

Optimal market making under partial information with general intensities

Luciano Campi and Diego Zabaljauregui*

Department of Statistics
London School of Economics and Political Science

Abstract

Starting from the Avellaneda–Stoikov framework [AS08], we consider a market maker who wants to optimally set bid/ask quotes over a finite time horizon, to maximize her expected utility. The intensities of the orders she receives depend not only on the spreads she quotes, but also on unobservable factors modelled by a hidden Markov chain. We tackle this stochastic control problem under partial information with a model that unifies and generalizes many existing ones under full information, combining several risk metrics and constraints, and using general decreasing intensity functionals. We use stochastic filtering, control and piecewise-deterministic Markov processes theory, to reduce the dimensionality of the problem and characterize the reduced value function as the unique continuous viscosity solution of its dynamic programming equation. We then solve the analogous full information problem and compare the results numerically through a concrete example. We show that the optimal full information spreads are biased when the exact market regime is unknown, and the market maker needs to adjust for additional regime uncertainty in terms of P&L sensitivity and observed order flow volatility. This effect becomes higher, the longer the waiting time in between orders.

Keywords: Market making, High-frequency trading, Algorithmic trading, Stochastic optimal control, Hidden Markov model, Stochastic filtering, Viscosity solutions, Piecewise-deterministic Markov processes.

Introduction

Given a financial market, a market maker (MM) can be understood as someone who provides liquidity for a certain asset. That is, she (almost) continuously posts bid/ask quotes for the asset, in the hope to profit from the bid-ask spread. In choosing how to do so, the MM faces a complicated problem on several levels; namely: the instantaneous margin/volume trade-off (the further away she quotes from the “fair price”, the less she gets executed, and vice-versa), adverse price movements and inventory risk (exposure), execution costs, and many others.

An increasingly popular mathematical approach to the MM problem, both in academia and in practice, is by means of stochastic optimal control. In particular, a line of research has focused on the modelling framework proposed by Avellaneda and Stoikov [AS08] (rooted in turn in [HS81]). Although widely motivated in the literature by order-driven markets such as equity markets, the shape of the limit order book is not explicitly taken into account. Thus, the framework is more easily understood and applied to over-the-counter (OTC) quote-driven

*Corresponding Author. E-mail: d.zabaljauregui@lse.ac.uk

The authors would like to thank Álvaro Cartea and Katia Colaneri for helpful discussions on the subject.

markets such as the foreign exchange (FX) market, and we therefore choose to present it in this setting.

In this framework, the MM gives firm bid-ask quotes during a finite time interval by choosing bid/ask spreads with respect to a certain *reference price*.¹ Depending on the market, this price could be for example an aggregated mid-price or a dealer-to-dealer price, and it is frequently assumed to behave as an arithmetic Brownian motion. To explicitly model the margin/volume trade-off, the probability that the MM gets executed decays as a function of the corresponding spread. More precisely, it is assumed that the MM receives market orders according to counting processes of stochastic intensity. Most typically in the mathematical literature, intensities are assumed to decay exponentially on the spreads (mainly for tractability reasons). The goal of the MM is to find a strategy that allows her to maximize her expected terminal utility, which is taken as a constant absolute risk aversion (CARA) utility. The problem is then translated into a deterministic one: solving the associated Hamilton–Jacobi–Bellman (HJB) equation, which is a partial-integro differential equation (PIDE) for the MM’s value function, and ultimately retrieving the optimal strategy in feedback form.

Within the described framework, a lot variants have been put forward and extensively studied. In [BL14, GLFT12], Bayraktar and Ludkovski, and Guéant, Lehalle and Fernandez-Tapia, apply a one-trading-side version to optimal liquidation in the risk-neutral and risk averse contexts, respectively. The introduction of a constraint on the inventory in [GLFT13], allowed the authors to rigorously solve the original problem of [AS08] with exponential intensities, by means of a verification theorem. This constraint has been widely used moving forward, with the exception of [FL12, FL13], where strategies are derived without a verification theorem.

Guéant and Lehalle continued to actively contribute to the area. In [GL15], they revisit the risk averse optimal liquidation problem for general intensities satisfying a certain ordinary differential inequality, and in [Gué17] the same is done for market making.² This assumption had been firstly introduced in [BL14, Sect.5.3] under risk-neutrality.

Other prolific contributors in the field are Cartea, Jaimungal and their coauthors [CJ15, CJP15, CJR14], who introduced a quadratic running penalty on the inventory to manage the “accumulated” inventory risk for an otherwise risk-neutral MM. Constraints on the MM’s spreads have also been considered in some of them. It is worth noting that [Gué17] shows how the two seemingly different subclasses of models (risk averse and “risk-neutral” with running penalty) can ultimately be characterized by a unique system of ordinary differential equations (ODEs) when considering two appropriate ansatz.

A relevant issue in practice that has not been included in the previous models is the following one. Empirical evidence (see, e.g., [CJ13]) suggests that liquidity taken by clients depends not only on the quoted spreads but also on other unobservable factors. Indeed, the confluence of factors such as market sentiment towards the asset and the competition with other market makers also affects the intensities at which the MM receives orders. This complicated effect can be modelled in a simplified fashion by making the intensities depend as well on a hidden finite-state Markov chain, effectively reflecting the regime or state of the market (different levels from very slow to very active). To the best of our knowledge, this has only been briefly done in the Avellaneda–Stoikov framework by proposing an approximation for the optimal strategy with exponential intensities [CJ13, Sect.5.1] or studying a simple two-states version with power-law intensities [BL14, Sect.5.4]. Both of these papers deal only with the “risk-neutral” (possibly penalized) case and make the unrealistic assumption that the current market regime is known by the MM.³

When the state of the market is unknown, the problem becomes significantly more challeng-

¹Also referred to by some authors as *micro-price* [CJP15] or *efficient price* [DRR13] in the martingale case.

²The latter paper also considers the multi-asset case.

³[BL09, Sect.3.3] (arXiv version) and [CJ19] also study optimal trade execution under partial information, albeit with uncontrolled intensities.

ing for several reasons:

- It becomes a combined problem of stochastic control and filtering. The MM needs to dynamically make her best possible prediction of the market regime (or more precisely, its distribution), based on the information she has (i.e., the orders she has received so far and the evolution of the reference price), and adjust her spreads accordingly. This prediction is known as a filter.
- The associated HJB PIDE has higher dimension and more non-linearities.
- The standard approach used in all of the previously cited papers relies on reducing the HJB PIDE to a system of ODEs by means of an ansatz, proving that such a system has a classical (smooth) solution, and recovering the value function via a verification theorem. Under partial information however, the reduced HJB equation is still a complicated PIDE such that, in general, a classical solution may not exist (or it may be too difficult to prove otherwise). Hence, the ansatz argument breaks down.
- As a consequence, one needs to resort to the concept of viscosity solutions [FS06]. In addition, the numerical resolution unfaillingly becomes a lot more involved than for simple systems of ODEs.
- On a technical level, the construction of the model is not straightforward. The MM needs to adjust her strategy based on her observable information, such as the arriving orders, but this flow of information is in turn affected by the MM's actions.

In this paper, we solve the problem of the MM under partial information. First, we start by unifying in one single formulation all the modelling features described so far. This allows us to simultaneously tackle all the models at once with a single approach, while generalising them at the same time. Indeed, our formulation allows for the interaction of any CARA utility (whether risk-neutral or risk averse) with running inventory penalty, terminal execution cost, inventory constraints and spread constraints. Further, motivated by practitioners' needs, we strongly generalize the intensity shapes to any continuous, decreasing to zero functions, adding modelling flexibility.⁴ We even allow for the inventory to be unconstrained when no penalties are present. (This scenario is considered mainly for the sake of completeness and comparison, at almost no extra cost.)

Secondly, we let the intensities depend on a k -dimensional hidden Markov chain. Following [CEFS16a], we use a weak formulation⁵ to construct a well-defined model (with exogenous information) and solve the filtering problem by means of the reference probability approach of stochastic filtering [Bré81, Chpt.VI]. The rigorous setting of the model and its full characterization are carried out in Section 1, while the filtering problem is solved in Section 2.

The optimization problem of the MM is then reformulated in terms of the usual state variables together with the k -dimensional observable distribution of the Markov chain (Section 3). At this point, the problem is too involved both analytically and numerically. We proceed by showing that, formally, there is an ansatz for the value function that reduces the dimensionality of the problem. However, as the verification approach is not valid any more, we explicitly find the guessed decomposition in terms of the value function of a new, diffusion-free, problem (Theorem 3.1.1). We then characterize the reduced value function as the unique continuous viscosity solution of its formally derived equation (Theorem 3.1.4). The latter is done by harnessing results of piecewise-deterministic Markov process (PDMPs) [CEFS16a, DF99], as first defined in [Dav84].

⁴This is done at the expense of renouncing to uniqueness in the optimal strategy. We also assume the decay to be "fast enough" in certain cases.

⁵That is, with controlled probability measures.

Thirdly, we solve the idealized problem of a MM with full information, who can observe the Markov chain (Section 4). We show that a similar ansatz and the standard approach work out in this case. We prove that the MM's value function is a classical solution of its HJB equation via a general verification theorem (Theorem 4.4.1) and we recover the well known strategies for one regime as particular cases.

Finally, we compare the optimal strategies under full and partial information through a concrete example, by numerical analysis (Section 5). In particular, we show that the optimal full information spreads are biased when the exact regime is unknown, and using them becomes suboptimal. We interpret the adjustment needed in terms of observable order flow volatility and sensitivity of the expected profit to observable regime changes; and we show how this effect becomes higher, the longer the waiting time in between orders (leading to higher uncertainty for the MM).

1 Setting and main assumptions

1.1 Preliminaries on the probability spaces

We start by setting up our framework in an abstract fashion, deferring the question of existence of a model to Proposition 1.4.1, and a characterization of all such models to Proposition 1.4.2. As mentioned in the Introduction, we seek to construct a model under a weak formulation, so that the information flow remains exogenous (i.e., unaffected by the MM's actions).

Let $T > 0$ be a given finite horizon and $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T})$ a right-continuous filtered measurable space with $\mathcal{F}_T = \mathcal{F}$. Suppose it supports three adapted stochastic processes W, N^+, N^- . (More assumptions to be added in the sequel.) Let $\mathbb{F}^{W, N^+, N^-} = (\sigma(W_u, N_u^+, N_u^- : 0 \leq u \leq t) \vee \mathcal{F}_0)_{0 \leq t \leq T}$ be the natural filtration of these processes enlarged ("completed") by \mathcal{F}_0 and define the set of *admissible spreads* as

$$\mathcal{A} := \{\delta : [0, T] \times \Omega \rightarrow \overline{(\delta_*, \delta^*)} : \delta \text{ is } \mathbb{F}^{W, N^+, N^-}\text{-predictable and bounded}\}, \quad (1.1.1)$$

where $-\infty \leq \delta_* < \delta^* \leq +\infty$ are fixed constants and $\overline{(a, b)} = \mathbb{R} \cap [a, b]$ denotes the closure of the interval (a, b) in \mathbb{R} , for any $-\infty \leq a < b \leq +\infty$. Note that for $\delta_* = -\infty$ and $\delta^* = +\infty$, the admissible spreads are not uniformly bounded. The self-imposed constraints δ_*, δ^* for the MM's spreads can be taken to be different for the bid and the ask if wanted, without any additional effort.

Consider on the former space a family $(\mathbb{P}^\alpha)_{\alpha \in \mathcal{A}^2}$, $\alpha = (\delta^-, \delta^+)$, of equivalent probability measures such that the sigma algebra generated by their null sets is \mathcal{F}_0 . Note that for each α , $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}^\alpha)$ is under the usual conditions and \mathcal{F}_0 is the completed trivial sigma algebra. We refer to \mathbb{P}^α as the *physical* or *historical* probability given the *admissible strategy* (or *control*) α , and we write \mathbb{E}^α for the expectation under \mathbb{P}^α . We will drop the " α " from the notation when there is no room for ambiguity.

Henceforth, all the processes (resp. properties) considered are supposed to be defined (resp. hold) on the space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}^\alpha)$ for all $\alpha \in \mathcal{A}^2$, unless otherwise stated. For example, a "Brownian motion" is a process on $[0, T] \times \Omega$ that is an $(\mathbb{F}, \mathbb{P}^\alpha)$ -Brownian motion for all α . All subfiltrations are understood to have been augmented to satisfy the usual conditions (which in the case of the natural filtrations of Feller processes, it amounts simply to completing them). Càdlàg versions of the processes are used whenever available.

1.2 Description of the model

A market maker in a quote-driven market gives binding bid/ask quotes S^-, S^+ resp. for a certain financial asset, during the time interval $[0, T]$. She receives unitary size market orders⁶

⁶Any constant size would work in the same way, but this convention simplifies the notation. See, e.g., [Gué17] for some formulas with arbitrary constant size under full information.

to buy/sell according to counting processes N^-, N^+ starting at 0, with a.s. no common jumps and stochastic intensities λ^-, λ^+ resp. (see [Bré81, p.27, Def. D7]). Each intensity at time t depends on the corresponding spread $\delta_t^\pm := \pm(S_t^\pm - S_t)$ between the MM's quotes and a reference price S . S may be interpreted, e.g., as an aggregated mid-price or a dealer-to-dealer price, depending on the market. We assume S is a Brownian motion with drift, of the form

$$S_t = s_0 + \mu t + \sigma W_t, \quad (1.2.1)$$

where W is a Wiener process, $s_0, \mu \in \mathbb{R}$ and $\sigma \geq 0$. Note that the MM can fully specify her quote S_t^\pm by choosing the spread δ_t^\pm , since $S_t^\pm = S_t \pm \delta_t^\pm$. We assume $\delta^-, \delta^+ \in \mathcal{A}$. In particular, the MM chooses her spreads based on the observation of W , N^- and N^+ . We refer to $\mathbb{F}^{N^-, N^+, W}$ as the *observable filtration*.

In addition, the intensities also depend on a hidden Markov chain Y with state space $\{1, \dots, k\}$, initial distribution μ_0 and deterministic generator matrix $Q = (q_t^{ij})_{1 \leq i, j \leq k}$. We assume Q is a continuous, stable and conservative $Q(t)$ -matrix, i.e.:

Assumptions 1.2.1 (Generator matrix). For all $t \in [0, T]$ and $1 \leq i \leq k$:

$$0 \leq q_t^{ij} < +\infty \text{ for all } 1 \leq j \leq k, j \neq i, \quad \sum_{j=1}^k q_t^{ij} = 0 \quad \text{and} \quad Q \in C[0, T].$$

The previous are standing assumptions in the literature of Markov chains and they guarantee in particular the existence of a unique (up to evanescence) càdlàg version of Y , and the existence of the predictable intensity kernel for its jump measure, as given in (1.3.1). In addition, we suppose that

$$Y \text{ has a.s. no common jumps with } N^- \text{ and } N^+. \quad (1.2.2)$$

In this model, Y_t represents the *market regime* at time t and results from the interaction of a range of different factors, such as market sentiment towards the asset and varying levels of competition with other liquidity providers. These effects are almost never explicitly modelled in the mathematical literature of market making. To this end, it is natural to assume that Y is not directly observable by the MM who can only see N^-, N^+ and W . (As often assumed in filtering and control with hidden Markov chain models, the parameters k, Q , and μ_0 of Y are known to the MM, who needs to model/estimate them in practice.)

We define the MM's inventory process

$$N_t := n_0 + N_t^- - N_t^+, \quad (1.2.3)$$

and the cash account process

$$X_t^{\delta^-, \delta^+} := x_0 + \int_0^t (S_u + \delta_u^+) dN_u^+ - \int_0^t (S_u - \delta_u^-) dN_u^-, \quad (1.2.4)$$

for some fixed initial values $n_0 \in \mathbb{Z}$ and $x_0 \in \mathbb{R}$. The MM's preferences at terminal time are represented by a CARA utility function:

$$U_\gamma(c) = \frac{1 - e^{-\gamma c}}{\gamma}, \text{ for } c \in \mathbb{R} \text{ and } \gamma > 0, \text{ and } U_0 = Id.$$

Note that

$$\mathbb{F}^{N^-, N^+} = \mathbb{F}^N \quad \text{and} \quad \mathbb{F}^{N^-, N^+, W} = \mathbb{F}^{N, W},$$

since N^-, N^+ have a.s. no common jumps.

We shall make some very natural modelling assumptions on the intensities. These generalize those in [BL14, Gué17, GL15] to a context with several market regimes, without the need of any conditions on the derivatives or even smoothness.

Assumptions 1.2.2 (Orders intensities). *There exist functions $\Lambda_i^\pm : (\overline{\delta_*}, \delta^*) \rightarrow (0, +\infty)$ for $i = 1, \dots, k$ and $N^* \in \mathbb{N} \cup \{+\infty\}$ with $-N^* \leq n_0 \leq N^*$, such that:*

(i) $\lambda_t^\pm = \lambda_t^{\pm, \alpha} = \Lambda_{Y_{t-}}^\pm(\delta_t^\pm) \mathbb{1}_{\{\mp N_{t-} < N^*\}}$ if $\alpha = (\delta^-, \delta^+)$.

(ii) Λ_i^\pm is continuous, decreasing (not necessarily strictly) and $\lim_{\delta \rightarrow +\infty} \delta^{\mathbb{1}_{\{\gamma=0\}}} \Lambda_i^\pm(\delta) = 0$ if $\delta^* = +\infty$, for all $1 \leq i \leq k$.

Notation. For the sake of readability, we will often write $\mathbb{1}_t^\pm$ instead of $\mathbb{1}_{\{\mp N_{t-} < N^*\}}$.

N^* is a self-imposed constraint on the MM's inventory (whichever its sign) as originally proposed in [GLFT13]. When $N^* = \infty$ we are in the unconstrained context, while a finite N^* effectively means that the MM will not buy (resp. sell) whenever $N_{t-} = N^*$ (resp. $N_{t-} = -N^*$). In practice, the MM could achieve this either by abstaining from quoting or by quoting with an excessively large spread that prevents any transaction (i.e., a “stub” or “placeholder quote”).⁷

Remark 1.2.3. The last assumption states in particular that the intensities should decrease to zero when the spreads grow arbitrarily large. Furthermore, when $\gamma = 0$ we require that they decay faster than $1/\delta$. Loosely speaking, this states that for any $\gamma \geq 0$ the “expected” utility of the MM's instantaneous margin, $F_i^\pm(\delta) := \Lambda_i(\delta)U_\gamma(\delta)$, should vanish for “stub quotes”. This is a reasonable assumption in practice, as the opposite could lead to unrealistic optimal strategies such as continuously quoting “infinite spreads”.

Remark. Note that the above intensities are predictable, and since any admissible spreads δ^-, δ^+ are bounded, λ^-, λ^+ are in turn bounded. Accordingly, N^-, N^+, N and X are non-explosive (see, e.g., [Bré81, p.27 T8]). Furthermore, for any constant $\lambda^* > 0$ such that $\lambda^- + \lambda^+ \leq \lambda^*$, it is easy to see that for all $p \in \mathbb{N}$, $t \in [0, T]$ and $r \geq 0$,

$$\mathbb{E}^\alpha[|N_t|^p] \leq \mathbb{E}^\alpha[(n_0 + N_t^- + N_t^+)^p] \leq \mathbb{E}^\alpha[(n_0 + M)^p], \quad \mathcal{M}_{|N_t|}^\alpha(r) \leq \mathcal{M}_{n_0 + N_t^- + N_t^+}^\alpha(r) \leq \mathcal{M}_{n_0 + M}^\alpha(r), \quad (1.2.5)$$

where $M \sim \text{Poisson}(\lambda^*T)$ and \mathcal{M}_R^α is the moment generating function of a random variable R .⁸

The following are some examples which we shall come to back in Section 4.5, when revisiting the standard assumptions in the literature. We remark that exponential, power-law and logistic intensities are the explicit ones most commonly used in the mathematical literature, for tractability reasons.

