

HIGH FREQUENCY MARKET MAKING

RENÉ CARMONA AND KEVIN WEBSTER

ABSTRACT. Since they were authorized by the U.S. Security and Exchange Commission in 1998, electronic exchanges have boomed, and by 2010 high frequency trading accounted for over 70% of equity trades in the US. Such markets are thought to increase liquidity because of the presence of market makers, who are willing to trade as counterparties at any time, in exchange for a fee, the bid-ask spread. In this paper, we propose an equilibrium model showing how such market makers provide liquidity. The model relies on a codebook for client trades, the implied alpha. After solving the individual clients optimization problems and identifying their implied alphas, we frame the market maker stochastic optimization problem as a stochastic control problem with an infinite dimensional control variable. Assuming either identical time horizons for all the clients, or a stochastic partial differential equation model for their beliefs, we solve the market maker problem and derive tractable formulas for the optimal strategy and the resulting limit-order book dynamics.

1. INTRODUCTION

Electronic exchanges play an increasingly important role in financial markets and market microstructure became the key to understanding them. Here market microstructure is understood as *the study of the trading mechanisms used for financial securities* [11]. High frequency trading strategies depend very strongly on these mechanisms which in turn, vary from market to market. Three main themes were proposed to unite the common features of most of these markets [11].

- (1) The first theme is the limit-order book, where agents can post trading intentions at prices above (ask) or below (bid) the mid-price, which is regarded as the current *fair price* for the security. These posts, called *limit-orders*, can then lead to executions if they are paired with a matching order. In this case, the initial agent, commonly known as the liquidity provider, gains a small amount of money, having sold at the ask or bought at the bid. It is said in that case that the liquidity taker has paid a *liquidity fee* for the right to trade immediately.
- (2) The second theme is adverse selection due to asymmetric sources of information. It is based on the principle that liquidity takers choose which limit-orders are actually executed, and when. Moreover, since liquidity providers publicly announce their intentions by posting their limit-orders, and these two market rules create an *information bias* in favor of liquidity takers that compensates for the *liquidity fee*.
- (3) The last theme is statistical predictions. All agents on electronic markets attempt to predict prices on a short time scale. Such predictions are possible both because of the existence of private information and the actual market mechanisms that create short-range price correlations. While statistical predictions are used by all traders on electronic markets, the objectives of these agents vary widely. Some make their profits from these predictions, while others just reduce their transaction costs by *trading smartly*.

Date: October 12, 2012.

Key words and phrases. Limit-order book, Transaction costs, Stochastic partial differential equations, Pontryagin maximum principle, Market making.

After the liquidity crisis of 2008, liquidity became the central object of interest for market microstructure. Liquidity is rarely defined precisely, although intuitively, it should quantify how difficult it is to engage in a trade. Commonly accepted measures of liquidity are the bid-ask spread¹ and the volume present in a limit-order book. We tie these two quantities together by modeling what we will call the liquidity cost curve, or simply cost curve. We define this to be the total fee paid by the liquidity taker as a function of the volume asked for. If the liquidity taker only executes limit-orders at the best ask or bid, the fee will be equal to the bid-ask spread times the volume. Some clients, however, might wish to go beyond the best bid or ask and take more volume from the market. The marginal fee then becomes a function of the volume.

Trade volumes are difficult to model, in part because they are so dependent on liquidity takers' decisions and their execution strategies. Empirically, trade volumes present very heavy tails and are strongly autocorrelated ([6]): there are frequent outliers, buys tend to follow buys, and sells tend to follow sells, which is what makes them hard to evaluate statistically. Limit-orders, on the other hand, are mean-reverting and have been widely studied in the literature (see, for instance, [6, 7, 19]). The interplay between trend-following liquidity takers and mean-reverting liquidity providers is what makes the market reach the critical state of diffusion, as seen in [5]. From a modeling stand point, all these attributes are what makes a direct statistical description of trade volumes so difficult, which is why the introduction of a codebook is desirable.

1.1. Market making. Market makers are a special class of liquidity providers. They essentially act as a scaled down version of the market itself, always providing limit-orders on both sides of the mid-price. How these limit-orders are placed in terms of volume, distribution, and distance from the mid-price, determines the pricing strategy. We can therefore define the liquidity curve associated with a market maker's pricing strategy. The agents trading with a market maker are considered as his *clients*. Market making is a purely passive strategy that essentially corresponds to the service of providing liquidity to the market against a fee. This strategy makes money as long as the pricing accurately anticipates future price variations. The first step towards that objective is to find a model for client volumes that is consistent with price dynamics.

There are two schools of thoughts on market making models. The first focuses on inventory risk. There, the market maker has a preferred inventory position and prices according to his risk aversion to diverging inventory levels. Some of the first models of this type are [10, 4], which can be found in [11]. The second, initiated by [14], focuses on adverse selection, usually distinguishing between informed and noise trades. This paper belongs to the second line of thought.

According to our stylized definition of a market maker, the latter does not possess a view on the market. Clients, on the other hand, have views on the market and this leads to trading needs. When this view is short term, then the client has a statistical edge on the rest of the market and tries to push that edge to make a short term profit. But even when the view is long term, the client will still attempt to minimize transaction costs by predicting prices on a short time scale. Long-term strategies and liquidity constraints can therefore be modeled as noise around optimal short term execution strategies. The intuition behind this is that long term constraints, while the main motives behind most trades, will not necessarily dictate their execution. This is the case because there are increasingly more layers between the entity that formulates the long term strategy (for example investment banks or hedge funds) and the one that actually executes it (for example brokers or execution engines). But only the latter actually has a short term view on the market, which is what

¹The bid-ask spread is defined as the price difference between the lowest ask and the highest bid in the limit-order book.

the market maker is interested in. Optimal executions based on short term beliefs and associated trade volumes form therefore the cornerstone of the upcoming analysis and will link prices to trades.

1.2. Alpha. Alpha is the term often used to signify that a client has a directional view on the market. In an adverse selection model the aim of any market maker is to discover this view and hedge against it. Most papers (see for example [2, 5, 18, 11]) introduce a notion of market impact or response function, to study the relationship between trades and prices. Trade size is often – though not always – included, and a lot of work has been done to fit various curves to the response function. A large number of papers on execution strategies ([1, 2, 3, 16]) rely heavily on this notion of market impact. So do most practitioners. Essentially, alpha and market impact are the same quantity, even though the stories told to justify these concepts are quite different. In the market impact literature, the premise is that a trade *causes* a price movement, while alpha is viewed as an attempt at *predicting* the movement. Response function is a more neutral term that allows for both interpretations.

