the limit order book. This is more complex but gives an explanation of the price impact in terms of fundamental quantities.

Portfolio liquidation strategies are executed at a high trading frequency. Hence a sound market impact model should be consistent with key stylized facts of high frequency data as discussed for

^{*}School of Mathematics, University of Leeds, LS2 9JT, Leeds, UK. K.Colaneri@leeds.ac.uk

[†]Institute for Statistics and Mathematics, Vienna University of Economics and Business (WU), Welthandelsplatz 1, 1020 Vienna, Austria, zehra.eksi@wu.ac.at

[‡]Institute for Statistics and Mathematics, Vienna University of Economics and Business (WU), ruediger.frey@wu.ac.at, corresponding author

[§]Seminar for Applied Mathematics and RiskLab Switzerland, ETH Zurich, Rämistrasse 101, 8092 Zurich, Switzerland, michaela.szolgyenyi@math.ethz.ch

instance by Cartea et al. [19] or Cont [24]. First, on very fine time scales the bid price of an asset is best described by a pure jump process, since in reality prices move on a discrete grid defined by the tick size. Second, the order flow is clustered in time: there are random periods with a lot of buy orders or with a lot of sell orders, interspersed by quieter times with less trading activity. Cont [24] attributes this to the fact that many observed orders are components of a larger parent order that is executed in small blocks. A further reason for the clustering in the inter-event times are random fluctuations in the arrival rate of new information, see, e.g. Andersen [5]. Third, the distribution of returns over short time intervals is strongly non-Gaussian but has heavy tails and a large mass around zero; to a certain extent this is a consequence of the first two stylized facts. Finally, there is permanent price impact, that is the implementation of a liquidation strategy pushes prices downwards.

To capture these stylized facts we model the bid price as marked point process with Markov switching whose local characteristics (intensity and jump size distribution) depend on the trader's current liquidation rate ν_t and on the value Y_t of a finite state Markov chain Y . The fact that the local characteristics depend on ν_t is used to model permanent price impact. Markov switching allows us to reproduce the observed clustering in the order flow. Our framework encompasses models with a high intensity of downward jumps in one state of Y and a high intensity of upward jumps in another state of Y and models where inter-event times are given by a mixture of exponential distributions. We view the process Y as an abstract modelling device that generates clustering and assume therefore that Y is unobservable by the trader. This is consistent with the fact that economic sources for clustering such as the trading activity of other investors are not directly observable. Markov modulated marked point processes with partial information (without price impact) were considered previously in the statistical modelling of high frequency data, see for instance Zeng [45], Cvitanic et al. [26], or Cartea and Jaimungal [17]; however, we are the first to study optimal liquidation in such a setting.

The first step in the analysis of a control problem with partial information is to derive an equivalent problem under full information via stochastic filtering. Hence we have to determine the dynamics of the conditional distribution of Y_t given the bid price history up to time t . Note that this provides a further rationale for modelling the bid price as a marked point process: the strong non-normality of short-period returns implies that it is very problematic to use high frequency data as input for the numerical solution of the filtering equations in the classical setup where observations are modelled as a Brownian motion with drift, as the resulting filters become extremely unstable. Instead one should take the structure of the observation process seriously and work in a point process model. We use the reference probability approach to derive the filtering equations for our model. In this way we circumvent the issue that the information available to the investor depends on her liquidation strategy. We end up with a control problem whose state process X consists of the stock price, the inventory level, and the filter process. We provide a detailed mathematical analysis of this problem. The form of the asset return dynamics implies that X is a piecewise deterministic Markov process (PDMP) so that we rely on control theory for PDMPs; a general introduction to this theory is given in Davis [29] or in Bäuerle and Rieder [12]. We establish the dynamic programming equation for the value function and we derive conditions on the data of the problem that guarantee the continuity of the value function. This requires a careful analysis of the behaviour of the value function close to the boundary of the state space. As a further step we characterize the value function as the unique continuous viscosity solution of the Hamilton-Jacobi-Bellman (HJB) partial integro-differential equation associated with the problem and we give an example showing that in general the HJB equation does not admit a classical solution. Moreover, we prove a novel comparison theorem

for the HJB equation which is valid in more general PDMP setups. A comparison principle is necessary to ensure the convergence of numerical schemes to the value function, see Barles and Souganidis [9].

The paper closes with a section on applications. We discuss properties of the optimal liquidation rate and of the expected liquidation profit and we use a finite difference approximation of the HJB equation to analyze the influence of the temporary and permanent price impact parameters on the form of the optimal liquidation rate. Among others, we find that for certain parameter constellations the optimal strategy displays a surprising gambling behaviour of the trader that cannot be guessed upfront and we give an economic interpretation that is based on the form of the HJB equation. Moreover, we study the additional liquidation profit from the use of a filtering model, and we report results from a small calibration study that provides further support for our model.

We continue with a brief discussion of the existing literature. Starting with market impact models, the first contribution is Bertsimas and Lo [15] who analyze the optimal portfolio execution problem for a risk-neutral agent in a model with linear and purely permanent price impact where the fundamental price follows an arithmetic random walk. This model has been generalized by Almgren and Chriss [2] who consider also risk aversion and temporary price impact. Since then, market impact models have been extensively studied. Important contributions include He and Mamaysky [38], Schied and Schöneborn [44], Schied [43], Ankirchner et al. [6], Guo and Zervos [37]. Recently, Cayé and Muhle-Karbe [20] studied an extension of the Almgren Chriss model with a self-exciting temporary price impact. All these models work in a (discretized) diffusion framework.

In the order book literature on the other hand, a few contributions based on point process models exist. Bayraktar and Ludkovski [13] analyze the optimal portfolio execution problem in a model with discrete order flow represented by a Poisson process with observable intensity. The price impact is purely temporary and is represented in terms of a cost function. Bäuerle and Rieder [11] consider the same setting with a standard Poisson process and solve the cost minimization problem by using tools from the control theory of PDMPs. A further order book model based on point process methodology is Bayraktar and Ludkovski [14]. There it is assumed that the trader uses limit orders and that she can control the intensity of the order flow by choosing the spread at which she is willing to trade. Additional contributions based on diffusion models are Alfonsi et al. [1], Obizhaeva and Wang [41], Cartea and Jaimungal [18]. For a detailed overview we refer to the surveys Gökay et al. [36], Gatheral and Schied [35] or Cartea et al. [19]. From a methodological point of view our analysis is also related to the literature on expected utility maximization or hedging for pure jump process such as Bäuerle and Rieder [10] or Kirch and Runggaldier [40]. Important contributions to the control theory of PDMPs include Davis [29], Dempster and Ye [31], Almudevar [4], Forwick et al. [33], Bäuerle and Rieder [11], Costa and Dufour [25]. Viscosity solutions for PDMP control problems were previously considered in Davis and Farid [30].

The outline of the paper is the following. In Section 2, we introduce our model, the main assumptions and the optimization problem. In Section 3, we derive the filtering equations for our model. Section 4 contains the mathematical analysis of the optimization problem via PDMP techniques. In Section 5 we provide a viscosity solution characterization of the value function. Finally, in Section 6, we present the results of our numerical experiments. The appendix contains additional proofs.

2 The Model

Throughout we work on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{P})$, where the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}$ satisfies the usual conditions. Here \mathbb{F} is the global filtration, i.e. all considered processes are \mathbb{F} -adapted, and \mathbf{P} is the historical probability measure. We consider a trader who wants to liquidate $w_0 > 0$ units of a given security (referred to as the *stock* in the sequel) over the period $[0, T]$ for a given time horizon T . We denote the bid price process by $S = (S_t)_{0 \leq t \leq T}$ and \mathbb{F}^S is the filtration generated by S . In what follows, we assume that \mathbb{F}^S satisfies the usual conditions. We assume that the trader sells the shares at a nonnegative \mathbb{F}^S -adapted rate $\nu = (\nu_t)_{0 \leq t \leq T}$ such that for every $t \in [0, T]$, $\nu_t \in [0, \nu^{\max}]$ for a given positive constant ν^{\max} . Hence her inventory, i.e. the amount of shares she holds at time $t \in [0, T]$, is given by

$$W_t = w_0 - \int_0^t \nu_u du, \quad t \in [0, T]. \quad (2.1)$$

Modelling the inventory as an absolutely continuous process corresponds to the situation where the trader is frequently submitting small sell orders. By taking ν to be nonnegative, we confine the trader to pure selling strategies; the motivation for imposing the upper bound ν^{\max} on the liquidation rate is discussed in Section 2.2 below. The goal of the trader is to maximize the expected revenue from her trading strategy. We assume that the implementation of the liquidation strategy generates temporary and permanent price impact, where permanent price impact is the impact of trading on the dynamics of S and temporary price impact is the impact of trading on the execution price of the current trade.

2.1 Dynamics of the bid price. In order to reproduce stylized facts of high frequency data such as the path structure of asset prices and the clustering of the order flow, we model the bid price as a Markov-modulated geometric finite activity pure jump process. Let $Y = (Y_t)_{0 \leq t \leq T}$ be a continuous-time finite-state Markov chain on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{P})$ with state space $\mathcal{E} = \{e_1, e_2, \dots, e_K\}$ (e_k is k -th unit vector in \mathbb{R}^K), generator matrix $Q = (q^{ij})_{i,j=1,\dots,K}$ and initial distribution $\pi_0 = (\pi_0^1, \dots, \pi_0^K)$. We assume that the bid price has the dynamics

$$dS_t = S_{t-} dR_t, \quad S_0 = s \in (0, \infty), \quad (2.2)$$

where the *return process* $R = (R_t)_{0 \leq t \leq T}$ is a finite activity pure jump process. We assume that $\Delta R_t := R_t - R_{t-} > -1$ so that S is strictly positive. Denote by μ^R the random measure associated with R , defined by

$$\mu^R(dt, dz) := \sum_{u \geq 0, \Delta R_u \neq 0} \delta_{\{u, \Delta R_u\}}(dt, dz),$$

and by $\eta^{\mathbf{P}}$ the (\mathbb{F}, \mathbf{P}) -dual predictable projection (or compensating random measure) of μ^R . We assume that $\eta^{\mathbf{P}}$ is absolutely continuous and of the form $\eta^{\mathbf{P}}(t, Y_{t-}, \nu_{t-}; dz)dt$, for a finite measure $\eta^{\mathbf{P}}(t, e, \nu; dz)$ on \mathbb{R} . Moreover, we assume that the processes R and Y have no common jumps, so that R and Y are orthogonal, $[R, Y]_t \equiv 0$ for all $t \in [0, T]$, \mathbf{P} -a.s.

The measure $\eta^{\mathbf{P}}(t, e, \nu; dz)$ is a crucial quantity as it determines the law of the bid price with respect to filtration \mathbb{F} under \mathbf{P} . The fact that $\eta^{\mathbf{P}}$ depends on the current liquidation rate serves to model permanent price impact; the dependence of $\eta^{\mathbf{P}}$ on Y_{t-} can be used to reproduce the clustering in inter-event durations observed in high frequency data and to model the feedback effect from the trading activity of the rest of the market. Finally, time-dependence of $\eta^{\mathbf{P}}$ can be used to model the strong intra-day seasonality patterns observed for high frequency data. These

aspects are explained in more detail in Example 2.3 below. Now we turn to the semimartingale decomposition of the bid price with respect to the full information filtration \mathbb{F} . Denote for all $(t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, \nu^{\max}]$, the mean of $\eta^{\mathbf{P}}$ by

$$\bar{\eta}^{\mathbf{P}}(t, e, \nu) := \int_{\mathbb{R}} z \eta^{\mathbf{P}}(t, e, \nu; dz); \quad (2.3)$$

$\bar{\eta}^{\mathbf{P}}(t, e, \nu)$ exists under Assumption 2.1 below. Fix some liquidation strategy ν . Then the martingale part M^R of the return process is given by $M_t^R = R_t - \int_0^t \bar{\eta}^{\mathbf{P}}(s, Y_{s-}, \nu_{s-}) ds$, for all $t \in [0, T]$, and the \mathbb{F} -semimartingale decomposition of S equals

$$S_t = S_0 + \int_0^t S_{s-} dM_s^R + \int_0^t S_{s-} \bar{\eta}^{\mathbf{P}}(s, Y_{s-}, \nu_{s-}) ds, \quad t \in [0, T].$$

It is well-known that the semimartingale decomposition of S with respect to the trader's filtration \mathbb{F}^S is obtained by projecting the process $\bar{\eta}^{\mathbf{P}}(t, Y_{t-}, \nu_{t-})$ onto \mathbb{F}^S . In the sequel we assume that for all $(t, e) \in [0, T] \times \mathcal{E}$, the mapping $\nu \mapsto \bar{\eta}^{\mathbf{P}}(t, e, \nu)$ is decreasing on $[0, \infty)$, that is selling pushes the price down on average. Furthermore, we make the following regularity assumption.

Assumption 2.1 (Properties of $\eta^{\mathbf{P}}$). There is a deterministic finite measure $\eta^{\mathbf{Q}}$ on \mathbb{R} whose support, denoted by $\text{supp}(\eta)$, is a compact subset of $(-1, \infty)$, such that for all $(t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, \infty)$ the measure $\eta^{\mathbf{P}}(t, e, \nu; dz)$ is equivalent to $\eta^{\mathbf{Q}}(dz)$. Furthermore, for every $\nu^{\max} < \infty$ there is some constant $M > 0$ such that

$$M^{-1} < \frac{d\eta^{\mathbf{P}}(t, e, \nu)}{d\eta^{\mathbf{Q}}}(z) < M \text{ for all } (t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, \nu^{\max}]. \quad (2.4)$$

The assumption implies that for every ν^{\max} there is a $\lambda^{\max} < \infty$ such that

$$\sup\{\eta^{\mathbf{P}}(t, e, \nu; \mathbb{R}) : (t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, \nu^{\max}]\} \leq \lambda^{\max}; \quad (2.5)$$

in particular the counting process associated to the jumps of S is \mathbf{P} -nonexplosive. Moreover, it provides a sufficient condition for the existence of a *reference probability measure*, i.e. a probability measure \mathbf{Q} equivalent to \mathbf{P} on (Ω, \mathcal{F}_T) , such that under \mathbf{Q} , μ^R is a Poisson random measure with intensity measure $\eta^{\mathbf{Q}}(dz)$, independent of Y and ν . This is needed in the analysis of the filtering problem of the trader in Section 3. Note that the equivalence of $\eta^{\mathbf{P}}$ and $\eta^{\mathbf{Q}}$ implies that for all $(t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, \infty)$ the support of $\eta^{\mathbf{P}}$ is equal to $\text{supp}(\eta)$. The assumption that $\text{supp}(\eta)$ is compact is not restrictive, since in reality the bid price moves only by a few ticks at a time.

The following examples serve to illustrate our framework; they will be taken up in our numerical experiments in Section 6.

Example 2.2. Consider the case where the return process R follows a bivariate point process, i.e. there are two possible jump sizes, $\Delta R \in \{-\theta, \theta\}$ for some $\theta > 0$. In this example we assume that the dynamics of S is independent of Y and t . Moreover, the intensity λ^+ of an upward jump is constant and equal to $c^{\text{up}} > 0$, and the intensity λ^- of a downward jump depends on the rate of trading and is given by $\lambda^-(\nu) = c^{\text{down}}(1 + a\nu)$ for constants $c^{\text{down}}, a > 0$. Note that, with this choice of λ^- , the intensity of a downward jump in S is linearly increasing in the liquidation rate ν . The function $\bar{\eta}^{\mathbf{P}}$ from (2.3) is independent of t and e and linearly decreasing in ν ; it is given by $\bar{\eta}^{\mathbf{P}}(\nu) = \theta(c^{\text{up}} - c^{\text{down}}(1 + a\nu))$. Linear models for the permanent price impact are frequently considered in the literature as they have theoretical and empirical advantages; see for instance Almgren et al. [3] or Gatheral and Schied [35].

Example 2.3. Now we generalize Example 2.2 and allow $\eta^{\mathbf{P}}$ to depend on the state process Y . We consider a two-state Markov chain Y with the state space $\mathcal{E} = \{e_1, e_2\}$ and we assume that e_1 is a ‘good’ state and e_2 a ‘bad’ state in the following sense: in state e_1 the intensity of an upward move of the stock is larger than in state e_2 ; the intensity of a downward move on the other hand is larger in state e_2 than in e_1 . We therefore choose constants $c_1^{\text{up}} > c_2^{\text{up}} > 0$, $c_2^{\text{down}} > c_1^{\text{down}} > 0$ and a price impact parameter $a > 0$ and we set for $i = 1, 2$,

$$\lambda^+(e_i, \nu) = (c_1^{\text{up}}, c_2^{\text{up}})e_i \quad \text{and} \quad \lambda^-(e_i, \nu) = (1 + a\nu)(c_1^{\text{down}}, c_2^{\text{down}})e_i.$$

Then, $\eta^{\mathbf{P}}(e_i, \nu, dz) = \lambda^+(e_i, \nu)\delta_{\{\theta\}}(dz) + \lambda^-(e_i, \nu)\delta_{\{-\theta\}}(dz)$, for $i = 1, 2$. Since $c_1^{\text{up}} > c_2^{\text{up}}$, in state e_1 one has on average more buy orders; this might represent a scenario where another trader is executing a large buy program. Similarly, since $c_2^{\text{down}} > c_1^{\text{down}}$, there are on average more sell orders in state e_2 , for instance because another trader is executing a large sell program. The form of $\eta^{\mathbf{P}}$ implies that the permanent price impact is linear and proportional to the intensity of a downward move and hence larger in the ‘bad’ state e_2 than in the good state e_1 .