Examples 1.2.4. 1. $\Lambda_i^\pm(\delta) = a_i^\pm e^{-b_i^\pm \delta}$, with $a_i^\pm, b_i^\pm > 0$, $-\infty \leq \delta_* < \delta^* \leq +\infty$.

2. $\Lambda_i^\pm(\delta) = \frac{a_i^\pm}{1 + c_i^\pm e^{b_i^\pm \delta}}$, with $a_i^\pm, b_i^\pm, c_i^\pm > 0$, $-\infty \leq \delta_* < \delta^* \leq +\infty$.

3. $\Lambda_i^\pm(\delta) = a_i^\pm (\pi/2 + \arctan(-b_i^\pm \delta + c_i^\pm))$, with $a_i^\pm, b_i^\pm > 0$, $c_i^\pm \in \mathbb{R}$, $-\infty \leq \delta_* < \delta^* \leq +\infty$.

4. $\Lambda_i^\pm(\delta) = a_i^\pm (b_i^\pm \delta + c_i^\pm)^{-d_i^\pm}$, with $a_i^\pm, b_i^\pm > 0$, $d_i^\pm > \mathbb{1}_{\{\gamma=0\}}$, $c_i^\pm \in \mathbb{R}$, $\max_i \{-c_i^\pm/b_i^\pm\} < \delta_* < \delta^* \leq +\infty$.

To fix ideas, consider Example 1 and let us observe that if Y represents decreasing levels of competition for the MM, the values of a_i^\pm (resp. b_i^\pm) should increase (resp. decrease) as Y increases. The same is true if Y represents increasing levels of positive sentiment (or “bullishness”) towards the asset. In practice, Y will result from the combination of these effects and many more.

⁷For simplicity and without loss of generality, only the first alternative is formalized in our model (as done in [Gué17, GL15, GLFT13]) and this is reflected in the admissible spreads being real-valued. We could also allow for different constraints depending on the sign of the inventory, but we refrain from this to simplify the notation.

⁸By possibly enlarging the space, one can consider a counting process Z with no common jumps with N^-, N^+ and stochastic intensity $\lambda^* - \lambda^- - \lambda^+ \geq 0$. Then the process $Z + N^- + N^+$ is a Poisson with intensity λ^* that dominates $N^+ + N^-$. The claim follows immediately.

Suppose the MM can hedge any remaining inventory at time T at the reference price S_T minus a certain execution cost. Neglecting discounting between 0 and T , we consider the optimization problem faced by the MM who tries to maximize the expected utility of her terminal *penalized profit and loss* (P&L),

$$\sup_{\delta^-, \delta^+ \in \mathcal{A}} \mathbb{E}^{\delta^-, \delta^+} \left[U_\gamma \left(X_T^{\delta^-, \delta^+} + S_T N_T - \ell(N_T) - \frac{1}{2} \sigma^2 \zeta \int_0^T N_t^2 dt \right) \right], \quad (1.2.6)$$

where:

- (i) $\zeta \geq 0$. When $\zeta > 0$ we are including in the model a *running penalty*, as firstly done in [CJ15], for the MM to further control her *accumulated inventory risk*. This is the same as subtracting the variance of the mark-to-market value of the inventory, weighted by ζ .
- (ii) $\ell : \mathbb{R} \rightarrow [0, +\infty)$ represents the final execution cost and is increasing on $[0, +\infty)$, decreasing on $(-\infty, 0]$ and $\ell(0) = 0$. (And usually convex in practice.)
- (iii) If $N^* = +\infty$, we set $\gamma = \zeta = 0 \equiv \ell$.

The last restriction on the parameters states that the only case we will consider in which the inventory could be arbitrary large (whichever its sign) is that of a completely risk-neutral MM with negligible costs. The risk averse case is far more challenging and has not been treated in full mathematical detail even in the complete information case (see discussion in [GLFT13]). The model given by (iii) is of secondary interest in practice, yet, it will allow us to further understand the general problem and methodology and to have a more holistic view at almost no extra cost.

Remark. Previous works in optimal market making do not consider both penalty and CARA utility (with $\gamma > 0$) in unison, and instead treat the two families of models separately [Gué17]. Besides the obvious interest in unifying the approach and generalizing existing models, allowing for $\gamma > 0$ and $\zeta > 0$ simultaneously adds flexibility to the risk managing capabilities of the MM. Indeed, in [Gué17] the author derives some HJB-type systems of ODEs for each problem with a unique risk aversion parameter $\bar{\gamma}$, and then relates them by the introduction of an auxiliary parameter $0 \leq \bar{\xi} \leq \bar{\gamma}$. The later is afterwards interpreted as measuring risk aversion to non-execution risk and it is carried forward to the formulae of the “optimal strategies”. However, such strategies are not shown to result from any particular optimization problem for $0 < \bar{\xi} < \bar{\gamma}$. By looking at the equations of the full information version of our model (Section 4) with a single market regime, one can immediately see that the equations in [Gué17] are obtained via the reparametrization $\gamma = \bar{\xi}$ and $\zeta = \bar{\gamma} - \bar{\xi}$. Thus, the formulation in (1.2.6) allows us, in particular, to establish the optimality of the strategies in [Gué17] for any value of non-execution risk aversion and, in general, to differentiate between aversion to different types of risks. We can say that, in our model, γ represents aversion to all types of risks, while the penalty ζ is used to further increase the aversion to risks other than non-execution risk.

1.3 Jump measures

It will be useful in the sequel to understand which are the jump measures and intensity kernels involved in our model. Let $\mu^N(dt, dz)$ (or simply μ^N) be the jump (or counting) measure of N (see def. in [Bré81, p.234] or [JS02, p.69]). We use a similar notation for the jump measures of Y and the pair (Y, N) . These are all finite random measures since the processes are non-explosive for any $\delta^+, \delta^- \in \mathcal{A}$ (see (1.2.5)), and they admit predictable intensity kernels (see [Bré81, p.235 D2]). If $m_{z_0}(dz)$ is the Dirac measure at a point $z_0 \in \mathbb{R}$, then the $(\mathbb{P}^\alpha, \mathbb{F})$ -predictable intensity kernels of μ^N, μ^Y and $\mu^{(Y, N)}$ are respectively given by:

$$\begin{aligned}
\eta_t^{\alpha, N}(dz) &= \eta^N(Y_{t-}, \delta_t^-, \delta_t^+, N_{t-}, dz) \\
&= \lambda_t^- m_1(dz) + \lambda_t^+ m_{-1}(dz) \\
&= \Lambda_{Y_{t-}}^-(\delta_t^-) \mathbb{1}_{\{N_{t-} < N^*\}} m_1(dz) + \Lambda_{Y_{t-}}^+(\delta_t^+) \mathbb{1}_{\{-N_{t-} < N^*\}} m_{-1}(dz), \\
\eta_t^Y(dh) &= \eta^Y(Y_{t-}, dh) = \sum_{j \neq Y_{t-}} q_t^{Y_{t-}, j} m_{j-Y_{t-}}(dh)
\end{aligned} \tag{1.3.1}$$

and

$$\eta_t^{\alpha, (Y, N)}(dh, dz) = \eta_t^{\alpha, N}(dz) \otimes m_0(dh) + \eta_t^Y(dh) \otimes m_0(dz). \tag{1.3.2}$$

We denote by $\bar{\mu}_\alpha^N(dt, dz)$, $\bar{\mu}_\alpha^Y(dt, dh)$ and $\bar{\mu}_\alpha^{(Y, N)}(dt, dh, dz)$ the corresponding $(\mathbb{F}, \mathbb{P}^\alpha)$ -compensated measures.

Remark. Note that equation (1.3.2) is a consequence of assuming that Y, N have a.s. no common jumps (which in turn is equivalent to (1.2.2)). Indeed, it is easy to see that Y, N have a.s. no common jumps if and only if

$$\mu^{(Y, N)}(dt, dh, dz) = \mu^N(dt, dz) \otimes m_0(dh) + \mu^Y(dt, dh) \otimes m_0(dz), \tag{1.3.3}$$

which implies (1.3.2).

1.4 Construction and characterization of the model

In this section, we rigorously construct and characterize our model. On top of the mathematical interest in a well defined framework, the results and techniques presented here will be used throughout the paper, when tackling both the filtering and the optimization problems. We start by proving the existence of a model under our assumptions (Proposition 1.4.1). We construct it by means of a reference probability (an overview of which can be found, e.g., in [Br681, Chpt.VI]) and, in particular, the Girsanov theorem for point processes.

For the following proposition, let us consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{Q})$ under the usual conditions, with $\mathcal{F} = \mathcal{F}_T$ and \mathcal{F}_0 the completed trivial sigma algebra. We refer to \mathbb{Q} as the *reference probability*. Suppose this space supports a two-dimensional counting process (N^-, N^+) (see def. in [Br681, Chapt.II]), a Wiener process W and a Markov chain Y with finite state space $\{1, \dots, k\}$, initial distribution μ_0 and time-dependent generator matrix $Q = (q_t^{ij})_{1 \leq i, j \leq k}$ satisfying Assumptions 1.2.1, such that N^-, N^+, Y verify (1.2.2). We set $N = N^- - N^+$ as before and suppose N^\pm has intensity $\mathbb{1}_t^\pm = \mathbb{1}_{\{\mp N_{t-} < N^*\}}$, where $N^* \in \mathbb{N} \cup \{+\infty\}$ is given.⁹ Let $\Lambda_i^\pm : (\delta_*, \delta^*) \rightarrow \mathbb{R}$ for $i = 1, \dots, k$ be functions under the Assumptions 1.2.2 (ii), for given $-\infty \leq \delta_* < \delta^* \leq +\infty$ and $N^* \in \mathbb{N} \cup \{+\infty\}$ with $-N^* \leq n_0 \leq N^*$, and define \mathcal{A} as in (1.1.1).

Proposition 1.4.1. *Let $\alpha = (\delta^-, \delta^+) \in \mathcal{A}^2$ and define the process Z^α as the stochastic exponential*

$$Z^\alpha := \mathcal{E} \left(\int_0^\cdot (\Lambda_{Y_{u-}}^-(\delta_u^-) - 1) (dN_u^- - \mathbb{1}_u^- du) + (\Lambda_{Y_{u-}}^+(\delta_u^+) - 1) (dN_u^+ - \mathbb{1}_u^+ du) \right). \text{ Then,}$$

- (i) Z^α is a strictly positive uniformly integrable (\mathbb{F}, \mathbb{Q}) -martingale. In particular, $\mathbb{E}^\mathbb{Q}[Z_T^\alpha] = 1$.
- (ii) \mathbb{P}^α defined by $\frac{d\mathbb{P}^\alpha}{d\mathbb{Q}} = Z_T^\alpha$ is an equivalent probability measure such that, for $(\mathbb{F}, \mathbb{P}^\alpha)$, N^\pm is a counting process with intensity $\lambda_t^\pm = \Lambda_{Y_{t-}}^\pm(\delta_t^\pm) \mathbb{1}_t^\pm$, W is a Wiener process and Y is a Markov chain with state space $\{1, \dots, k\}$, initial distribution μ_0 and generator matrix Q . (That is, Assumptions 1.2.1, 1.2.2 and (1.2.2) are all verified.)

⁹Such a simpler model can be constructed for example as a product of canonical spaces, with the existence of the counting processes with the right intensities proved in [JP82, Thm.24 and Cor.31]. The finite-dimensional result has the same proof as in one dimension, starting from independent Poisson measures.

Proof. See Appendix A. □

Reciprocally, suppose we start with a family of filtered probability spaces $\{(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P}^\alpha)\}_{\alpha \in \mathcal{A}^2}$ as in Section 1.1, supporting a counting process (N^-, N^+) and a Markov chain Y satisfying condition (1.2.2), together with a Wiener process W . Suppose also that Y has finite state space $\{1, \dots, k\}$, initial distribution μ_0 and time-dependent generator matrix Q , and that our Assumptions 1.2.1 and 1.2.2 are in place. We set $N = N^- - N^+$ as before. We would like to characterize any such model in terms of a reference probability as in Proposition 1.4.1 by an inverse change of measure. However, we only claim the uniqueness of the reference probability on \mathbb{F}^N . We will come back to this result in the sequel.

Proposition 1.4.2. *Let $\alpha = (\delta^-, \delta^+) \in \mathcal{A}^2$ and define*

$$\bar{Z}^\alpha := \mathcal{E} \left(\int_0^\cdot \left(1/\Lambda_{Y_{u-}}^-(\delta_u^-) - 1 \right) (dN_u^- - \lambda_u^- du) + \left(1/\Lambda_{Y_{u-}}^+(\delta_u^+) - 1 \right) (dN_u^+ - \lambda_u^+ du) \right). \text{ Then,}$$

- (i) \bar{Z}^α is a strictly positive uniformly integrable $(\mathbb{F}, \mathbb{P}^\alpha)$ -martingale. In particular, $\mathbb{E}^\alpha[\bar{Z}_T] = 1$.
- (ii) \mathbb{Q}^α defined by $\frac{d\mathbb{Q}^\alpha}{d\mathbb{P}^\alpha} = \bar{Z}_T^\alpha$ is a probability measure equivalent to \mathbb{P}^α such that, for $(\mathbb{F}, \mathbb{Q}^\alpha)$, N^\pm is a counting process with intensity $\mathbb{1}_t^\pm$, W is a Wiener process and Y is a Markov chain with state space $\{1, \dots, k\}$, initial distribution μ_0 and generator matrix Q . Furthermore, Y, N, W are independent.
- (iii) If we define Z^α as in Lemma 1.4.1 under \mathbb{Q}^α , then $Z^\alpha = 1/\bar{Z}^\alpha$ (i.e., the two changes of measure are inverse of each other).
- (iv) For all $\tilde{\alpha} \in \mathcal{A}^2$, $\mathbb{Q}^{\tilde{\alpha}} \equiv \mathbb{Q}^\alpha$ on \mathbb{F}_T^N .

Proof. See Appendix A. □

2 Filtering problem

Since the MM cannot directly observe all the information in \mathbb{F} but only $\mathbb{F}^{N,W}$ (in particular, she cannot observe Y), in order to solve the optimization problem (1.2.6) under partial information we want to reduce it first to an equivalent one under full information. Throughout this section we work under \mathbb{P}^α with $\alpha = (\delta^-, \delta^+) \in \mathcal{A}^2$ fixed. We sometimes omit α from the notation for simplicity.

Let us consider the optional projections

$$\Pi^{\alpha,i} := {}^\circ(\mathbb{1}_{\{Y_t=i\}})^{(\mathbb{P}^\alpha, \mathbb{F}^{N,W})}, \quad 1 \leq i \leq k, \quad \text{and} \quad \Pi^\alpha := (\Pi^{\alpha,1}, \dots, \Pi^{\alpha,k})^{10}$$

That is, Π^α is the unique càdlàg version of the conditional distribution of Y given the observable information: $\Pi_t^{\alpha,i} = \mathbb{P}^\alpha(Y_t = i | \mathcal{F}_t^{N,W})$. It reflects the MM's beliefs on the state of the market as she updates them through time.

At any point in time, the MM typically observes orders arriving at rates which do not seem to match exactly those of any particular regime, but some average ones instead. We want to characterize the observable (that is, the $\mathbb{F}^{N,W}$ -) predictable intensities of N^-, N^+ and μ^N in terms of Π^α . Loosely speaking, the observable intensity of, say N^+ , is obtained projecting: $\mathbb{E}^\alpha[\lambda_t^+ | \mathcal{F}_t^{W,N}]$ [Bré81, p.32 Comment and Pseudo-Proof of T14]. This yields an average of the different regime intensities weighted by the observable probabilities of the market states. The

¹⁰For any càdlàg bounded process M (not necessarily adapted) on a filtered probability space $(\Sigma, \mathcal{H}, \mathbb{H}, \mathbb{P})$ satisfying the usual conditions, the *optional projection* of M on \mathbb{H} is the unique càdlàg process ${}^\circ M^{(\mathbb{P}, \mathbb{H})}$ such that ${}^\circ M_t^{(\mathbb{P}, \mathbb{H})} = \mathbb{E}^\mathbb{P}[M_t | \mathcal{H}_t]$ a.s. for each t . Its existence is guaranteed by the Optional Projection Theorem (see, e.g., [JYC09, p.264] or [Nik06, p.357-358]).

only technical difficulty is that of finding a predictable version of such process. ${}^o(\lambda^+)$ has the desired projective property but it is not predictable in general. ${}^o(\lambda^+)_{t-}$, on the other hand, is predictable but does not normally enjoy the projective property. In fact, the process we are looking for is “in between” these two.

Proposition 2.1.1 (Observable intensities). *The $(\mathbb{P}^\alpha, \mathbb{F}^{N,W})$ -predictable intensities of N^\pm and μ^N , resp., are*

$$\widehat{\lambda}_t^\pm{}^\alpha := \mathbb{1}_t^\pm \sum_{i=1}^k \Pi_{u-}^{\alpha,i} \Lambda_i^\pm(\delta_u^\pm) \quad \text{and} \quad \widehat{\eta}_t^{\alpha,N}(dz) := \widehat{\lambda}_t^-{}^\alpha m_1(dz) + \widehat{\lambda}_t^+{}^\alpha m_{-1}(dz).$$

Notation. We set $\widehat{\Lambda}^\pm(\pi, \delta) = \sum_{i=1}^k \pi^i \Lambda_i^\pm(\delta)$; thus, $\widehat{\lambda}_t^\pm{}^\alpha = \mathbb{1}_t^\pm \widehat{\Lambda}^\pm(\Pi_{u-}^\alpha, \delta_u^\pm)$.

Proof. See Appendix A. □

We give now the *filtering (or Kushner-Stratonovich) equations* for the observable distribution of Y . These are coupled stochastic differential equations (SDEs) governing the dynamics of Π^α . They allow the MM to dynamically and efficiently update an existing estimate Π_t^α upon receiving new information, without having to recompute it from scratch. In a nutshell, given a side of the market and the spread quoted, an order arrival modifies the observed probability of regime i by a multiplicative factor. This factor is given by the percentage difference between the orders intensity of regime i and the observed one. In between orders however, the probabilities change continuously and in a more subtle way, depending not only on the different levels of liquidity but also on the transition rates between them.

Notation. We denote by $\Delta \subset \mathbb{R}^k$ the $(k-1)$ -simplex (i.e., $\Delta = \{\pi \in \mathbb{R}^k : 0 \leq \pi^i \leq 1 \text{ for all } i, \text{ and } \sum_i \pi^i = 1\}$) and by Δ° its interior relative to the hyperplane $\{\pi \in \mathbb{R}^k : \sum_i \pi^i = 1\}$ (i.e., $\Delta^\circ = \{\pi \in \Delta : 0 < \pi^i < 1 \text{ for all } i\}$).

Proposition 2.1.2 (Observable distribution of Y). *The process $\Pi^\alpha = (\Pi^{\alpha,1}, \dots, \Pi^{\alpha,k})$ is the unique strong solution of the constrained system of SDEs*

$$\begin{aligned} d\Pi_t^i &= \sum_{j=1}^k \left(q_t^{ji} \Pi_t^j + \Pi_t^i \Pi_t^j \left(\mathbb{1}_t^- (\Lambda_j^- - \Lambda_i^-)(\delta_t^-) + \mathbb{1}_t^+ (\Lambda_j^+ - \Lambda_i^+)(\delta_t^+) \right) \right) dt \\ &+ \Pi_{t-}^i \left(\frac{\Lambda_i^-(\delta_t^-)}{\widehat{\Lambda}^-(\Pi_{u-}^\alpha, \delta_u^-)} - 1 \right) dN_{t-}^- + \Pi_{t-}^i \left(\frac{\Lambda_i^+(\delta_t^+)}{\widehat{\Lambda}^+(\Pi_{u-}^\alpha, \delta_u^+)} - 1 \right) dN_{t-}^+, \end{aligned} \tag{2.1.1}$$

subject to $\Pi_0 = \mu_0$ and $\Pi_t \in \Delta$ for all $t \in [0, T]$ a.s.