In this paper, we present a simple model for client decisions that proves, under very general assumptions, a systematic relationship between trade sizes and short term expected price variations. It states that marginal costs should equal expected price variations. This is an intuitive result given that price variations are the marginal gains of a risk-neutral client. In this model, each client has a short term view on the market and trades optimally according to this view: his optimal execution strategy depends upon his belief. The theoretical notion of *implied alpha* appears quite naturally from the stochastic optimal control problem that defines the client's model.

On the market maker side, our model tries to capture the fact that market makers do not act on a personal view of the market. The market maker uses the notion of client implied alpha to infer an approximate price process from his clients' behaviors, essentially aggregating the views of his clients to form a probability distribution on the price. The optimal order book strategy then replicates this distribution, with a corrective term that takes into account the profitability of trades, that is, the trade-off between spread and volume.

1.3. Results. The thrust of this paper is to propose a framework in which a market making strategy appears endogenously. This framework is based on

- (1) a result to imply a market view from client trades under a simple yet robust model of client behavior;
- (2) a procedure for the market maker to infer his own view on the market from that of his clients;
- (3) a profitability function that measures how profitable a posted volume on the order book is likely to be;
- (4) a penalizing term taking into account possible feedback effects of the market maker's decision.

The last part will be motivated by the idea that, just as the market maker implies information from his clients' trades, clients can infer information from his order book. Our hypotheses are chosen, and our results are derived, in order to make the market making problem tractable. Once the theoretical framework is put in place, the market maker optimal control problem is formulated and solved, leading to an explicit market making strategy.

The story told by the model is closest to that of the informed trade literature [14], although one major difference is that there is no clear cut distinction between informed and noise traders. Our conservative market maker assumes that all of his clients' implied alphas carry information. Moreover, the existence of at least one informed trader underpins the estimation formula used by the market maker. The paper also relates strongly to both the market impact [2, 5, 18, 11] and optimal execution [1, 3, 2, 16] literatures. A cost function c_t and an order book γ_t'' are derived

endogenously from the premises of the model. We hope that the proposed implied alpha codebook will find applications in other limit-order book models. Finally, while we ignore inventory risk to make our market maker risk-neutral, our approach still relates to utility function and inventory models such as [4, 10].

The results of the paper are organized in four sections. First, the setup for the model is given, which includes a methodology for modeling heterogeneous beliefs and transaction costs. Second, a simple client model is presented and solved, leading to a relationship between trades and alphas which motivates the market maker's choice for a codebook. Third, the market maker model is built step by step, starting from the client model and working through a series of results and approximations to lead to a reasonable control problem. Last, the market maker's problem is solved and the dynamics of the order book are determined analytically. For the sake of definiteness, we focus on a tractable example for which we compute the two-humped limit order book shape. An appendix is added at the end of the paper to provide the proofs of two technical results which, had they been included in the text, could have distracted from the main objective of the modeling challenge.

2. SETUP OF THE MODEL

In this section, all the elements of the model are presented.

2.1. Heterogeneous beliefs on the price. Consider n clients and one market maker who interact on an electronic exchange, with n very large. We will denote by $i \in \{1 \dots n\}$ a client index, and $k \in \{0 \dots n\}$ a generic index, with $k = 0$ corresponding to the market maker. We first introduce the following setting for the model:

- (1) a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ representing the “real life” filtration and probability measure. The filtration is generated by a d -dimensional \mathbb{P} Wiener process W_t .
- (2) a different filtration and measure $((\mathcal{F}_t^k)_{t \geq 0}, \mathbb{P}^k)$ for each of the agents. Assume furthermore that $\mathcal{F}_t^k \subset \mathcal{F}_t$, that $\mathbb{P}^k|_{\mathcal{F}_t^k}$ and $\mathbb{P}|_{\mathcal{F}_t^k}$ are equivalent and that $\mathbb{P}^0|_{\mathcal{F}_t^0} = \mathbb{P}|_{\mathcal{F}_t^0}$.
- (3) a d^k -dimensional \mathbb{P} Wiener process W_t^k that generates the filtration $(\mathcal{F}_t^k)_{t \geq 0}$.
- (4) a price process p_t which is an Itô process adapted to *all* the filtrations $((\mathcal{F}_t^k)_{t \geq 0})_{k=0 \dots n}$.
- (5) the drift and volatility of p_t grow at most polynomially in t under all probability measures.

where the last hypothesis must be understood in the a.s. and L^2 sense.

Let

$$dp_t = a_t dt + \sigma_t dW_t \quad (2.1)$$

be the Itô decomposition of p_t under $((\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. Hypotheses 2 and 3 imply that there exists an $(\mathcal{F}_t^k)_{t \geq 0}$ adapted process r_t^k such that

$$W_t^k = B_t^k + \int_0^t r_s^k ds \quad (2.2)$$

for some $\mathbb{P}^k|_{\mathcal{F}_t^k}$ Wiener process B_t^k and with $r_t^0 = 0$. Because W^k is d_k dimensional, so are r_t^k and B_t^k . Furthermore, by the martingale representation theorem, given that W_t^k is a \mathbb{P} martingale, there exists an $(\mathcal{F}_t)_{t \geq 0}$ adapted, $d_k \times d$ dimensional matrix Σ_t^k such that

$$dW_t^k = \Sigma_t^k dW_t \quad (2.3)$$

Finally, agent k has the following view on the market under $((\mathcal{F}_t^k)_{t \geq 0}, \mathbb{P}^k)$:

$$dp_t = a_t^k dt + \sigma_t^k dB_t^k \quad (2.4)$$

with $\sigma_t = \sigma_t^k \Sigma_t^k$ and $a_t^k = a_t + \sigma_t^k r_t^k$. In particular, σ_t needs to live in the intersection of the images of all the $(\Sigma_t^k)^T$ for p_t to be adapted to all the filtrations. Note that the d^k is allowed to differ from one agent to another, in which case the σ_t^k must be of different dimensions and hence differ.

In conclusion, all the agents have views on the price process that potentially conflict with each other's probability measure and even filtration, but can be compared coherently within the larger probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$.