Note that within our general setup this example could be enhanced in a number of ways. For instance, the transition intensities c_i^{up} and c_i^{down} and the liquidity parameter a could be made time dependent to reflect the fact that on most markets trading activity during the day is U -shaped with more trades occurring at the beginning and the end of a day than in the middle. Moreover, one could introduce an additional state where the market is moving sideways, or one could consider the case where the liquidity parameter a depends on Y .

Remark 2.4 (Calibration.) We briefly discuss a potential approach for parameter estimation in our setup. For $\nu \equiv 0$ the model is a hidden Markov model with point process observation. It is therefore natural to use the expectation maximization (EM) methodology for Markov modulated point processes as described in Elliott and Malcolm [32] or in Damian et al. [27] to estimate the generator matrix of Y and parameters of the compensator $\eta^{\mathbf{P}}$. A numerical case study with simulated and real data in the context of Example 2.3 is given in Section 6.3. Using proprietary data on the performance of equity sales, Almgren et al. [3] find empirical support for a linear permanent price impact function; the parameter a can be estimated by regressing price changes on trading volume.

2.2 The optimization problem. In this section we specify the ingredients of the traders optimization problem in detail.

Liquidation strategies. We assume that the state process Y is not directly observable by the trader. Instead, she observes the price process S and knows the model parameters, so that information available to her is carried by filtration \mathbb{F}^S or, equivalently, by the filtration generated by the return process R . Hence we assume that the trader uses only liquidation strategies that are \mathbb{F}^S -adapted. Moreover we impose a bound on the maximal speed of trading: we fix some constant $\nu^{\text{max}} > w_0/T$ and we call a liquidation strategy ν *admissible* if ν is \mathbb{F}^S -adapted and if $\nu_t \in [0, \nu^{\text{max}}]$ for all $t \in [0, T]$ \mathbf{P} -a.s. Note that the condition $\nu^{\text{max}} > w_0/T$ ensures that it is feasible for the trader to liquidate the whole inventory over the period $[0, T]$.

The assumption of a *bounded* liquidation rate merits a discussion. From a mathematical point of view a bound on the liquidation rate facilitates the application of results for the control of piecewise deterministic Markov processes, since in this theory it is typically assumed that the strategies take values in a compact control space. Moreover, without this assumption the

viscosity solution characterization of the value function (see Theorem 5.3 below) does not hold. A counterexample is given in Section 5.2, where we show that for unbounded liquidation rate the value function is a strict supersolution of the corresponding dynamic programming equation, cf. Remark 5.5. Finally, the upper bound on ν_t ensures that under Assumption 2.1 for every admissible strategy ν a return process R with compensating measure $\eta^{\mathbf{P}}(t, Y_{t-}, \nu_{t-}; dz)dt$ (and hence the bid price process (2.2)) exists.

From a financial point of view an upper bound on ν_t is reasonable, as trading at infinite speed would correspond to large block transactions; allowing such transactions at some time point $t < T$ would require an explicit model for market resiliency. It is however not clear how to determine ν^{\max} empirically. In Proposition 2.5 below we therefore show that $J^{*,m}$, the optimal liquidation value if the trader uses \mathbb{F}^S -adapted strategies with $\nu_t \leq m$ for all t , is bounded independently of m . The sequence $\{J^{*,m}\}_{m \in \mathbb{N}}$ is obviously increasing, since a higher m means that the trader can optimize over a larger set of strategies. Hence, $\{J^{*,m}\}_{m \in \mathbb{N}}$ is Cauchy. This implies that optimal proceeds from liquidation are nearly independent of the precise numerical value chosen for ν^{\max} . In order to further support this argument we present results of numerical experiments in the framework of Example 2.3. Table 1 displays the value function $J^{*,m}$ for varying ν^{\max} (expressed as multiple of the initial inventory w_0) and for fixed $w_0 = 6000$. The value grows in ν^{\max} , but for $\nu^{\max} > 2w_0$, the additional gain is small. For details on the numerical analysis we refer to Section 6.

ν^{\max}	w_0/T	$2w_0/T$	$3w_0/T$	$5w_0/T$	$7w_0/T$	$10w_0/T$
$V'(0, w_0, \pi^1; \nu^{\max})$	5648.99	5749.14	5750.14	5750.3	5750.33	5750.34

Table 1: The expected proceeds from liquidation for varying ν^{\max} and fixed $w_0 = 6000$. Details on the numerical methodology are given in Section 6.

Objective of the trader. To account for the case where not all shares have been sold prior to time T we specify the liquidation value of the remaining share position W_T . This liquidation value is of the form $h(W_T)S_T$ for some increasing, continuous and concave function h with $h(w) \leq w$ and $h(0) = 0$. For instance, the choice $h(w) = \frac{w}{1+\vartheta w}$ for some $\vartheta > 0$ models the situation where the liquidation value that is strictly smaller than the book value, reflecting the limited liquidity of the market for the stock.¹ For $\vartheta \rightarrow \infty$ we obtain the limit $h(w) \equiv 0$; this models the situation where a block transaction at the terminal date is prohibitively expensive. We model the temporary price impact by a nonnegative, continuous increasing function f , so that the proceeds from liquidation are given for every $t \in [0, T]$ by $\int_0^t \nu_s S_s (1 - f(\nu_s)) ds$. For instance, Almgren et al. [3] propose a power function of the form $f(\nu) = c_f \nu^\zeta$ and they estimate the coefficient $\zeta \approx 0.6$.

Now we define the time τ as the minimum of the first time the inventory is completely liquidated and the horizon T :

$$\tau := \inf\{t \geq 0 : W_t \leq 0\} \wedge T. \quad (2.6)$$

¹Note that when considering a block transaction at the horizon date T we do not need to model market resiliency or permanent price impact as the model ‘ends’ at T .

Denote by ρ the (subjective) discount rate of the trader. Consider an admissible strategy ν and denote the corresponding bid price by S^ν . The expected discounted value of the proceeds generated by the liquidation strategy ν is equal to

$$J(\nu) = \mathbb{E} \left(\int_0^\tau e^{-\rho u} \nu_u S_u^\nu (1 - f(\nu_u)) du + e^{-\rho \tau} S_\tau^\nu h(W_\tau) \right). \quad (2.7)$$

The trader wants to maximize (2.7) over all admissible strategies; the corresponding optimal value is denoted by J^* , or, if we want to emphasize the dependence on the upper bound on the liquidation rate, by $J^{*, \nu^{\max}}$. Note that the form of the objective function in (2.7) implies that the trader is risk neutral. Risk neutrality is frequently assumed in the literature on optimal order execution, see for instance Bertsimas and Lo [15]. Moreover, in our setup the risk the trader may take is limited as we consider pure selling strategies and as the time period $[0, T]$ is fairly short.

The next proposition shows that, if the trader uses \mathbb{F}^S -adapted strategies with $\nu_t \leq m$ for all $t \in [0, T]$, the optimal liquidation value $J^{*, m}$ is bounded independently of m .

Proposition 2.5. *Suppose that Assumption 2.1 holds and that the function $(t, e, \nu) \rightarrow \bar{\eta}^{\mathbf{P}}(t, e, \nu)$ from (2.3) is decreasing in ν , and set*

$$\bar{\eta} = 0 \vee \sup\{\bar{\eta}^{\mathbf{P}}(t, e, 0) - \rho : t \in [0, T], e \in \mathcal{E}\}.$$

Then $\sup_{m>0} J^{*, m} \leq w_0 S_0 e^{\bar{\eta} T}$.

Note that the upper bound on J^* corresponds to the liquidation value of the inventory in a frictionless model where the expected value of the bid price grows at the maximum rate $\bar{\eta} + \rho$.

Proof. Fix some \mathbb{F}^S -adapted strategy ν with values in $[0, m]$ and let $\tilde{S}_t^\nu = e^{-\rho t} S_t^\nu$. Since $W_t = w_0 - \int_0^t \nu_s ds$ we get by partial integration that

$$\int_0^\tau \nu_s \tilde{S}_s^\nu ds = - \int_0^\tau \tilde{S}_s^\nu dW_s = S_0 w_0 - \tilde{S}_\tau^\nu W_\tau + \int_0^\tau W_s d\tilde{S}_s^\nu.$$

Since $h(w) \leq w$ and $f(\nu) \geq 0$ we thus get that

$$\int_0^\tau \nu_u \tilde{S}_u^\nu (1 - f(\nu_u)) du + \tilde{S}_\tau^\nu h(W_\tau) \leq \int_0^\tau \nu_u \tilde{S}_u^\nu du + \tilde{S}_\tau^\nu W_\tau = S_0 w_0 + \int_0^\tau W_u d\tilde{S}_u^\nu.$$

Now $\int_0^\tau W_u d\tilde{S}_u^\nu = \int_0^\tau W_u \tilde{S}_u^\nu dM_u^R + \int_0^\tau W_u \tilde{S}_u^\nu (\bar{\eta}^{\mathbf{P}}(u, Y_{u-}, \nu_{u-}) - \rho) du$. Moreover, $\int_0^{\wedge \tau} W_u \tilde{S}_u^\nu dM_u^R$ is a true martingale: as $0 \leq W_u \leq w_0$, a similar argument as in the proof of Lemma A.1 shows that this process is of integrable quadratic variation. Since $\bar{\eta}^{\mathbf{P}}(u, Y_{u-}, \nu_{u-}) - \rho \leq \bar{\eta}$, $\tau \leq T$ and $W_u \leq w_0$, we get

$$\begin{aligned} J(\nu) &\leq S_0 w_0 + \mathbb{E} \left(\int_0^\tau W_u \tilde{S}_u^\nu (\bar{\eta}^{\mathbf{P}}(u, Y_{u-}, \nu_{u-}) - \rho) du \right) \\ &\leq S_0 w_0 + \mathbb{E} \left(\int_0^T w_0 \tilde{S}_u^\nu \bar{\eta} du \right). \end{aligned} \quad (2.8)$$

Next we show that $\mathbb{E}(\tilde{S}_t^\nu) \leq S_0 e^{\bar{\eta} t}$. To this end, note that by Lemma A.1, $\int_0^\cdot \nu_{s-} dM_s^R$ is a true martingale so that

$$\mathbb{E}(\tilde{S}_t^\nu) = S_0 + \mathbb{E} \left(\int_0^t \tilde{S}_u^\nu (\bar{\eta}^{\mathbf{P}}(u, Y_{u-}, \nu_{u-}) - \rho) du \right) \leq S_0 + \bar{\eta} \int_0^t \mathbb{E}(\tilde{S}_u^\nu) du,$$

and the claim follows from the Gronwall inequality. Using (2.8) we finally get that $J(\nu) \leq S_0 w_0 (1 + \int_0^T \bar{\eta} e^{\bar{\eta} u} du) = S_0 w_0 e^{\bar{\eta} T}$, and hence the result. \square

3 Partial Information and Filtering

Considering \mathbb{F}^S -adapted investment strategies results in an optimal control problem under partial information. The standard approach to dealing with such problems is to introduce the filter for the Markov chain as additional state variable of the control problem. In this section we therefore derive the filtering equations for our model. Filtering for point process observations is for instance considered in Ceci and Gerardi [23], Elliott and Malcolm [32], Frey and Schmidt [34], Ceci and Colaneri [21, 22]. This literature is mostly based on the innovations approach. In this paper, instead, we address the filtering problem via the *reference probability approach*. This methodology relies on the existence of an equivalent probability measure such that the observation process is driven by a random measure with dual predictable projection independent of the Markov chain, see for instance Brémaud [16, Chapter 6]. The reference probability approach permits us to overcome the difficulties caused by the fact that the observation process S is affected by the liquidation strategy chosen by the trader.

3.1 Reference probability. We start from a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{Q})$ that supports a Markov chain Y with state space \mathcal{E} and generator matrix Q , and an *independent* Poisson random measure μ^R with compensator $\eta^{\mathbf{Q}}(dz)dt$ as in Assumption 2.1.2; \mathbf{Q} is known as the *reference probability* measure. Note that under \mathbf{Q} the dynamics of S and R are independent of the liquidation strategy ν so that the filtration \mathbb{F}^S is exogenously given. Moreover, the independence of Y and μ^R implies that R and Y have no common jumps. For $(t, e, \nu, z) \in [0, T] \times \mathcal{E} \times [0, \nu^{\max}] \times \text{supp}(\eta)$, we define the function β by

$$\beta(t, e, \nu, z) := \frac{d\eta^{\mathbf{P}}(t, e, \nu; dz)}{d\eta^{\mathbf{Q}}(dz)}(z) - 1,$$

i.e. $\beta(t, e, \nu, z) + 1$ is the Radon-Nikodym derivative of the measure $\eta^{\mathbf{P}}(t, e, \nu; dz)$ with respect to $\eta^{\mathbf{Q}}(dz)$.

Fix now an admissible liquidation strategy ν and define for $t \in [0, T]$ the stochastic exponential \tilde{Z} by

$$\tilde{Z}_t = 1 + \int_0^t \int_{\mathbb{R}} \tilde{Z}_{s-} \beta(s, Y_{s-}, \nu_{s-}, z) (\mu^R(ds, dz) - \eta^{\mathbf{Q}}(dz)ds). \quad (3.1)$$

Then we have the following result.

Lemma 3.1. *Let Assumption 2.1 prevail. Then the process \tilde{Z} is a strictly positive martingale with $\mathbb{E}^{\mathbf{Q}}(\tilde{Z}_T) = 1$. Define a measure \mathbf{P} on \mathcal{F}_T by setting $\frac{d\mathbf{P}}{d\mathbf{Q}}|_{\mathcal{F}_T} = \tilde{Z}_T$. Then \mathbf{P} and \mathbf{Q} are equivalent and, under \mathbf{P} , the random measure μ^R has the compensator $\eta^{\mathbf{P}}$.*

The proof of the lemma is postponed to Appendix A.

3.2 Filtering equations. For a function $f: \mathcal{E} \rightarrow \mathbb{R}$, we introduce the filter $\pi(f)$ as the optional projection of the process $f(Y)$ on the filtration \mathbb{F}^S , i.e. $\pi(f)$ is a càdlàg process such that for all $t \in [0, T]$, it holds that $\pi_t(f) = \mathbb{E}(f(Y_t) | \mathcal{F}_t^S)$. Note that $f(Y_t) = \langle \mathbf{f}, Y_t \rangle$ for all $t \in [0, T]$, where $\langle \cdot, \cdot \rangle$ denotes the scalar product on \mathbb{R}^K and $\mathbf{f}_i = f(e_i)$, $i \in \{1, \dots, K\}$, so that functions of the Markov chain can be identified with K -vectors. Let for all $t \in [0, T]$ and $i \in \{1, \dots, K\}$, $\pi_t^i := \mathbb{E}(\mathbf{1}_{\{Y_t=e_i\}} | \mathcal{F}_t^S)$. Then, we can represent the filter as

$$\pi_t(f) = \sum_{i=1}^K \mathbf{f}_i \pi_t^i = \langle \mathbf{f}, \pi_t \rangle, \quad t \in [0, T].$$

The objective of this section is to derive the dynamics of the process $\boldsymbol{\pi} = (\pi^1, \dots, \pi^K)$. To this end, we first observe that by the Kallianpur-Striebel formula we have $\pi_t(f) := \frac{p_t(f)}{p_t(1)}$ for all $t \in [0, T]$, where $p(f)$ denotes the *unnormalized* version of the filter, which is defined by

$$p_t(f) := \mathbb{E}^{\mathbf{Q}} \left(\tilde{Z}_t \langle \mathbf{f}, Y_t \rangle \mid \mathcal{F}_t^S \right), \quad t \in [0, T]. \quad (3.2)$$

The dynamics of $p(f)$ is given in the next theorem.

Theorem 3.2 (The Zakai equation). *Suppose Assumption 2.1 holds and let $f: \mathcal{E} \rightarrow \mathbb{R}$. Then, for all $t \in [0, T]$, the unnormalized filter (3.2) solves the equation:*

$$p_t(f) = \pi_0(f) + \int_0^t p_s(Qf) ds + \int_0^t \int_{\mathbb{R}} p_{s-}(\beta(z)f) (\mu^R(ds, dz) - \eta_s^{\mathbf{Q}}(dz) ds), \quad (3.3)$$

where $p_{t-}(\beta(z)f) = \mathbb{E}^{\mathbf{Q}} \left(f(Y_{t-}) \tilde{Z}_{t-} \beta(t, Y_{t-}, \nu_{t-}, z) \mid \mathcal{F}_t^S \right)$ and $p_t(Qf) = \mathbb{E}^{\mathbf{Q}} \left(\tilde{Z}_t \langle Q\mathbf{f}, Y_t \rangle \mid \mathcal{F}_t^S \right)$.