Proof. See Appendix A. □

Remark 2.1.3. Consider the identification $\Delta \simeq [0, 1]^{k-1}$ (resp. $\Delta^\circ \simeq (0, 1)^{k-1}$) obtained by the substitution $\pi^k = 1 - \sum_{j < k} \pi^j$ (where the choice of the k -th coordinate over the rest is completely arbitrary). Then the constrained system of SDEs (2.1.1) for $(\Pi^{\alpha,1}, \dots, \Pi^{\alpha,k})$ becomes an “unconstrained” system for $(\Pi^{\alpha,1}, \dots, \Pi^{\alpha,k-1}) \in [0, 1]^{k-1}$. Henceforth, we shall use this identification whenever convenient.

We finish this section with a short lemma. It states that the conditional distribution Π^α can never reach the relative border of the simplex Δ , provided it starts from the relative interior. This amounts to saying that all regimes have some positive probability at time zero.

Lemma 2.1.4. *If $\mu_0 \in \Delta^\circ$, then $\Pi_t^\alpha \in \Delta^\circ$ for all $0 \leq t \leq T$ a.s.*

Proof. See Appendix A. □

In light of the previous lemma, we will assume from here onwards that

$$\mu_0 \in \Delta^\circ, \quad (2.1.2)$$

and therefore work with Δ° instead of Δ .

3 Value function and HJB equation

In this section, we tackle the control problem of the MM with the filter as an additional state variable. We define the MM's value function and we aim to characterize it by means of an HJB equation.

Let $t \in [0, T]$. We consider our model “starting at t ” instead of 0. Whenever a process is defined from time “ t^- onwards” (i.e., from time t onwards and decreasing its left-limit value at t) this implicitly means it is constant on $[0, t)$. In particular, we use this convention for all integrals (stochastic or not) of the form \int_t^\cdot and (with a slight abuse of notation) for the processes $W - W_t$ and $N^\pm - N_{t^-}^\pm$. We work with t^- instead of t due to the jumps of the processes N^-, N^+ . However, since N^-, N^+ are quasi-left continuous,¹¹ for most intended purposes one can drop the left limit with no harm. We define $\mathbb{F}^{t, W, N} = (\sigma(W_r - W_t, N_r - N_{t^-} : t \leq r \leq u \vee t) \vee \mathcal{F}_0)_u$ and $\mathbb{F}^{t, N}$ analogously.

Let $s, x \in \mathbb{R}$, $n \in \mathbb{Z} \cap [-N^*, N^*]$ and $\pi \in \Delta^\circ \subset \mathbb{R}^k$. The set \mathcal{A}_t of admissible spreads starting at t is the set of $\delta \in \mathcal{A}$ which are independent of $\mathcal{F}_{t^-}^{N, W}$ (equivalently, the $\delta \in \mathcal{A}$ which are $\mathbb{F}^{t, W, N}$ -predictable). Consider for each $\alpha = (\delta^-, \delta^+) \in \mathcal{A}_t^2$ the processes $S^{t, s}, X^{\alpha, t, x, s}, N^{t, n}, \Pi^{\alpha, t, n, \pi}$ defined pathwise, outside some set $A \in \mathcal{F}_0$, by (1.2.1), (1.2.4), (1.2.3), (2.1.1) resp., replacing the initial conditions s_0, x_0, n_0, μ_0 at time 0 by s, x, n, π resp. at time t^- . We remark that $\mathbb{F}^{t, W, N} = \mathbb{F}^{t, W, N^{t, n}}$ (since $N^{t, n} = n + N^- - N^+ - (N_{t^-}^- - N_{t^-}^+)$) and all the processes defined in this section are adapted to this filtration. We further assume there exists a family of “physical” probabilities $(\mathbb{P}^{\alpha, t, n, \pi})_{\alpha \in \mathcal{A}_t^2}$ such that their null sets generate \mathcal{F}_0 , and for $(\mathbb{P}^{\alpha, t, n, \pi}, \mathbb{F}^{t, W, N})$ it holds that $W - W_t$ is a Wiener process and $N^\pm - N_{t^-}^\pm$ has predictable intensity $\widehat{\lambda}^{\pm, \alpha, t, n, \pi}$, as defined in Proposition 2.1.1 in terms of $N^{t, n}$ and $\Pi^{\alpha, t, n, \pi}$.

We define the penalized P&L from t to T (see (1.2.6) for parameter restrictions) as

$$P_{t, T}^{\alpha, s, x, n} := X_T^{\alpha, t, x, s} + S_T^{t, s} N_T^{t, n} - \ell(N_T^{t, n}) - \frac{1}{2} \sigma^2 \zeta \int_t^T (N_u^{t, n})^2 du, \quad (3.1.1)$$

and the *value function* of problem (1.2.6) as

$$V(t, s, x, n, \pi) := \sup_{\alpha \in \mathcal{A}_t^2} \mathbb{E}^{\alpha, t, n, \pi} \left[U_\gamma \left(P_{t, T}^{\alpha, s, x, n} \right) \right]. \quad (3.1.2)$$

Our goal is to compute optimal or “close to optimal” strategies.¹² The Dynamic Programming Principle and Ito's Lemma allow us to formally derive (see, e.g., [Bou07]) the *Hamilton–Jacobi–Bellman* (or *dynamic programming*) partial-integro differential equation associated to V :

$$\begin{aligned} 0 = & v_t + \mu v_s + \frac{1}{2} \sigma^2 v_{ss} + \sum_{i, j=1}^k q_t^{ji} \pi^j v_{\pi^i} + \frac{1}{2} \sigma^2 \zeta n^2 (\gamma v - 1) \\ & + \mathbb{1}_{\{n < N^*\}} \sup_{\delta^- \in (\delta_*, \delta^*)} \left\{ \sum_{i, j=1}^k (\Lambda_j^- - \Lambda_i^-) (\delta^-) \pi^j \pi^i v_{\pi^i} + d_{\delta^-}^-(v) \sum_{i=1}^k \pi^i \Lambda_i^-(\delta^-) \right\} \\ & + \mathbb{1}_{\{-n < N^*\}} \sup_{\delta^+ \in (\delta_*, \delta^*)} \left\{ \sum_{i, j=1}^k (\Lambda_j^+ - \Lambda_i^+) (\delta^+) \pi^j \pi^i v_{\pi^i} + d_{\delta^+}^+(v) \sum_{i=1}^k \pi^i \Lambda_i^+(\delta^+) \right\}, \end{aligned} \quad (3.1.3)$$

¹¹For example, because they are increasing càdlàg processes admitting continuous compensators for (any one of) the physical probabilities [JS02, p.70 Prop.1.19 or p.77 Prop.2.9].

¹²By “close to optimal” we mean that for each $\varepsilon > 0$ there exists a strategy such that the supremum in (3.1.2) is attained up to ε .

with terminal condition $v(T, s, x, n, \pi) = U_\gamma(x + sn - \ell(n))$, where:

$$d_\delta^\pm(v)(t, s, x, n, \pi) = v\left(t, s, x \pm (s \pm \delta), n \mp 1, \frac{1}{\sum_{j=1}^k \pi^j \Lambda_j^\pm(\delta)} (\pi^1 \Lambda_1^\pm(\delta), \dots, \pi^k \Lambda_k^\pm(\delta))\right) - v(t, s, x, n, \pi),$$

and we convene the following:

Notation. (i) The derivatives with respect to $\pi = (\pi^1, \dots, \pi^k)$ should be understood via the identification of Remark 2.1.3. (ii) Although it is not meaningful to evaluate v on the inventories $\pm N^* \pm 1$, this only happens in equation (3.1.3) when the corresponding term vanishes. This slight abuse of notation can be found throughout previous works and we will be using it as well.

Equation (3.1.3) can also be seen as a coupled system of PIDEs indexed in $n \in \mathbb{Z} \cap [-N^*, N^*]$. (We will talk about system of equations or simply “equation” indistinctly). Nonlinearity aside, (3.1.3) is rather complex, in particular due to being of second order, high-dimensional and with derivatives in almost all of these dimensions. Tackling it directly (either analytically or numerically) is utterly challenging. Consequently, it has become common practice for optimal market making and optimal liquidation models *à la* Avellaneda–Stoikov [AS08] to propose an ansatz for the solution [AS08, BL14, CDJ17, CJ15, CJR14, FL12, FL13, Gué17, GL15, GLFT13]. This approach however, relies heavily on the existence of a classical solution for the resulting simplified equation, so that the ansatz is ultimately proved valid by a suitable verification theorem. (See Section 4 for more details.) When the simplified equation does not admit (or cannot be guaranteed to admit) a classical solution, and a viscosity approach needs to be used instead, the previous argument breaks down.

If we attempted to solve our problem by the standard approach, a plausible ansatz for the value function could be

$$V(t, s, x, n, \pi) = U_\gamma(x + sn + \Theta(t, n, \pi)). \quad (3.1.4)$$

Formal substitution yields the following equation for Θ :

$$\begin{aligned} 0 = & \theta_t + \mu n - \frac{1}{2} \sigma^2 n^2 (\gamma + \zeta) + \sum_{i,j=1}^k q_t^{j,i} \pi^j \theta_{\pi^i} \\ & + \mathbb{1}_{\{n < N^*\}} \sup_{\delta^- \in (\delta_*, \delta^*)} \left\{ \sum_{i,j=1}^k (\Lambda_j^- - \Lambda_i^-)(\delta^-) \pi^j \pi^i \theta_{\pi^i} + U_\gamma(\delta^- + d_{\delta^-}^-(\theta)) \sum_{i=1}^k \pi^i \Lambda_i^-(\delta^-) \right\} \\ & + \mathbb{1}_{\{-n < N^*\}} \sup_{\delta^+ \in (\delta_*, \delta^*)} \left\{ \sum_{i,j=1}^k (\Lambda_j^+ - \Lambda_i^+)(\delta^+) \pi^j \pi^i \theta_{\pi^i} + U_\gamma(\delta^+ + d_{\delta^+}^+(\theta)) \sum_{i=1}^k \pi^i \Lambda_i^+(\delta^+) \right\}, \end{aligned} \quad (3.1.5)$$

with terminal condition $\theta(T, n, \pi) = -\ell(n)$, where:

$$d_\delta^\pm(\theta)(t, n, \pi) = \theta\left(t, n \mp 1, \frac{1}{\sum_{j=1}^k \pi^j \Lambda_j^\pm(\delta)} (\pi^1 \Lambda_1^\pm(\delta), \dots, \pi^k \Lambda_k^\pm(\delta))\right) - \theta(t, n, \pi).$$

The new system of PIDEs is of first order and no longer depends on the variables s and x (there is no diffusion anymore). This is a considerable simplification; one that will permit effective numerical solution in Section 5. But it is not good enough for us to assert existence of a classical solution. Notwithstanding, we are able to rigorously prove the decomposition (3.1.4) and explicitly find Θ as a new “value function” (Theorem 3.1.1). When the control

space is compact, this ultimately allows us to characterize Θ as the unique solution of the terminal condition PIDE (3.1.5) in the viscosity sense (Theorem 3.1.4), further simplified in the unconstrained inventory case. These two theorems constitute the main mathematical results of this paper. They allow us to safely postulate reasonable candidates for optimal (or ϵ -optimal) strategies for the MM, i.e., those given by spreads that (at least approximately) realize the suprema in (3.1.5).

Theorem 3.1.1. *There exists a unique function $\Theta : [0, T] \times (\mathbb{Z} \cap [-N^*, N^*]) \times \Delta^\circ \rightarrow \mathbb{R}$ such that the decomposition (3.1.4) holds true. Furthermore, there exists a family of equivalent probability measures $\tilde{\mathbb{P}}^{\alpha, t, n, \pi} \sim \mathbb{P}^{\alpha, t, n, \pi}$, $\alpha = (\delta^-, \delta^+) \in \mathcal{A}_t^2$, such that*

(i) $\Pi^{\alpha, t, n, \pi}$ is the unique strong solution of (2.1.1) with initial condition (t^-, π) under $\tilde{\mathbb{P}}^{\alpha, t, n, \pi}$.

(ii) $\mu^{N^{t, n}}$ has $(\tilde{\mathbb{P}}^{\alpha, t, n, \pi}, \mathbb{F}^{t, W, N})$ -predictable intensity kernel

$$\tilde{\eta}_u^N(dz) := e^{-\gamma \delta_u^-} \widehat{\lambda}_u^-{}^{\alpha, t, n, \pi} m_1(dz) + e^{-\gamma \delta_u^+} \widehat{\lambda}_u^+{}^{\alpha, t, n, \pi} m_{-1}(dz).$$

(iii) $\Theta = U_\gamma^{-1} \circ \Psi = -\frac{1}{\gamma} \log(1 - \gamma \Psi)$ with

$$\Psi(t, n, \pi) := \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \tilde{\mathbb{E}}^{\alpha, t, n, \pi} \left[U_\gamma \left(\tilde{P}_{t, T}^{\alpha, n, \pi} \right) \right], \quad (3.1.6)$$

where $\tilde{P}_{t, T}^{\alpha, n, \pi} := \int_t^T \{ U_\gamma(\delta_u^-) \widehat{\lambda}_u^-{}^{\alpha, t, n, \pi} + U_\gamma(\delta_u^+) \widehat{\lambda}_u^+{}^{\alpha, t, n, \pi} + \mu N_u^{t, n} - \frac{1}{2} \sigma^2 (\gamma + \zeta) (N_u^{t, n})^2 \} du - \ell(N_T^{t, n})$ and $\tilde{\mathcal{A}}_t := \{ \delta \in \mathcal{A} : \delta \text{ is } \mathbb{F}^{t, N^{t, n}}\text{-predictable} \}$.

Proof. For shortness, we make an abuse of notation and omit (t, n, π) from the probability measures and expectations. Let us start by proving (3.1.4) and finding $(\tilde{\mathbb{P}}^\alpha)_{\alpha \in \mathcal{A}_t^2}$ with the desired properties. Using integration by parts we can re-write the penalized P&L (3.1.1) as

$$\begin{aligned} P_{t, T}^{\alpha, s, x, n} &= x + sn + \int_t^T \left\{ \mu N_u^{t, n} - \frac{1}{2} \sigma^2 (\gamma + \zeta) (N_u^{t, n})^2 \right\} du - \ell(N_T^{t, n}) \\ &+ \sigma \int_t^T N_u^{t, n} dW_u + \frac{1}{2} \sigma^2 \gamma \int_t^T (N_u^{t, n})^2 du + \int_t^T \delta_u^- dN_u^- + \delta_u^+ dN_u^+ \\ &=: x + sn + \bar{P}_{t, T}^{\alpha, n}, \end{aligned} \quad (3.1.7)$$

Consider first the case $\gamma = 0$. The integrals with respect to W, N^-, N^+ all have bounded integrands (and predictable for N^-, N^+), except in the case of unconstrained inventory: $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$ (see (1.2.6)). Regardless, we still have $\mathbb{E}^\alpha \left[\int_t^T (N_u^{t, n})^2 du \right] = \int_t^T \mathbb{E}^\alpha \left[(N_u^{t, n})^2 \right] du < +\infty$ by (1.2.5). Choosing $\tilde{\mathbb{P}}^\alpha := \mathbb{P}^\alpha$ the conclusion follows by taking expectation and by Propositions 2.1.1 and 2.1.2.

Consider now $\gamma > 0$. Hence, we are in the case $|N| \leq N^* < +\infty$. We define

$$A^\alpha := \mathcal{E} \left(-\gamma \sigma \int_t^\cdot N_u^{t, n} dW_u \right) = \exp \left(-\gamma \sigma \int_t^\cdot N_u^{t, n} dW_u - \frac{1}{2} \sigma^2 \gamma^2 \int_t^\cdot (N_u^{t, n})^2 du \right).$$

By Novikov's condition, A^α is a strictly positive uniformly integrable (UI) martingale with $\mathbb{E}^\alpha[A_T] = 1$, and therefore defines an equivalent probability measure $\mathbb{A}^\alpha \sim \mathbb{P}^\alpha$ via $\frac{d\mathbb{A}^\alpha}{d\mathbb{P}^\alpha} = A_T^\alpha$. Note that the Girsanov–Meyer Theorem. [Pro04, p.132 Thm.35] ensures the $\mathbb{F}^{t, W, N}$ -intensities of $N^\pm - N_t^\pm$ remain the same when changing to \mathbb{A}^α . Let us set

$$\begin{aligned} B^\alpha &:= \mathcal{E} \left(-\gamma \int_t^\cdot U_\gamma(\delta_u^-) d\bar{N}_u^-{}^\alpha + U_\gamma(\delta_u^+) d\bar{N}_u^+{}^\alpha \right) \\ &= \mathcal{E} \left(\int_t^\cdot (e^{-\gamma \delta_u^-} - 1) d\bar{N}_u^-{}^\alpha + (e^{-\gamma \delta_u^+} - 1) d\bar{N}_u^+{}^\alpha \right), \end{aligned}$$

where $\overline{N_u^\pm}^\alpha$ denote the corresponding $(\mathbb{P}^\alpha, \mathbb{F}^{t,W,N})$ -compensated (or equivalently, $(\mathbb{A}^\alpha, \mathbb{F}^{t,W,N})$ -compensated) processes. By the same arguments of Propositions 1.4.1 and 1.4.2, B^α is a strictly positive UI martingale with $\mathbb{E}^{\mathbb{A}^\alpha}[B_T^\alpha] = 1$ and defines an equivalent probability measure $\tilde{\mathbb{P}}^\alpha \sim \mathbb{A}^\alpha \sim \mathbb{P}^\alpha$ via $\frac{d\tilde{\mathbb{P}}^\alpha}{d\mathbb{A}^\alpha} = B_T^\alpha$, such that (ii) holds true. Note that (i) is also trivially verified due to the equivalence of the probability measures.

Suppose for the time being that Ψ is defined as in (3.1.6) but taking supremum over the whole set of admissible controls $\mathcal{A}_t^2 \supseteq \tilde{\mathcal{A}}_t^2$ instead. We will see afterwards that this makes no difference. To see (3.1.4), observe that the identity $U_\gamma(a+b) = U_\gamma(b)e^{-\gamma a} + U_\gamma(a)$ and (3.1.7) yield $U_\gamma(P_{t,T}^{\alpha,s,x,n}) = U_\gamma(\overline{P}_{t,T}^{\alpha,n})e^{-\gamma(x+sn)} + U_\gamma(x+sn)$, giving

$$V(t, s, x, n, \pi) = \sup_{\alpha \in \mathcal{A}_t^2} \mathbb{E}^\alpha \left[U_\gamma(\overline{P}_{t,T}^{\alpha,n,\pi}) \right] e^{-\gamma(x+sn)} + U_\gamma(x+sn).$$

On the other hand, by the same identity,

$$U_\gamma(x+sn+\Theta) = (U_\gamma \circ \Theta)e^{-\gamma(x+sn)} + U_\gamma(x+sn) = \Psi e^{-\gamma(x+sn)} + U_\gamma(x+sn).$$

As a consequence, (3.1.4) is equivalent to the equality $\Psi(t, n, \pi) = \sup_{\alpha \in \mathcal{A}_t^2} \mathbb{E}^\alpha \left[U_\gamma(\overline{P}_{t,T}^{\alpha,n}) \right]$. We check instead the stronger statement

$$\mathbb{E}^\alpha \left[\exp \left(-\gamma \overline{P}_{t,T}^{\alpha,n} \right) \right] = \tilde{\mathbb{E}}^\alpha \left[\exp \left(-\gamma \tilde{P}_{t,T}^{\alpha,n} \right) \right], \text{ for all } \alpha \in \mathcal{A}_t^2. \quad (3.1.8)$$

Using the explicit exponential formula (see (A.1.1)) and by straightforward computations:

$$\begin{aligned} \exp \left(-\gamma \overline{P}_{t,T}^{\alpha,n} \right) &= \exp \left(-\gamma \left(\int_t^T \left\{ \mu N_u^{t,n} - \frac{1}{2} \sigma^2 (\gamma + \zeta) (N_u^{t,n})^2 \right\} du - \ell(N_T^{t,n}) \right) \right) \\ &\quad \times A_T^\alpha \exp \left(-\gamma \int_t^T \delta_u^- dN_u^- + \delta_u^+ dN_u^+ \right) \\ &= \exp \left(-\gamma \tilde{P}_{t,T}^{\alpha,n} \right) A_T^\alpha \exp \left(\gamma \int_t^T \left\{ U_\gamma(\delta_u^-) \widehat{\lambda}_u^{-,\alpha,t,n,\pi} + U_\gamma(\delta_u^+) \widehat{\lambda}_u^{+,\alpha,t,n,\pi} \right\} du \right) \\ &\quad \times \prod_{\substack{t \leq u \leq T: \\ \Delta N_u^- \neq 0}} \exp \left(-\gamma \delta_u^- \right) \prod_{\substack{t \leq u \leq T: \\ \Delta N_u^+ \neq 0}} \exp \left(-\gamma \delta_u^+ \right) \\ &= \exp \left(-\gamma \tilde{P}_{t,T}^{\alpha,n} \right) A_T^\alpha B_T^\alpha, \end{aligned}$$

which yields (3.1.8) after taking \mathbb{P}^α -expectation.