Example 2.1. Consider the case where you have three Wiener processes W^1, W^2 and W^3 . Let $p_t = W_t^1 + W_t^2 + W_t^3$, $\mathcal{F}_t^i = \sigma\left((W_s^i)_{s \leq t}, (p_s)_{s \leq t}\right)$ and $r_t^i = f(W_t^i)$. You then have that:

$$\begin{aligned} dp_t &= dW_t^1 + dW_t^2 + dW_t^3 \text{ under } \mathbb{P} \\ &= f(W_t^1)dt + dB_t^{1,1} + dB_t^{1,2} \text{ under } \mathbb{P}^1 \\ &= f(W_t^2)dt + dB_t^{2,1} + dB_t^{2,2} \text{ under } \mathbb{P}^2 \end{aligned}$$

where the two last decompositions are not adapted with respect to the other filtration, but the price process is nevertheless adapted to both filtrations.

2.2. Transaction costs. We now allow trades among agents. All clients have a cumulative position L^i in the asset, which starts off at 0 at the beginning of the trading period. For the market maker, we rescale all the quantities according to the number of clients, hence L_t^0 is his average cumulative position per client. Clients control their position through its first derivative, l_t^i but incur transaction costs. The market maker, on the other hand, has no direct control over his position, but receives the liquidity fee c_t . To be precise, the market maker announces to his clients a cost function $l \mapsto c_t(l)$, which denotes the price of trading at speed l at moment t . A client then chooses her preferred trading volume l_t^i , and pays $c_t(l_t^i)$ in total transaction fees. We make the following hypotheses on these two processes:

- (1) The trade volume process l_t^i for each client is adapted to $(\mathcal{F}_t^i)_{t \geq 0}$ and $(\mathcal{F}_t^0)_{t \geq 0}$.
- (2) The cost function process c_t is adapted to all filtrations.
- (3) Marginal costs are defined: c_t' is almost everywhere continuous.
- (4) Clients may choose not to trade, $c_t(0) = 0$ and the mid-price is well defined at p_t , $c_t'(0) = 0$.
- (5) Marginal costs increase with volume: c_t is convex.
- (6) There is a fixed amount of liquidity the market maker offers.

Hypotheses 1 and 2 describe the information the different agents have access to. 3-5 are intuitive properties that the cost function must verify to make sense in terms of transaction costs. Finally, the hypothesis 6 is left purposefully vague, because it is best expressed with notation we introduce now.

2.2.1. Duality and link to the order book. This section introduces a change of variable that is both mathematically convenient and will carry deeper financial meaning when the model is solved. The new variable γ is defined as follows:

$$\gamma_t(\alpha) = \sup_{l \in \text{supp}(c_t)} (\alpha l - c_t(l)) \tag{2.5}$$

Under the assumptions $c_t \in C^1$, convex and $c_t(0) = 0$, this is equivalent to defining $\gamma_t' = (c_t')^{-1}$ and $\gamma_t(0) = 0$. In that case, γ_t is simply the Legendre transform of c_t . The following properties can be derived:

- (1) The Legendre transform maps the space of convex functions onto itself. It is its own inverse. In particular, γ_t'' , where the second derivative has to be understood in the sense of distributions, is a positive, finite measure.
- (2) $\forall \alpha \in \mathbb{R}, \forall l \in \mathbb{R}, \gamma_t(\alpha) + c_t(l) \geq \alpha l$ (Fenchel's inequality), with equality for $\alpha^* = c_t'(l)$ or $l^* = \gamma_t'(\alpha)$.
- (3) γ_t'' is a description of the limit-order book of the market maker. Indeed, it represents marginal volumes as a function of marginal costs. The property γ_t'' being a positive, finite measure corresponds to the fact that the market maker can only post positive, finite volumes on the order book.

This means we can actually replace c_t by the second derivative of its dual γ_t'' as the control variable of the market maker. Similarly, α_t^i is the dual to the client's control variable l_t^i . Call γ_t'' the market maker's liquidity offer and α_t^i the client's implied alpha. The second terminology will be explained later. We therefore recast all the hypotheses with these new dual variables:

- (1) The implied alpha process α_t^i for each client is adapted only to $(\mathcal{F}_t^i)_{t \geq 0}$ and $(\mathcal{F}_t^0)_{t \geq 0}$.
- (2) The liquidity process γ_t'' for each market maker is adapted to all filtrations.
- (3) We have $l_t^i = \gamma_t'(\alpha_t^i)$ and $\alpha_t^i = c_t'(l_t^i)$.
- (4) Clients may choose not to trade and the order book is centered around the mid-price: $\gamma_t(0) = 0$ and $\gamma_t'(0) = 0$.
- (5) Only positive, finite volumes are posted on a market maker's liquidity offer: γ_t'' is a positive, finite measure.
- (6) The total mass of γ_t'' is fixed. For convenience sake we will renormalize it to one, making γ_t'' a probability measure.

As a consequence of 3-5, l_t^i is bounded. Furthermore, γ_t'' therefore lives on the space of finite measures, which is a complete, separable metric space under the Lévy-Prokhorov metric and a convex set.

2.3. Putting it all together. We summarize the model. First a public price process p_t is given. Assume it is exogenously given to every one, that is, all the agents consider that they have no impact on the price under their probability measure, and only try to estimate its future movements. Then, a public liquidity offer, which can both be followed either by the cost function c_t or its dual, γ_t'' is announced. Clients pick their trade volumes l_t^i , or equivalently, their implied alphas α_t^i . We furthermore have the following state variable equations:

$$\begin{cases} dL_t^i &= l_t^i dt \\ &= \gamma_t'(\alpha_t^i) dt \\ dL_t^0 &= -\frac{1}{n} \sum_i \gamma_t'(\alpha_t^i) dt \end{cases} \quad (2.6)$$

since γ_t' is bounded, L_t^k has at most linear growth in t .

Finally, we write an infinite-horizon objective function for each agent, using distinct time scales β^k for each of them. Cumulative positions appreciate or depreciate at every moment by dp_t and client i pays the liquidity fee $c_t(l_t^i)$ to the market maker when readjusting her portfolio. We also rescale his objective function according to the number of clients he has.

$$\begin{cases} J^i &= \mathbb{E}_{\mathbb{P}^i} \left[\int_0^\infty e^{-\beta^i t} (L_t^i dp_t - c_t(l_t^i) dt) \right] \\ &= \mathbb{E}_{\mathbb{P}^i} \left[\int_0^\infty e^{-\beta^i t} (L_t^i dp_t - \alpha_t^i \gamma_t'(\alpha_t^i) dt + \gamma_t(\alpha_t^i) dt) \right] \\ J^0 &= \mathbb{E}_{\mathbb{P}} \left[\int_0^\infty e^{-\beta^0 t} (L_t^0 dp_t + \frac{1}{n} \sum_i (\alpha_t^i \gamma_t'(\alpha_t^i) - \gamma_t(\alpha_t^i)) dt) \right] \end{cases} \quad (2.7)$$

The first term of each objective function is bounded by the linear growth of L_t^k and the polynomial growth of the drift and volatility of p_t under \mathbb{P}^k .