We now provide the general idea of the proof, details are given in Appendix A. Consider the process \tilde{Z} defined in (3.1) and some function $f: \mathcal{E} \rightarrow \mathbb{R}$. Then by Itô's formula the product $\tilde{Z}_t f(Y_t)$ has the following (\mathbf{Q}, \mathbb{F}) -semimartingale decomposition

$$\begin{aligned} \tilde{Z}_t f(Y_t) &= f(Y_0) + \int_0^t \tilde{Z}_s \langle Q\mathbf{f}, Y_t \rangle ds + \int_0^t \tilde{Z}_s dM_s^f \\ &\quad + \int_0^t \tilde{Z}_s f(Y_s) \int_{\mathbb{R}} \beta(s, Y_{s-}, \nu_{s-}, z) (\mu^R(ds, dz) - \eta_s^{\mathbf{Q}}(dz) ds), \end{aligned}$$

where $M^f = (M^f)_{t \in [0, T]}$ is the true (\mathbb{F}, \mathbf{Q}) -martingale appearing in the semimartingale decomposition of $f(Y)$. Taking the conditional expectation with respect to \mathcal{F}_t^S yields the result, since it can be shown that $\mathbb{E}^{\mathbf{Q}} \left(\int_0^t \tilde{Z}_s dM_s^f \mid \mathcal{F}_t^S \right) = 0$.

We introduce the notation

$$\pi_{t-}(\eta^{\mathbf{P}}(dz)) := \sum_{i=1}^K \pi_{t-}^i \eta^{\mathbf{P}}(t, e_i, \nu_t, dz), \quad t \in [0, T].$$

By applying [16, Ch. II, Theorem T14] it is easy to see that $\pi_{t-}(\eta^{\mathbf{P}}(dz)) dt$ provides the $(\mathbb{F}^S, \mathbf{P})$ -dual predictable projection of the measure μ^R . The next proposition provides the dynamics of the conditional state probabilities.

Proposition 3.3. *The process $\boldsymbol{\pi}$ solves the*

$$\pi_t^i = \pi_0^i + \int_0^t \sum_{j=1}^K q^{ji} \pi_s^j ds + \int_0^t \int_{\mathbb{R}} \pi_{s-}^i u^i(s, \nu_{s-}, \pi_s, z) (\mu^R(ds, dz) - \pi_{s-}(\eta^{\mathbf{P}}(dz)) ds), \quad (3.4)$$

for every $t \in [0, T]$ and $1 \leq i \leq K$, where $u^i(t, \nu, \pi, z) := \frac{(d\eta^{\mathbf{P}}(t, e_i, \nu)/d\eta^{\mathbf{Q}})(z)}{\sum_{j=1}^K \pi^j (d\eta^{\mathbf{P}}(t, e_j, \nu)/d\eta^{\mathbf{Q}})(z)} - 1$.

Proof. By the Kallianpur-Striebel formula we have that $\pi_t(f) := \frac{p_t(f)}{p_t(1)}$, for every $t \in [0, T]$. Then, by (3.3) and Itô formula we get the dynamics of the normalized filter $\pi(f)$. The claimed result is obtained by setting $f(Y_t) = \mathbf{1}_{\{Y_t = e_i\}}$, for every $i \in \{1, \dots, K\}$. \square

Note that the filtering equation (3.4) does not depend on the particular choice of $\eta^{\mathbf{Q}}$.

Filter equations for Example 2.3. In the following we give the dynamics of π for Example 2.3. For a two-state Markov chain it is sufficient to specify the dynamics of $\pi = \pi^1$, since $\pi^2 = 1 - \pi^1$. Define two point processes by $N_t^{\text{up}} = \sum_{T_n \leq t} 1_{\{\Delta R_{T_n} = \theta\}}$ and $N_t^{\text{down}} = \sum_{T_n \leq t} 1_{\{\Delta R_{T_n} = -\theta\}}$, for all $t \in [0, T]$, that count the upward and the downward jumps of the return process. It is easily seen that for every $(\nu, \pi, z) \in [0, \nu^{\max}] \times [0, 1] \times \{-\theta, \theta\}$, the function u^1 is given by

$$u^1(\nu, \pi, z) = \frac{\lambda^+(e_1, \nu)}{\pi \lambda^+(e_1, \nu) + (1 - \pi) \lambda^+(e_2, \nu)} 1_{\{z = \theta\}} + \frac{\lambda^-(e_1, \nu)}{\pi \lambda^-(e_1, \nu) + (1 - \pi) \lambda^-(e_2, \nu)} 1_{\{z = -\theta\}}. \quad (3.5)$$

By Corollary 3.3 we then get the following equation for $\pi_t = \pi_t^1$:

$$\begin{aligned} d\pi_t &= (q^{11}\pi_t + q^{21}(1 - \pi_t)) dt \\ &+ \pi_t(1 - \pi_t) ((\lambda^+(e_1, \nu_t) + \lambda^-(e_1, \nu_t)) - (\lambda^+(e_2, \nu_t) + \lambda^-(e_2, \nu_t))) dt \\ &+ \pi_{t-} \left(\frac{\lambda^+(e_1, \nu_t)}{\pi_{t-} \lambda^+(e_1, \nu_t) + (1 - \pi_{t-}) \lambda^+(e_2, \nu_t)} - 1 \right) dN_t^{\text{up}} \\ &+ \pi_{t-} \left(\frac{\lambda^-(e_1, \nu_t)}{\pi_{t-} \lambda^-(e_1, \nu_t) + (1 - \pi_{t-}) \lambda^-(e_2, \nu_t)} - 1 \right) dN_t^{\text{down}}. \end{aligned} \quad (3.6)$$

4 Control Problem I: Analysis via PDMPs

We begin with a brief overview of our analysis of the control problem (2.7). In Proposition 4.3 below we show that the Kushner-Stratonovich equation (3.4) has a unique solution. Then standard arguments ensure that the original control problem under incomplete information is equivalent to a control problem under complete information with state process equal to the $(K+2)$ -dimensional process $X := (W, S, \pi)$. This process is a PDMP in the sense of Davis [29], that is a trajectory of X consists of a deterministic part which solves an ordinary differential equation (ODE), interspersed by random jumps. Therefore, to solve the optimal liquidation problem we apply control theory for PDMPs. This theory is based on the observation that a control problem for a PDMP is discrete in time: loosely speaking, at every jump-time of the process one chooses a control policy to be followed up to the next jump time or until maturity. Therefore, one can identify the control problem for the PDMP with a control problem for a discrete-time, infinite-horizon Markov decision model (MDM). Using this connection we show that the value function of the optimal liquidation problem is continuous and that is the unique solution of the dynamic programming or optimality equation for the MDM. These results are the basis for the viscosity-solution characterization of the value function in Section 5.

4.1 Optimal liquidation as a control problem for a PDMP. From the viewpoint of the trader endowed with the filtration \mathbb{F}^S , the state of the economic system at time $t \in [0, T]$ is given by $X_t = (W_t, S_t, \pi_t)$. Since it is more convenient to work with autonomous Markov processes we include time into the state and define $\tilde{X}_t := (t, X_t)$. The state space of \tilde{X} is $\tilde{\mathcal{X}} = [0, T] \times \mathcal{X}$ where $\mathcal{X} = [0, w_0] \times \mathbb{R}^+ \times \mathcal{S}^K$ with \mathcal{S}^K being the K -dimensional simplex. Let ν be the liquidation strategy followed by the trader. It follows from (2.1), (3.4), and from the fact that the bid price is a pure jump process that between jump times the state process follows the ODE $d\tilde{X}_t = g(\tilde{X}_t, \nu_t) dt$, where the vector field $g(\tilde{x}, \nu) \in \mathbb{R}^{K+3}$ is given by $g^1(\tilde{x}, \nu) = 1$, $g^2(\tilde{x}, \nu) = -\nu$, $g^3(\tilde{x}, \nu) = 0$, and for $k = 1, \dots, K$,

$$g^{k+3}(\tilde{x}, \nu) = \sum_{j=1}^K q^{jk} \pi^j - \pi^k \sum_{j=1}^K \pi^j \int_{\mathbb{R}} u^k(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_j, \nu, dz).$$

For our analysis we need the following regularity property of g .

Lemma 4.1. *Under Assumption 2.1, the function g is Lipschitz continuous in \tilde{x} uniformly in $(t, \nu) \in [0, T] \times [0, \nu^{\max}]$; the Lipschitz constant is denoted by K_g .*

The proof is postponed to Appendix B.

The jump rate of the state process \tilde{X} is given by $\lambda(\tilde{X}_{t-}, \nu_{t-})$, $t \in (0, T]$, where for every $(\tilde{x}, \nu) \in \tilde{\mathcal{X}} \times [0, \nu^{\max}]$,

$$\lambda(\tilde{x}, \nu) = \lambda(t, w, s, \pi, \nu) := \sum_{j=1}^K \pi^j \eta^{\mathbf{P}}(t, e_j, \nu, \mathbb{R}).$$

Next, we identify the transition kernel $Q_{\tilde{X}}$ that governs the jumps of \tilde{X} . Denote by $\{T_n\}_{n \in \mathbb{N}}$ the sequence of jump times of \tilde{X} . It follows from (3.4) that for any measurable function $f: \tilde{\mathcal{X}} \rightarrow \mathbb{R}^+$,

$$Q_{\tilde{X}} f(\tilde{x}, \nu) := E(f(\tilde{X}_{T_n}) \mid T_n = t, X_{T_n-} = x, \nu_{T_n-} = \nu) = \frac{1}{\lambda(\tilde{x}, \nu)} \bar{Q}_{\tilde{X}} f(\tilde{x}, \nu),$$

where the unnormalized kernel $\bar{Q}_{\tilde{X}}$ is given by

$$\bar{Q}_{\tilde{X}} f(\tilde{x}, \nu) = \sum_{j=1}^K \pi^j \int_{\mathbb{R}} f(t, w, s(1+z), \pi^1(1+u^1), \dots, \pi^K(1+u^K)) \eta^{\mathbf{P}}(t, e_j, \nu, dz).$$

Here u^i is short for $u^i(t, \nu, \pi, z)$. Summarizing, \tilde{X} is a PDMP with characteristics given by the vector field g , the jump rate λ and the transition kernel $Q_{\tilde{X}}$.

It is standard in control theory for PDMPs to work with so-called open-loop controls. In the current context this means that the trader chooses at each jump time $T_n < \tau$ a liquidation policy ν^n to be followed up to $T_{n+1} \wedge \tau$. This policy may depend on the state $\tilde{X}_{T_n} = (T_n, X_{T_n})$.

Definition 4.2. Denote by \mathcal{A} the set of measurable mappings $\alpha: [0, T] \rightarrow [0, \nu^{\max}]$. An *admissible open loop liquidation strategy* is a sequence of mappings $\{\nu^n\}_{n \in \mathbb{N}}$ with $\nu^n: \tilde{\mathcal{X}} \rightarrow \mathcal{A}$; the liquidation rate at time t is given by $\nu_t = \sum_{n=0}^{\infty} \mathbf{1}_{(T_n \wedge \tau, T_{n+1} \wedge \tau]}(t) \nu^n(t - T_n, \tilde{X}_{T_n})$.

It follows from Brémaud [16, Theorem T34, Appendix A2] that an admissible strategy as defined in Section 2.2 is of the form given in Definition 4.2, but for $\mathcal{F}_{T_n}^S$ measurable mappings $\nu^n: \Omega \rightarrow \mathcal{A}$ for every $n \in \mathbb{N}$, that ν^n may depend on the entire history of the system. General results for Markov decision models (see Bäuerle and Rieder [12, Theorem 2.2.3]) show that the expected profit of the trader stays the same if instead we consider the smaller class of admissible open loop strategies, so that we may restrict ourselves to this class.

Proposition 4.3. *Let Assumption 2.1 hold. For every admissible liquidation strategy $\{\nu^n\}_{n \in \mathbb{N}}$ and every initial value \tilde{x} , a unique PDMP with characteristics g , λ , and $Q_{\tilde{X}}$ as above exists. In particular the Kushner-Stratonovic equation (3.4) has a unique solution.*

Proof. Lemma 4.1 implies that for $\alpha \in \mathcal{A}$ the ODE $d\tilde{X}_t = g(\tilde{X}_t, \alpha_t)dt$ has a unique solution so that between jumps the state process is well-defined. At any jump time T_n , \tilde{X}_{T_n} is uniquely defined in terms of observable data $(T_n, \Delta R_{T_n})$. Moreover, since the jump intensity is bounded by λ^{\max} , jump times cannot accumulate. \square

Denote by $\mathbf{P}_{(t,x)}^{\{\nu^n\}}$ (equiv. $\mathbf{P}_{\tilde{x}}^{\{\nu^n\}}$) the law of the state process provided that $X_t = x \in \mathcal{X}$ and that the trader uses the open-loop strategy $\{\nu^n\}_{n \in \mathbb{N}}$. The reward function associated to an admissible liquidation strategy $\{\nu^n\}_{n \in \mathbb{N}}$ is defined by

$$V(t, x, \{\nu^n\}_{n \in \mathbb{N}}) = \mathbb{E}_{(t,x)}^{\{\nu^n\}} \left(\int_t^\tau e^{-\rho(u-t)} \nu_u S_u (1 - f(\nu_u)) du + e^{\rho(\tau-t)} h(W_\tau) S_\tau \right),$$

and the value function of the liquidation problem under partial information is

$$V(t, x) = \sup \{V(t, x, \{\nu^n\}_{n \in \mathbb{N}}) : \{\nu^n\}_{n \in \mathbb{N}} \text{ admissible liquidation strategy} \}. \quad (4.1)$$

Remark 4.4. Note that the compensator $\eta^\mathbf{P}$ and the dynamics of the filter π are independent of the current bid price s , and that the payoff of a liquidation strategy $\{\nu^n\}_{n \in \mathbb{N}}$ is positively homogeneous in s . This implies that the reward and the value function of the liquidation problem are positively homogeneous in s and, in particular, $V(t, w, s, \pi) = sV(t, w, 1, \pi)$.

4.2 Associated Markov decision model. The optimization problem in (4.1) is discrete in time since the control policy is chosen at the discrete time points T_n , $n \in \mathbb{N}$, and the value of the state process at these time points forms a discrete-time Markov chain (for $T_n < \tau$). Hence (4.1) can be rewritten as a control problem in an infinite horizon Markov decision model. The *state process* of the MDM is given by the sequence $\{L_n\}_{n \in \mathbb{N}}$ of random variables with

$$L_n = \tilde{X}_{T_n} \text{ for } T_n < \tau \text{ and } L_n = \bar{\Delta} \text{ for } T_n \geq \tau, \quad n \in \mathbb{N},$$

where $\bar{\Delta}$ is the cemetery state. In order to derive the transition kernel of the sequence $\{L_n\}_{n \in \mathbb{N}}$ and the reward function of the MDM, we introduce some notation. For a function $\alpha \in \mathcal{A}$ we denote by $\tilde{\varphi}_t^\alpha(\tilde{x})$ or by $\tilde{\varphi}_t(\alpha, \tilde{x})$ the flow of the initial value problem $\frac{d}{dt} \tilde{X}(t) = g(\tilde{X}(t), \alpha_t)$ with initial condition $\tilde{X}(0) = \tilde{x}$. Whenever we want to make the dependence on time explicit we write $\tilde{\varphi}^\alpha$ in the form (t, φ^α) . Moreover, we define the function λ_u^α by

$$\lambda_u^\alpha(\tilde{x}) = \lambda(\tilde{\varphi}_u^\alpha(\tilde{x}), \alpha_u) = \lambda((t+u, \varphi_u^\alpha), \alpha_u) \quad u \in [0, T-t],$$

and we let $\Lambda_u^\alpha(\tilde{x}) = \int_0^u \lambda_v^\alpha(\tilde{x}) dv$.