It remains to see that

$$\Psi(t, n, \pi) := \sup_{\alpha \in \mathcal{A}_t^2} \tilde{\mathbb{E}}^\alpha \left[U_\gamma \left(\tilde{P}_{t,T}^{\alpha,n,\pi} \right) \right] = \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \tilde{\mathbb{E}}^\alpha \left[U_\gamma \left(\tilde{P}_{t,T}^{\alpha,n,\pi} \right) \right] =: \tilde{\Psi}(t, n, \pi).$$

Clearly $\Psi \geq \tilde{\Psi}$. Let us check $\Psi \leq \tilde{\Psi}$. As done in Proposition 1.4.2 we can define a family of “reference” equivalent probability measures $\tilde{\mathbb{Q}}^\alpha \sim \tilde{\mathbb{P}}^\alpha$, $\alpha \in \mathcal{A}_t^2$, such that for $(\tilde{\mathbb{Q}}^\alpha, \mathbb{F}^{t,W,N})$ it holds: $W - W_t$ is a Wiener process independent of the counting process $(N^- - N_t^-, N^+ - N_t^+)$ and $N^\pm - N_t^\pm$ has predictable intensity $\mathbb{1}_u^{\pm,t,n} := \mathbb{1}_{\{\mp N_u^{t,n} < N^*\}}$ (in particular, its law does not depend on α). Furthermore, the inverse change of measure is given by $d\tilde{\mathbb{P}}^\alpha/d\tilde{\mathbb{Q}}^\alpha = \tilde{Z}_T^\alpha$ with

$$\tilde{Z}^\alpha := \mathcal{E} \left(\int_t^\cdot \left(\widehat{\Lambda}^-(\Pi_{u^-}^{\alpha,t,n,\pi}, \delta_u^-) - 1 \right) (dN_u^- - \mathbb{1}_u^{-,t,n} du) + \left(\widehat{\Lambda}^+(\Pi_{u^-}^{\alpha,t,n,\pi}, \delta_u^+) - 1 \right) (dN_u^+ - \mathbb{1}_u^{+,t,n} du) \right).$$

Let us fix $\alpha \in \mathcal{A}_t^2$. Denote by $\mathcal{D} = \mathcal{D}([t, T], \mathbb{R})$ the Skorokhod space of càdlàg functions with its usual sigma algebra and by $\mathbb{P}^W, \mathbb{P}^N$ the laws (or pushforward measures) induced on \mathcal{D} by $W - W_t, N^{t,n}$ resp. when starting from \mathbb{Q}^α . These laws do not depend on α and characterize the joint law of $(W - W_t, N^{t,n})$ on \mathcal{D}^2 as $\mathbb{P}^W \otimes \mathbb{P}^N$, due to the independence of the two processes. Since α is $\mathbb{F}^{t,W,N}$ -predictable, by a monotone class argument one can show there exists a jointly measurable process $f : [t, T] \times \mathcal{D}^2 \rightarrow \mathbb{R}$ such that $\alpha_u = f_u(W - W_t, N^{t,n})$ and for \mathbb{P}^W -almost every $w \in \mathcal{D}$, the process $\tilde{\alpha}_u := f_u(w, N^{t,n})$ is in $\tilde{\mathcal{A}}_t^2$. Note also that we can write $\tilde{Z}_T^\alpha U_\gamma(\tilde{P}_{t,T}^{\alpha,n,\pi}) = g(\alpha, N^{t,n}) = g(f(W - W_t, N^{t,n}), N^{t,n})$ for some function g . By Fubini's theorem,

$$\tilde{\mathbb{E}}^\alpha \left[U_\gamma(\tilde{P}_{t,T}^{\alpha,n,\pi}) \right] = \tilde{\mathbb{E}}^{\mathbb{Q}^\alpha} \left[\tilde{Z}_T^\alpha U_\gamma(\tilde{P}_{t,T}^{\alpha,n,\pi}) \right] = \int \mathbb{E}^{\mathbb{P}^N} [g(f(w, \cdot), \cdot)] d\mathbb{P}^W(w) \leq \int \tilde{\Psi} d\mathbb{P}^W(w) = \tilde{\Psi}.$$

Since $\alpha \in \mathcal{A}_t^2$ was arbitrary, we conclude that $\Psi = \tilde{\Psi}$. \square

Just as it occurs under full information (see Section 4), for a fully risk-neutral MM with negligible costs (i.e. $N^* = +\infty$, $\gamma = \zeta = 0 \equiv \ell$), Θ can be further decomposed. Note, from their definition, that in this case $\widehat{\lambda}^{\pm, \alpha, t, n, \pi}$ and $\Pi^{\alpha, t, n, \pi}$ do not depend on n . As it was proved in Theorem 3.1.1, when $\gamma = 0$ the family $(\tilde{\mathbb{P}}^{\alpha, t, n})$ can be taken as the original physical probabilities, and these do not depend on n either. The following corollary is now immediate.

Corollary 3.1.2. *If $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$ then*

$$V(t, s, x, n, \pi) = x + sn + \mu n(T - t) + \Phi(t, \pi),$$

$$\text{with } \Phi(t, \pi) = \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \mathbb{E}^{t, \pi} \left[\int_t^T \left\{ \delta_u^- \widehat{\lambda}_u^{-, \alpha, t, \pi} + \delta_u^+ \widehat{\lambda}_u^{+, \alpha, t, \pi} + \mu(\widehat{\lambda}_u^{-, \alpha, t, \pi} - \widehat{\lambda}_u^{+, \alpha, t, \pi}) \right\} du \right].$$

In the context of the previous corollary, formal substitution in (3.1.3) or (3.1.5) yields the following PIDE for Φ :

$$\begin{aligned} 0 &= \phi_t + \sum_{i,j=1}^k q_t^{ji} \pi^j \phi_{\pi^i} \\ &+ \sup_{\delta^- \in (\delta_*, \delta^*)} \left\{ \sum_{i,j=1}^k (\Lambda_j^- - \Lambda_i^-) (\delta^-) \pi^j \pi^i \phi_{\pi^i} + (\delta^- + \mu(T - t) + d_{\delta^-}^-(\phi)) \sum_{i=1}^k \pi^i \Lambda_i^-(\delta^-) \right\} \\ &+ \sup_{\delta^+ \in (\delta_*, \delta^*)} \left\{ \sum_{i,j=1}^k (\Lambda_j^+ - \Lambda_i^+) (\delta^+) \pi^j \pi^i \phi_{\pi^i} + (\delta^+ - \mu(T - t) + d_{\delta^+}^+(\phi)) \sum_{i=1}^k \pi^i \Lambda_i^+(\delta^+) \right\}, \end{aligned} \quad (3.1.9)$$

with terminal condition $\phi(T, \pi) = 0$, where:

$$d_{\delta^\pm}^\pm(\phi)(t, \pi) = \phi\left(t, \frac{1}{\sum_{j=1}^k \pi^j \Lambda_j^\pm(\delta)} (\pi^1 \Lambda_1^\pm(\delta), \dots, \pi^k \Lambda_k^\pm(\delta))\right) - \phi(t, \pi).$$

Remark 3.1.3. Besides its theoretical motivation, Theorem 3.1.1 is interesting in itself. It enables the use of probabilistic and PDMPs numerical techniques when the dimension of the problem renders PDEs methods prohibitive. Furthermore, as readily seen in the proof, the decomposition of the expected P&L utility holds true for any admissible strategy. This means, in particular, that one can efficiently simulate the proceeds of any strategy without the need to simulate the diffusive state variables (see Section 5.2).

We now want to prove that Θ (resp. Φ) is the unique continuous viscosity solution of the terminal condition PIDE (3.1.5) (resp. (3.1.9)). (See, e.g., [Son88, Def.2.1] for the relevant definition, or more in general [DF99, Def.7.3], recalling that in our case we have no boundary

conditions other than that at terminal time.) A complication inevitably arises, as classical viscosity techniques [Bou07, FS06, ØS09] cannot be applied directly to weak formulation models such as ours. However, the decomposition of Theorem 3.1.1 (resp. Corollary 3.1.2) not only reduces the dimensionality of the problem, but also states that the MM may neglect the diffusion component of the state process altogether, focusing solely on the time-space state variable $(u, N^{t,n}, \Pi^{\alpha,t,n,\pi})$ (resp. $(u, \Pi^{\alpha,t,n,\pi})$). This is a PDMP as introduced in [Dav84] (detailed treatments also found in [BR11, Dav93]). Using results from PDMPs theory, our continuous-time problem is identified with a control problem for a discrete-time Markov decision model, as in [BR09, BR10, BR11, CEFS16a], and linked again with viscosity solutions of HJB PIDEs as in [CEFS16a, DF99].

An inevitable drawback is that the PDMPs approach relies on the use of the so-called *randomized* (or *relaxed*) controls and requires the control space to be compact. Hence, for the following theorem we will assume $-\infty < \delta_* < \delta^* < +\infty$. Assuming a uniform lower constraint $\delta_* > -\infty$ is hardly a problem. On the contrary, $\delta_* = 0$ (or even some small positive number) is the most meaningful in practice, as negative spreads imply the MM is willing to offer her clients better prices than the reference price S . ($\delta_* = -\infty$ is motivated in the literature by mathematical convenience rather than modelling accuracy.) A uniform upper bound $\delta^* < +\infty$, on the other hand, is harder to assess a priori. Fortunately, in most situations encountered in practice, the unconstrained optimization will yield bounded optimal spreads nonetheless, and the MM can dispense with δ_*, δ^* if she wishes to do so (see Section 5 for an example).

Theorem 3.1.4. *Assume $-\infty < \delta_* < \delta^* < \infty$. For $N^* < \infty$ (resp. $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$), let Θ (resp. Φ) be as in Theorem 3.1.1 (resp. Corollary 3.1.2). Then Θ (resp. Φ) is the unique continuous viscosity solution of the terminal condition PIDE (3.1.5) (resp. (3.1.9)).*

Proof. See Appendix A. □

4 Full information

In this section we consider the idealized case of a MM with full information. We assume the MM has inside information in such a way that she can observe the full filtration \mathbb{F} , and in particular she can observe Y . For example, if Y represents different levels of competition amongst liquidity providers, this would practically mean the MM has information regarding her competitors' quotes. We will see that in this case the value function turns out to be a regular, classical solution, of its HJB equation. Afterwards, we will compare the results with those obtained in the more realistic setting of partial information.

4.1 Dimensionality reduction and a general system of ODEs

We consider problem (1.2.6) but under full information, with the set of admissible spreads:

$$\mathcal{U} := \{\delta : [0, T] \rightarrow \overline{(\delta_*, \delta^*)} : \delta \text{ is } \mathbb{F}\text{-predictable and bounded from below}\}.$$

Note that in this section, and when $\delta^* = +\infty$, we do not assume upper-boundedness of the spreads. Let $t \in [0, T]$, $s, x \in \mathbb{R}$ and $(n, i) \in \mathcal{I} := (\mathbb{Z} \cap [-N^*, N^*]) \times \{1, \dots, k\}$. Consider the processes $S^{t,s}, X^{\alpha,t,x,s}, N^{t,n}, P^{\alpha,s,x,n}$ as defined in Section 3 and $Y^{t,i}$ a Markov chain with deterministic generator matrix Q , state space $\{1, \dots, k\}$ and such that $Y_t^{t,i} = i$. We assume the physical probabilities $\mathbb{P}^{\alpha,t,n,i}$ are defined for every $\alpha \in \mathcal{U}^2$ and that Assumptions 1.2.1, 1.2.2 and (1.2.2) (starting at time t^-) are still in place. The value function in this case is

$$V(t, s, x, n, i) := \sup_{\alpha \in \mathcal{U}_t^2} \mathbb{E}^{t,n,i} \left[U_\gamma(P_{t,T}^{\alpha,s,x,n}) \right], \quad (4.1.1)$$

By means of the Dynamic Programming Principle and Ito's Lemma one can formally derive the following HJB PIDE for V :

$$\begin{aligned}
0 &= v_t(t, s, x, n, i) + \mu v_s(t, s, x, n, i) + \frac{1}{2} \sigma^2 v_{ss}(t, s, x, n, i) \\
&+ \frac{1}{2} \sigma^2 \zeta n^2 (\gamma v(t, s, x, n, i) - 1) + \sum_{j=1}^k q_t^{ij} v(t, s, x, n, j) \\
&+ \mathbb{1}_{\{n < N^*\}} \sup_{\delta^- \in (\overline{\delta_*}, \delta^*)} \Lambda_i^-(\delta^-) \left(v(t, s, x - (s - \delta^-), n + 1, i) - v(t, s, x, n, i) \right) \\
&+ \mathbb{1}_{\{-n < N^*\}} \sup_{\delta^+ \in (\overline{\delta_*}, \delta^*)} \Lambda_i^+(\delta^+) \left(v(t, s, x + (s + \delta^+), n - 1, i) - v(t, s, x, n, i) \right),
\end{aligned} \tag{4.1.2}$$

with terminal condition $v(T, s, x, n, i) = U_\gamma(x + sn - \ell(n))$.

In this new context, instead of formally proving a decomposition of V as in Theorem 3.1.1, it is more straightforward to propose an ansatz and ultimately prove it valid with a verification theorem. (This is the standard approach used in the Avellaneda–Stoikov framework.) Let us consider an ansatz for the value function analogous to those used for the one regime case:

$$V(t, s, x, n, i) = U_\gamma(x + sn + \Theta(t, n, i)),$$

for some function $\Theta : [0, T] \times \mathcal{I} \rightarrow \mathbb{R}, C^1$ in time. Substituting in (4.1.2) and using Assumptions 1.2.1, we see that Θ must satisfy a system of ODEs indexed in $(n, i) \in \mathcal{I}$:

$$\begin{aligned}
0 &= \theta_t(t, n, i) + \mu n - \frac{1}{2} \sigma^2 n^2 (\zeta + \gamma) + \sum_{j \neq i} q_t^{ij} U_\gamma(\theta(t, n, j) - \theta(t, n, i)) \\
&+ \mathbb{1}_{\{n < N^*\}} H_i^-(\theta(t, n + 1, i) - \theta(t, n, i)) + \mathbb{1}_{\{-n < N^*\}} H_i^+(\theta(t, n - 1, i) - \theta(t, n, i)),
\end{aligned} \tag{4.1.3}$$

with terminal condition $\theta(T, n, i) = -\ell(n)$, where:

1. $H_i^\pm(d) := \sup_{\delta \in (\overline{\delta_*}, \delta^*)} h_i^\pm(\delta, d)$, for $d \in \mathbb{R}$.
2. $h_i^\pm(\delta, d) := \Lambda_i^\pm(\delta) U_\gamma(\delta + d)$, for $d \in \mathbb{R}, \delta \in (\overline{\delta_*}, \delta^*)$.

The original problem is simplified in this way, both by the dimension of the state variable and by the complexity of the equations, provided we can show that problem (4.1.3) admits a solution. With this aim in mind, let us prove first the following property of the Hamiltonian functions H_1^\pm, \dots, H_k^\pm .

Lemma 4.1.1. *For all $1 \leq i \leq k$, it holds:*

- (i) *For each compact $K \subset \mathbb{R}$, there exists $[a, b] \subseteq (\overline{\delta_*}, \delta^*)$ such that $H_i^\pm(d) = \max_{\delta \in [a, b]} h_i^\pm(\delta, d)$, for all $d \in K$.*
- (ii) *H_i^\pm is locally Lipschitz.*

Proof. Fix $1 \leq i \leq k$ and let $K \subset \mathbb{R}$ be a compact set. We verify first that (i) is a consequence of Assumptions 1.2.2. Let $C > 0$ such that $|d| \leq C$ for all $d \in K$. If $\delta_* = -\infty$, then we can take any $a < \min\{-C, \delta^*\}$ since $h_i^\pm(\delta, d) < h_i^\pm(a, d)$ for all $\delta < a$ and $d \in K$. On the other hand, if $\delta^* = +\infty$, take some $c > \max\{\delta_*, 2C\}$. It holds that $h_i^\pm(c, d) \geq \Lambda_i^\pm(c) U_\gamma(C) =: \varepsilon > 0$ and we can choose $b > c$ such that $h_i^\pm(\delta, d) \leq h_i^\pm(\delta, C) < \varepsilon$ for all $\delta \geq b$ and $d \in K$. Replacing supremum by maximum is now immediate due to the continuity of $h_i^\pm(\cdot, d)$ on $[a, b]$ for all d .

(ii) is routinely verified using that the family $\{h_i^\pm(\delta, \cdot)\}_{\delta \in [a, b]}$ is equi-Lipschitz on K . \square

We want to prove now that the Cauchy problem (4.1.3) admits a unique global classical solution θ which is C^1 in time. To this purpose, we will treat the cases of the finite system ($N^* < \infty$) and infinite system ($N^* = \infty$) of equations separately.

4.2 Constrained inventory ODEs

For $N^* < \infty$ we are dealing with a finite system of ODEs. We know that under certain regularity conditions the Cauchy problem (4.1.3) is guaranteed to have a classical solution on some neighbourhood $(\tau, T] \subset [0, T]$ of T . Nonetheless, it is not always the case that such a local solution can be extended to a global one on $[0, T]$. Following [Gué17, GL15], we start by proving a comparison principle for (4.1.3) that will allow us, in particular, to show the existence of a global solution. The argument used is standard for comparison principles of HJB-type equations.

Proposition 4.2.1 (Comparison Principle). *Let $I \subseteq [0, T]$ be an interval containing T and let $\bar{\theta}, \underline{\theta} : I \times \mathcal{I} \rightarrow \mathbb{R}$ be classical (C^1 with respect to time) super- and sub-solutions resp. of (4.1.3). That is,*

$$\begin{aligned} 0 \leq & -\bar{\theta}_t(t, n, i) - \mu n + \frac{1}{2}\sigma^2 n^2(\zeta + \gamma) - \sum_{j \neq i} q_t^{ij} U_\gamma(\bar{\theta}(t, n, j) - \bar{\theta}(t, n, i)) \\ & - \mathbb{1}_{\{n < N^*\}} H_i^-(\bar{\theta}(t, n+1, i) - \bar{\theta}(t, n, i)) - \mathbb{1}_{\{-n < N^*\}} H_i^+(\bar{\theta}(t, n-1, i) - \bar{\theta}(t, n, i)), \end{aligned} \quad (4.2.1)$$

$$\begin{aligned} 0 \geq & -\underline{\theta}_t(t, n, i) - \mu n + \frac{1}{2}\sigma^2 n^2(\zeta + \gamma) - \sum_{j \neq i} q_t^{ij} U_\gamma(\underline{\theta}(t, n, j) - \underline{\theta}(t, n, i)) \\ & - \mathbb{1}_{\{n < N^*\}} H_i^-(\underline{\theta}(t, n+1, i) - \underline{\theta}(t, n, i)) - \mathbb{1}_{\{-n < N^*\}} H_i^+(\underline{\theta}(t, n-1, i) - \underline{\theta}(t, n, i)) \end{aligned} \quad (4.2.2)$$

and

$$\underline{\theta}(T, \cdot, \cdot) \leq -\ell \leq \bar{\theta}(T, \cdot, \cdot). \quad (4.2.3)$$

Then

$$\underline{\theta} \leq \bar{\theta}.$$

Proof. See Appendix A. □

We can now prove the existence and uniqueness of a classical global solution to the Cauchy problem (4.1.3).