3. THE CLIENT CONTROL PROBLEM

We first solve the problem from the clients' perspective, in order to know what drives trades. The results of this simple model will be used in the next section as a codebook for a more powerful market making model. The section concludes with a small robustness analysis on the underlying hypothesis of the model.

3.1. Solving the control problem. In this section, we solve the control problem for one generic client. Because her decisions have no impact on either p_t or c_t , it will not affect any of the other clients' decisions. Her control variable is not even adapted to their filtration. Similarly, her own control problem is not affected by the other clients' decision. This allows us to drop the index i in this subsection. Let $\left(\tilde{\mathbb{P}}, \left(\tilde{F}_t\right)_{t \geq 0}\right)$ denote that client's probability measure and filtration and \tilde{W} her Wiener process. Remember that the client tries to solve the following control problem:

- An admissible control l_t is a stochastic process adapted to $\left(\tilde{F}_t\right)_{t \geq 0}$ that lives in the support of c_t . Because of the hypothesis $\langle \gamma_t'', 1 \rangle = 1$, we have that the support of c_t is included in $[-1, 1]$, which means that the control set A is bounded.
- The state variable L_t verifies the dynamics:

$$dL_t = l_t dt \quad (3.1)$$

- The objective function is

$$\sup_l \mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_0^\infty e^{-\beta t} (L_t dp_t - c_t(l_t) dt) \right] \quad (3.2)$$

We assume β large enough for the problem to be well defined.

Using the notation introduced in the previous section, we have the central result:

Theorem 3.1. (*Implied alpha*)

A client who trades optimally will follow the relationship

$$\alpha_t = \mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_t^\infty e^{-\beta(s-t)} dp_s \middle| \tilde{\mathcal{F}}_t \right] \quad (3.3)$$

where α_t is defined by the codebook $c'_t(l_t) = \alpha_t$.

Proof. The aim now is to apply the Pontryagin maximum principle to the above control problem. The gain function is not in standard form, but a simple integration by parts solves this issue.

$$\int_0^\infty e^{-\beta t} L_t dp_t = [e^{-\beta t} L_t p_t]_0^\infty - \int_0^\infty e^{-\beta t} l_t p_t dt + \int_0^\infty e^{-\beta t} L_t \beta p_t dt. \quad (3.4)$$

The first term is equal to 0 because we assume $L_0 = 0$ and $\lim_{t \rightarrow \infty} e^{-\beta t} L_t p_t = 0$ by the linear growth of L_t and the polynomial growth of p_t . The client therefore maximizes:

$$\mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_0^\infty e^{-\beta t} ((\beta L_t - l_t) p_t - c_t(l_t)) dt \right] \quad (3.5)$$

We write out the generalized Hamiltonian of the system:

$$\mathcal{H}(t, \omega, L, l, Y) = lY + e^{-\beta t} ((\beta L - l)p_t - c_t(l)) \quad (3.6)$$

The generalized Hamiltonian is linear in L and concave in l by convexity of c , and therefore overall concave in (L, l) .

p_t and c_t are exogenously defined Itô process. Then, the backward equation

$$-dY_t = \beta p_t e^{-\beta t} dt - Z_t \cdot d\tilde{W}_t \quad (3.7)$$

has a unique solution

$$Y_t = \mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_t^\infty \beta p_s e^{-\beta s} ds \middle| \tilde{\mathcal{F}}_t \right] \quad (3.8)$$

By polynomial growth of the volatility of p_t , the Z term of the Backward Stochastic Differential Equation (BSDE from now on) satisfies the growth condition (A.6).

Therefore, the candidate optimal control is l_t^* verifying $e^{\beta t} Y_t - p_t = c'_t(l_t^*)$.

This determines $l_t = \gamma'_t(e^{\beta t} Y_t - p_t)$. Therefore the forward equation of L_t has a unique solution.

Finally, the quantity

$$c'_t(l_t^*) = e^{\beta t} Y_t - p_t = \mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_t^\infty \beta p_s e^{-\beta(s-t)} ds - p_t \middle| \tilde{\mathcal{F}}_t \right] \quad (3.9)$$

$$= \mathbb{E}_{\tilde{\mathbb{P}}} \left[\int_t^\infty e^{-\beta(s-t)} dp_s \middle| \tilde{\mathcal{F}}_t \right] \quad (3.10)$$

can be seen as the “implied alpha” of the deal, that is, the difference between the projected value into the future and the current value of p . \square

While it makes intuitive sense², this a non-trivial result. Indeed, c_t and hence l_t depend on the market maker’s decision, and yet $\alpha_t = c'_t(l_t)$ becomes a quantity that is “intrinsic” to the client. It is independent of the market maker’s pricing and only depends on the client’s view on the market. It intuitively represents the client’s price estimator and summarizes her beliefs on the price dynamics. Note that the discount factor now becomes the time-scale of her prediction. We define the quantity the client tries to predict as the *realized alpha* over the time scale $\frac{1}{\beta}$:

$$\alpha_t^r = \int_t^\infty e^{-\beta(s-t)} dp_s. \quad (3.11)$$

This coincides with what practitioners refer to as the ‘alpha’ of a client.

3.2. Dynamics of the implied alpha. (3.3) can be rewritten in terms of Itô dynamics:

$$d\alpha_t = \beta \alpha_t dt - dp_t + e^{\beta t} Z_t d\tilde{W}_t \quad (3.12)$$

$$= \beta \alpha_t dt - dp_t + \theta_t d\tilde{W}_t \quad (3.13)$$

with $\theta_t = e^{\beta t} Z_t$ therefore being the volatility of the estimation. All these equations summarize the link between optimal client volumes and price dynamics *under the client’s probability measure and filtration*.

Three things can be noted.

- (1) The drift of the implied alpha is the result of two opposing forces. On the one hand, the self-correlation term guarantees a certain coherence in the client’s decisions over the time-scale β^{-1} . On the other hand, the implied alpha, through the $-dp_t$ term, automatically takes into account the last price variation to recenter the estimation.

²All the formula says is that marginal costs equal expected marginal gains.