Next we take a closer look at the *boundary* of $\tilde{\mathcal{X}}$. First note that the process π takes values in the hyperplane $\mathcal{H}^K = \{x \in \mathbb{R}^K : \sum_{i=1}^K x_i = 1\}$, so that $\tilde{\mathcal{X}}$ is contained in the set $\mathcal{H} = \mathbb{R}^3 \times \mathcal{H}^K$, which is a hyperplane of \mathbb{R}^{K+3} . When considering the boundary or the interior of the state space we always refer to the *relative* boundary or the *relative* interior with respect to \mathcal{H} . Of particular interest to us is the *active boundary* Γ of the state space, that is the part of the boundary of $\tilde{\mathcal{X}}$ which can be reached by the flow $\tilde{\varphi}^\alpha(\tilde{x})$ starting in an interior point $\tilde{x} \in \text{int}(\tilde{\mathcal{X}})$. The boundary of $\tilde{\mathcal{X}}$ can only be reached if $w = 0$, if $t = T$, or if the filter process reaches the boundary of the K -dimensional simplex. The latter is not possible: indeed, if $\pi_0^i > 0$, then $\pi_t^i > 0$ for all $t \in [0, T]$, since there is a positive probability that the Markov chain has not changed its state and since the conditional distribution of Y_t given \mathcal{F}_t^S is equivalent to the unconditional distribution of Y_t by the Kallianpur-Striebel formula. Hence the active boundary equals $\Gamma = \Gamma_1 \cup \Gamma_2$, where

$$\Gamma_1 = [0, T] \times \{0\} \times (0, \infty) \times \mathcal{S}_0^K \quad \text{and} \quad \Gamma_2 = \{T\} \times [0, w_0] \times (0, \infty) \times \mathcal{S}_0^K, \quad (4.2)$$

and where \mathcal{S}_0^K is the interior of \mathcal{S}^K , i.e. $\mathcal{S}_0^K := \{x \in \mathcal{S}^K : x_i > 0 \text{ for all } i\}$. In (4.2) Γ_1 is the *lateral* part of the active boundary corresponding to an inventory level equal to zero, and Γ_2

is the terminal boundary corresponding to the exit from the state space at maturity T . In the sequel we denote the first exit time of the flow $\tilde{\varphi}^\alpha(\tilde{x})$ from $\tilde{\mathcal{X}}$ by

$$\tau^\varphi = \tau^\varphi(\tilde{x}, \alpha) = \inf\{u \geq 0 : \tilde{\varphi}_u^\alpha(\tilde{x}) \in \Gamma\}.$$

Notice that the stopping time τ defined in (2.6) corresponds to the first time the state process \tilde{X} reaches the active boundary Γ .

Using similar arguments as in Bäuerle and Rieder [12, Section 8.2] or in Davis [29, Section 44], it is easily seen that the *transition kernel* Q_L of the sequence $\{L_n\}_{n \in \mathbb{N}}$ is given by

$$Q_L f((t, x), \alpha) = \int_0^{\tau^\varphi(\tilde{x})} e^{-\Lambda_u^\alpha(\tilde{x})} \overline{Q}_{\tilde{X}} f(u+t, \varphi_u(\tilde{x}), \alpha_u) du + e^{-\Lambda_{\tau^\varphi}^\alpha(\tilde{x})} f(\bar{\Delta});$$

we omit the details. Moreover, since the cemetery state is absorbing, $Q_L \mathbf{1}_{\{\bar{\Delta}\}}(\bar{\Delta}, \alpha) = 1$. Finally we define the *one-period reward function* $r: \tilde{\mathcal{X}} \times \mathcal{A} \rightarrow \mathbb{R}^+$ by

$$r(\tilde{x}, \alpha) = \int_0^{\tau^\varphi} e^{-\rho u} e^{-\Lambda_u^\alpha(\tilde{x})} \alpha_u s(1 - f(\alpha_u)) du + e^{-\rho \tau^\varphi} e^{-\Lambda_{\tau^\varphi}^\alpha(\tilde{x})} h(w_{\tau^\varphi}) s, \quad (4.3)$$

and w_{τ^φ} the inventory-component of $\tilde{\varphi}^\alpha$, and we set $r(\bar{\Delta}) = 0$. For an admissible strategy $\{\nu^n\}_{n \in \mathbb{N}}$ we set $J_\infty^{\{\nu^n\}}(\tilde{x}) = \mathbb{E}_{\tilde{x}}^{\{\nu^n\}} \left(\sum_{n=0}^{\infty} r(L_n, \nu^n(L_n)) \right)$, and

$$J_\infty(\tilde{x}) := \sup \left\{ J_\infty^{\{\nu^n\}}(\tilde{x}) : \{\nu^n\}_{n \in \mathbb{N}} \text{ admissible liquidation strategy} \right\}. \quad (4.4)$$

The next lemma shows that the MDM with transition kernel Q_L and one-period reward $r(L, \alpha)$ is equivalent to the optimization problem (4.1).

Lemma 4.5. *For every admissible strategy $\{\nu^n\}$ it holds that $V^{\{\nu^n\}} = J_\infty^{\{\nu^n\}}$. Hence $V = J_\infty$, and the control problems (4.1) and (4.4) are equivalent.*

The proof is similar to the proof of Davis [29, Theorem 44.9] and is therefore omitted.

4.3 The Bellman equation. In this section we study the Bellman equation for the value function V . Define for $\alpha \in \mathcal{A}$ and a measurable function $v: \tilde{\mathcal{X}} \rightarrow \mathbb{R}^+$ the function $\mathcal{L}v(\cdot, \alpha)$ by

$$\mathcal{L}v(\tilde{x}, \alpha) = r(\tilde{x}, \alpha) + Q_L v(\tilde{x}, \alpha), \quad \tilde{x} \in \tilde{\mathcal{X}}.$$

The *maximal reward operator* \mathcal{T} is then given by $\mathcal{T}v(\tilde{x}) = \sup_{\alpha \in \mathcal{A}} \mathcal{L}v(\tilde{x}, \alpha)$. Since the one-period reward function is nonnegative we have a so-called *positive* MDM and it follows from Bäuerle and Rieder [12, Theorem 7.4.3] that the value function satisfies the so-called *Bellman* or *optimality equation*

$$V(\tilde{x}) = \mathcal{T}V(\tilde{x}), \quad \tilde{x} \in \tilde{\mathcal{X}},$$

that is V is a fixed point of the operator \mathcal{T} . In order to characterize V as viscosity solution of the HJB equation associated with the PDMP \tilde{X} (see Section 5) we need a stronger result. We want to show: i) that the value function V is the unique fixed point of \mathcal{T} in a suitable function class \mathcal{M} ; ii) that for a starting point $v^0 \in \mathcal{M}$ iterations of the form $v^{n+1} = \mathcal{T}v^n$, $n \in \mathbb{N}$, converge to V ; and iii) that V is continuous on $\tilde{\mathcal{X}}$.

Points i) and ii) follow from the next lemma.

Lemma 4.6. Define for $\gamma > 0$, the function $b: \tilde{\mathcal{X}} \cup \{\bar{\Delta}\} \rightarrow \mathbb{R}^+$ by $b(\tilde{x}) = b(t, x) := swe^{\gamma(T-t)}$, $\tilde{x} \in \tilde{\mathcal{X}}$, and $b(\bar{\Delta}) = 0$. Then under Assumption 2.1, b is a bounding function for the MDM with transition kernel Q_L and reward function r , that is there are constants c_r, c_Q such that for all $(\tilde{x}, \alpha) \in \tilde{\mathcal{X}} \times \mathcal{A}$.

$$|r(\tilde{x}, \alpha)| \leq c_r b(\tilde{x}) \quad \text{and} \quad Q_L b(\tilde{x}, \alpha) \leq c_Q b(\tilde{x}).$$

Moreover, for γ sufficiently large it holds that $c_Q < 1$, that is the MDM is contracting.

The proof is postponed to Appendix B. In the sequel we denote by \mathcal{B}_b the set of functions

$$\mathcal{B}_b := \{v: \tilde{\mathcal{X}} \rightarrow \mathbb{R} \text{ such that } \sup_{\tilde{x} \in \tilde{\mathcal{X}}} |v(\tilde{x})/b(\tilde{x})| < \infty\},$$

and we define for $v \in \mathcal{B}_b$ the norm $\|v\|_b = \sup_{\tilde{x} \in \tilde{\mathcal{X}}} |v(\tilde{x})/b(\tilde{x})|$. Then the following holds, see Bäuerle and Rieder [12, Section 7.3]: a) $(\mathcal{B}_b, \|\cdot\|_b)$ is a Banach space; b) $\mathcal{T}(\mathcal{B}_b) \subset \mathcal{B}_b$; c) $\|\mathcal{T}v - \mathcal{T}u\|_b \leq c_Q \|v - u\|_b$.

If the MDM is contracting, the maximal reward operator is a contraction on $(\mathcal{B}_b, \|\cdot\|_b)$, and the value function is an element of \mathcal{B}_b . Banach's fixed point theorem thus gives properties i.) and ii.) above with $\mathcal{M} = \mathcal{B}_b$. In order to establish property iii.) (continuity of V) we observe that the set

$$\mathcal{C}_b := \{v \in \mathcal{B}_b: v \text{ is continuous}\}$$

is a closed subset of $(\mathcal{B}_b, \|\cdot\|_b)$, see Bäuerle and Rieder [12, Section 7.3]. Moreover, we show in Proposition 4.8 that under certain continuity conditions (see Assumptions 2.1 and 4.7), \mathcal{T} maps \mathcal{C}_b into itself. Hence it follows from Banach's fixed point theorem that $V \in \mathcal{C}_b$.

Assumption 4.7. 1. The measure $\eta^j(t, \nu; dz)$ for $j \in \{1, \dots, K\}$ is continuous in the weak topology, i.e. for all bounded and continuous ϕ , the mapping $(t, \nu) \mapsto \int_{\mathbb{R}} \phi(z) \eta^j(dz)$ is continuous on $[0, T] \times [0, \nu^{\max}]$.

2. For the functions w^j introduced in Corollary 3.3 the following holds: for any sequence $\{(t^n, \nu^n, \pi^n)\}_{n \in \mathbb{N}}$ with $(t^n, \nu^n, \pi^n) \in [0, T] \times [0, \nu^{\max}] \times \mathcal{S}^K$ for every $n \in \mathbb{N}$, such that $(t^n, \nu^n, \pi^n) \xrightarrow{n \rightarrow \infty} (t, \nu, \pi)$, one has

$$\lim_{n \rightarrow \infty} \sup_{z \in \text{supp}(\eta)} |w^j(t^n, \nu^n, \pi^n, z) - w^j(t, \nu, \pi, z)| = 0.$$

Proposition 4.8. Suppose that Assumptions 2.1 and 4.7 hold and let $v \in \mathcal{C}_b$. Then $\mathcal{T}v \in \mathcal{C}_b$.

Proof. Consider some sequence $\tilde{x}_n \rightarrow \tilde{x}$ for $n \rightarrow \infty$. Since $|\mathcal{T}v(\tilde{x}_n) - \mathcal{T}v(\tilde{x})| \leq \sup_{\alpha \in \mathcal{A}} |\mathcal{L}v(\tilde{x}_n, \alpha) - \mathcal{L}v(\tilde{x}, \alpha)|$, it suffices to estimate the difference $\sup_{\alpha \in \mathcal{A}} |\mathcal{L}v(\tilde{x}_n, \alpha) - \mathcal{L}v(\tilde{x}, \alpha)|$. First, note that by the Lipschitz continuity of g , established in Lemma 4.1, we have

$$|\tilde{\varphi}_t^\alpha(\tilde{x}_n) - \tilde{\varphi}_t^\alpha(\tilde{x})| \leq |\tilde{x}_n - \tilde{x}| + K_g \int_0^t |\tilde{\varphi}_u^\alpha(\tilde{x}_n) - \tilde{\varphi}_u^\alpha(\tilde{x})| du.$$

Gronwall inequality hence yields that

$$\sup_{t \in [0, T], \alpha \in \mathcal{A}} |\tilde{\varphi}_t^\alpha(\tilde{x}_n) - \tilde{\varphi}_t^\alpha(\tilde{x})| \leq |\tilde{x}_n - \tilde{x}| e^{K_g T}, \quad (4.5)$$

and thus uniform convergence for $n \rightarrow \infty$ of the flow $\tilde{\varphi}^\alpha(\tilde{x}_n)$ to $\tilde{\varphi}^\alpha(\tilde{x})$. This does however not imply that τ^{φ^n} , the entrance time of $\tilde{\varphi}^\alpha(\tilde{x}_n)$ into the active boundary of the state space, converges to τ^φ for $n \rightarrow \infty$. To deal with this issue we distinguish two cases:

Case 1. The flow $\tilde{\varphi}^\alpha(\tilde{x})$ exits the state space $\tilde{\mathcal{X}}$ at the terminal boundary Γ_2 (see (4.2)). This implies that $\tau^\varphi = T - t$ and that the inventory level w_u is strictly positive for $u < T - t$. We therefore conclude from (4.5) that τ^{φ_n} converges to $T - t$. Under Assumptions 2.1 and 4.7 the uniform convergence $\lim_{n \rightarrow \infty} \sup_{\alpha \in \mathcal{A}} |\mathcal{L}v(\tilde{x}_n, \alpha) - \mathcal{L}v(\tilde{x}, \alpha)| = 0$ thus follows immediately using the definition of r and the continuity of the mapping $(\tilde{x}, \nu) \mapsto \bar{Q}v(\tilde{x}, \nu)$ established in Lemma B.1, see Appendix B.

Case 2. The flow $\tilde{\varphi}^\alpha(\tilde{x})$ exits $\tilde{\mathcal{X}}$ at the lateral boundary Γ_1 so that $w_{\tau^\varphi} = 0$. In that case (4.5) implies that $\liminf_{n \rightarrow \infty} \tau^{\varphi_n} \geq \tau^\varphi$; it is however possible that this inequality is strict. We first show continuity of the reward function for that case. We decompose $r(\tilde{x}_n, \alpha)$ as follows, setting $\rho = 0$ for simplicity:

$$r(\tilde{x}_n, \alpha) = s \int_0^{\tau^\varphi \wedge \tau^{\varphi_n}} e^{-\Lambda_u^\alpha(\tilde{x}_n)} \alpha_u (1 - f(\alpha_u)) du \quad (4.6)$$

$$+ s \int_{\tau^\varphi \wedge \tau^{\varphi_n}}^{\tau^{\varphi_n}} e^{-\Lambda_u^\alpha(\tilde{x}_n)} \alpha_u (1 - f(\alpha_u)) du + s e^{-\Lambda_{\tau^\varphi}^\alpha(\tilde{x}_n)} h(w_{\tau^\varphi}). \quad (4.7)$$

Now it follows from (4.5) that the integral in (4.6) converges for $n \rightarrow \infty$ to $r(\tilde{x}, \alpha)$ uniformly in $\alpha \in \mathcal{A}$. The terms in (4.7) are bounded from above by $sw_{\tau^\varphi \wedge \tau^{\varphi_n}}$; this can be shown via a similar partial integration argument as in the proof of Lemma 4.6. Moreover, $w_{\tau^\varphi \wedge \tau^{\varphi_n}}$ converges uniformly in $\alpha \in \mathcal{A}$ to $w_{\tau^\varphi} = 0$, so that (4.7) converges to zero. Next we turn to the transition kernel. We decompose $Q_L v$:

$$Q_L v(\tilde{x}_n, \alpha) = \int_0^{\tau^\varphi \wedge \tau^{\varphi_n}} e^{-\Lambda_u^\alpha(\tilde{x}_n)} \bar{Q}v(\tilde{\varphi}_u(\tilde{x}_n), \alpha_u) du + \int_{\tau^\varphi \wedge \tau^{\varphi_n}}^{\tau^{\varphi_n}} e^{-\Lambda_u^\alpha(\tilde{x}_n)} \bar{Q}v(\tilde{\varphi}_u^\alpha(\tilde{x}_n), \alpha_u) du.$$

For $n \rightarrow \infty$, the first integral converges to $Q_L v(\tilde{x}, \alpha)$ using (4.5) and the continuity of the mapping $(\tilde{x}, \nu) \mapsto \bar{Q}v(\tilde{x}, \nu)$ (Lemma B.1). To estimate the second term note that $\bar{Q}v(\tilde{x}, \nu) \leq \|v\|_b s w \lambda(\tilde{x}, \nu)$ (as $\frac{1}{\lambda} \bar{Q}$ is a probability transition kernel), so that the integral is bounded by

$$\|v\|_b s w_{\tau^\varphi \wedge \tau^{\varphi_n}} \int_{\tau^\varphi \wedge \tau^{\varphi_n}}^{\tau^{\varphi_n}} \lambda_u^\alpha e^{-\Lambda_u^\alpha(\tilde{x}_n)} du \leq \|v\|_b s w_{\tau^\varphi \wedge \tau^{\varphi_n}},$$

and the last term converges to zero for $n \rightarrow \infty$, uniformly in $\alpha \in \mathcal{A}$. \square

Remark 4.9. Note that existing continuity results for $\mathcal{L}v(\cdot, \alpha)$ such as Davis [29, Theorem 44.11] make the assumption that the flow φ^α reaches the active boundary at a uniform speed, independent of the chosen control. In order to ensure this hypothesis in our framework we would have to impose a strictly positive lower bound on the admissible liquidation rate. This is an economically implausible restriction of the strategy space which is why we prefer to rely on a direct argument.

We summarize the results of this section in the following theorem.