Theorem 4.2.2. *There exists a unique $\Theta : [0, T] \times \mathcal{I} \rightarrow \mathbb{R}$, C^1 in time, which (classically) solves the Cauchy problem (4.1.3).*

Proof. For each $1 \leq i \leq k$, Lemma 4.1.1 tells us that $H_i^\pm : \mathbb{R} \rightarrow \mathbb{R}$ is locally Lipschitz. We also assumed $Q : [0, T] \rightarrow \mathbb{R}^{k \times k}$ continuous (Assumptions 1.2.1). Hence, by the Cauchy-Lipschitz Theorem, the terminal condition system of ODEs (4.1.3) admits a unique C^1 local solution $(\Theta(\cdot, n, i))_{(n, i) \in \mathcal{I}}$, defined on some maximal interval $I \subset [0, T]$ containing T .

Suppose that $I \subsetneq [0, T]$. Then $I = (\tau, T]$ for some $0 \leq \tau < T$ and $\|\Theta(t)\| \rightarrow \infty$ as $t \searrow \tau$. We claim that Θ is actually bounded, resulting in a contradiction. Indeed, take

$$\bar{K} = \max_{(i, n) \in \mathcal{I}} \left\{ \left| \mu n - \frac{1}{2}\sigma^2 n^2(\zeta + \gamma) + \mathbb{1}_{\{n < N^*\}} H_i^-(0) + \mathbb{1}_{\{-n < N^*\}} H_i^+(0) \right| \right\}$$

and

$$\underline{K} = \ell(N^*),$$

and define $\bar{\theta} : I \times \mathcal{I} \rightarrow \mathbb{R}$ by $\bar{\theta}(t, n, i) = \bar{K}(T - t)$ and $\underline{\theta} := -\bar{\theta} - \underline{K}$. Then $\bar{\theta}$ (resp. $\underline{\theta}$) is a super- (resp. sub-) solution of the Cauchy problem (4.1.3). By the Comparison Principle (4.2.1),

$$-\bar{K}T - \underline{K} \leq \underline{\theta} \leq \Theta \leq \bar{\theta} \leq \bar{K}T,$$

proving that Θ is bounded. □

4.3 Unconstrained inventory ODEs

We consider now $N^* = +\infty$, for which (4.1.3) becomes an infinite system of ODEs. Recall that in this case we assumed $\gamma = \zeta = 0 \equiv \ell$, i.e., the MM is fully risk-neutral and has negligible costs. This allows us to further reduce the dimensionality of the state variable by the additional ansatz

$$\Theta(t, n, i) = \mu n(T - t) + \Phi_i(t), \quad (4.3.1)$$

for some $\Phi = (\Phi_i)_{i=1}^k \in C^1([0, T], \mathbb{R}^k)$. Substituting in (4.1.3), we get that Φ must solve the finite linear system of ODEs

$$\begin{cases} \phi'(t) = -Q(t)\phi(t) + b(t) \\ \phi(T) = 0, \end{cases} \quad (4.3.2)$$

with $b_i(t) = -H_i^-(\mu(T - t)) - H_i^+(-\mu(T - t))$, $i = 1, \dots, k$. By continuity of Q and b (see Lemma 4.1.1), the previous system is known to have a unique global solution $\Phi \in C^1([0, T], \mathbb{R}^k)$ which can be computed by the variation of parameters method. Straightforward verification now gives the following:

Proposition 4.3.1. *Let $\Phi \in C^1([0, T], \mathbb{R}^k)$ be the unique solution of (4.3.2). Then $\Theta : [0, T] \times \mathcal{I} \rightarrow \mathbb{R}$ such that $\Theta(t, n, i) = \mu n(T - t) + \Phi_i(t)$, is the unique, C^1 in time, (classical) solution of the Cauchy problem (4.1.3) with $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$.*

4.4 General Verification Theorem

We give now the complete solution for the general model under full information. For the next theorem we note that given the function Θ defined in Theorem 4.2.2 for $N^* < +\infty$ (resp. Proposition 4.3.1 for $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$) the difference or ‘‘jump’’ terms of equation (4.1.3) are bounded by continuity. That is, $\Theta(t, n \mp 1, i) - \Theta(t, n, i)$ is bounded on $[0, T] \times \mathcal{I}$ for $N^* < +\infty$ (resp. $\Theta(t, n \mp 1, i) - \Theta(t, n, i) = \mp \mu(T - t)$ is bounded on $[0, T]$).

Theorem 4.4.1 (Verification Theorem). *Let Θ be as in Theorem 4.2.2 for $N^* < +\infty$ and as in Proposition 4.3.1 for $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$. Then the value function in (4.1.1) is*

$$V(t, s, x, n, i) = U_\gamma(x + sn + \Theta(t, n, i)).$$

Furthermore, given $C > 0$ such that $|\Theta(t, n \mp 1, i) - \Theta(t, n, i)| \leq C$ for all $(t, n, i) \in [0, T] \times \mathcal{I}$ with $-N^* \leq n \mp 1 \leq N^*$, there exist Borel measurable functions $\bar{\delta}_1^\pm, \dots, \bar{\delta}_k^\pm : [-C, C] \rightarrow (\bar{\delta}_*, \bar{\delta}^*)$ such that $\bar{\delta}_i^\pm(d) \in \arg \max_{\delta \in (\bar{\delta}_*, \bar{\delta}^*)} h_i^\pm(\delta, d)$ for all $d \in [-C, C]$, $1 \leq i \leq k$; and for any such functions, the strategy $(\bar{\delta}_u^+, \bar{\delta}_u^-)$, with

$$\bar{\delta}_u^\pm := \bar{\delta}_{Y_{u^-}^{t,i}}^\pm \left(\Theta(u, N_{u^-}^{t,n} \mp 1, Y_{u^-}^{t,i}) - \Theta(u, N_{u^-}^{t,n}, Y_{u^-}^{t,i}) \right),$$

is optimal.

Notation. Note that we make a slight abuse of notation, writing $\bar{\delta}_i^\pm$ for the Borel functions in the theorem and $(\bar{\delta}_u^\pm)$ for the spread processes.

Proof. Let $1 \leq i \leq k$. We check first that we can choose a maximizer of $h_i^\pm(\cdot, d)$ in a measurable way with respect to d . Lemma 4.1.1 tells us that there exists $[a_i, b_i] \subseteq (\bar{\delta}_*, \bar{\delta}^*)$ such that $h_i^\pm(\delta, d)$ attains its maximum in $[a_i, b_i]$ for all $d \in [-C, C]$. Due to Assumptions 1.2.2 (ii), h_i^\pm is continuous and, in particular, a Carathéodory function [AB06, Def.4.50]. Thus, the Measurable Maximum Theorem [AB06, Thm.18.19] guarantees the existence of a Borel selector

of the arg max, $\overline{\delta_i^\pm} : [-C, C] \rightarrow [a_i, b_i] \subseteq \overline{(\delta_*, \delta^*)}$. (Note that the weak measurability assumption is trivially verified in this case.)

From here onwards, let $\overline{\delta_1^\pm}, \dots, \overline{\delta_k^\pm} : [-C, C] \rightarrow \overline{(\delta_*, \delta^*)}$ be some Borel selectors as above. By Lemma 4.1.1 again, these functions must be bounded from below, and the spread processes $(\overline{\delta_u^\pm})$ defined as in the theorem are clearly admissible.

We define now $\tilde{V}(t, s, x, n, i) := U_\gamma(x + ns + \Theta(t, n, i))$ and we want to show that $\tilde{V} = V$. Let us fix the initial time and values, $t \in [0, T]$, $s, x \in \mathbb{R}$ and $(n, i) \in \mathcal{I}$, and consider an arbitrary strategy $\alpha = (\delta_u^-, \delta_u^+) \in \mathcal{U}_t^2$. For shortness, we omit these initial conditions and strategy from the notation of the processes and the expectation. We denote by $\mathcal{S} := \{j - l : 1 \leq l, j \leq k, l \neq j\}$, the set of jump heights of Y , and

$$R_{t,v} := -\frac{1}{2}\sigma^2\zeta \int_t^v N_u^2 du, \quad Z_{t,v} := \exp(-\gamma R_{t,v}).$$

By the identity $U_\gamma(a+b) = U_\gamma(a)e^{-\gamma b} + U_\gamma(b)$, we can rewrite the utility of the MM's penalized P&L as

$$\begin{aligned} U_\gamma(P_{t,T}) &= U_\gamma(X_T + S_T N_T - \ell(N_T)) Z_{t,T} + U_\gamma(R_{t,T}) \\ &= \tilde{V}(T, S_T, X_T, N_T, Y_T) Z_{t,T} + U_\gamma(R_{t,T}). \end{aligned} \quad (4.4.1)$$

Using integration by parts and Ito's Lemma (recalling (1.2.2)), we re-express the last two terms as

$$U_\gamma(R_{t,T}) = -\frac{1}{2}\sigma^2\zeta \int_t^T Z_{t,u} N_u^2 du$$

and

$$\begin{aligned} &\tilde{V}(T, S_T, X_T, N_T, Y_T) Z_{t,T} \\ &= \tilde{V}(t, s, x, n, i) + \int_t^T \tilde{V}(u, S_u, X_u, N_u, Y_u) dZ_{t,u} + \int_t^T Z_{t,u} d\tilde{V}(u, S_u, X_u, N_u, Y_u) \\ &= \tilde{V}(t, s, x, n, i) + \frac{1}{2}\sigma^2\zeta\gamma \int_t^T V(u, S_u, X_u, N_u, Y_u) Z_{t,u} N_u^2 du + \int_t^T Z_{t,u} \tilde{V}_t(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du \\ &+ \mu \int_t^T Z_{t,u} \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du + \sigma \int_t^T Z_{t,u} \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) dW_u \\ &+ \frac{1}{2}\sigma^2 \int_t^T Z_{t,u} \tilde{V}_{ss}(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du \\ &+ \int_t^T Z_{t,u} (\tilde{V}(u, S_u, X_{u-} - (S_u - \delta_u^-), N_{u-} + 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) dN_{u-}^- \\ &+ \int_t^T Z_{t,u} (\tilde{V}(u, S_u, X_{u-} + S_u + \delta_u^+, N_{u-} - 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) dN_{u-}^+ \\ &+ \int_t^T Z_{t,u} \int_{\mathcal{S}} (\tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-} + h) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) \mu^Y(dt, dh), \end{aligned}$$

where μ^Y is the jump measure of Y . Substituting in (4.4.1),

$$\begin{aligned}
U_\gamma(P_{t,T}) &= \tilde{V}(t, s, x, n, i) + \sigma \int_t^T Z_{t,u} \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) dW_u \\
&+ \int_t^T Z_{t,u} \left\{ \tilde{V}_t(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du + \mu \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du \right. \\
&+ \frac{1}{2} \sigma^2 \tilde{V}_{ss}(u, S_u, X_{u-}, N_{u-}, Y_{u-}) du + \frac{1}{2} \sigma^2 \zeta N_u^2 (\gamma V(u, S_u, X_u, N_u, Y_u) - 1) du \\
&+ \int_S (\tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-} + h) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) \mu^Y(dt, dh) \\
&+ (\tilde{V}(u, S_u, X_{u-} - (S_u - \delta_u^-), N_{u-} + 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) dN_u^- \\
&\left. + (\tilde{V}(u, S_u, X_{u-} + S_u + \delta_u^+, N_{u-} - 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) dN_u^+ \right\}. \tag{4.4.2}
\end{aligned}$$

Next, we want to verify that the process $\left(\int_t^v Z_{t,u} \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) dW_u \right)_{t \leq v \leq T}$ is a zero mean martingale. Since $Z_{t,u}$ is bounded (either $N^* < \infty$ or $\gamma = 0$) it suffices to check

$$\int_t^T \mathbb{E} \left[\tilde{V}_s(u, S_u, X_u, N_u, Y_u)^2 \right] du < +\infty. \tag{4.4.3}$$

If $\gamma = 0$,

$$\tilde{V}_s(u, S_u, X_u, N_u, Y_u)^2 = N_u^2$$

and (4.4.3) is a consequence of (1.2.5).

If $\gamma \neq 0$ (hence $N^* < \infty$), let C be a lower bound for (δ_u^-) and (δ_u^+) . Integration by parts and Hölder inequality yield

$$\begin{aligned}
\mathbb{E} \left[\tilde{V}_s(u, S_u, X_u, N_u, Y_u)^2 \right] &= \gamma^2 \mathbb{E} \left[\exp(-2\gamma(X_u + S_u N_u + \theta(u, N_u, Y_u))) N_u^2 \right] \\
&= \gamma^2 \exp(-2\gamma(x + sn)) \\
&\times \mathbb{E} \left[\exp \left(-2\gamma \int_t^u \delta_w^- dN_w^- + \delta_w^+ dN_w^+ + \mu N_w dw + \sigma N_w dW_w + \theta(w, N_w, Y_w) \right) N_u^2 \right] \\
&\leq \gamma^2 N^{*2} \exp \left(2\gamma(x + sn + \mu N^*(T-t) + \|\theta\|_{L^\infty([t,T] \times \mathcal{I})} + 2\sigma^2 \gamma N^{*2}(T-t)) \right) \\
&\times \mathbb{E} \left[\exp \left(-4\gamma C(N_T^- + N_T^+ - N_{t^-}^- - N_{t^-}^+) \right) \right]^{\frac{1}{2}} \mathbb{E} \left[\exp \left(-4\gamma \sigma \int_t^T N_u dW_u - 8\gamma^2 \sigma^2 \int_t^T N_u^2 du \right) \right]^{\frac{1}{2}}
\end{aligned}$$

and (4.4.3) is again consequence of (1.2.5) and Novikov's condition (trivially satisfied).

In the same way, recalling that $(Z_{t,u})$ is bounded, (δ_u^\pm) is bounded from below and that Q is bounded by continuity, one can check that

$$\int_t^T \mathbb{E} \left[Z_{t,u} |\tilde{V}(u, S_u, X_{u-} \pm (S_u \pm \delta_u^\pm), N_{u-} \mp 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})| \lambda_u^\pm \right] du < +\infty$$

and

$$\int_t^T \mathbb{E} \left[Z_{t,u} \sum_{j \neq Y_{u-}} q_t^{Y_{u-}, j} |\tilde{V}(u, S_u, X_{u-}, N_{u-}, j) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})| \right] du < +\infty.$$

Taking expectation in (4.4.2), the Brownian term vanishes and integration with respect to dN^- , dN^+ and $d\mu^Y$ is replaced by integration with respect to their dual predictable projections (see, e.g., [Bré81, p.27 T8 and p.235 C4]). That is,

$$\begin{aligned}
\mathbb{E} \left[U_\gamma(P_{t,T}) \right] &= \tilde{V}(t, s, x, n, i) + \int_t^T \mathbb{E} \left[Z_{t,u} \left\{ \tilde{V}_t(u, S_u, X_{u-}, N_{u-}, Y_{u-}) \right. \right. \\
&+ \mu \tilde{V}_s(u, S_u, X_{u-}, N_{u-}, Y_{u-}) + \frac{1}{2} \sigma^2 \tilde{V}_{ss}(u, S_u, X_{u-}, N_{u-}, Y_{u-}) \\
&+ \frac{1}{2} \sigma^2 \zeta N_u^2 (\gamma \tilde{V}(u, S_u, X_u, N_u, Y_u) - 1) + \sum_{j=1}^k q_t^{Y_{u-}-j} \tilde{V}(u, S_u, X_{u-}, N_{u-}, j) \\
&+ \mathbb{1}_u^- \Lambda_{Y_{u-}}^- (\delta_u^-) (\tilde{V}(u, S_u, X_{u-} - (S_u - \delta_u^-), N_{u-} + 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) \\
&+ \left. \left. \mathbb{1}_u^+ \Lambda_{Y_{u-}}^+ (\delta_u^+) (\tilde{V}(u, S_u, X_{u-} + S_u + \delta_u^+, N_{u-} - 1, Y_{u-}) - \tilde{V}(u, S_u, X_{u-}, N_{u-}, Y_{u-})) \right\} \right] du \\
&\leq \tilde{V}(t, s, x, n, i),
\end{aligned}$$

as \tilde{V} solves (4.1.2), with the equality attained for $(\delta_u^-) = (\overline{\delta_u^-})$ and $(\delta_u^+) = (\overline{\delta_u^+})$ by definition. We conclude that $V = \tilde{V}$ and that the pair $(\overline{\delta_u^-}, \overline{\delta_u^+})$ is optimal. \square

Remark. For $N^* < \infty$, since the MM will not buy (resp. sell) whenever (N_{u-}) hits N^* (resp. $-N^*$), the value of $(\overline{\delta_u^-})$ (resp. $(\overline{\delta_u^+})$) at these stopping times is essentially irrelevant. From a strict mathematical perspective, the only constraint is that whichever the value we choose, the process needs to remain admissible.

4.5 Computing the optimal spreads: some particular cases

We have shown in Theorem 4.4.1 that the optimal spreads for the full information problem (4.1.1) can be computed in feedback form in terms of $(t, n, i) = (t, N_{t-}, Y_{t-})$ at each time $t \in [0, T]$. Practically, this means finding Θ (typically, numerically) that solves the terminal condition system of ODEs (4.1.3), and finding the spreads by maximization:

$$\overline{\delta}^\pm(t, n, i) \in \arg \max_{\delta \in (\delta_*, \delta^*)} \Lambda_i^\pm(\delta) U_\gamma(\delta + \Theta(t, n \mp 1, i) - \Theta(t, n, i)). \quad (4.5.1)$$

(See the proof of Lemma 4.1.1 to see how to reduce the maximization to a compact domain in the cases of $\delta_* = -\infty$ or $\delta^* = +\infty$.) The functions in (4.5.1) may admit multiple maximizers in general. In [BL14, Gué17, GL15] stronger assumptions are imposed on the orders intensities which guarantee, in particular, the uniqueness of the maximizers. We now extend these assumptions to our context and give the corresponding characterization of the spreads. Henceforth, Assumption 1.2.2 (ii) is replaced by the following:

Assumptions 4.5.1. $\Lambda_i^\pm \in C^2(\overline{(\delta_*, \delta^*)})$, $\Lambda_i^{\pm'} < 0$ and $\Lambda_i^\pm \Lambda_i^{\pm''} < c(\Lambda_i^{\pm'})^2$ for some $0 < c < 2$, for all $1 \leq i \leq k$.¹³

Remark. The intensity functions of Examples 1.2.4 1, 2 and 3 all verify these stronger assumptions, while Example 4 only does so for $d_i^\pm > 1$. As a matter of fact, by solving the differential inequality in Assumptions 4.5.1, one sees that any Λ_i^\pm within this new framework has a strict upper bound of the form of Example 4. In particular, for $\delta^* = +\infty$, $\lim_{\delta \rightarrow +\infty} \delta \Lambda_i^\pm(\delta) = 0$ is necessarily satisfied.

Under Assumptions 4.5.1, the maximization of (4.5.1) (for fixed (t, n, i)) can be replaced essentially by solving a contractive fixed point equation and flooring and capping the results at

¹³Lateral derivatives are considered on the domain border. The strict inequality $c < 2$ is needed when $\gamma = 0$ and spreads are unconstrained.

δ_* and δ^* respectively. To make this precise, even when the intensity functions are not defined beyond $(\overline{\delta_*}, \overline{\delta^*})$, let us set

$$d_{i*}^{\pm} := \begin{cases} -U_{\gamma}^{-1} \left(\frac{\Lambda_i^{\pm}(\delta_*)}{\Lambda_i^{\pm'}(\delta_*)} \right) - \delta_* & \text{if } \delta_* > -\infty \\ +\infty & \text{if } \delta_* = -\infty, \end{cases} \quad d_i^{\pm*} := \begin{cases} -U_{\gamma}^{-1} \left(\frac{\Lambda_i^{\pm}(\delta^*)}{\Lambda_i^{\pm'}(\delta^*)} \right) - \delta^* & \text{if } \delta^* < +\infty \\ -\infty & \text{if } \delta^* = +\infty. \end{cases}$$

The following result shows how, up to constraints, the optimal spreads are given by a first term that maximizes the “expected” utility of the MM’s instantaneous margin (see Remark 1.2.3), plus an additional risk adjustment taking into account the inventory held and the prospect of the market shifting. Note how the first term depends on $\Lambda_i^{\pm'}/\Lambda_i^{\pm}$, the percentage sensitivity of the liquidity to spread changes.