- (2) θ_t is a measure of intelligence of a client over the price process p_t , given that in the limit where $\theta_t = 0$, a client has a perfect view on the market. Conversely, clients with a big θ_t will have a higher variance on their price estimator, and can at the limit be considered as “noise” traders.
- (3) We can write the dynamics of α_t under \mathbb{P} to obtain:

$$d\alpha_t = \beta\alpha_t dt - dp_t + \theta_t \Sigma_t dW_t + \theta_t r_t dt \quad (3.14)$$

which provides the link between trade and price dynamics. Note that, while under $\tilde{\mathbb{P}}$, α_t is intrinsic to the client, under \mathbb{P} , α_t may depend on the market maker’s decisions. In words, this means that while the market maker cannot influence the client’s decision under her own view of the market, he can affect that view itself.

3.2.1. The cost of information. Given the above result, it is clear that a market maker has a privileged position on the market: he catches a glimpse of everyone’s belief on the price. We can now give an interpretation of γ beyond the fact that γ'' represents the order book:

Under a martingale measure of p_t , $c_t(l_t)$ represents the cost the client pays to the market maker for the liquidity l_t , and this is the conservative interpretation of transaction costs. However, $\gamma_t(\alpha_t)$ represents the cost the market maker pays *under the client’s view of the market*. Under $\tilde{\mathbb{P}}$, it is as if the market maker pays the client for information on the price process.

This gives us a good intuition about the market maker’s strategy: he collects information from each client, pricing them according to the current beliefs and how much the new information brings to him: $\gamma(\alpha)$ is essentially how much the market maker is willing to pay for a prediction of strength α , assuming that the client is correct.

3.3. Robustness analysis. In this section we define the notion of a ‘cash-insensitive’ agent and generalize the implied alpha relationship to the utility function case. This illustrates the robustness and limits of the proposed codebook.

Define the state variables

$$\begin{cases} dL_t &= l_t dt \\ dK_t &= -(p_t l_t + c_t(l_t)) dt \end{cases} \quad (3.15)$$

The second variable represents the client’s *cash position* at time t . If we define her wealth process as $X_t = p_t L_t + K_t$ then it verifies the standard dynamics

$$dX_t = L_t dp_t \quad (3.16)$$

The proposed decomposition has a clear economic meaning: the agent calculates her wealth by adding her cash K_t and the marked-to-the-mid value of her asset position L_t . This means that a ‘cash-insensitive’ objective function would be of the standard form

$$J = \mathbb{E}[U(X_\tau, p_\tau)] = \mathbb{E}[U(p_\tau L_\tau + K_\tau, p_\tau)] \quad (3.17)$$

with τ a stopping time adapted to the client’s filtration and U her utility function. The special form guarantees that the client does not differentiate between wealth in cash and wealth in the asset. An agent concerned with the liquidity of the asset would not be ‘cash-insensitive’.

One could generalize the utility framework to ‘cash-sensitive’ agents by considering a utility function of the form $U(L_\tau, K_\tau, p_\tau)$. In this case, the below result would not hold. The new Hamiltonian becomes:

$$\mathcal{H}(t, \omega, L, K, l, Y_L, Y_K) = lY_L - (p_t l + c_t(l))Y_K \quad (3.18)$$

where the dual variables satisfy in the 'cash-insensitive' case the BSDEs

$$\begin{aligned} Y_{L,t} &= \mathbb{E}[\partial_X U(X_\tau, p_\tau) p_\tau | \mathcal{F}_t] \\ Y_{K,t} &= \mathbb{E}[\partial_X U(X_\tau, p_\tau) | \mathcal{F}_t] \end{aligned}$$

The optimal execution strategy therefore verifies:

$$c'_t(l_t^*) = \mathbb{E} \left[\frac{Y_{K,\tau}}{Y_{K,t}} p_\tau \middle| \mathcal{F}_t \right] - p_t \quad (3.19)$$

which can be rewritten in a manner very similar to (3.3):

$$\alpha_t = \mathbb{E}_{\mathbb{Q}} [p_\tau | \mathcal{F}_t] - p_t \quad (3.20)$$

where $\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{Y_{K,\tau}}{\mathbb{E}[Y_{K,\tau}]}$ is a legitimate change of measure if $\partial_X U(X_\tau, p_\tau)$ is positive and integrable. In the case where τ is exponential and independent of the price process, we simply recover (3.3) under a different probability measure. Given that we assume all the clients to have differing probability measures anyway, we can without loss of generality stick to (3.3).

The above computation is somewhat formal, but can be made rigorous by giving explicit integrability assumptions on τ such that the growth condition (A.2) is verified. This is in particular the case when τ is independent of p and has exponential moments, or if τ is bounded.

4. REWORKING THE MARKET MAKER'S <ODEL

In this section, we work under the measure \mathbb{P} and often drop the superscript 0 referring to the market maker. Our goal is to provide an optimal market making strategy. Because of its complexity, the problem cannot be solved in full generality, and we propose a set of approximations which we justify on financial grounds.

The next four subsections identify a succession of simplifications guided by intuition based on the behavior of a typical market maker:

- (1) First, he should 'not hold a view on the market'. Mathematically speaking, this means that in his model for the price, the main explanatory variables are his client's beliefs. The error associated to this first simplification is proved to be small, though a function of the market maker's control.
- (2) Second, he should model how his clients' views might evolve into the future. A straightforward system of correlated Ornstein-Uhlenbeck processes is proposed to serve this purpose. This will be used to define an approximate objective function for the market maker.
- (3) The third approximation is made for mathematical convenience: we assume that the market maker has an infinite number of clients. This leads to the previous models becoming SPDEs.
- (4) Finally, a placeholder function is proposed to model the source of error identified in the first subsection.