Theorem 4.10. *Suppose that Assumptions 2.1 and 4.7 hold. Then the value function V is continuous on $\tilde{\mathcal{X}}$ and satisfies the boundary conditions $V(\tilde{x}) = 0$ for \tilde{x} in the lateral boundary Γ_1 and $V(T, x) = sh(w)$. Moreover, V is the unique solution of the Bellman or optimality equation $V = \tilde{\mathcal{T}}V$ in \mathcal{B}_b .*

5 Control Problem II: Viscosity Solutions

In this section we show that the value function is a viscosity solution of the standard HJB equation associated with the controlled Markov process (W, π) and we derive a comparison principle for

that equation. These results are crucial to ensure the convergence of suitable numerical schemes for the HJB equation and thus for the numerical solution of the optimal liquidation problem. In Section 5.2 we provide an example which shows that in general the HJB equation does not admit a classical solution.

5.1 Viscosity solution characterization. As a first step we write down the Bellman equation and we use the positive homogeneity of V in the bid price (see Remark 4.4) to eliminate s from the set of state variables. Define $\tilde{\mathcal{Y}} = [0, T] \times [0, w_0] \times \mathcal{S}^K$ and denote by $\text{int } \tilde{\mathcal{Y}}$ and $\partial \tilde{\mathcal{Y}}$ the *relative* interior and the *relative* boundary of $\tilde{\mathcal{Y}}$ with respect to the hyperplane $\mathbb{R}^2 \times \mathcal{H}^K$. For $\tilde{y} \in \tilde{\mathcal{Y}}$ we set

$$V'(\tilde{y}) = V'(t, w, \pi) := V(t, w, 1, \pi),$$

so that the value function satisfies the relation $V(\tilde{x}) = sV'(\tilde{y})$. For $\nu \in [0, \nu^{\max}]$, $\tilde{y} \in \tilde{\mathcal{Y}}$, and any measurable function $\Psi: \tilde{\mathcal{Y}} \rightarrow \mathbb{R}^+$, define

$$\bar{Q}'\Psi(\tilde{y}, \nu) := \sum_{j=1}^K \pi^j \int_{\mathbb{R}} (1+z)\Psi(t, w, (\pi^i(1+u^i(t, \pi, \nu, z)))_{i=1, \dots, K}) \eta^{\mathbf{P}}(t, e_j, \nu, dz)$$

and note that $\bar{Q}V(\tilde{x}, \nu) = s\bar{Q}'V'(\tilde{y}, \nu)$. From now on we denote by $\tilde{\varphi}_u^\alpha(\tilde{y})$ the flow of the vector field g with price component g^3 omitted, and we write τ^φ for the first time this flow reaches the active boundary of $\tilde{\mathcal{Y}}$ given by $\Gamma := [0, T] \times \{0\} \times \mathcal{S}_0^K \cup \{T\} \times [0, w_0] \times \mathcal{S}_0^K$ of $\tilde{\mathcal{Y}}$.

By positive homogeneity, the Bellman equation for V reduces to the following optimality equation for V' :

$$V'(\tilde{y}) = \sup_{\alpha \in A} \left\{ \int_0^{\tau^\varphi} e^{-(\rho u + \Lambda_u^\alpha(\tilde{y}))} (\alpha_u(1 - f(\alpha_u)) + \bar{Q}'V'(\tilde{\varphi}_u^\alpha(\tilde{y}), \alpha_u)) du + e^{-(\rho\tau^\varphi + \Lambda_{\tau^\varphi}^\alpha(\tilde{y}))} h(w_{\tau^\varphi}) \right\}. \quad (5.1)$$

For $\Psi: \tilde{\mathcal{Y}} \rightarrow \mathbb{R}^+$ bounded, define the function $\ell^\Psi: \tilde{\mathcal{Y}} \times [0, \nu^{\max}] \rightarrow \mathbb{R}^+$ and the operator \mathcal{T}' by

$$\ell^\Psi(\tilde{y}, \nu) = \nu(1 - f(\nu)) + \bar{Q}'\Psi(\tilde{y}, \nu), \quad (5.2)$$

$$\mathcal{T}'\Psi(\tilde{y}) = \sup_{\alpha \in A} \left\{ \int_0^{\tau^\varphi} e^{-(\rho u + \Lambda_u^\alpha(\tilde{y}))} \ell^\Psi(\tilde{\varphi}_u^\alpha(\tilde{y}), \alpha_u) du + e^{-(\rho\tau^\varphi + \Lambda_{\tau^\varphi}^\alpha(\tilde{y}))} h(w_{\tau^\varphi}) \right\}. \quad (5.3)$$

Note that for fixed Ψ , $v^\Psi := \mathcal{T}'\Psi$ is the value function of a deterministic exit-time optimal control problem with instantaneous reward ℓ^Ψ and boundary value h . Viscosity solutions for this problem are studied extensively in Barles [8]. Moreover, the optimality equation (5.1) for V' can be written as the fixed point equation $V' = \mathcal{T}'V'$. Davis and Farid [30] observed that this can be used to obtain a viscosity solution characterization of the value function in a PDMP control problem, and we now explain how this idea applies in our framework. Define for $\Psi: \tilde{\mathcal{Y}} \rightarrow \mathbb{R}^+$ the function $F_\Psi: \tilde{\mathcal{Y}} \times \mathbb{R}^+ \times \mathbb{R}^{K+2} \rightarrow \mathbb{R}$ by

$$F_\Psi(\tilde{y}, v, p) = - \sup \left\{ -(\rho + \lambda(\tilde{y}, \nu))v + g(\tilde{y}, \nu)'p + \ell^\Psi(\tilde{y}, \nu) : \nu \in [0, \nu^{\max}] \right\}.$$

The dynamic programming equation associated with the control problem (5.3) is

$$F_\Psi(\tilde{y}, v^\Psi(\tilde{y}), \nabla v^\Psi(\tilde{y})) = 0 \text{ for } \tilde{y} \in \text{int } \tilde{\mathcal{Y}}, \quad v^\Psi(\tilde{y}) = h(\tilde{y}) \text{ for } \tilde{y} \in \partial \tilde{\mathcal{Y}}. \quad (5.4)$$

Moreover, since $V' = \mathcal{T}'V'$, we expect that V' solves in a suitable sense the equation

$$F_{V'}(\tilde{y}, V'(\tilde{y}), \nabla V'(\tilde{y})) = 0, \text{ for } \tilde{y} \in \text{int } \tilde{\mathcal{Y}}, \quad V'(\tilde{y}) = h(\tilde{y}) \text{ for } \tilde{y} \in \partial \tilde{\mathcal{Y}}. \quad (5.5)$$

Remark 5.1. Notice that, equations (5.4) and (5.5) differs in the sense that in (5.4) the function F_Ψ enters with Ψ fixed, whereas in (5.5) one works with the function $F_{V'}$. This reflects the fact that control problem (5.3) associated to equation (5.4) has an exogenously given running cost, while in the optimization problem (5.1), leading to equation (5.5), function V' is the solution of a fixed point equation, and therefore the running cost is endogenous.

There are two issues with equations (5.4) and (5.5): v^Ψ and V' are typically not \mathcal{C}^1 functions, and the value of these functions on the non-active part $\partial\tilde{\mathcal{Y}} \setminus \Gamma$ of the boundary is determined endogenously. Following Barles [8] we therefore work with the following notion of viscosity solutions.

Definition 5.2. 1. A bounded upper semi-continuous (u.s.c.) function v on $\tilde{\mathcal{Y}}$ is a *viscosity subsolution* of (5.4), if for all $\phi \in \mathcal{C}^1(\tilde{\mathcal{Y}})$ and all local maxima $\tilde{y}_0 \in \tilde{\mathcal{Y}}$ of $v - \phi$ one has

$$\begin{aligned} F_\Psi(\tilde{y}_0, v(\tilde{y}_0), \nabla\phi(\tilde{y}_0)) &\leq 0 \text{ for } \tilde{y}_0 \in \text{int}\tilde{\mathcal{Y}}, \\ \min \{F_\Psi(\tilde{y}_0, v(\tilde{y}_0), \nabla\phi(\tilde{y}_0)), v(\tilde{y}_0) - h(\tilde{y}_0)\} &\leq 0 \text{ for } \tilde{y}_0 \in \partial\tilde{\mathcal{Y}}. \end{aligned} \quad (5.6)$$

A bounded lower semi-continuous (l.s.c.) function u on $\tilde{\mathcal{Y}}$ is a *viscosity supersolution* of (5.4), if for all $\phi \in \mathcal{C}^1(\tilde{\mathcal{Y}})$ and all local minima $\tilde{y}_0 \in \tilde{\mathcal{Y}}$ of $u - \phi$ one has

$$\begin{aligned} F_\Psi(\tilde{y}_0, u(\tilde{y}_0), \nabla\phi(\tilde{y}_0)) &\geq 0 \text{ for } \tilde{y}_0 \in \text{int}\tilde{\mathcal{Y}}, \\ \max \{F_\Psi(\tilde{y}_0, u(\tilde{y}_0), \nabla\phi(\tilde{y}_0)), u(\tilde{y}_0) - h(\tilde{y}_0)\} &\geq 0 \text{ for } \tilde{y}_0 \in \partial\tilde{\mathcal{Y}}. \end{aligned} \quad (5.7)$$

A *viscosity solution* v^Ψ of (5.4) is either a continuous function on $\tilde{\mathcal{Y}}$ that is both a sub and a supersolution of (5.4), or a bounded function with u.s.c. and l.s.c. envelopes that are a sub and a supersolution of (5.4).

2. A bounded u.s.c. function v on $\tilde{\mathcal{Y}}$ is a *viscosity subsolution* of (5.5), if the relation (5.6) holds for $F = F_v$. Similarly, a bounded l.s.c. function u on $\tilde{\mathcal{Y}}$ is a *viscosity supersolution* of (5.4), if (5.7) holds for $F = F_u$. Finally, V' is a *viscosity solution* of (5.5), if it is both a sub and a supersolution of that equation.

Note that Definition 5.2 allows for the case that $v^\Psi(\tilde{y}_0) \neq h(\tilde{y}_0)$ for certain boundary points $\tilde{y}_0 \in \partial\tilde{\mathcal{Y}}$. In particular, if $F_\Psi(\tilde{y}_0, v^\Psi(\tilde{y}_0), \nabla v^\Psi(\tilde{y}_0)) = 0$ in the viscosity sense, (5.6) and (5.7) hold irrespectively of the value of $h(\tilde{y}_0)$.

Theorem 5.3. *Suppose that Assumptions 2.1 and 4.7 hold. Then the value function V' is a continuous viscosity solution of (5.5) in $\tilde{\mathcal{Y}}$. Moreover, a comparison principle holds for (5.5): if $v \geq 0$ is a subsolution and $u \geq 0$ a supersolution of (5.5) such that $v(\tilde{y})/w$ and $u(\tilde{y})/w$ are bounded on $\tilde{\mathcal{Y}}$ and such that $v = u = h$ on the active boundary Γ of $\tilde{\mathcal{Y}}$, then $v \leq u$ on $\text{int}\tilde{\mathcal{Y}}$. It follows that V' is the only continuous viscosity solution of (5.5).*

Proof. First, by Theorem 4.10, V' is continuous. Moreover, Barles [8, Theorem 5.2] implies that V' is a viscosity solution of (5.4) with $\Psi = V'$ and hence of equation (5.5).

Next we prove the comparison principle. In order to establish the inequality $v \leq u$ we use an inductive argument based on the monotonicity of \mathcal{T}' and on a comparison result for (5.4). Let $u_0 := u$ and define $u_1 = \mathcal{T}u_0$. It follows from Barles [8, Theorem 5.2] that u_1 is a viscosity solution of (5.4) with $\Psi = u_0$. Moreover, $u_1(\tilde{y})/w$ is bounded on $\tilde{\mathcal{Y}}$ so that $u_1 = h$ on Γ . Since u_0 is a supersolution of (5.5) it is also a supersolution of (5.4) with $\Psi = u_0$. Barles [8, Theorem

5.7] gives the inequality $u_1 \leq u_0$ on $\text{int } \tilde{\mathcal{Y}}$, since the functions u^+ and u^- defined in that theorem coincide in our case. Define now inductively $u_n = \mathcal{T}'u_{n-1}$, and suppose that $u_n \leq u_{n-1}$. Then, using the monotonicity of \mathcal{T}' , we have

$$u_{n+1} = \mathcal{T}'u_n \leq \mathcal{T}'u_{n-1} = u_n.$$

This proves that $u_{n+1} \leq u_n$ for every $n \in \mathbb{N}$. Moreover, as explained in Section 4.3, the sequence $\{u_n\}_{n \in \mathbb{N}}$ converges to V' , so that $u_n \geq V'$ for all n . In the same way we can construct a sequence of functions $\{v^n\}$ with $v_0 = v$ such that $v^n \uparrow V'$, and we conclude that $v \leq V' \leq u$. The remaining statements are clear. \square

Remark 5.4. Note that the results in Davis and Farid [30] do not apply directly in our case since their assumptions on the behaviour of the vector field g on the lateral boundary are not satisfied in our model, see also the discussion in Remark 4.9. Moreover, Davis and Farid [30] do not give a comparison principle for (5.5).

Finally, we write the dynamic programming equation (5.5) explicitly. To this end, we use the fact that $\lambda(\tilde{y}, \nu) = \sum_{k=1}^K \pi^k \eta^{\mathbf{P}}(t, e_k, \nu, \mathbb{R})$, the definition of g , and the definition of $l^{V'}$ in (5.2) to obtain

$$0 = \frac{\partial V'}{\partial t}(t, w, \pi) + \sup \{H(\nu, t, w, \pi, V', \nabla V') : \nu \in [0, \nu^{\max}]\}, \quad \text{with} \quad (5.8)$$

$$\begin{aligned} H(\nu, t, w, \pi, V', \nabla V') &= -\rho V' + \nu(1 - f(\nu)) - \nu \frac{\partial V'}{\partial w}(t, w, \pi) \\ &+ \sum_{k,j=1}^K \frac{\partial V'}{\partial \pi^k}(t, w, \pi) \pi^j \left(q^{jk} - \pi^k \int_{\mathbb{R}} u^k(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_j, \nu, dz) \right) \\ &+ \sum_{j=1}^K \pi^j \int_{\mathbb{R}} \Delta V'(t, w, \pi, z) \eta^{\mathbf{P}}(t, e_j, \nu, dz), \end{aligned} \quad (5.9)$$

and $\Delta V'(t, w, \pi, z) := (1+z)V'(t, w, (\pi^i(1+u^i(t, \nu, \pi, z)))_{i=1, \dots, K}) - V'(t, w, \pi)$. This equation coincides with the standard HJB equation associated with the controlled Markov process (W, π) . The advantage of using viscosity solution theory is that we are able to give a mathematical meaning to this equation even if V' is merely continuous. This is relevant in our context. Indeed, in the next section we present a simple example where V' is not \mathcal{C}^1 .

5.2 A counterexample. We now give an example within a setup where the value function is a viscosity solution of the dynamic programming equation but not a classical solution. Precisely, we work in the context of Example 2.2 with linear permanent price impact and deterministic compensator $\eta^{\mathbf{P}}$. For simplicity we let $\rho = 0$, $s = 1$, $h(w) \equiv 0$, $f(\nu) \equiv 0$ (zero terminal liquidation value and no temporary price impact). Moreover, we assume that $c^{\text{up}} < c^{\text{down}}$. The function $\bar{\eta}^{\mathbf{P}}$ from (2.3) is thus given by $\bar{\eta}^{\mathbf{P}}(\nu) := \theta(c^{\text{up}} - c^{\text{down}}(1 + a\nu))$ and $\bar{\eta}^{\mathbf{P}}(\nu) < 0$ for $\nu > 0$. It follows that S^ν is a supermartingale for any admissible ν , and we conjecture that it is optimal to sell as fast as possible to reduce the loss due to the falling bid price. Denote by $\tau(w) := w/\nu^{\max}$ the minimal time necessary to liquidate the inventory w . The optimal strategy is thus given by $\nu_t^* = \nu^{\max} \mathbf{1}_{[0, \tau(w_0) \wedge T]}(t)$. Moreover, for $t < \tau(w_0) \wedge T$ one has $\bar{\eta}^{\mathbf{P}}(\nu_t) = \bar{\eta}^{\mathbf{P}}(\nu^{\max})$ and $\mathbb{E}(S_t^*) = \exp(t\bar{\eta}^{\mathbf{P}}(\nu^{\max}))$. Hence we get that

$$J(\nu^*) = \int_0^{\tau(w_0) \wedge T} \nu^{\max} \exp(u \bar{\eta}^{\mathbf{P}}(\nu^{\max})) du.$$

Solving this integral we get the following candidate for the value function

$$V'(t, w) := \frac{\nu^{\max}}{\bar{\eta}^{\mathbf{P}}(\nu^{\max})} \left\{ \exp(\bar{\eta}^{\mathbf{P}}(\nu^{\max})(\tau(w) \wedge (T-t))) - 1 \right\}, \quad (t, w) \in [0, T] \times [0, \nu^{\max}]. \quad (5.10)$$

In order to verify that V' is in fact the value function we show that V' is a viscosity solution of the HJB equation 5.8. In the current setting this equation becomes

$$-\frac{\partial V'}{\partial t} - \sup \left\{ \nu - \nu \frac{\partial V'}{\partial w} - \bar{\eta}^{\mathbf{P}}(\nu) V' : \nu \in [0, \nu^{\max}] \right\} = 0. \quad (5.11)$$

First note that V' satisfies the correct terminal and boundary conditions. Define the set

$$G := \{(t, w) \in [0, T] \times [0, w_0] : \tau(w) = (T-t)\}.$$

The function V' is \mathcal{C}^1 on $[0, T] \times [0, w_0] \setminus G$, and it is a classical solution of (5.11) on this set. However V' is not differentiable on G and hence not a classical solution everywhere.