Proposition 4.5.2. *Under Assumptions 4.5.1, the functions $\overline{\delta}_1^{\pm}, \dots, \overline{\delta}_k^{\pm}$ of Theorem 4.4.1 are uniquely characterized by*

$$\overline{\delta}_i^{\pm}(d) = \delta^* \text{ if } d < d_{i*}^{\pm*}, \quad \overline{\delta}_i^{\pm}(d) = \delta_* \text{ if } d > d_{i*}^{\pm} \quad (4.5.2)$$

and if $d \in (\overline{d_{i*}^{\pm*}}, \overline{d_{i*}^{\pm}})$, then $\overline{\delta}_i^{\pm}(d)$ is the unique solution of the fixed point equation

$$\overline{\delta}_i^{\pm}(d) = -U_{\gamma}^{-1} \left(\frac{\Lambda_i^{\pm}(\overline{\delta}_i^{\pm}(d))}{\Lambda_i^{\pm'}(\overline{\delta}_i^{\pm}(d))} \right) - d. \quad (4.5.3)$$

Additionally, the Hamiltonians of equation (4.1.3) can be expressed as

$$H_i^{\pm}(d) = \begin{cases} h_i^{\pm}(\delta^*, d) & \text{if } d \leq d_{i*}^{\pm*} \\ -\frac{\Lambda_i^{\pm 2}(\widehat{\delta}^{\pm}(d))}{\Lambda_i^{\pm'}(\widehat{\delta}^{\pm}(d)) - \gamma \Lambda_i^{\pm}(\widehat{\delta}^{\pm}(d))} & \text{if } d_{i*}^{\pm*} < d < d_{i*}^{\pm} \\ h_i^{\pm}(\delta_*, d) & \text{if } d \geq d_{i*}^{\pm}. \end{cases}$$

Proof. For each $1 \leq i \leq k$, define

$$f_i^{\pm}(\delta, d) = \delta + U_{\gamma}^{-1} \left(\frac{\Lambda_i^{\pm}(\delta)}{\Lambda_i^{\pm'}(\delta)} \right) + d, \quad \text{with } \delta \in (\overline{\delta_*}, \overline{\delta^*}), \quad d \in \mathbb{R}.$$

By Assumptions 4.5.1 and straightforward computations (see computations, e.g., in [Gué17, Lemma 3.1]), one verifies that $\text{sgn} \left(\frac{\partial h_i^{\pm}}{\partial \delta} \right) = -\text{sgn}(f_i^{\pm})$ and f_i^{\pm} is strictly increasing on each variable.¹⁴ It follows that for any $d < d_{i*}^{\pm*}$ and $\delta < \delta^*$ (resp. $d > d_{i*}^{\pm}$ and $\delta > \delta_*$), $f_i^{\pm}(\delta, d) < f_i^{\pm}(\delta^*, d_{i*}^{\pm*}) = 0$ (resp. $f_i^{\pm}(\delta, d) > f_i^{\pm}(\delta_*, d_{i*}^{\pm}) = 0$), which proves (4.5.2).

In the same way, if $d \in (d_{i*}^{\pm*}, d_{i*}^{\pm})$ and $-\infty < \delta_* < \delta^* < +\infty$, then $f_i^{\pm}(\delta_*, d) < 0$ and $f_i^{\pm}(\delta^*, d) > 0$. Hence, the continuity of $f_i^{\pm}(\cdot, d)$ implies (4.5.3) (the uniqueness of the solution being due to the strict monotonicity of $f_i^{\pm}(\cdot, d)$). The cases of unconstrained spreads are proved as in [Gué17].

Lastly, the cases $d = d_{i*}^{\pm*} > -\infty$ and $d = d_{i*}^{\pm} < +\infty$ follow in the same manner; and the new expressions for the Hamiltonians are immediate, putting $H_i^{\pm}(d) = h_i^{\pm}(\overline{\delta}_i^{\pm}(d), d)$. \square

As done in [Gué17, Lemma 3.1], (4.5.3) can be replaced by the explicit formula

$$\overline{\delta}_i^{\pm}(d) = (\Lambda_i^{\pm})^{-1} \left(\gamma H_i^{\pm}(d) - H_i^{\pm'}(d) \right),$$

¹⁴ sgn denotes the sign function, with $\text{sgn}(0) := 0$.

but the computation of $(\Lambda_i^\pm)^{-1}$, H_i^\pm and $H_i^{\pm'}$ still has to be carried out numerically, in general. We state now a few simplifications that arise in some particular cases, i.e., when the intensities are exponential or when $N^* = \infty$, $\gamma = \zeta = 0 \equiv \ell$. These are all derived by straightforward substitution. We refer to [Gué17, Sect.4] for some asymptotic approximations when $t \ll T$ in the one regime case, with unconstrained spreads and $N^* < \infty$.

Corollary 4.5.3. *If $N^* = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$, then*

$$\overline{\delta^\pm}(t, i) = \delta^* \text{ if } \mp \mu(T - t) < d_i^{\pm*}, \quad \overline{\delta_i^\pm}(d) = \delta_* \text{ if } \mp \mu(T - t) > d_{i*}^\pm, \quad (4.5.4)$$

and if $\mp \mu(T - t) \in (\overline{d_i^{\pm*}}, \overline{d_{i*}^\pm})$, then $\overline{\delta^\pm}(t, i)$ is the unique solution of the fixed point equation

$$\overline{\delta^\pm}(t, i) = -\frac{\Lambda_i^\pm(\overline{\delta^\pm}(t, i))}{\Lambda_i^{\pm'}(\overline{\delta^\pm}(t, i))} \pm \mu(T - t). \quad (4.5.5)$$

If additionally $\Lambda_i^\pm(\delta) = a_i^\pm e^{-b_i^\pm \delta}$, with $a_i^\pm, b_i^\pm > 0$ for each $1 \leq i \leq k$, then

$$\overline{\delta^\pm}(t, i) = \delta_* \vee \left(\frac{1}{b_i^\pm} \pm \mu(T - t) \right) \wedge \delta^*.$$

In other words, if $N^* = \infty$ and $\gamma = \zeta = 0 \equiv \ell$, then solving (4.1.3) is no longer necessary and one only needs to solve the fixed point equations (4.5.5). Further, the MM's spreads do not depend on n . This is to be expected, since in this case she is neutral to all types of inventory risks and the terminal execution cost is neglected. However, she does need to re-adjust for changes in the regime as these impact on her probability of getting orders. If additionally, the orders intensities are exponential, then the spreads are computed straightforwardly avoiding the need of any numerical scheme. Note also how this simple model makes evident a second component of the optimal spreads through equation (4.5.5): a drift adjustment by which the MM takes into account the overall tendency of the asset's price.

Now we look at the results obtained for exponential intensities in general, and in particular for $N^* < +\infty$. Equation (4.5.3) becomes an explicit formula in this case, as the percentage liquidity sensitivities, $\Lambda_i^{\pm'}/\Lambda_i^\pm$, are constant.

Corollary 4.5.4. *If $\Lambda_i^\pm(\delta) = a_i^\pm e^{-b_i^\pm \delta}$, with $a_i^\pm, b_i^\pm > 0$ for each $1 \leq i \leq k$, then*

$$\overline{\delta^\pm}(t, n, i) = \delta_* \vee \left(-U_\gamma^{-1} \left(-1/b_i^\pm \right) - (\Theta(t, n \mp 1, i) - \Theta(t, n, i)) \right) \wedge \delta^*.$$

This means that for exponential intensities in general, it is no longer necessary to solve any fixed point equations but only the system of ODEs (4.1.3). The latter still needs to be solved numerically in general. (Alternatively, see [Gué17] for some asymptotic approximations.)

Remark. Apart from the case $N^* = +\infty$, $\gamma = \zeta = 0 \equiv \ell$, a noteworthy scenario in which (4.1.3) can be further simplified to a linear system of ODEs (with constant coefficients) is when $k = 1$, $N^* < +\infty$, $\delta_* = -\infty$, $\delta^* = +\infty$ and $\Lambda^\pm(\delta) = ae^{-b\delta}$, with $a, b > 0$. The reduction is achieved via the transformation $\Theta(t, n) = \frac{1}{b} \log \Psi(t, n)$. See [GLFT13] for more details.

5 Numerical analysis

In this section we present our numerical results, focusing on the difference in optimal behaviours between the partial and full information frameworks, and the intuition behind the filter. To exemplify our findings in a concrete manner, while keeping the presentation as simple as possible, we will assume throughout this section that:

- The MM has risk aversion parameter $\gamma = 0$.
- There are only two possible states: state 1 represents a “bad” regime with low liquidity taken by clients, and state 2 a “good” regime with high liquidity.
- The transition rate matrix Q is constant.
- The intensities are symmetric, proportional and exponential, i.e., $\Lambda_1^\pm(\delta) = \Lambda_1(\delta) = ae^{-b\delta}$ and $\Lambda_2^\pm(\delta) = \Lambda_2(\delta) = m\Lambda_1(\delta)$, with $a, b > 0$, $m > 1$.

The latter assumption allows us to perform the optimizations in equation (3.1.5) analytically. Practically, proportional intensities mean that while there is more active trading on the good regime, the way in which the clients react to movements in the spreads remains unaffected. As in [CJ15], we will allow for the three type of penalties to manage inventory risk: constraints on the maximum long and short positions, accumulated inventory penalty and a quadratic (possibly negligible) terminal penalty (or cost) for the MM. That is,

- $N^* < +\infty$, $\zeta \geq 0$ and $\ell(n) = cn^2$, for some $c \geq 0$.

Let us write $\pi := \pi_1$ for the conditional probability of being in the bad regime given the observable information. Note that π is a scalar in this section, since $\pi_2 = 1 - \pi_1$ neglects the need of the additional variable. The PIDE (3.1.5) at the point (t, n, π) reads:

$$0 = \theta_t + \mu n - \frac{1}{2}\sigma^2 n^2 \zeta + 2\hat{q}(\pi)\theta_\pi + \frac{a}{b}\hat{m}(\pi) \left(\mathbb{1}_{\{n < N^*\}} e^{-b\widehat{\delta}^-} + \mathbb{1}_{\{n > -N^*\}} e^{-b\widehat{\delta}^+} \right), \quad (5.0.1)$$

with terminal condition $\theta(T, n, \pi) = -cn^2$ and partial information optimal spreads given by

$$\widehat{\delta}^\pm(t, n, \pi) = \frac{1}{b} - 2\frac{w(\pi)}{\hat{m}(\pi)}\theta_\pi(t, n, \pi) - (\theta(t, n \mp 1, \pi/\hat{m}(\pi)) - \theta(t, n, \pi)), \quad \text{where:} \quad (5.0.2)$$

- (i) $\hat{q}(\pi) = q^{11}\pi + q^{21}(1 - \pi)$ is the observable transition rate to the bad regime.
- (ii) $\hat{m}(\pi) = \pi + (1 - \pi)m$ is the observable intensity increase from the bad regime (as a ratio).
- (iii) $w(\pi) = (m - 1)\pi(1 - \pi)$ is the observable variance, of the square root, of the percentage intensity increase from the bad regime; i.e., a measure of observable order flow volatility.

The previous equations are valid for $\delta_* \ll 0 \ll \delta^*$, or more precisely, for δ_*, δ^* such that $\delta_* \leq \widehat{\delta}^\pm \leq \delta^*$ holds true over the whole domain. Otherwise, one needs to floor and cap the optimal spreads and change the Hamiltonians accordingly, as done in Proposition 4.5.2.

A finite differences scheme of simple implementation to solve (5.0.1) consists of reversing time and using an explicit upwind Euler scheme over a uniform grid for (t, π) and $n = -N^*, -N^* + 1, \dots, N^*$. The terms of the form $\theta(t, n \mp 1, \pi/\hat{m}(\pi))$, where $\pi/\hat{m}(\pi)$ will typically fall outside of the grid, can be approximated by linear interpolation as in [DF14]. The limiting equations can be used for $\pi \rightarrow 0^+$ and $\pi \rightarrow 1^-$ provided that $-q^{11} > 0$, which we assume henceforth. Such a scheme can be shown to be consistent, stable and monotone under an appropriate CFL condition. In light of equation (3.1.5) satisfying a comparison principle (see [CEFS16a, Thm.5.3] and why it holds for our model in the proof of Theorem 3.1.4), we know the scheme converges to the unique continuous viscosity solution [BS91], and we can recover the expected penalized P&L of the MM (Theorems 3.1.1 and 3.1.4). The optimal strategy to be followed by the MM is then given in feedback form by $(\widehat{\delta}^-(t, N_{t^-}, \Pi_{t^-}), \widehat{\delta}^+(t, N_{t^-}, \Pi_{t^-}))$.¹⁵

¹⁵As previously mentioned, these are technically only candidates for optimal (or ϵ -optimal) strategies. We do not rigorously prove their optimality character here, but merely note that well known results of discrete time dynamic programming, together with the convergence of the discrete solutions of (5.0.1) to the analytical one, suggest that they can be safely used as such.

We will focus our attention on the optimal ask spread, with analogous observations holding for the bid spread. The parameter values in Table 1 were used for all the experiments presented in this section. For those present in the classical one regime models, we have chosen values used in previous works (in the $\gamma = 0$ case [CJ15]) to make the comparison clearer. In particular, the value of c will be taken as either 0 or 0.01, to be further specified in each experiment. The time horizon will always be the one displayed on the corresponding axis. Note that we work in a symmetric market, which justifies analyzing one side only.

Parameter	μ	σ	ζ	a	b	N^*	m	$-q^{11} = q^{21}$	$-\delta_* = \delta^*$
Value	0	0.1	0.1	2	25	3	5	5	10

Table 1: Standing parameter values for numerical tests presented.

5.1 Comparing full and partial information optimal strategies

Under the standing assumptions of this section, the full information equation (4.1.3), for a function $\tilde{\theta}$, becomes:

$$0 = \tilde{\theta}_t + \mu n - \frac{1}{2}\sigma^2 n^2 \zeta + \tilde{q}_i \left(\tilde{\theta}(t, n, 1) - \tilde{\theta}(t, n, 2) \right) + \frac{a}{b} \tilde{m}_i \left(\mathbb{1}_{\{n < N^*\}} e^{-b\tilde{\delta}^-} + \mathbb{1}_{\{n > -N^*\}} e^{-b\tilde{\delta}^+} \right), \quad (5.1.1)$$

with terminal condition $\tilde{\theta}(T, n, i) = -cn^2$ and full information optimal spreads given by

$$\tilde{\delta}^\pm(t, n, i) = \frac{1}{b} - \left(\tilde{\theta}(t, n \mp 1, i) - \tilde{\theta}(t, n, i) \right), \quad \text{where:} \quad (5.1.2)$$

- (i) $\tilde{q}_i = q^{11} \mathbb{1}_{\{1\}}(i) + q^{21} \mathbb{1}_{\{2\}}(i)$ is the effective transition rate to the bad regime.
- (ii) $\tilde{m}_i = \mathbb{1}_{\{1\}}(i) + \mathbb{1}_{\{2\}}(i)m$ is the effective intensity increase from the bad regime (as a ratio).

The previous equations are valid for $\delta_* \ll 0 \ll \delta^*$, as in the partial information case.

Although similar, equation (5.1.1) for the bad regime $i = 1$ (resp. good regime $i = 2$) is not the limiting equation of (5.0.1) for $\pi \rightarrow 1^-$ (resp. $\pi \rightarrow 0^+$). Indeed, a MM with full information can expect to make a larger profit. Thus, in general, $\tilde{\theta} > \theta$ even in these extreme cases.¹⁶ However, the corresponding optimal strategies do (at least approximately) agree, as Figures 1 and 2 show.

Figure 1 shows the optimal ask spread under partial information as a function of time, for different inventory levels (solid lines). It also displays how it changes for different filter values ($\pi = 0, 0.6$ and 1 from left to right) and how it compares to the optimal ask spread under full information (dashed lines) in the good regime (left) and bad regime (right). The inventory levels increase from top to bottom in all cases, and the terminal execution cost is neglected ($c = 0$). We have chosen to display $\pi = 0.6$ due to the effect of the filter being more pronounced around this value; it results in an optimal spread which is, for $t \ll T$, between 6% and 30% higher (depending on the inventory level) than the corresponding ones under full information (a sizeable difference in practice). We will come back to this in Section 5.2. All other values show similar intermediate behaviours. Recall that when the inventory reaches the minimum $n = -N^* = -3$, the MM will not sell any more (either abstaining from quoting or giving a “stub” quote) until her inventory increases again, which is why there is no spread plotted for this position. We see that some of the features already present in one regime models are preserved for two regimes, both under full and partial information; namely:

¹⁶Our numerical findings were indeed consistent with this intuitive statement.




Figure 1: Optimal ask spread as a function of time, under full and partial information (dashed and solid lines resp.), and for different inventory and filter values. Inventory levels increase from top to bottom in all cases. Terminal execution cost is neglected ($c = 0$).

- For $t \ll T$ the spread does not depend on time (approximately) and its asymptotic value becomes higher as the inventory decreases to the lower constraint $n = -N^*$. Indeed, for flat or short positions, any additional ask order will increase inventory risk, moving it closer to the minimum allowed and raising the accumulated inventory penalty. The MM manages this risk by increasing the spread, thus demanding a higher instantaneous profit and reducing the probability of getting executed. Similarly for long positions, the higher the exposure the lower the spread, as the MM seeks to unwind her inventory.
- In the case of negligible terminal cost, the spread converges to a terminal value independent of the inventory. This is the optimal spread of a fully risk-neutral MM with negligible costs (Corollary 4.5.3), who only maximizes “expected instantaneous margin”. (This value is given by $\widehat{\delta}^+(T, \cdot, \cdot) = 1/b$, the reciprocal percentage sensitivity of the liquidity to spread changes; cf. Proposition 4.5.2). The reason being that as the time horizon approaches, it becomes less and less likely for the MM to get executed again. The penalty accumulated between t and T vanishes, and the risk of reaching the inventory constraints diminishes. The MM therefore takes a bit more risk and makes a last attempt at increasing her expected P&L, either by increasing execution probability (for null or short positions) or increasing her instantaneous profit if executed (for long positions).

Nevertheless, the following differences can be observed:

- Under full information (and for symmetric markets), the spreads are always symmetric in the inventory with respect the risk-neutral value, i.e., $1/b - \widehat{\delta}^+(\cdot, n, \cdot) = \widehat{\delta}^+(\cdot, 1-n, \cdot) - 1/b$, for $0 < n \leq N^*$. However, under partial information this symmetry is broken, as the MM

increases her spreads. This holds in fact for any $0 < \pi < 1$ and it is suggested by the comparison of equations (5.0.2) and (5.1.2), since $\theta_\pi < 0$ (an increasing probability of being in the bad regime lowers the expected P&L; see argument of [BL09, Lemma 3.3]). As a result, the full information spreads are downward biased when the exact regime is unknown. Loosely speaking, this bias is approximately proportional to the product of θ_π (sensitivity of the expected P&L to observable regime changes) and $w(\pi)$ (observable order flow volatility), and inversely proportional to the observable intensity increase $\hat{m}(\pi)$. Intuitively, a partially informed MM faces not only the risk of the market regime shifting but also uncertainty on the current state, and must increase her spreads accordingly. This is done by considering the cost of any observable changes on the P&L and the fluctuations in the order flow, and discounting by liquidity increases.