4.1. Approximate Price Process. As the market maker should not hold a view on the market, we refrain from directly modeling p_t under the market maker's measure and filtration. Instead, the market maker constructs his model from the client's implied alphas. This is done in the following fashion:

Equation (3.14) can be turned around to describe price dynamics using the implied alpha of client i under \mathbb{P} :

$$dp_t = \beta^i \alpha_t^i dt - d\alpha_t^i + \theta_t^i \Sigma_t^i dW_t + \theta_t^i r_t^i dt \quad (4.1)$$

and this equation holds true for all i . Notice that there is no contradiction with the uniqueness of the Itô decomposition: these representations correspond to the same Itô process rewritten in terms

of the variable α_t^i . Using a sequence of positive and constant weights λ^i averaging to 1 (i.e. such that $(1/n) \sum_{i=1}^n \lambda^i = 1$), we obtain:

$$dp_t = \frac{1}{n} \sum_{i=1}^n \lambda^i (\beta^i \alpha_t^i dt - d\alpha_t^i) + \left(\frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i \Sigma_t^i \right) dW_t + \left(\frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i r_t^i \right) dt \quad (4.2)$$

Again, this equation is just a reformulation of the Itô decomposition of p_t . The advantage of that particular representation is that the first term is adapted to $(\mathcal{F}_t^0)_{t \geq 0}$. Hence, if he also knows (or rather, chooses) the weights $(\lambda^i)_{i=1 \dots n}$, then the market maker can follow the first term of this decomposition *in real time*. We introduce the special notation p_t^λ for this term and we refer to it as his price estimator: So:

$$dp_t^\lambda = \frac{1}{n} \sum_{i=1}^n \lambda^i (\beta^i \alpha_t^i dt - d\alpha_t^i). \quad (4.3)$$

The remainder which we denote ϵ_t^λ includes quantities that are unknown to him since

$$d\epsilon^\lambda = \left(\frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i \Sigma_t^i \right) dW_t + \left(\frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i r_t^i \right) dt \quad (4.4)$$

If we replace p by p^λ in the market maker's problem, the only difference appears in the objective function, which now contains an extra term. We refer to it as the error term:

$$\text{err} = \mathbb{E} \left[\int_0^\infty e^{-\beta t} L_t \left(\frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i r_t^i \right) dt \right] \quad (4.5)$$

on the market maker's objective function. Next we define

$$(\sigma^i)^2 = \mathbb{E} \left[\int_0^\infty e^{-\beta t} |\theta_t^i|^2 dt \right]. \quad (4.6)$$

This quantity is a measure of how 'noisy' client i is, and can be estimated by fitting the expression for the implied alpha to historical data. To be more specific, recall that under the client measure and filtration,

$$\alpha_t^i = \mathbb{E}_{P^i} \left[\int_t^\infty e^{-\beta^i(s-t)} dp_s \middle| \mathcal{F}_t^i \right] \quad (4.7)$$

and that θ_t^i is the error on this estimation. Hence, the smaller θ_t^i , the closer the implied alpha is to the realized one, which means that the client is particularly well informed. Note that in the finance literature, clients are often called informed if their market impact function (their average alpha) is unusually large. Our notion of intelligence of the price process does not coincide with this practice, as one can have at the same time a systematically small though correct alpha. In this paper, an informed trader is a trader for whom the implied and realized alphas nearly coincide, whereas for a noise trader, relationship (3.3) has much higher variance (due, for example, to liquidity concerns, a strongly non-linear utility function, or a poor filtration). The β^i and σ^i can be estimated from historical data by regressing the implied alpha against the realized one, that is for example by solving the least squares regression problem

$$\inf_{\beta > 0} \frac{1}{N} \sum_{k=1}^N \left(\alpha_{t_k}^i - \int_{t_k}^\infty e^{-\beta(s-t_k)} dp_s \right)^2 \quad (4.8)$$

where the t_k are the times in the past at which client i traded.

A skilled market maker can therefore construct his approximate price process by choosing a λ which puts most of its weight on such intelligent clients and little weight on *noise traders* with large σ^i 's.

Using Cauchy-Schwartz's inequality and the particular choice of weights

$$\lambda^i = \frac{n(\sigma^i)^{-2}}{\sum_j (\sigma^j)^{-2}}, \quad (4.9)$$

and assuming L_t is uniformly bounded by a constant \bar{L} , we obtain:

$$\begin{aligned} |\text{err}|^2 &\leq \beta^{-1}(\bar{L})^2 \mathbb{E} \left| \int_0^\infty \beta e^{-\beta t} \frac{1}{n} \sum_{i=1}^n \lambda^i \theta_t^i r_t^i dt \right|^2 \\ &\leq \beta(\bar{L})^2 \frac{1}{n} \sum_{i=1}^n \int_0^\infty e^{-\beta t} \mathbb{E} |\lambda^i \theta_t^i|^2 \frac{1}{n} \int_0^\infty e^{-\beta t} \sum_{i=1}^n \mathbb{E} |r_t^i|^2 dt \\ &\leq \beta I^{-1}(\bar{L})^2 \int_0^\infty e^{-\beta t} \frac{1}{n} \sum_{i=1}^n \mathbb{E} |r_t^i|^2 dt \end{aligned}$$

where

$$I^{-1} = \min_\lambda \frac{1}{n} \sum_{i=1}^n (\lambda^i \sigma^i)^2 = \left(\frac{1}{n} \sum_{i=1}^n (\sigma^i)^{-2} \right)^{-1}. \quad (4.10)$$

I can therefore be seen as a measure of the aggregate intelligence the clients have over the price process. Note that it suffices for *one* σ^i to be of order ϵ for I to be of order ϵ^{-1} . By Girsanov's theorem, we then have that

$$\mathbb{E} |r_t^i|^2 = -2 \frac{d}{dt} \mathbb{E} \log \frac{d\mathbb{P}^i|_{\mathcal{F}_t^i}}{d\mathbb{P}|_{\mathcal{F}_t^i}}, \quad (4.11)$$

and finally,

$$|\text{err}|^2 \leq 2\beta^2 I^{-1}(\bar{L})^2 \int_0^\infty e^{-\beta t} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \log \frac{d\mathbb{P}^i|_{\mathcal{F}_t^i}}{d\mathbb{P}|_{\mathcal{F}_t^i}} dt \quad (4.12)$$

Two important remarks are in order at this point. First, as long as at least *one* agent is well informed, I^{-1} is small. Second, the dependence of p and α^i upon the market maker's control γ'' is hidden in the Radon Nykodym derivative $\frac{d\mathbb{P}^i|_{\mathcal{F}_t^i}}{d\mathbb{P}|_{\mathcal{F}_t^i}}$. To understand why it is (unfortunately) reasonable to assume that this term may be strongly dependent upon γ'' is because a client can use the information on the order book γ'' publicly available as one of the sources of information he uses to form his probability measure on the price. This problem will be addressed in the last subsection.