Fix some point $(\bar{t}, \bar{w}) \in G$. In order to show that V' is a viscosity solution of the (5.8) we need to verify the subsolution property in this point. (For the supersolution property there is nothing to show as there is no \mathcal{C}^1 -function ϕ such that $V' - \phi$ has a local minimum in (\bar{t}, \bar{w}) .) Consider $\phi \in \mathcal{C}^1$ such that $V' - \phi$ has a local maximum in (\bar{t}, \bar{w}) . By considering the left and right derivatives of the functions $t \mapsto (V' - \phi)(t, \bar{w})$ respectively $w \mapsto (V' - \phi)(\bar{t}, w)$ we get the following inequalities for the partial derivatives of ϕ

$$-\nu^{\max} e^{\bar{\eta}^{\mathbf{P}}(\nu^{\max})(T-\bar{t})} \leq \frac{\partial \phi}{\partial t}(\bar{t}, \bar{w}) \leq 0 \quad \text{and} \quad 0 \leq \frac{\partial \phi}{\partial w}(\bar{t}, \bar{w}) \leq \exp(\bar{\eta}^{\mathbf{P}}(\nu^{\max})\tau(\bar{w})).$$

Moreover, it holds on G that $V'(t, w) = \frac{\nu^{\max}}{\bar{\eta}^{\mathbf{P}}(\nu^{\max})} \left\{ \exp(\bar{\eta}^{\mathbf{P}}(\nu^{\max})(T-t)) - 1 \right\}$. As $w = \nu^{\max}(T-t)$ on G , differentiating with respect to t gives that

$$\left(\frac{\partial \phi}{\partial t} - \nu^{\max} \frac{\partial \phi}{\partial w} \right)(\bar{t}, \bar{w}) = -\nu^{\max} \exp(\bar{\eta}^{\mathbf{P}}(\nu^{\max})(T-\bar{t})). \quad (5.12)$$

Applying the inequalities for $\frac{\partial \phi}{\partial w}$ we get that

$$\sup \left\{ \nu - \nu \frac{\partial \phi}{\partial w} + \bar{\eta}^{\mathbf{P}}(\nu) V' : \nu \in [0, \nu^{\max}] \right\} = \nu^{\max} \left(-\frac{\partial \phi}{\partial w} + e^{\bar{\eta}^{\mathbf{P}}(\nu^{\max})(T-t)} \right).$$

Using (5.12) this gives $-\frac{\partial \phi}{\partial t} - \sup \left\{ \nu - \nu \frac{\partial \phi}{\partial w} + \bar{\eta}^{\mathbf{P}}(\nu) V' : \nu \in [0, \nu^{\max}] \right\} = 0$ and hence the subsolution property.

Remark 5.5. It can be shown that for $\nu^{\max} \rightarrow \infty$ the value function V' from (5.10) converges to $V'^{\infty}(t, w) := -\frac{1}{\theta e^{\text{down}_a}} (\exp(-w\theta a c^{\text{down}}) - 1)$ and that V'^{∞} is a strict (classical) supersolution of equation (5.11). Hence V'^{∞} is the value function of the optimal liquidation problem for $\nu^{\max} = \infty$, and we conclude that without an upper bound on the liquidation rate the viscosity-solution characterization from Theorem 5.3 does not hold in general.

6 Examples and numerical results

In this section we study the optimal liquidation rate and the expected liquidation profit in our model. For concreteness we work in the framework of Example 2.3, that is the example where $\eta^{\mathbf{P}}$ depends on the liquidation strategy as well as on a two-state Markov chain. We focus on two different research questions: (i) the influence of model parameters on the form of the optimal liquidation rate; (ii) the additional liquidation profit from the use of stochastic filtering and a comparison to classical approaches. Moreover, we report the results of a small calibration study.

Numerical method. Since the HJB equation cannot be solved analytically, we resort to numerical methods. We apply an explicit finite difference scheme to solve the HJB equation and to compute the corresponding liquidation strategy. First, we turn the HJB equation into an initial value problem via time reversion. Given a time discretization $0 = t_0 < \dots < t_k < \dots < t_m = T$ we set $V'_{t_0} = h$, and given V'_{t_k} , we approximate the liquidation strategy as follows:

$$\nu_{t_k}^*(w, \pi) := \operatorname{argmax}_{\nu \in [0, \nu^{\max}]} H(\nu, t_k, w, \pi, V'_{t_k}, \nabla^{\text{disc}} V'_{t_k}), \quad (6.1)$$

where ∇^{disc} is the gradient operator with derivatives replaced by suitable finite differences. In a slight abuse of language we refer to $\nu_{t_k}^*$ from (6.1) as the optimal liquidation rate. With this we obtain the next time iterate of the value function,

$$V'_{t_{k+1}} = V'_{t_k} + (t_{k+1} - t_k) H(\nu_{t_k}^*, t_k, w, \pi, V'_{t_k}, \nabla^{\text{disc}} V'_{t_k}).$$

Since the comparison principle holds, as shown in Theorem 5.3, and the value function is the unique viscosity solution of our HJB equation, we get convergence of the proposed procedure to the value function by similar arguments as in Barles and Souganidis [9], Dang and Forsyth [28]; details are omitted.

6.1 Optimal liquidation rate. We start by computing the optimal liquidation rate $\nu_{t_k}^*$ for Example 2.3, assuming that the temporary price impact is of the form $f(\nu) = c_f \nu^\varsigma$ for $\varsigma > 0$. Since $\pi_t^1 + \pi_t^2 = 1$ for all $t \in [0, T]$, we can eliminate the process π^2 from the set of state variables. In the sequel we denote by π_t the conditional probability of being in the good state e_1 at time t and by $V'(t, w, \pi)$ the value function evaluated at the point $(t, w, (\pi, 1 - \pi))$. To compute $\nu_{t_k}^*$ we substitute the functions u_i given in (3.5) and the dynamics of the process $(\pi_t)_{t \in [0, T]}$ from (3.6) into the general HJB equation (5.8). Denote by

$$\pi_t^{\text{post}} = \frac{\pi_t c_1^{\text{down}}}{\pi_t c_1^{\text{down}} + (1 - \pi_t) c_2^{\text{down}}}, \quad t \in [0, T],$$

the updated (posterior) probability of state e_1 given that a downward jump occurs at t . Moreover, denote the discretized partial derivatives of V' appearing in (6.1) by $\frac{\delta V'}{\delta w}$ and $\frac{\delta V'}{\delta \pi}$. Substitution into (5.9) leads to

$$\nu_{t_k}^* = \operatorname{argmax}_{\nu \in [0, \nu^{\max}]} \{ \nu(1 - c_f \nu^\varsigma) - \nu C(t_k, w, \pi) \}, \quad \text{with} \quad (6.2)$$

$$\begin{aligned} C(t_k, w, \pi) &= \frac{\delta V'}{\delta w}(t_k, w, \pi) + \frac{\delta V'}{\delta \pi}(t_k, w, \pi) \pi(1 - \pi) a(c_1^{\text{down}} - c_2^{\text{down}}) \\ &\quad - \{ (1 - \theta) V'(t_k, w, \pi^{\text{post}}) - V'(t_k, w, \pi) \} (\pi c_1^{\text{down}} + (1 - \pi) c_2^{\text{down}}) a. \end{aligned} \quad (6.3)$$

Maximizing (6.2) with respect to ν , we get that $\nu_{t_k}^* = 0$ if $C(t_k, w, \pi) > 1$; for $C(t_k, w, \pi) \leq 1$ one has $\nu_{t_k}^* = \tilde{\nu}^* \wedge \nu^{\max}$, where $\tilde{\nu}^*$ solves the equation

$$1 - c_f(\varsigma + 1)\nu^\varsigma = C(t_k, w, \pi). \quad (6.4)$$

This characterization of $\nu_{t_k}^*$ is very intuitive: $1 - c_f(\varsigma + 1)\nu^\varsigma$ gives the marginal liquidation benefit due to an increase in ν and $C(t_k, w, \pi)$ can be viewed as marginal cost of an increase in ν (see below). For $C(t_k, w, \pi) \leq 1$, $\tilde{\nu}^*$ is found by equating marginal benefit and marginal cost; for $C(t_k, w, \pi) > 1$ the marginal benefit is smaller than the marginal cost for all $\nu \geq 0$ and $\nu_{t_k}^* = 0$.

The optimal liquidation rate $\nu_{t_k}^*$ is thus determined by the marginal cost $C(t_k, w, \pi)$, and we now give an economic interpretation of the terms in (6.3). First, $\frac{\delta V'}{\delta w}$ is a marginal opportunity

cost, since selling inventory reduces the amount that can be liquidated in the future. Moreover, it holds that

$$-((1 - \theta)V'(t_k, w, \pi^{\text{post}}) - V'(t_k, w, \pi)) = \theta V'(t_k, w, \pi^{\text{post}}) - (V'(t_k, w, \pi^{\text{post}}) - V'(t_k, w, \pi)).$$

The term $\theta V'(t_k, w, \pi^{\text{post}})$ gives the reduction in the expected liquidation value due to a downward jump in the return process, and $(\pi c_1^{\text{down}} + (1 - \pi)c_2^{\text{down}})a$ is the marginal increase in the intensity of a downward jump, so that the term

$$\theta V(t_k, w, \pi^{\text{post}})(\pi c_1^{\text{down}} + (1 - \pi)c_2^{\text{down}})a \quad (6.5)$$

measures the marginal cost due to permanent price impact; in the sequel we refer to (6.5) as *illiquidity cost*. Finally, note that $\pi^{\text{post}} - \pi = \frac{\pi(1-\pi)(c_1^{\text{down}} - c_2^{\text{down}})}{\pi c_1^{\text{down}} + (1-\pi)c_2^{\text{down}}}$. Hence the remaining terms in (6.3) are equal to

$$-\left(V'(t_k, w, \pi^{\text{post}}) - V'(t_k, w, \pi) - \frac{\delta V'}{\delta \pi}(t_k, w, \pi)(\pi^{\text{post}} - \pi)\right)a(\pi c_1^{\text{down}} + (1 - \pi)c_2^{\text{down}}). \quad (6.6)$$

Simulations indicate that V' is convex in π ; this is quite natural as it implies that uncertainty about the true state reduces the optimal liquidation value. It follows that (6.6) is negative which leads to an increase in the optimal liquidation rate (6.4). Since $\pi^{\text{post}} - \pi$ is largest for $\pi \approx 0.5$, this effect is most pronounced if the investor is uncertain about the true state. Hence (6.6) can be viewed as an uncertainty correction that makes the trader sell faster if he is uncertain about the true state.

Numerical analysis and varying price impact parameters. To gain further insight into the structure of the optimal liquidation rate we resort to numerical experiments. We work with the parameter set given in Table 2. Moreover, we set the liquidation value $h(w) \equiv 0$, that is we assume that block transactions at the horizon date are prohibitively expensive. Without loss of generality we set $s = 1$, so that the expected liquidation profit is equal to V' .

w_0	ν^{max}	T	ρ	θ	$c_1^{\text{up}}, c_2^{\text{down}}$	$c_1^{\text{down}}, c_2^{\text{up}}$	a	ς	q^{12}	q^{21}
6000	9000	2 days	0.00005	0.001	1000	900	7×10^{-6}	0.6	4	4

Table 2: Parameter values used in numerical experiments.

First, we discuss the form of the optimal liquidation rate for varying size of the temporary price impact, that is for varying c_f , keeping the permanent price impact parameter a constant at the moderate value $a = 7 \times 10^{-6}$. Figure 1 shows the liquidation rate at $t = 0$ for the cases of no, intermediate, and large temporary price impact as a function of w and π . The figure is a contour plot: white areas correspond to $\nu_0 = 0$; black areas correspond to selling at maximum speed ($\nu_0 = 18000$); grey areas correspond to selling at a moderate speed, see also the color bars below the graphs. Comparing the graphs we see that for higher temporary price impact (high c_f) the trader tends to trade more evenly over the state space to keep the cost due to the temporary price impact small. Note that for $c_f \rightarrow 0$ the liquidation strategy converges to a *bang-bang type strategy*. The optimal policy is then characterized by two regions: a *sell* region, where the trader sells at the maximum speed, and a *wait* region, where she does not sell at all. This reaction of $\nu_{t_k}^*$ to variations in c_f can also be derived theoretically by inspection of (6.4).




Figure 1: Contour plot of the liquidation policy as a function of w (abscissa) and π (ordinate) for $c_f = 5 \times 10^{-11}$ (left), $c_f = 5 \times 10^{-5}$ (middle), and $c_f = 10^{-5}$ (right) and $t = 0$ for Example 2.3.

Now we study the impact of the permanent price impact a on the form of the optimal liquidation rate. Figure 1 shows that for moderate a the liquidation rate is decreasing in π and increasing in the inventory level. The situation changes when the permanent price impact becomes large. The left plot in Figure 2 depicts the sell and wait regions under partial information in dependence of the inventory level w and the filter probability π for $a = 7 \times 10^{-5}$, and essentially without temporary price impact. For this value of a the sell region forms a band from low values of w and π to high values of w and π . In particular, for large w and small π there is a *gambling region* where the trader does not sell, even if a small value of π means that the bid price is trending downward (recall that π gives the probability that Y is in the good state). In the presence of a temporary price impact (right plot of Figure 2) the qualitative behaviour of the liquidation rate is similar to the case without temporary price impact, but the transition from waiting to selling at the maximum rate is smooth.




Figure 2: Contour plot of the liquidation policy as a function of w (abscissa) and π (ordinate) for $c_f = 5 \times 10^{-11}$ (left) and $c_f = 10^{-5}$ (right) for $a = 7 \times 10^{-5}$ and $t = 0$ for Example 2.3.

The observed form of $\nu_{t_k}^*$ has the following explanation. Our numerical experiments show that for the chosen parameter values V' is almost linear in π , so that the uncertainty correction (6.6) is negligible. Hence the liquidation rate $\nu_{t_k}^*$ is determined by the interplay of the opportunity cost $\frac{\delta V'}{\delta w}(t_k, w, \pi)$ and of the illiquidity cost (6.5). We found that the opportunity cost is increasing in π . This is very intuitive: in the good state the investor expects an increase in the expected bid price which makes additional inventory more valuable. Moreover, we found that $\frac{\delta V'}{\delta w}(t_k, w, \pi)$ is decreasing in w , that is the optimal liquidation problem has decreasing returns to scale. The illiquidity cost has the opposite monotonicity behaviour: it is increasing in w (as it is proportional to $V'(t_k, w, \pi^{\text{post}})$) and, for the given parameters, decreasing in π . Now for small values of a the opportunity cost dominates the illiquidity cost for all (w, π) and $C(t_k, w, \pi)$ is increasing in π .

and decreasing in w . By (6.4), the liquidation rate is thus decreasing in π and increasing in w , which is in line with the monotonicity behaviour observed in Figure 1. If a is large the situation is more involved. The opportunity cost dominates for small w , leading to a liquidation rate that is decreasing in π . For large w the illiquidity cost dominates, C is decreasing in π , and the optimal liquidation rate is increasing in π . For w large enough this effect is strong enough to generate the unexpected gambling region observed in Figure 2.

Impact of other model components. In reality the support of $\eta^{\mathbf{P}}$ is larger than $\{-\theta, \theta\}$ as the price may jump by more than one tick. Hence it is important to test the sensitivity of $\nu_{t_k}^*$ with respect to the precise form of the support. To this end, we computed the optimal strategy for a different parameter set $\tilde{\theta}, \tilde{c}_i^{\text{up}}, \tilde{c}_i^{\text{down}}, i = 1, 2$ with $\tilde{\theta} = 2\theta$ and $\tilde{c}_i^{\text{up}} = 0.5c_i^{\text{up}}, \tilde{c}_i^{\text{down}} = 0.5c_i^{\text{down}}, i = 1, 2$. Note that for the new parameters the support of $\eta^{\mathbf{P}}$ is different but the expected return of the bid price in each of the two states is the same. We found that the liquidation value and the optimal strategy were nearly identical to the original case. This shows that our approach is quite robust with respect to the exact form of the support of $\eta^{\mathbf{P}}$ and justifies the use of a simple model with only two possible values for the jump size of R .