- The good and bad regimes differ in the rate of order flow at the reference price (i.e., the amplitude parameters a and am of the orders intensities). Consequently, one could erroneously think that the partial information strategy should essentially be a one regime strategy for some intermediate parameter, such as $a\hat{m}(\pi)$. However, in the full information framework, an increase in the liquidity taken by clients results in the asymptotic spreads (i.e., for $t \ll T$) moving closer to the risk-neutral value (see [CJ15, GLFT13]). On the other hand, this is not true any more in the partial information case, due to the regime risk adjustment previously mentioned that shifts the spreads.
- Lastly, we remark that there is a distinguished change in behaviour close to expiry, with the spreads overshooting above the risk-neutral value for some inventories, before approaching it again. Intuitively, the overshooting is due to the added regime risk adjustment, and the convergence is due to the vanishing of additional inventory risks. (Recall that the MM's main concern becomes her terminal hedging cost, should there be one, and this does not depend on the market regime.) A somewhat similar effect can be found under full information in asymmetric markets [CJ15].

Figure 2 serves the same comparison as Figure 1, but including in this case a terminal execution cost. The observations remain mostly the same, except for a change in the terminal behaviour of the spreads, just as in the classical one regime case. Here, the spreads diverge from the risk-neutral no-costs value instead of approaching it, as the MM makes a last instant attempt to reduce her hedging cost. We remark that in both examples, a higher penalty ζ increases the difference between the full information optimal spreads of the two regimes, but it does not qualitatively change the observations made.

Figure 3 shows the partial information optimal ask spread as a function of time and filter, for three different inventory positions: extreme short (left), flat (center) and extreme long (right), and no terminal penalty. We can see that the spread is concave on the filter, and the maximum concavity is reached for a null inventory and $\pi \approx 0.6$, gradually decreasing as the position becomes short or long. We had observed already in equation (5.0.2) that the partially informed MM increases her spreads to manage higher regime risk, and that this correction is approximately proportional to the expected P&L sensitivity, θ_π , times the observed order flow volatility $w(\pi) = (m - 1)\pi(1 - \pi)$. Heuristically, it seems reasonable that the change in spread should resemble the concavity of $w(\pi)$ for a risk averse MM, and that the cost of regime uncertainty is the highest for a flat position (since deviations from it increase in turn price exposure). The question of why the maximum variation is reached around $\pi \approx 0.6$ is once again differed to Section 5.2.




Figure 2: Optimal ask spread as a function of time, under full and partial information (dashed and solid lines resp.), and for different inventory and filter values. Inventory levels increase from top to bottom in all cases. Terminal execution cost: $\ell(n) = -cn^2$, with $c = 0.01$.

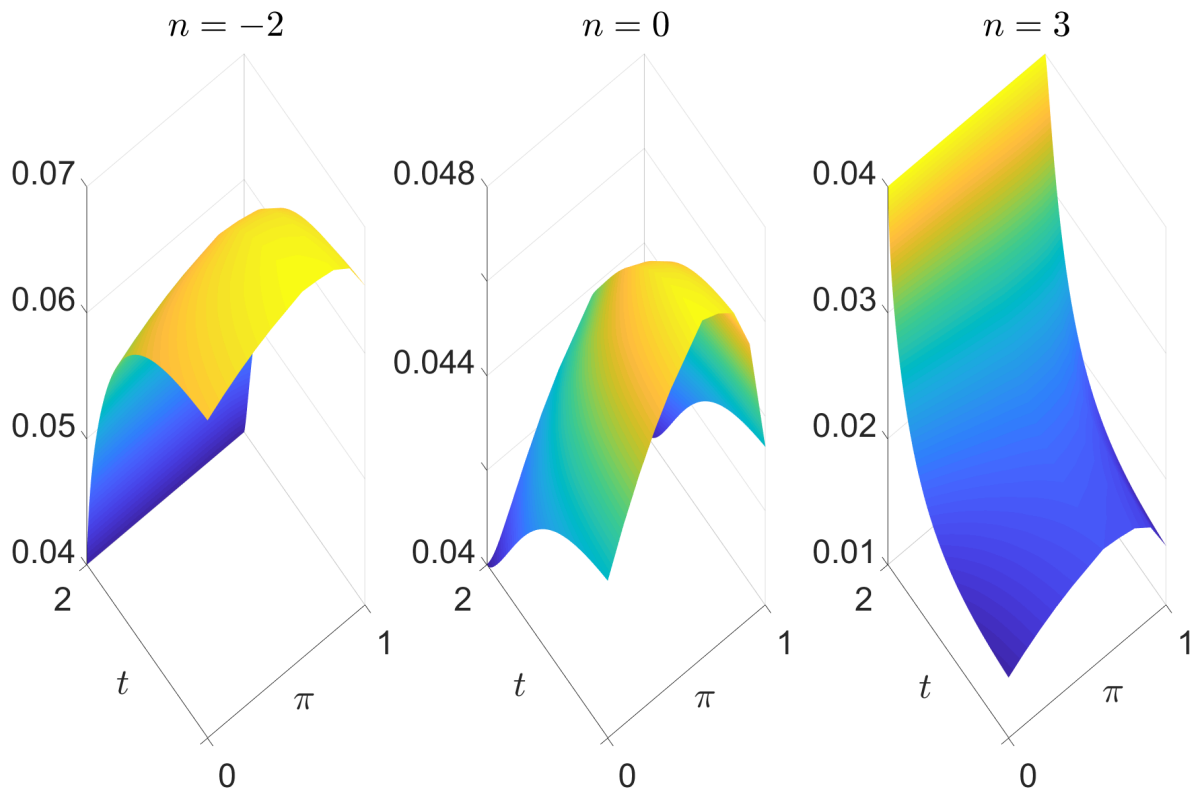


Figure 3: Optimal partial information ask spread as a function of time and filter (observable probability of bad regime), for different inventory levels. Terminal execution cost is neglected ($c = 0$).

5.2 Sample paths and a closer look at the filter

There are different ways to simulate point processes with stochastic intensity. A classical approach of simple implementation is the so-called *thinning method* [Oga81]. This method is particularly well suited to our framework, as it can be combined with the filtering theory developed in Section 2 by making use of the observable intensity of N (see Theorem 3.1.1 and [GKM11, Alg.3.2]). The interested reader is referred to [LTT18] for optimizations in terms of the thinning bound.

Figure 4 shows four sample paths resulting from an optimal behaviour of the MM with incomplete information. They were obtained by jointly simulating the inventory N (middle) starting from $n_0 = 0$, the filter $\Pi = \Pi^{\alpha,1}$ (bottom) starting from $\pi_0 = 0.5$, and the optimal strategy $\alpha = (\widehat{\delta}^-(t, N_{t-}, \Pi_{t-}), \widehat{\delta}^+(t, N_{t-}, \Pi_{t-}))$ (top), for $c = 0$.




Figure 4: Sample paths of optimal partial information bid/ask spreads (first), inventory (second) and filter (third), with $c = 0$, $n_0 = 0$ and $\pi_0 = 0.5$.

We want to analyze the behaviour of Π . Let us recall that Π_t represents the probability of being in the bad regime (a slow market with low liquidity) given the information observed by the MM at time t . The MM makes her assessment based on the orders (both buy and sell) that she received so far, which is the same as looking at the evolution of her inventory. If we take a close look at what happens between any two consecutive orders, we see that there is a pattern that repeats itself. As time passes without the MM receiving any order, she deems more likely that the market is in the bad regime, and so Π_t increases. If enough time goes by in this way, then Π_t gets asymptotically closer to some preponderant “equilibrium” value which is around 0.6. (Recall that this is the same value we have encountered before as having the greatest effect on the spreads.) But as soon as an order arrives, the MM revisits her probabilities, as being in the good regime of high liquidity seems a lot more likely. This is why we see that Π_t jumps downwards with each trade (i.e., at the jump times of the inventory).

So what makes $\pi \approx 0.6$ so special? If we look at the SDE (2.1.1) governing the filter dynamics, it now reads:

$$d\Pi_t = \left(\hat{q}(\Pi_t) + w(\Pi_t)a(\mathbb{1}_t^- e^{-b\hat{\delta}_t^-} + \mathbb{1}_t^+ e^{-b\hat{\delta}_t^+}) \right) dt + \Pi_{t-} \left(\frac{1}{\hat{m}(\Pi_{t-})} - 1 \right) (dN_t^- + dN_t^+). \quad (5.2.1)$$

In between jumps, each path of Π_t evolves according to an ODE. If enough time passes by without another order arriving, we can consider the asymptotic behaviour for $t \rightarrow T$. For $c = 0$, this leads to

$$\frac{d\Pi_t}{dt} \approx \hat{q}(\Pi_t) + w(\Pi_t)\frac{a}{e}\beta, \quad \text{with } \beta = 1 \text{ or } \beta = 2,$$

depending on the current inventory level. The right-hand side of this ODE is a concave parabola in Π_t with a negative root and another one, π^* , between 0 and 1. For the parameter values in Table 1, it is either $\pi^* \approx 0.572$ for $\beta = 1$, or $\pi^* \approx 0.636$ for $\beta = 2$. Therefore, π^* is an attractor (i.e., a stable equilibrium) for the filter, which explains the tendency of Π_t towards 0.6 and the central role of this value. Indeed, as the filter approaches one of its attractors, the MM exhausts all the information she has and can make no further assessment on the state of the market until she is hit by another order. In a sense, this is when she faces the most uncertainty, and she increases her spreads accordingly to manage regime risk.

A Additional proofs

Proof of Proposition 1.4.1. N^-, N^+ are finite variation (FV) processes with no common jumps, hence $[N^-, N^+] = 0$ \mathbb{Q} -a.s. Therefore, the multiplicativity of the stochastic exponential [JS02, p.138] and the exponential formula for FV processes [Br681, p.337 T4] yield the explicit expression $Z^\alpha = Z^{\alpha,-} Z^{\alpha,+}$, with

$$Z_t^{\alpha,\pm} = \exp \left(\int_0^t (1 - \Lambda_{Y_{u-}}^\pm(\delta_u^\pm)) \mathbb{1}_u^\pm du \right) \prod_{\substack{u \leq t: \\ \Delta N_u^\pm \neq 0}} \Lambda_{Y_{u-}}^\pm(\delta_u^\pm). \quad (\text{A.1.1})$$

Then (i) follows from the strict positivity and boundedness of $\Lambda_i^\pm(\delta^\pm)$ for $1 \leq i \leq k$, as in [Br681, p.168 T4] with terminal time T . The only difference is that in our case the \mathbb{Q} -intensity of N^\pm is $\mathbb{1}_{\{\mp N_{t-} < N^*\}} \leq 1$ instead of 1. The proof remains the same though, simply recalling that the \mathbb{Q} -moment generating function of N_T^\pm is dominated by the one of a standard Poisson random variable (see (1.2.5)). Uniform integrability is immediate.

(i) guarantees that \mathbb{P}^α , as defined in (ii), is an equivalent probability measure. The shape of the \mathbb{P}^α -intensities of N^-, N^+ is due to [Br681, p.166 T3].

The fact that W is still a Wiener process under \mathbb{P}^α is a consequence of the Girsanov–Meyer Theorem [Pro04, p.132 Thm.35] and Levy’s characterization Theorem.

As for Y , note first that its initial distribution does not change as $Z_0^\alpha = 1$, and the same goes for its infinitesimal generator operator. To see this, consider the \mathbb{Q} -generator operator $\mathcal{A}_t^Y : \mathbb{R}^E \rightarrow \mathbb{R}^E$, $\mathcal{A}_t^Y f(i) = \sum_{j=1}^k q_t^{ij} f(j)$. We know the process $M_t = f(Y_t) - f(Y_0) - \int_0^t \mathcal{A}_u^Y f(Y_u) du$ is a \mathbb{Q} -local martingale. Once again, by the Girsanov–Meyer Theorem., M is also a \mathbb{P}^α -local martingale ($[M, Z] = 0$ since Y, N^-, N^+ are FV processes satisfying (1.2.2)). Moreover, being bounded, M is a true \mathbb{P}^α -martingale and Y solves, for $(\mathbb{F}, \mathbb{P}^\alpha)$, a well-posed martingale problem for (\mathcal{A}_t^Y, μ_0) . This implies Y is a \mathbb{P}^α -Markov chain with a uniquely determined law [EK09, p.184 Thm.4.2].¹⁷ That Q is also the \mathbb{P}^α -generator matrix follows from uniqueness.

(1.2.2) clearly remains unchanged under an equivalent change of probability measure. \square

¹⁷Although our martingale problem is non-homogeneous in time, (Q_t) is deterministic, so this does not represent a problem.

Proof of Proposition 1.4.2. (i) and (ii) are proved just as in Proposition 1.4.1 (the processes $1/\Lambda_i^\pm(\delta^\pm)$ are strictly positive and bounded for $1 \leq i \leq k$). The \mathbb{Q}^α -independence of Y, N, W is a consequence of [JYC09, p.543 Lem.9.5.4.1]. That is, the $(\mathbb{Q}^\alpha, \mathbb{F})$ -martingales $Y - \int_0^\cdot \sum_{j=1}^k q_{u^-}^{Y_{u^-} - j} j du$, $N - \int_0^\cdot (\mathbb{1}_u^- - \mathbb{1}_u^+) du$ and W have the predictable representation property with respect to \mathbb{F}^Y , \mathbb{F}^N and \mathbb{F}^W resp. (see [Bré81, p.239 T8] for compensated counting measures) and they are $(\mathbb{F}, \mathbb{Q}^\alpha)$ -orthogonal, which implies their independence.

The explicit expressions for the stochastic exponentials (see (A.1.1)) show straightforwardly that $Z^\alpha \bar{Z}^\alpha = 1$, proving (iii). (iv) is due to [Bré81, p.64 T8]. \square

Proof of Proposition 2.1.1. We prove it just for N^+ , the others being analogous, and we omit α for simplicity. It is clear that $\widehat{\lambda}^+$ is predictable. We need to check that for any $\mathbb{F}^{W, N}$ -predictable process $\psi \geq 0$,

$$\mathbb{E} \left[\int_0^T \psi_t \widehat{\lambda}_t^+ dt \right] = \mathbb{E} \left[\int_0^T \psi_t dN_t^+ \right].$$

For any $1 \leq i \leq k$, each path of the càdlàg processes Π^i and $\mathbb{1}_{\{Y_{\cdot} = i\}}$ has only countably many jumps. We can therefore interchange these processes and their left limits when integrating with respect to dt . By properties of the conditional expectation and Fubini's Theorem,

$$\begin{aligned} \mathbb{E} \left[\int_0^T \psi_t \widehat{\lambda}_t^+ dt \right] &= \mathbb{E} \left[\int_0^T \psi_t \mathbb{1}_t^+ \sum_{i=1}^k \Pi_{t-}^i \Lambda_i^+(\delta_t^+) dt \right] = \mathbb{E} \left[\int_0^T \psi_t \mathbb{1}_t^+ \sum_{i=1}^k \Pi_t^i \Lambda_i^+(\delta_t^+) dt \right] \\ &= \mathbb{E} \left[\int_0^T \psi_t \mathbb{1}_t^+ \sum_{i=1}^k \mathbb{E}[\mathbb{1}_{\{Y_t = i\}} | \mathcal{F}_t^{N, W}] \Lambda_i^+(\delta_t^+) dt \right] = \int_0^T \mathbb{E} \left[\mathbb{E} \left[\psi_t \mathbb{1}_t^+ \sum_{i=1}^k \mathbb{1}_{\{Y_t = i\}} \Lambda_i^+(\delta_t^+) | \mathcal{F}_t^{N, W} \right] \right] dt \\ &= \int_0^T \mathbb{E} \left[\psi_t \mathbb{1}_t^+ \sum_{i=1}^k \mathbb{1}_{\{Y_t = i\}} \Lambda_i^+(\delta_t^+) \right] dt = \mathbb{E} \left[\int_0^T \psi_t \mathbb{1}_t^+ \sum_{i=1}^k \mathbb{1}_{\{Y_{t-} = i\}} \Lambda_i^+(\delta_t^+) dt \right] = \mathbb{E} \left[\int_0^T \psi_t \lambda_t^+ dt \right] \\ &= \mathbb{E} \left[\int_0^T \psi_t dN_t^+ \right]. \end{aligned}$$

\square

Proof of Proposition 2.1.2. Let us check that Π^α solves (2.1.1). Clearly, the constraint and the initial condition are satisfied. The verification of the SDEs is due to [CEFS16a, Prop.3.3] (with more details in [CEFS16b, App.A, Lemma A.2 and Prop.3.3]), albeit some considerations need to be made.

On the one hand, the authors work with a pure jump model, with strategies adapted to the natural filtration of the driving jump process only (i.e., there is no diffusion) and constant generator matrix for the Markov chain. However, *mutatis mutandis* the former differences yield no major change in the proofs.

On the other hand, the main assumption of the authors, [CEFS16a, Asm.2.1], postulates the existence of some deterministic measure $\tilde{\eta}^N(dz)$ on \mathbb{R} with compact support¹⁸ such that for all $i \in E$, $\delta^+, \delta^- \in (\delta_*, \delta^*)$, $n \in [-N^*, N^*] \cap \mathbb{Z}$, the measure $\eta^N(i, \delta^+, \delta^-, n, dz)$ is equivalent to $\tilde{\eta}^N(dz)$ and the Radon-Nikodym derivative $d\eta^N(i, \delta^+, \delta^-, n, \cdot) / d\tilde{\eta}^N$ is uniformly bounded and bounded away from zero $\tilde{\eta}^N(dz)$ -a.s.

¹⁸In [CEFS16a] the support is assumed to be a subset of $(-1, \infty)$, but this is only for a “return (or yield) process” as in their case.

Since the spread processes are fixed and bounded, we can assume without loss of generality that δ_* is finite. Setting $\tilde{\eta}^N(dz) := m_1(dz) + m_{-1}(dz)$ we see straightforwardly that the two measures are equivalent with derivative

$$\frac{d\eta^N(i, \delta^-, \delta^+, n, \cdot)}{d\tilde{\eta}^N}(z) = \Lambda_i^-(\delta^-) \mathbb{1}_{\{n < N^*\}} \mathbb{1}_{\{z=1\}} + \Lambda_i^+(\delta^+) \mathbb{1}_{\{-n < N^*\}} \mathbb{1}_{\{z=-1\}}, \quad (\text{A.1.2})$$

uniformly bounded by $\Lambda_i^-(\delta_*) + \Lambda_i^+(\delta_*)$. However, our model allows for $d\eta^N/d\tilde{\eta}^N(z) = 0$, which is a consequence of having vanishing intensities λ^\pm . This poses no issue nonetheless, as $d\eta^N/d\tilde{\eta}^N(z) > 0$ is only used in [CEFS16a, Prop.3.3] to guarantee $\bar{Z}^\alpha > 0$ (see Proposition 1.4.2). This condition, also satisfied in our model,¹⁹ allows to go from the physical probability \mathbb{P}^α to a reference probability \mathbb{Q}^α and backwards.

We turn now to the proof of uniqueness. We remark first that the jump height coefficients in (2.1.1) will typically not be Lipschitz (classical results for SDEs such as [Pro04, p.253 Thm.7] cannot be applied) and the paths of the spreads need not be continuous between the jump times of N (ruling out the most classical results of ODEs [CL55]). Nevertheless, we can still follow a pathwise ODEs approach.