4.2. Approximate Objective Function. In this subsection, we take a crucial methodological step. Instead of maximizing

$$J = \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left(L_t dp_t + \frac{1}{n} \sum_i \alpha_t^i \gamma_t'(\alpha_t^i) dt - \frac{1}{n} \sum_i \gamma_t(\alpha_t^i) dt \right) \right] \quad (4.13)$$

we assume that the market maker maximizes the *approximate* objective function

$$J^\lambda = \int_0^\infty e^{-\beta t} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[L_t \lambda^i (\beta^i - \beta) \alpha_t^i + \left(\alpha_t^i - \frac{1}{n} \sum_{j=1}^n \lambda^j \alpha_t^j \right) \gamma_t'(\alpha_t^i) - \gamma_t(\alpha_t^i) \right] dt \quad (4.14)$$

subject to a constraint of the form $|\text{err}|^2 \leq C$ for some constant $C > 0$. The new objective function was obtained by replacing dp_t by dp_t^λ and integrating by parts. Because of our choice of the form of the approximate price (4.3), the market maker's objective function does not depend upon p_t anymore, and *he only needs to model the client belief distribution*. This is consistent with the intuition that a market maker should not hold a view on the market. Rather, he should model the behavior of his clients with respect to each other, and price according to the information they provide. This is exactly the approach used in what follows.

However, we still need to propose a model for the α_t^i . For reasons of tractability we choose them as correlated Ornstein-Uhlenbeck processes:

$$d\alpha_t^i = -\rho\alpha_t^i dt + \sigma dM_t^i + \nu dW_t \quad (4.15)$$

with $\rho > 0$ and the M^i Wiener processes which are independent of each other and of W . $\sqrt{\sigma^2 + \nu^2}$ is the overall volatility level of a client, and ν the volatility that is due to some common information amongst clients (for example, the movement of the midprice). The next step of our strategy is to introduce a penalization term to account for the possible feedback effects hidden inside the error term introduced when replacing the objective function by its approximation. But first, we take the limit $n \rightarrow \infty$ to identify effective equations providing informative approximation to the properties of the original system comprising *finitely many clients*.

4.3. Infinitely Many Clients. Assume that the number of clients of the market maker is large enough to justify an approximation in the asymptotic regime n large. This will greatly improve the tractability of the model by allowing us to work in function spaces and rely on stochastic calculus tools to solve the model. The mainstay of this subsection is the Stochastic Partial Differential Equation (SPDE for short):

$$dv_t = \left(\frac{1}{2}(\sigma^2 + \nu^2)\Delta v_t + \rho \nabla(\text{id } v_t) \right) dt - \nu \nabla v_t dW_t \quad (4.16)$$

describing the dynamics of an infinite dimensional measure valued Ornstein-Uhlenbeck process. The following lemma links this macroscopic SPDE to our microscopic Ornstein-Uhlenbeck model for the implied alphas.

Proposition 4.1. *If $(\epsilon^i)_{i \geq 1}$ is a sequence of random variables such that $(\alpha_0^i, \epsilon^i)_{i \geq 1}$ is an iid sequence independent of W and the sequence $(M^i)_{i \geq 1}$, and such that the joint distribution, say m , of all the couples (α_0^i, ϵ^i) satisfies:*

$$\int (|\alpha|^p + |\epsilon|^2) m(d\alpha, d\epsilon) < \infty, \quad (4.17)$$

for all $p > 0$, then for each $t \geq 0$, the limit

$$\nu_t = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \epsilon^i \delta_{\alpha_t^i} \quad (4.18)$$

exists almost surely in the sense of weak convergence of measures, almost surely for every $t \geq 0$ it holds:

$$\int |\alpha|^p d\nu_t(\alpha) < \infty \quad (4.19)$$

for every $p > 0$, and the measure valued process $(\nu_t)_{t \geq 0}$ is a weak solution of the SPDE (4.16) in the sense that for any twice continuously differentiable function f (i.e. $f \in C^2$) such as f and its two derivatives have at most polynomial growth, we have that

$$d\langle f, \nu_t \rangle = \left\langle \frac{1}{2}(\sigma^2 + \nu^2)\Delta f - \rho \text{id} \nabla f, \nu_t \right\rangle dt + \nu \langle \nabla f, \nu_t \rangle dW_t \quad (4.20)$$

Furthermore, for each $t > 0$, the measure ν_t possesses an L^2 density almost surely.

Proof. See appendix. \square

The papers [13, 12] provide similar results for a more general class of microscopic models, including existence and uniqueness of the solution of the SPDE (4.16). However, given the simple form of the dynamics chosen in our particular model, we can provide an *explicit* form for the solution and detailed properties on the nature of its tails. Note that the correlation between client beliefs is crucial in having a fully stochastic model, given that for $\nu = 0$, ν_t only satisfies a deterministic partial differential equation.

We shall use four measure valued solutions of the above SPDE. In each case, the weights ϵ^i are explicit functions of the parameters σ^i and β^i introduced earlier. Because the values of these parameters appear as outcomes of statistical estimation procedures in practice, assuming that they are random and satisfy some form of ergodicity is not restrictive³. To construct these measures we assume that $(\alpha_0^i)_{i \geq 1}$ is an iid sequence of random variables whose common distribution has finite moments of all orders. The first of our four measures is obtained by choosing $\epsilon^i \equiv 1$ for all $i \geq 1$. Then:

$$\mu_t = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta_{\alpha_0^i}. \quad (4.21)$$

Next we assume that $\{(\sigma^i)^{-2}\}_{i \geq 1}$ is a sequence of positive random variables of order 1 such that $\{(\alpha_0^i, (\sigma^i)^{-2})\}_{i \geq 1}$ is an iid sequence independent of W and the sequence $(M^i)_{i \geq 1}$, so by choosing $\epsilon^i = (\sigma^i)^{-2}$ for all $i \geq 1$ we can define:

$$I_t = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (\sigma^i)^{-2} \delta_{\alpha_0^i}. \quad (4.22)$$

We shall also assume that the number $I = \mathbb{E}[(\sigma^i)^{-2}]$ is finite and strictly positive, and for each $t \geq 0$ we define the non-negative measure λ_t by:

$$\lambda_t = I_0^{-1} I_t. \quad (4.23)$$

Finally, we assume that $(\beta^i)_{i \geq 1}$ is a sequence of bounded random variables such that $\{(\alpha_0^i, (\sigma^i)^{-2}, \beta^i)\}_{i \geq 1}$ is an iid sequence independent of W and the sequence $(M^i)_{i \geq 1}$, so by choosing $\epsilon^i = \beta^i (\sigma^i)^{-2}$ for all $i \geq 1$ we can define:

$$\beta_t = I^{-1} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \beta^i (\sigma^i)^{-2} \delta_{\alpha_0^i}. \quad (4.24)$$

We now make a few remarks on the properties shared by essentially all the solutions $(\nu_t)_{t \geq 0}$ of the SPDE (4.16). For the sake of definiteness, we shall assume that $(\nu_t)_{t \geq 0}$ is a non-negative measure valued process solving this SPDE.