6.2 Gain from filtering and comparison to classical approaches. In this section we compare the expected proceeds of the optimal liquidation rate to the expected proceeds of a trader who mistakenly uses a model with deterministic $\eta^{\mathbf{P}}$ as in Example 2.2. We use the following parameters for the deterministic model: $c^{\text{up}} = 0.5c_1^{\text{up}} + 0.5c_2^{\text{up}}, c^{\text{down}} = 0.5c_1^{\text{down}} + 0.5c_2^{\text{down}}$, that is the trader ignores regime switching but works with the stationary distribution of the Markov chain throughout. To compute the resulting liquidation rate $\nu_{t_k}^{*,\text{det}}$, we consider the HJB equation for the value function $V^{\prime,\text{det}}$ for Example 2.2, which is given by

$$\frac{\partial V^{\prime,\text{det}}}{\partial t}(-\rho V^{\prime,\text{det}} + \sup_{\nu \in [0, \nu^{\text{max}}]} \left\{ \nu(1 - c_f \nu^\varsigma) - \nu \frac{\partial V^{\prime,\text{det}}}{\partial w} - \bar{\eta}^{\mathbf{P}}(\nu) V^{\prime,\text{det}}(t, w) \right\}) = 0, \quad (6.7)$$

where, for the given model specification, $\bar{\eta}^{\mathbf{P}}(\nu) = \theta c^{\text{down}} a$. Then $\nu_{t_k}^{*,\text{det}}$ is the maximizer in (6.7) (with partial derivatives replaced by finite differences) and depends only on time and inventory level.

Numerical results of the performance comparisons show that, for a time horizon of two days, the expected gain from the use of filtering is € 113.38. For a longer liquidation horizon of four days the effect becomes even stronger with € 167.35. This shows that the additional complexity of using a filtering model may be worthwhile.

Remark 6.1 (Comparison to Almgren and Chriss [2]). It is interesting that the optimal liquidation rate ν^{*,det,t_k} is identical to the optimal rate in a geometric version of the well-known model of Almgren and Chriss [2], referred to as *geometric AC-model* in the sequel. In particular, the performance comparison applies also to the case where the investor uses this classical model. In the geometric AC-model it is assumed that the bid price has dynamics

$$dS_t^\nu = \bar{\eta}^{\mathbf{P}}(\nu_t) S_t^\nu dt + \sigma S_t^\nu dB_t, \quad (6.8)$$

for a Brownian motion B . By standard arguments the HJB equation for the value function V^{AC} of the optimal liquidation problem in the geometric AC-model is

$$\frac{\partial V^{\text{AC}}}{\partial t} - \rho V^{\text{AC}} + \sup_{\nu \in [0, \nu^{\text{max}}]} \left\{ s\nu(1 - c_f \nu^\varsigma) - \nu \frac{\partial V^{\text{AC}}}{\partial w} - \bar{\eta}^{\mathbf{P}}(\nu) s \frac{\partial V^{\text{AC}}}{\partial s} + \frac{1}{2} \sigma^2 s^2 \frac{\partial^2 V^{\text{AC}}}{\partial s^2} \right\} = 0.$$

Moreover, since V^{AC} is homogeneous in s , $V^{\text{AC}}(t, s, w) = sV'^{\text{AC}}(t, w)$. It follows that $\frac{\partial^2 V^{\text{AC}}}{\partial s^2} = 0$, and the HJB equation for V'^{AC} reduces to (6.7). Hence the optimal liquidation rate in the geometric AC model and in the jump-model with deterministic compensator coincide. The equivalence between the jump model and the geometric AC-model (6.8) holds only for the case where the compensator is deterministic: a model of the form (6.8) with drift driven by an unobservable Markov chain would lead to a diffusion equation for the filter and hence to a control problem for diffusion processes.

6.3 Model calibration. Finally we report the results of a small calibration study. We used a robust version of the EM algorithm to estimate the parameters of the bid price dynamics for the model specification from Example 2.3; see Damian et al. [27] for details on the methodology. First, in order to test the performance of the algorithm we ran a study with simulated data for two different parameter sets. In set 1 we use the parameters from Table 2; in set 2 we work with $c_1^{\text{up}} = c_2^{\text{up}} = c_1^{\text{down}} = c_2^{\text{down}} = 1000$, that is we consider a situation without Markov switching in the true data-generating process. However, the EM algorithm allows for different parameters in the two states. Hence this is a test, if the EM methodology points out spurious regime changes and trading opportunities which are not really in the data. The outcome of this exercise is presented in Figure 3, where we plot the hidden trajectory of Y together with the filter estimate \hat{Y} generated from the simulated data using the estimated model parameters. We see that in the left plot the filter nicely picks up the regime change, in the right plot the estimate \hat{Y}_t is close to 1.5 throughout, that is the estimated model correctly indicates that there is no Markov switching in the data. Finally we applied the algorithm to bid price data from the share price of Google, sampled at a frequency of one second. The EM estimates are $\hat{c}_1^{\text{up}} = 2128$, $\hat{c}_2^{\text{up}} = 1751$, $\hat{c}_1^{\text{down}} = 1769$, $\hat{c}_2^{\text{down}} = 1888$, which shows the same qualitative behaviour as the values used in our simulation study. A trajectory of the ensuing filter is given in Fig 4.

One would need an extensive empirical study to confirm and refine these results, but this is beyond the scope of the present paper.




Figure 3: A trajectory of the Markov chain Y (dashed) and of the corresponding filter \hat{Y} (straight line) computed using the parameter estimates from the EM algorithm as input. Left plot: results for parameter set 1 (with Markov switching); right plot: results for parameter set 2 (no Markov switching) In the graphs state e_1 (e_2) is represented by the value 1 (the value 2), and $\hat{Y}_t = \pi_t 1 + (1 - \pi_t) 2$. The estimated parameters for parameter set 1 are as follows: $\hat{c}_1^{\text{up}} = 993$; $\hat{c}_2^{\text{up}} = 875$; $\hat{c}_1^{\text{down}} = 842$; $\hat{c}_2^{\text{down}} = 960$. For parameter set 2 we obtained $\hat{c}_1^{\text{up}} = 940$; $\hat{c}_2^{\text{up}} = 941$; $\hat{c}_1^{\text{down}} = 9445$; $\hat{c}_2^{\text{down}} = 957$.




Figure 4: Trajectory of \widehat{Y} computed from the Google share price on 2012-06-21, sampled at a frequency of one second. (Data are from the LOBSTER database, see <https://lobsterdata.com>)

Acknowledgements

A part of this article was written while K. Colaneri was affiliated with the Department of Economics, University of Perugia. Support by the Vienna Science and Technology Fund (WWTF) through project MA14-031 is gratefully acknowledged. M. Szölgényi is supported by the AXA Research Fund grant “Numerical Methods for Stochastic Differential Equations with Irregular Coefficients with Applications in Risk Theory and Mathematical Finance”. A part of this article was written while M. Szölgényi was affiliated with the Institute of Statistics and Mathematics, Vienna University of Economics and Business. The authors are grateful for the excellent research assistance by Camilla Damian.

A Setup and filtering: proofs and additional results

Lemma A.1. *Suppose that Assumption 2.1 holds. Fix $m > w_0/T$ and consider some \mathbb{F}^S -adapted strategy ν with values in $[0, m]$. Define*

$$C := 0 \vee \sup \left\{ \int_{\mathbb{R}} (z^2 + 2z) \eta^{\mathbf{P}}(t, e, \nu, dz) : (t, e, \nu) \in [0, T] \times \mathcal{E} \times [0, m] \right\}.$$

Then $C < \infty$, $\mathbb{E}((S_t^\nu)^2) \leq S_0^2 e^{Ct}$, and $(\int_0^t S_{s-}^\nu dM_s^R)_{0 \leq t \leq T}$ is a true martingale.

Proof. To ease the notation we write S_t for S_t^ν . We begin with the bound on S_t^2 . First note that C is finite by Assumption 2.1. At a jump time T_n of R it holds that $S_{T_n} = S_{T_n-}(1 + \Delta R_{T_n})$ and therefore

$$S_{T_n}^2 - S_{T_n-}^2 = S_{T_n-}^2 \Delta R_{T_n}^2 + 2S_{T_n-}^2 \Delta R_{T_n}.$$

Hence $S_t^2 = S_0^2 + \int_0^t \int_{\mathbb{R}} S_{s-}^2 (z^2 + 2z) \mu^R(dz, ds)$ and we get

$$\begin{aligned} \mathbb{E}(S_t^2) &= S_0^2 + \mathbb{E} \left(\int_0^t \int_{\mathbb{R}} S_s^2 (z^2 + 2z) \eta^{\mathbf{P}}(s, Y_{s-}, \nu_{s-}, dz) ds \right) \\ &\leq S_0^2 + C \int_0^t \mathbb{E}(S_s^2) ds, \end{aligned}$$

so that $\mathbb{E}((S_t^\nu)^2) \leq S_0^2 e^{Ct}$ by the Gronwall inequality. To show that $\int_0^\cdot S_{s-} dM_s^R$ is a true martingale we show that this process has integrable quadratic variation. Since $[\int_0^\cdot S_{s-} dM_s^R]_t =$

$\int_0^t \int_{\mathbb{R}} S_s^2 z^2 \mu^R(dz, ds)$, we have

$$\mathbb{E} \left(\left[\int_0^\cdot S_{s-} dM_s^R \right]_t \right) = \mathbb{E} \left(\int_0^t S_s^2 \int_{\mathbb{R}} z^2 \eta^{\mathbf{P}}(s, Y_{s-}, \nu_{s-} dz) ds \right) \leq S_0^2 \tilde{C} \int_0^t e^{C_s} ds,$$

for every $t \in [0, T]$, where $\tilde{C} = \sup \{ \int_{\mathbb{R}} z^2 \eta^{\mathbf{P}}(t, e, \nu, dz) : t \in [0, T], e \in \mathcal{E}, \nu \in [0, m] \}$ is finite by Assumption 2.1. \square

Lemma 3.1. Conditions (2.4) and (2.5) imply that \tilde{Z} is a true martingale, see Protter and Shimbo [42]. Moreover, $\beta(t, Y_{t-}, \nu_{t-}, z) > -1$, since $(d\eta_t^{\mathbf{P}}(t, e_i, \nu; dz)/d\eta_t^{\mathbf{Q}}(dz))(z) > 0$ by assumption. This implies that $\tilde{Z}_T > 0$, and hence the equivalence of \mathbf{P} and \mathbf{Q} . The Girsanov theorem for random measures (see [16, VIII, Theorem T10]) shows that under \mathbf{P} , $\mu^R(dt, dz)$ has the predictable compensator $(\beta(t, Y_{t-}, \nu_t, z) + 1)\eta_t^{\mathbf{Q}}(dz)dt$. By definition of β this is equal to $\eta^P(t, Y_{t-}, \nu_t, dz)dt$. Moreover, \tilde{Z} and Y are orthogonal, since R and Y have no common jumps, so that the law of Y is the same under \mathbf{P} and under \mathbf{Q} . \square

Theorem 3.2. Our derivation parallels the proof of [7, Theorem 3.24], which deals with the classical case where the observation process is a Brownian motion with drift. Recall that for a function $f: \mathcal{E} \rightarrow \mathbb{R}$ the semimartingale decomposition of $f(Y_t)$ is given by $f(Y_t) = f(Y_0) + \int_0^t \langle Q\mathbf{f}, Y_s \rangle ds + M_t^f$, where M^f is a true (\mathbb{F}, \mathbf{Q}) -martingale. Define the process $\tilde{Z}^\epsilon = (\tilde{Z}_t^\epsilon)_{t \in [0, T]}$ by

$$\tilde{Z}_t^\epsilon := \frac{\tilde{Z}_t}{1 + \epsilon \tilde{Z}_t},$$

and note that $\tilde{Z}_t^\epsilon < 1/\epsilon$ for every $t \in [0, T]$. Now we compute $\tilde{Z}^\epsilon f(Y)$. Notice that $[\tilde{Z}^\epsilon, Y]_t = 0$ for every $t \in [0, T]$, as R and Y have no common jumps. Hence, from Itô's product rule we get

$$\begin{aligned} d(\tilde{Z}_t^\epsilon f(Y_t)) &= \tilde{Z}_{t-}^\epsilon \langle Q\mathbf{f}, Y_t \rangle dt + \tilde{Z}_{t-}^\epsilon dM_t^f - f(Y_{t-}) \tilde{Z}_{t-}^\epsilon \int_{\mathbb{R}} \frac{\beta(t, Y_{t-}, \nu_{t-}, z)}{1 + \epsilon \tilde{Z}_{t-}^\epsilon} \eta_t^{\mathbf{Q}}(dz) dt \\ &\quad + f(Y_{t-}) \tilde{Z}_{t-}^\epsilon \int_{\mathbb{R}} \frac{\beta(t, Y_{t-}, \nu, z)}{1 + \epsilon \tilde{Z}_{t-}^\epsilon (1 + \beta(t, Y_{t-}, \nu_{t-}, z))} \mu^R(dt, dz). \end{aligned} \quad (\text{A.1})$$

Next we show that $\mathbb{E}^{\mathbf{Q}} \left(\int_0^t \tilde{Z}_{s-}^\epsilon dM_s^f \mid \mathcal{F}_t^S \right) = 0$. By the definition of conditional expectation, this is equivalent to $\mathbb{E}^{\mathbf{Q}} \left(H \int_0^t \tilde{Z}_{s-}^\epsilon dM_s^f \right) = 0$ for every bounded, \mathcal{F}_t^S -measurable random variable H . Define an $(\mathbb{F}^S, \mathbf{Q})$ -martingale by $H_u = \mathbb{E}^{\mathbf{Q}}(H \mid \mathcal{F}_u^S)$, $0 \leq u \leq t \leq T$, and note that $H = H_t$. By the martingale representation theorem for random measures, see, e.g., [39, Ch. III, Theorem 4.37] or [16, Ch. VIII, Theorem T8], we get that there is a bounded \mathbb{F}^S -predictable random function ϕ such that

$$H_t = H_0 + \int_0^t \int_{\mathbb{R}} \phi(s, z) (\mu^R(ds, dz) - \eta^{\mathbf{Q}}(dz) ds), \quad t \in [0, T].$$

Now, applying the Itô product rule and using that $[M^f, H]_t = [Y, R]_t = 0$ for every $t \in [0, T]$, we obtain

$$H_t \int_0^t \tilde{Z}_{s-}^\epsilon dM_s^f = \int_0^t H_{s-} \tilde{Z}_{s-}^\epsilon dM_s^f + \int_0^t \int_{\mathbb{R}} \left(\int_0^s \tilde{Z}_{u-}^\epsilon dM_u^f \right) \phi(s, z) (\mu^R(ds, dz) - \eta^{\mathbf{Q}}(dz) ds).$$

Both integrals on the right hand side of the above representation are martingales. This follows from the finite-state property of the Markov chain Y and the boundedness of \tilde{Z}^ϵ and H . Hence, taking the expectation we get that $\mathbb{E}^{\mathbf{Q}}\left(H \int_0^t \tilde{Z}_{s-}^\epsilon dM_s^f\right) = 0$ as claimed.