Let us fix a path and verify uniqueness inductively on the intervals $[\tau_m, \tau_{m+1})$, where $\tau_0 := 0$ and $0 < \tau_1 < \tau_2 < \dots < \tau_M = T$ are the jump times of N (including the terminal time T even if there is no jump at that point). Set $A_t^{ji} = \mathbb{1}_t^-(\Lambda_j^- - \Lambda_i^-)(\delta_t^-) + \mathbb{1}_t^+(\Lambda_j^+ - \Lambda_i^+)(\delta_t^+)$. Then A^{ji} is bounded for all $1 \leq i, j \leq k$. Now observe that any càdlàg process $\tilde{\Pi}$, solving the constrained system of SDEs (2.1.1), must solve pathwise for $m = 0, 1, \dots, M-1$ the following system of ODEs in integral form, for $t \in [\tau_m, \tau_{m+1})$:

$$\tilde{\Pi}_t^i = R_m^i(\tilde{\Pi}_{\tau_m^-}) + \int_{\tau_m}^t \sum_{j=1}^k \left(q_u^{ji} \tilde{\Pi}_u^j + \tilde{\Pi}_u^i \tilde{\Pi}_u^j A_u^{ji} \right) du, \quad (\text{A.1.3})$$

with $R_m^i(\tilde{\Pi}_{\tau_m^-}) := \tilde{\Pi}_{\tau_m^-}^i \Lambda_i^\pm(\delta_{\tau_m}^\pm) / \sum_{j=1}^k \tilde{\Pi}_{\tau_m^-}^j \Lambda_j^\pm(\delta_{\tau_m}^\pm)$ if $\Delta N_{\tau_m} = \mp 1$, $m > 0$, and $R_0^i(\tilde{\Pi}_{\tau_0^-}) = \mu_0^i$. Using that $\tilde{\Pi} \in \Delta$ is bounded, and elementary algebra of bounded Lipschitz functions, it follows that $f^i : [\tau_m, \tau_{m+1}) \times \Delta \rightarrow \mathbb{R}$, defined by $f^i(u, \pi) = \sum_{j=1}^k (q_u^{ji} \pi^j + \pi^i \pi^j A_u^{ji})$ is Lipschitz in π , uniformly in u . Let K be the maximum Lipschitz constant of f^l for $1 \leq l \leq k$ and suppose $\Pi_{\tau_m^-}^\alpha = \tilde{\Pi}_{\tau_m^-}$ (clearly satisfied for $m = 0$). Then (A.1.3) yields $\|\Pi_t^\alpha - \tilde{\Pi}_t\| \leq K \int_{\tau_m}^t \|\Pi_u^\alpha - \tilde{\Pi}_u\| du$, implying $\Pi_t^\alpha = \tilde{\Pi}_t$ on $[\tau_m, \tau_{m+1})$ by Grönwall's inequality. As a consequence, the equality on $[0, T)$ follows by induction. It must clearly hold at time T as well, either by continuity or (if there is a jump) because $\Pi_T^\alpha = R_M(\Pi_{\tau_M^-}^\alpha) = R_M(\tilde{\Pi}_{\tau_M^-}) = \tilde{\Pi}_T$. \square

Proof of Proposition 2.1.4. We want to show that $\Pi_t^{\alpha, i} > 0$ for all $1 \leq i \leq k$, $0 \leq t \leq T$ a.s. We proceed by induction on the jump times of N , for each path, as at the end of the proof of Proposition 2.1.2. Using the same notations, let $1 \leq i \leq k$ and suppose $\Pi_{\tau_m^-}^{\alpha, j} > 0$ for all $1 \leq j \leq k$ (satisfied for $m = 0$ by assumption). Then (A.1.3), Assumptions 1.2.1 and the fact that $A^{ii} \equiv 0$ by definition, show that $\Pi^{\alpha, i}$ is absolutely continuous on the interval $[\tau_m, \tau_{m+1})$ and satisfies

$$\begin{aligned} (\Pi_t^{\alpha, i})' &= \sum_{j=1}^k \left(q_t^{ji} \Pi_t^{\alpha, j} + \Pi_t^{\alpha, i} \Pi_t^{\alpha, j} A_t^{ji} \right) = \Pi_t^{\alpha, i} \left(q_t^{ii} + \sum_{j \neq i} \Pi_t^{\alpha, j} A_t^{ji} \right) + \sum_{j \neq i} q_t^{ji} \Pi_t^{\alpha, j} \\ &\geq \Pi_t^{\alpha, i} \left(q_t^{ii} + \sum_{j \neq i} \Pi_t^{\alpha, j} A_t^{ji} \right), \end{aligned} \quad (\text{A.1.4})$$

¹⁹This was ultimately a consequence of the decomposition in Assumptions 1.2.2 (i), that allowed for the vanishing factors of the intensities to be secluded as the reference probability intensities.

for dt -a.e. $t \in [\tau_m, \tau_{m+1})$, subject to $\Pi_{\tau_m}^{\alpha,i} = R_m^i(\Pi_{\tau_m^-}^\alpha) > 0$. Let us set $s := \sup([\tau_m, \tau_{m+1}) \cap \{t \in [0, T] : \Pi_t^{\alpha,i} > 0\})$. We need to prove that $s = \tau_{m+1}$. By the continuity of $\Pi^{\alpha,i}$, it must be $s > \tau_m$. Consequently, (A.1.4) and the absolute continuity of $\log \Pi^{\alpha,i}$ on $[\tau_m, t] \subset [\tau_m, s)$ yield

$$\Pi_t^{\alpha,i} \geq R_m^i(\Pi_{\tau_m^-}^\alpha) \exp\left(\int_{\tau_m}^t (q_u^{ii} + \sum_{j \neq i} \Pi_u^{\alpha,j} A_u^{ji}) du\right) \quad \text{for } dt\text{-a.e. } t \in [\tau_m, s).$$

If it were $s < \tau_{m+1}$, the continuity of $\Pi^{\alpha,i}$ again and the former inequality would imply $0 = \Pi_s^{\alpha,i} > 0$, proving by contradiction that $s = \tau_{m+1}$ and $\Pi^{\alpha,i}$ is positive on the whole interval $[\tau_m, \tau_{m+1})$. Positivity on $[0, T)$ now follows by induction, and it must clearly hold at time T as well, as either there is a jump or we can reason as we just did. \square

Proof of Theorem 3.1.4. *Case $N^* < +\infty$ and $\gamma > 0$:* Let us write $\Psi = \frac{1+\Upsilon}{\gamma}$, with

$$\Upsilon(t, n, \pi) = \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \tilde{\mathbb{E}}^{\alpha, t, n, \pi} \left[\exp\left(-\gamma(\tilde{P}_{t,T}^{\alpha, n, \pi} + \ell(N_T^{t, n}))\right) \left(-\exp(\gamma \ell(N_T^{t, n}))\right) \right].$$

Then Υ can be regarded as the value function of an optimization problem in the standard Bolza-Lagrange formulation, i.e.,

$$\Upsilon(a) = \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \tilde{\mathbb{E}}^{\alpha, a} \left[\int_t^T D_{t,u}^{\alpha, a} f(u, A_u^{\alpha, a}, \alpha_u) du + D_{t,T}^{\alpha, a} g(A_T^{\alpha, a}) \right],$$

where the state variable $A_u^{\alpha, a} = (u, N_u^{t, n}, \Pi_u^{\alpha, t, n, \pi})$ is a PDMP with bounded state space $[0, T] \times (\mathbb{Z} \cap [-N^*, N^*]) \times \Delta^\circ$ and initial condition $a = (t, n, \pi)$, $D_{t,u}^{\alpha, a} = \exp\left(-\int_t^u \rho(u, A_u^{\alpha, a}, \alpha_u) du\right)$ is the discount factor and f, g, ρ are bounded functions. (These functions are bounded thanks to the control space being bounded.) The continuity of Υ can be proved now in the same way as in [CEFS16a, Thm.4.10] albeit in a more straightforward manner. This is due to the boundedness of f, g, ρ , the fact that $\Pi^{\alpha, t, n, \pi}$ never visits the relative border of the simplex Δ , and that there is no exit time of the state space other than the terminal time. Assumptions [CEFS16a, Asm.4.7] are clearly verified in our model, and [CEFS16a, Asm.2.1] has already been accounted for at the beginning of the proof of Proposition 2.1.2.²⁰ We remark that in our case the bounding function (see [CEFS16a, Lem.4.6]) can be taken simply as $b(t, n, \pi) = \exp(\eta(T-t))$, for an $\eta > 0$ large enough to prove contractiveness.

Having proved the continuity, the same proof of [CEFS16a, Thm.5.3] (or [DF99, Thm.7.5]) shows that Υ is the unique continuous viscosity solution of its standard HJB equation. We remark once again that our case is simpler, in that Υ is bounded and we do not have any boundary conditions other than the terminal time condition. In particular, there is no need for additional assumptions on the growth of Υ .

Finally, the result for Θ is obtained via the two increasing diffeomorphic transformations $\Psi = \frac{1+\Upsilon}{\gamma}$ and $\Theta = U_\gamma^{-1} \circ \Psi$.

Case $N^ < +\infty$ and $\gamma = 0$:* The only difference with the previous case is that the Bolza-Lagrange representation of the problem is obtained directly, since

$$\Theta(t, n, \pi) = \sup_{\alpha \in \tilde{\mathcal{A}}_t^2} \mathbb{E}^{\alpha, t, n, \pi} \left[\int_t^T \left\{ \delta_u^- \widehat{\lambda}_u^{\alpha, t, n, \pi} + \delta_u^+ \widehat{\lambda}_u^{\alpha, t, n, \pi} + \mu N_u^{t, n} - \frac{1}{2} \sigma^2 \zeta(N_u^{t, n})^2 \right\} du - \ell(N_T^{t, n}) \right],$$

with no need for any transformation.

Case $N^ = +\infty$ and $\gamma = \zeta = 0 \equiv \ell$:* The same as the latter case but working instead with the state variable $(u, \Pi^{\alpha, t, n, \pi})$ and the value function Φ . \square

²⁰An additional detail now is that [CEFS16a, Asm.2.1] is also used in [CEFS16a, Lem.4.1], which states that the drift coefficient in equation (2.1.1) is Lipschitz in the state variable, uniform in time and control. This is routinely verified in our case, under our new assumption: $-\infty < \delta_* < \delta^* < \infty$.

Proof of Proposition 4.2.1. Suppose first $I = [\tau, T]$ for some $0 \leq \tau < T$ and let $\varepsilon > 0$. Since $N^* < \infty$, there exists $(t_\varepsilon, n_\varepsilon, i_\varepsilon) \in [\tau, T] \times \mathcal{I}$ such that

$$\bar{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon) - \underline{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon) + \varepsilon(T - t_\varepsilon) = \min_{(t, n, i) \in [\tau, T] \times \mathcal{I}} \bar{\theta}(t, n, i) - \underline{\theta}(t, n, i) + \varepsilon(T - t). \quad (\text{A.1.5})$$

If $t_\varepsilon < T$, then we must have

$$\bar{\theta}_t(t_\varepsilon, n_\varepsilon, i_\varepsilon) - \underline{\theta}_t(t_\varepsilon, n_\varepsilon, i_\varepsilon) \geq \varepsilon.$$

Let us see that the left-hand side is non-positive. By (4.2.1) and (4.2.2),

$$\begin{aligned} \bar{\theta}_t(t_\varepsilon, n_\varepsilon, i_\varepsilon) - \underline{\theta}_t(t_\varepsilon, n_\varepsilon, i_\varepsilon) &\leq \sum_{j \neq i_\varepsilon} q_t^{i_\varepsilon j} \left(U_\gamma(\underline{\theta}(t_\varepsilon, n_\varepsilon, j) - \underline{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) - U_\gamma(\bar{\theta}(t_\varepsilon, n_\varepsilon, j) - \bar{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) \right) \\ &+ \mathbb{1}_{\{n < N^*\}} \left(H_i^- (\underline{\theta}(t_\varepsilon, n_\varepsilon + 1, i_\varepsilon) - \underline{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) - H_i^- (\bar{\theta}(t_\varepsilon, n_\varepsilon + 1, i_\varepsilon) - \bar{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) \right) \\ &+ \mathbb{1}_{\{-n < N^*\}} \left(H_i^+ (\underline{\theta}(t_\varepsilon, n_\varepsilon - 1, i_\varepsilon) - \underline{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) - H_i^+ (\bar{\theta}(t_\varepsilon, n_\varepsilon - 1, i_\varepsilon) - \bar{\theta}(t_\varepsilon, n_\varepsilon, i_\varepsilon)) \right). \end{aligned}$$

H_i^\pm increasing (resp. U_γ increasing) and (A.1.5) imply that the last two terms (resp. the first one) are non-positive. We must have then that $t_\varepsilon = T$, and due to (A.1.5) and (4.2.3) for all $(t, n, i) \in [0, T] \times \mathcal{I}$:

$$\begin{aligned} \bar{\theta}(t, n, i) - \underline{\theta}(t, n, i) + \varepsilon(T - t) &\geq \bar{\theta}(T, n_\varepsilon, i_\varepsilon) - \underline{\theta}(T, n_\varepsilon, i_\varepsilon) + \varepsilon(T - T) \geq 0 \\ \bar{\theta}(t, n, i) &\geq \underline{\theta}(t, n, i) - \varepsilon(T - t). \end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, we obtain the desired result.

The case $I = (\tau, T]$ is now a consequence of comparing $\underline{\theta}$ and $\bar{\theta}$ on intervals of the form $[t_n, T] \subseteq (\tau, T]$ with $t_n \searrow \tau$. \square

References

- [AB06] C.D. Aliprantis and K.C. Border, *Infinite dimensional analysis: a hitchhiker's guide*, 3rd ed., Springer, 2006.
- [AS08] M. Avellaneda and S. Stoikov, *High-frequency trading in a limit order book*, Quantitative Finance **8** (2008), no. 3, 217–224.
- [BL09] E. Bayraktar and M. Ludkovski, *Optimal trade execution in illiquid markets*, arXiv preprint arXiv:0902.2516v1 (2009).
- [BL14] ———, *Liquidation in limit order books with controlled intensity*, Mathematical Finance **24** (2014), no. 4, 627–650.
- [Bou07] B. Bouchard, *Introduction to stochastic control of mixed diffusion processes, viscosity solutions and applications in finance and insurance*, Lecture Notes Preprint (2007).
- [BR09] N. Bäuerle and U. Rieder, *MDP algorithms for portfolio optimization problems in pure jump markets*, Finance and Stochastics **13** (2009), no. 4, 591–611.
- [BR10] ———, *Optimal control of piecewise deterministic Markov processes with finite time horizon*, Modern trends in controlled stochastic processes: theory and applications **123** (2010), 143.
- [BR11] ———, *Markov decision processes with applications to finance*, Springer Science & Business Media, 2011.

- [Bré81] P. Brémaud, *Point processes and queues: martingale dynamics*, vol. 50, Springer, 1981.
- [BS91] G. Barles and P.E. Souganidis, *Convergence of approximation schemes for fully nonlinear second order equations*, *Asymptotic analysis* **4** (1991), no. 3, 271–283.
- [CDJ17] Á. Cartea, R. Donnelly, and S. Jaimungal, *Algorithmic trading with model uncertainty*, *SIAM Journal on Financial Mathematics* **8** (2017), no. 1, 635–671.
- [CEFS16a] K. Colaneri, Z. Eksi, R. Frey, and M. Szölgényi, *Optimal liquidation under partial information with price impact*, arXiv preprint arXiv:1606.05079v5 (2016).
- [CEFS16b] ———, *Shall i sell or shall i wait? Optimal liquidation under partial information with price impact*, arXiv preprint arXiv:1606.05079v3 (2016).
- [CJ13] Á. Cartea and S. Jaimungal, *Modelling asset prices for algorithmic and high-frequency trading*, *Applied Mathematical Finance* **20** (2013), no. 6, 512–547.
- [CJ15] ———, *Risk metrics and fine tuning of high-frequency trading strategies*, *Mathematical Finance* **25** (2015), no. 3, 576–611.
- [CJ19] Philippe Casgrain and Sebastian Jaimungal, *Trading algorithms with learning in latent alpha models*, *Mathematical Finance* **29** (2019), no. 3, 735–772.
- [CJP15] Á. Cartea, S. Jaimungal, and J. Penalva, *Algorithmic and high-frequency trading*, Cambridge University Press, 2015.
- [CJR14] Á. Cartea, S. Jaimungal, and J. Ricci, *Buy low, sell high: A high frequency trading perspective*, *SIAM Journal on Financial Mathematics* **5** (2014), no. 1, 415–444.
- [CL55] E.A. Coddington and N. Levinson, *Theory of ordinary differential equations*, Tata McGraw-Hill Education, 1955.
- [Dav84] M.H.A. Davis, *Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models*, *Journal of the Royal Statistical Society. Series B (Methodological)* (1984), 353–388.
- [Dav93] ———, *Markov models and optimization*, Chapman and Hall/CRC, 1993.
- [DF99] M.H.A. Davis and M. Farid, *Piecewise-deterministic processes and viscosity solutions*, *Stochastic analysis, control, optimization and applications*, Springer, 1999, pp. 249–268.
- [DF14] D.-M. Dang and P.A. Forsyth, *Continuous time mean-variance optimal portfolio allocation under jump diffusion: An numerical impulse control approach*, *Numerical Methods for Partial Differential Equations* **30** (2014), no. 2, 664–698.
- [DRR13] S. Delattre, C.Y. Robert, and M. Rosenbaum, *Estimating the efficient price from the order flow: a brownian cox process approach*, *Stochastic Processes and their Applications* **123** (2013), no. 7, 2603–2619.
- [EK09] S.N. Ethier and T.G. Kurtz, *Markov processes: characterization and convergence*, vol. 282, John Wiley & Sons, 2009.
- [FL12] P. Fodra and M. Labadie, *High-frequency market-making with inventory constraints and directional bets*, arXiv preprint arXiv:1206.4810 (2012).

- [FL13] ———, *High-frequency market-making for multi-dimensional Markov processes*, arXiv preprint arXiv:1303.7177 (2013).
- [FS06] W.H. Fleming and H.M. Soner, *Controlled Markov processes and viscosity solutions*, 2nd ed., vol. 25, Springer Science & Business Media, 2006.
- [GKM11] K. Giesecke, H. Kakavand, and M. Mousavi, *Exact simulation of point processes with stochastic intensities*, *Operations Research* **59** (2011), no. 5, 1233–1245.
- [GL15] O. Guéant and C.-A. Lehalle, *General intensity shapes in optimal liquidation*, *Mathematical Finance* **25** (2015), no. 3, 457–495.
- [GLFT12] O. Guéant, C.-A. Lehalle, and J. Fernandez-Tapia, *Optimal portfolio liquidation with limit orders*, *SIAM Journal on Financial Mathematics* **3** (2012), no. 1, 740–764.
- [GLFT13] ———, *Dealing with the inventory risk: a solution to the market making problem*, *Mathematics and Financial Economics* **7** (2013), no. 4, 477–507.
- [Gué17] O. Guéant, *Optimal market making*, *Applied Mathematical Finance* **24** (2017), no. 2, 112–154.
- [HS81] T. Ho and H.R. Stoll, *Optimal dealer pricing under transactions and return uncertainty*, *Journal of Financial Economics* **9** (1981), no. 1, 47–73.
- [JP82] J. Jacod and P. Protter, *Quelques remarques sur un nouveau type d'équations différentielles stochastiques*, *Séminaire de Probabilités XVI 1980/81*, Springer, 1982, pp. 447–458.
- [JS02] J. Jacod and A. Shiryaev, *Limit theorems for stochastic processes*, 2nd ed., vol. 288, Springer Science & Business Media, 2002.
- [JYC09] M. Jeanblanc, M. Yor, and M. Chesney, *Mathematical methods for financial markets*, Springer Science & Business Media, 2009.
- [LTT18] V. Lemaire, M. Thieullen, and N. Thomas, *Exact simulation of the jump times of a class of piecewise deterministic Markov processes*, *Journal of Scientific Computing* **75** (2018), no. 3, 1776–1807.
- [Nik06] A. Nikeghbali, *An essay on the general theory of stochastic processes*, *Probability Surveys* **3** (2006), 345–412.
- [Oga81] Y. Ogata, *On Lewis' simulation method for point processes*, *IEEE Transactions on Information Theory* **27** (1981), no. 1, 23–31.
- [ØS09] B.K. Øksendal and A. Sulem, *Applied stochastic control of jump diffusions*, 3rd ed., vol. 498, Springer, 2009.
- [Pro04] P.E. Protter, *Stochastic integration and differential equations*, 2nd ed., Springer, 2004.
- [Son88] H.M. Soner, *Optimal control of jump-Markov processes and viscosity solutions*, *Stochastic differential systems, stochastic control theory and applications*, Springer, 1988, pp. 501–511.