Now note that for $t \in [0, T]$ and a generic integrable \mathcal{F}_t -measurable random variable U it holds that

$$\mathbb{E}^{\mathbf{Q}}(U \mid \mathcal{F}_t^S) = \mathbb{E}^{\mathbf{Q}}(U \mid \mathcal{F}_T^S); \quad (\text{A.2})$$

this can be shown with similar arguments as in [7, Proposition 3.15]. Taking the conditional expectation from (A.1) and applying (A.2) and the Fubini theorem we get for every $t \in [0, T]$,

$$\begin{aligned} \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_t^\epsilon f(Y_t) \mid \mathcal{F}_t^S\right) &= \frac{\pi_0(f)}{1+\epsilon} + \int_0^t \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_{s-}^\epsilon \langle Q\mathbf{f}, Y_s \rangle \mid \mathcal{F}_T^S\right) ds \\ &+ \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-}^\epsilon \frac{\beta(s, Y_{s-}, \nu_{s-}, z)}{1 + \epsilon \tilde{Z}_{s-} (1 + \beta(s, Y_{s-}, \nu, z))} \mid \mathcal{F}_T^S\right) \mu^R(ds, dz) \\ &- \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-}^\epsilon \frac{\beta(s, Y_{s-}, \nu_{s-}, z)}{1 + \epsilon \tilde{Z}_{s-}} \mid \mathcal{F}_T^S\right) \eta_s^{\mathbf{Q}}(dz) ds. \end{aligned} \quad (\text{A.3})$$

Note that, for every $t \in [0, T]$, $\tilde{Z}_t^\epsilon < \tilde{Z}_t$ and that \tilde{Z}_t is integrable. Since β is bounded by assumption, by dominated convergence we get the following three limits

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_t^\epsilon f(Y_t) \mid \mathcal{F}_t^S\right) &= \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_t f(Y_t) \mid \mathcal{F}_t^S\right), \\ \lim_{\epsilon \rightarrow 0} \int_0^t \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_{s-}^\epsilon \langle Q\mathbf{f}, Y_s \rangle \mid \mathcal{F}_T^S\right) ds &= \int_0^t \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_{s-} \langle Q\mathbf{f}, Y_s \rangle \mid \mathcal{F}_T^S\right) ds, \\ \lim_{\epsilon \rightarrow 0} \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-}^\epsilon \frac{\beta(s, Y_{s-}, \nu_{s-}, z)}{1 + \epsilon \tilde{Z}_{s-}} \mid \mathcal{F}_T^S\right) \eta_s^{\mathbf{Q}}(dz) ds \\ &= \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-} \beta(s, Y_{s-}, \nu_{s-}, z) \mid \mathcal{F}_T^S\right) \eta_s^{\mathbf{Q}}(dz) ds. \end{aligned}$$

Finally we consider the integral with respect to $\mu^R(ds, dz)$ in (A.3). Let $\{T_n, Z_n\}$ be the sequence of jump times and the corresponding jump sizes of the process R . Denote by $n(t)$ the number of jumps up to time t , so that $T_{n(t)}$ is the last jump time before t . Then

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-}^\epsilon \frac{\beta(s, Y_{s-}, \nu, z)}{1 + \epsilon \tilde{Z}_{s-} (1 + \beta(s, Y_{s-}, \nu, z))} \mid \mathcal{F}_T^S\right) \mu^R(ds, dz) \\ &= \lim_{\epsilon \rightarrow 0} \sum_{n=1}^{n(t)} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{T_n^-}) \tilde{Z}_{T_n^-}^\epsilon \frac{\beta(T_n, Y_{T_n^-}, \nu, \Delta R_{T_n})}{1 + \epsilon \tilde{Z}_{T_n^-} (1 + \beta(T_n, Y_{T_n^-}, \nu_{T_n^-}, \Delta R_{T_n}))} \mid \mathcal{F}_T^S\right) \\ &= \sum_{n=1}^{n(t)} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{T_n^-}) \tilde{Z}_{T_n^-} \beta(T_n, Y_{T_n^-}, \nu_{T_n^-}, \Delta R_{T_n}) \mid \mathcal{F}_T^S\right) \\ &= \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-} \beta(s, Y_{s-}, \nu_{s-}, z) \mid \mathcal{F}_T^S\right) \mu^R(ds, dz). \end{aligned}$$

Assembling the previous results we obtain

$$\begin{aligned} \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_t f(Y_t) \mid \mathcal{F}_T^S\right) &= \pi_0(f) + \int_0^t \mathbb{E}^{\mathbf{Q}}\left(\tilde{Z}_{s-} \langle Q\mathbf{f}, Y_s \rangle \mid \mathcal{F}_T^S\right) ds \\ &+ \int_0^t \int_{\mathbb{R}} \mathbb{E}^{\mathbf{Q}}\left(f(Y_{s-}) \tilde{Z}_{s-} \beta(s, Y_{s-}, \nu_{s-}, z) \mid \mathcal{F}_T^S\right) (\mu^R(ds, dz) - \eta_s^{\mathbf{Q}}(dz)), \end{aligned}$$

and hence the claim of the theorem follows from (A.2). \square

B Optimization via MDMs: proofs and additional results

Lemma 4.1. To establish the claim we show that the first derivatives of the vector field g are bounded, uniformly in ν . The components of $\frac{\partial g}{\partial w}$ and $\frac{\partial g}{\partial s}$ are all 0, and, using Assumption 2.1, the nonzero components of $\frac{\partial g}{\partial \pi^i}$, $i = 1, \dots, K$, can be estimated as follows. For $i \neq k$,

$$\begin{aligned} \left| \frac{\partial g^{k+3}}{\partial \pi^i} \right| &= \left| q^{ik} - \pi^k \int_{\mathbb{R}} u^k(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_i, \nu, dz) - \pi^k \sum_{j=1}^K \pi^j \int_{\mathbb{R}} \frac{\partial u^k(t, \nu, \pi, z)}{\partial \pi^i} \eta^{\mathbf{P}}(t, e_j, \nu, dz) \right| \\ &< \max_{i,k} q^{ik} + \pi^k \int_{\mathbb{R}} u^k(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_i, \nu, dz) \\ &\quad + \pi^k \sum_{j=1}^K \pi^j \int_{\mathbb{R}} \frac{d\eta^{\mathbf{P}}(t, e_i, \nu) / d\eta_t^{\mathbf{Q}}(z) d\eta^{\mathbf{P}}(t, e_k, \nu) / d\eta_t^{\mathbf{Q}}(z)}{\left(\sum_{l=1}^K \pi^l d\eta^{\mathbf{P}}(t, e_l, \nu) / d\eta_t^{\mathbf{Q}}(z) \right)^2} \eta^{\mathbf{P}}(t, e_j, \nu, dz), \end{aligned}$$

and this is smaller than $\max_{i,k} q^{ik} + (M^4 + M^2)\lambda^{\max}$. For $i = k$ we get

$$\begin{aligned} \left| \frac{\partial g^{i+3}}{\partial \pi^i} \right| &= \left| q^{ii} - 2\pi^i \int_{\mathbb{R}} u^i(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_i, \nu, dz) - \sum_{j \neq i} \pi^j \int_{\mathbb{R}} u^i(t, \nu, \pi, z) \eta^{\mathbf{P}}(t, e_j, \nu, dz) \right. \\ &\quad \left. - \pi^i \sum_{j=1}^K \pi^j \int_{\mathbb{R}} \frac{\partial u^i(t, \nu, \pi, z)}{\partial \pi^i} \eta^{\mathbf{P}}(t, e_j, \nu, dz) \right| < \max_i q^{ii} (M^4 + 3M^2) \lambda^{\max}. \end{aligned}$$

□

Lemma 4.6. First we estimate the reward function introduced in (4.3). Since $f \geq 0$, $e^{-\rho t} \leq 1$, and $h(w) \leq w$, we get that $r(\tilde{x}, \alpha) \leq s \int_0^{\tau^\varphi} e^{-\Lambda_u^\alpha} \alpha_u du + s e^{-\Lambda_{\tau^\varphi}^\alpha} w_{\tau^\varphi}^\alpha$. Partial integration gives

$$\int_0^{\tau^\varphi} e^{-\Lambda_u^\alpha} \alpha_u du = \left[-w_u^\alpha e^{-\Lambda_u^\alpha} \right]_0^{\tau^\varphi} - \int_0^{\tau^\varphi} \lambda_u^\alpha e^{-\Lambda_u^\alpha} w_u^\alpha du \leq w - e^{-\Lambda_{\tau^\varphi}^\alpha} w_{\tau^\varphi}^\alpha,$$

and hence $r(\tilde{x}, \alpha) \leq sw$. Next we estimate $Q_L b(\tilde{x}, \alpha)$. Recall the definition of $\bar{\eta}^{\mathbf{P}}$ from (2.3) and let $c_\eta := \sup \{ \bar{\eta}^{\mathbf{P}}(t, e, 0) : (t, e) \in [0, T] \times \mathcal{E} \}$. It holds that

$$Q_L b(\tilde{x}, \alpha) = \int_0^{\tau^\varphi} e^{\gamma(T-(u+t))} e^{-\Lambda_u^\alpha} \sum_{j=1}^K \pi_j s w_u^\alpha (1 + \bar{\eta}^{\mathbf{P}}(t+u, e_j, \alpha_u)) du \leq s w e^{\gamma(T-t)} c_\eta \int_0^{\tau^\varphi} e^{-\gamma u} du,$$

where we have used that $w_u^\alpha \leq w$ and $e^{-\Lambda_u^\alpha} < 1$. The last term is bounded by $b(\tilde{x}) \frac{c_\eta}{\gamma}$, and the MDM is contracting for $\gamma > c_\eta$. □

The following lemma is needed in the proof of Proposition 4.8.

Lemma B.1. *Consider a function $v \in \mathcal{C}_b$. Then the mapping $(\tilde{x}, \nu) \mapsto \bar{Q}v(\tilde{x}, \nu)$ is continuous on $\tilde{\mathcal{X}} \times [0, \nu^{\max}]$.*

Proof. It suffices to show that for $j = 1, \dots, K$ the mapping

$$(t, w, s, \pi, \nu) \mapsto \int_{\mathbb{R}} v(t, s(1+z), \pi^1(1+u^1(t, \nu, \pi, z)), \dots, \pi^K(1+u^K(t, \nu, \pi, z))) \eta^j(t, \nu, dz)$$

is continuous on $\tilde{\mathcal{X}} \times [0, \nu^{\max}]$, where $\eta^j(t, \nu, dz) := \eta(t, e_j, \nu, dz)$. Consider a sequence with elements $(t_n, \nu_n, \pi_n) \xrightarrow{n \rightarrow \infty} (t, \nu, \pi)$. Note that, for sufficiently large n , the set $\{s^n(1+z) : z \in \text{supp}(\eta)\}$ is contained in a compact subset $[\underline{s}, \bar{s}] \subset (0, \infty)$. Moreover, v is uniformly continuous on the compact set $[0, T] \times [0, w_0] \times [\underline{s}, \bar{s}] \times \mathcal{S}^K \times [0, \nu^{\max}]$. Then, Assumption 4.7-(2) implies that the sequence $\{v^n\}$ with

$$v^n(z) := v(t_n, s_n(1+z), \pi_n^1(1+u^1(t_n, \nu_n, \pi_n, z)), \dots, \pi_n^K(1+u^K(t_n, \nu_n, \pi_n, z)))$$

converges uniformly in $z \in \text{supp}(\eta)$ to $v(z) := v(t, s, \pi, \nu, z)$. Hence the following estimate holds:

$$\begin{aligned} & \left| \int_{\text{supp}(\eta)} v^n(z) \eta^j(t_n, \nu_n, dz) - \int_{\text{supp}(\eta)} v(z) \eta^j(t, \nu, dz) \right| \\ & \leq \int_{\text{supp}(\eta)} |v^n(z) - v(z)| \eta^j(t_n, \nu_n, dz) + \left| \int_{\text{supp}(\eta)} v(z) \eta^j(t_n, \nu_n, dz) - \int_{\text{supp}(\eta)} v(z) \eta^j(t, \nu, dz) \right|. \quad (\text{B.1}) \end{aligned}$$

Finally, the first term in (B.1) can be estimated by $\lambda^{\max} \sup\{|v^n(z) - v(z)| : z \in \text{supp}(\eta)\}$, which converges to zero as v^n converges to v uniformly; the second term in (B.1) converges to zero by Assumption 4.7-(1) (continuity of the mapping $(t, \nu) \mapsto \eta^j(t, \nu, dz)$ in the weak topology). \square

References

- [1] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.
- [2] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3: 5–40, 2001.
- [3] R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. *Risk*, 18(5752):10, 2005.
- [4] A. Almudevar. A dynamic programming algorithm for the optimal control of piecewise deterministic Markov processes. *SIAM Journal on Control and Optimization*, 40(2):525–539, 2001.
- [5] T. G. Andersen. Return volatility and trading volume: An information flow interpretation of stochastic volatility. *The Journal of Finance*, 51(1):169–204, 1996.
- [6] S. Ankirchner, C. Blanchet-Scalliet, and A. Eyraud-Loisel. Optimal portfolio liquidation with additional information. *Mathematics and Financial Economics*, pages 1–14, 2015.
- [7] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*, volume 3. Springer, 2009.
- [8] G. Barles. *Solutions de Viscosité des Equations de Hamilton-Jacobi*. Springer Verlag, 1994.
- [9] G. Barles and P. E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4:271–283, 1991.
- [10] N. Bäuerle and U. Rieder. MDP algorithms for portfolio optimization problems in pure jump markets. *Finance and Stochastics*, 13(4):591–611, 2009.
- [11] N. Bäuerle and U. Rieder. Optimal control of piecewise deterministic Markov processes with finite time horizon. *Modern trends in controlled stochastic processes: theory and applications*, 123:143, 2010.

- [12] N. Bäuerle and U. Rieder. *Markov Decision Processes with Applications to Finance*. Springer Science & Business Media, 2011.
- [13] E. Bayraktar and M. Ludkovski. Optimal trade execution in illiquid markets. *Mathematical Finance*, 21(4):681–701, 2011.
- [14] E. Bayraktar and M. Ludkovski. Liquidation in limit order books with controlled intensity. *Mathematical Finance*, 24(4):627–650, 2014.
- [15] D. Bertsimas and A. W. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.
- [16] P. Brémaud. *Point Processes and Queues: Martingale Dynamics*. Springer Series in Statistics. Springer-Verlag, New York Heidelberg Berlin, 1981.
- [17] A. Cartea and S. Jaimungal. Modelling asset prices for algorithmic and high frequency trading. *Applied Mathematical Finance*, 20(6):512–547, 2013.
- [18] Á. Cartea and S. Jaimungal. Optimal execution with limit and market orders. *Quantitative Finance*, 15(8):1279–1291, 2015.
- [19] Á. Cartea, S. Jaimungal, and J. Penalva. *Algorithmic and high-frequency trading*. Cambridge University Press, 2015.
- [20] T. Cayé and J. Muhle-Karbe. Liquidation with self-exciting price impact. *Mathematics and Financial Economics*, 10(1):15–28, 2016.
- [21] C. Ceci and K. Colaneri. Nonlinear filtering for jump diffusion observations. *Advances in Applied Probability*, 44(3):678–701, 2012.
- [22] C. Ceci and K. Colaneri. The Zakai equation of nonlinear filtering for jump-diffusion observation: existence and uniqueness. *Applied Mathematics and Optimization*, 69(1):47–82, 2014.
- [23] C. Ceci and A. Gerardi. A model for high frequency data under partial information: a filtering approach. *International Journal of Theoretical and Applied Finance*, 9(4):1–22, 2006.
- [24] R. Cont. Statistical modeling of high-frequency financial data. *IEEE Signal Processing Magazine*, 28(5):16–25, 2011.
- [25] O. L. Costa and F. Dufour. *Continuous Average Control of Piecewise Deterministic Markov Processes*. Springer, 2013.
- [26] J. Cvitanic, B. Rozovskii, and I. Zaliapin. Numerical estimation of volatility values from discretely observed diffusion data. *Journal of Computational Finance*, 9(4):1, 2006.
- [27] C. Damian, Z. Eksi, and R. Frey. EM algorithm for Markov chains observed via Gaussian noise and point process information: Theory and numerical experiments. *Statistics and Risk Modelling*, 35:51–72, 2018.
- [28] D. M. Dang and P. A. Forsyth. Continuous Time Mean-Variance Optimal Portfolio Allocation Under Jump Diffusion: A Numerical Impulse Control Approach. *Numerical Methods for Partial Differential Equations*, 30:664–698, 2014.

- [29] M. H. A. Davis. *Markov Models & Optimization*, volume 49. CRC Press, 1993.
- [30] Mark H.A. Davis and M. Farid. Piecewise-deterministic processes and viscosity solutions. In *Stochastic analysis, control, optimization and applications*, pages 249–268. Springer, 1999.
- [31] M.A.H. Dempster and J.J. Ye. Necessary and sufficient optimality conditions for control of piecewise deterministic Markov processes. *Stochastics: An International Journal of Probability and Stochastic Processes*, 40(3-4):125–145, 1992.
- [32] R. J. Elliott and W. P. Malcolm. Discrete-time expectation maximization algorithms for markov-modulated poisson processes. *IEEE Transactions on Automatic Control*, 53(1):247–256, 2008.
- [33] L. Forwick, M. Schäl, and M. Schmitz. Piecewise deterministic Markov control processes with feedback controls and unbounded costs. *Acta Applicandae Mathematica*, 82(3):239–267, 2004.
- [34] R. Frey and T. Schmidt. Pricing and hedging of credit derivatives via the innovation approach to nonlinear filtering. *Finance and Stochastics*, 16(1):105–133, 2012.
- [35] J. Gatheral and A. Schied. Dynamical models of market impact and algorithms for order execution. In J.P. Jean-Pierre Fouque and J.A. Langsam, editors, *Handbook on Systemic Risk*, pages 579–599. 2013.
- [36] S. Gökay, A. F. Roch, and H. M. Soner. *Liquidity Models in Continuous and Discrete Time*. Springer, 2011.
- [37] X. Guo and M. Zervos. Optimal execution with multiplicative price impact. *SIAM Journal on Financial Mathematics*, 6(1):281–306, 2015.
- [38] H. He and H. Mamaysky. Dynamic trading policies with price impact. *Journal of Economic Dynamics and Control*, 29(5):891–930, 2005.
- [39] J. Jacod and A.N. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer, 2nd edition, 2003.
- [40] M. Kirch and W. Runggaldier. Efficient hedging when asset prices follow a geometric Poisson process with unknown intensities. *SIAM Journal on Control and Optimization*, 43:1174–1195, 2004.
- [41] A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1):1–32, 2013.
- [42] P. Protter and K. Shimbo. No arbitrage and general semimartingales. *Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz*, 4:267–283, 2008.
- [43] A. Schied. Robust strategies for optimal order execution in the Almgren–Chriss framework. *Applied Mathematical Finance*, 20(3):264–286, 2013.
- [44] A. Schied and T. Schöneborn. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, 13(2):181–204, 2009.
- [45] Y. Zeng. A partially observed model for micromovement of asset prices with Bayes estimation via filtering. *Mathematical Finance*, 13:411–444, 2003.