(1) The total mass $\langle 1, \nu_t \rangle$ of the measure ν_t is constant over time. Indeed, using the constant test function $f \equiv 1$ in (4.16) we see that

$$d \langle 1, \nu_t \rangle = 0 \quad (4.25)$$

In particular, the intelligence assumption $I = \langle 1, I_0 \rangle \geq \epsilon^{-1}$ is conserved over time.

(2) If we use the test function $f = id$ in (4.16) where the identity function id is defined by $id(\alpha) = \alpha$, then we see that the first moment of ν_t is itself an Ornstein-Uhlenbeck process mean reverting around 0 since:

$$d \langle id, \nu_t \rangle = -\rho \langle id, \nu_t \rangle dt + \nu \langle 1, \nu_0 \rangle dW_t. \quad (4.26)$$

³The fact that we have to enlarge the Brownian filtration at time $t = 0$ to *randomize* the coefficients does not impact the martingale representation theorem.

(3) Using $f = id^2$ we see that the second moment mean reverts around $(\sigma^2 + \nu^2) \langle 1, v_0 \rangle$ since:

$$d \langle id^2, v_t \rangle = (\langle \sigma^2 + \nu^2, v_0 \rangle - 2\rho \langle id^2, v_t \rangle) dt + 2\nu \langle id, v_t \rangle dW_t. \quad (4.27)$$

These Stochastic Differential Equations (SDEs for short) guarantee the existence of a constant C_1 such that:

$$\begin{aligned} \mathbb{E} \langle id^2, v_t \rangle &\leq e^{C_1 t} \\ \mathbb{E} \langle |id|, v_t \rangle &\leq \sqrt{\mathbb{E} \langle id^2, v_t \rangle} \leq e^{\frac{1}{2} C_1 t}. \end{aligned}$$

We shall use these estimates for the measures μ_t , λ_t and β_t . In fact similar estimates hold for moments of all order as can be proved by induction from (4.16). We do not give the details as they will not be used in what follows.

Coming back to the optimal control problem of the market maker, since γ_t'' belongs to a space of probability measures whenever the control $(\gamma_t)_{t \geq 0}$ is admissible, the following estimates hold:

$$\begin{aligned} \|\gamma_t'\|_\infty &\leq 1 \\ |\gamma_t(\alpha)| &\leq \alpha \\ |\langle \gamma_t', v_t \rangle| &\leq \langle 1, v_0 \rangle = O(1) \\ \mathbb{E} |\langle \gamma_t, v_t \rangle| &\leq \mathbb{E} \langle |id|, v_t \rangle = o(e^{\beta t}) \\ \mathbb{E} |\langle id \gamma_t', v_t \rangle| &\leq \mathbb{E} \langle |id|, v_t \rangle = o(e^{\beta t}) \end{aligned}$$

As an immediate consequence of the above remarks we have:

Corollary 4.2. *For any progressively measurable process $(\gamma_t)_{t \geq 0}$ such that γ_t'' is a probability measure, the state dynamic equation of the market maker:*

$$dL_t = -\langle \gamma_t', \mu_t \rangle dt \quad (4.28)$$

makes sense, and for sufficiently large β , the approximate objective function

$$J^\lambda = \int_0^\infty e^{-\beta t} \mathbb{E} [L_t \langle id, \beta_t \rangle + \langle -L_t \beta id + (id - \bar{\alpha}_t) \gamma_t' - \gamma_t, \mu_t \rangle] dt \quad (4.29)$$

where $\bar{\alpha}_t = \langle id, \lambda_t \rangle$, is well defined.

4.4. Modeling the Error Term. As argued in the first subsection, the approximation technique hinges on one hypothesis: the existence of at least one 'epsilon-intelligent' client. Furthermore, the clients probability measure (by which we mean the distribution of the clients implied alphas) is potentially a function of the market maker's control. This causes an undesirable nonlinear feedback effect which needs to be reined in to avoid explosion of the approximation error. In this subsection we propose a direct description of the error term, leaving open the question of how to derive it from a specific model of the clients probability measures.

Intuitively, the feedback effect corresponds to how much new information the order book shape reveals to the clients. The clients' beliefs at time t can be summarized by the probability measure μ_t and the order book by γ_t'' . A reasonable model for the error term is given by the expression

$$E = \epsilon \mathbb{E} \int_0^\infty e^{-\beta t} H(\gamma_t'' | \mu_t) dt \quad (4.30)$$

where $H(\nu | \mu)$ denotes the Kullback-Leibler distance (also known as relative entropy) defined by

$$H(\nu | \mu) = \begin{cases} \int f \left(\frac{d\nu}{d\mu} \right) d\mu, & \text{whenever } \nu \ll \mu; \\ \infty & \text{otherwise} \end{cases} \quad (4.31)$$

with $f(x) = x \log x$. Note that $H(\nu|\mu)$ is minimal for $\gamma_t'' = \mu_t$ by convexity of f . As explained in the introduction, we choose this particular distance for its intuitive interpretation and the fact that it leads to an explicit expression for the two hump order book shape endogenous to the model. However, the results hold for general strictly convex functions f , in which case the pseudo-distance H defined in (4.31) is known as the f-divergence between the measures μ and ν . See [8].

5. THE MARKET MAKER'S CONTROL PROBLEM

With all the pieces of our model in place, we now solve the market maker's control problem.

5.1. Model Summary. We consider a sequence $(\beta^i, (\sigma^i)^{-2})_{i \geq 1}$ of random variables such that the assumptions of Proposition 4.1 are satisfied with $\epsilon^i \equiv 1$, $\epsilon^i = (\bar{\sigma}^i)^{-2}$ and $\epsilon^1 = \beta^i (\sigma^i)^{-2}$ respectively, so that the measure-valued processes $(\mu_t)_{t \geq 0}$, $(\lambda_t)_{t \geq 0}$ and $(\beta_t)_{t \geq 0}$ constructed in (4.21), (4.23) and (4.24) are well defined.

As explained earlier, the market maker's control at time t is the convex function γ_t , and given our choice of penalizing terms, we expect its second derivative γ_t'' to be a probability measure absolutely continuous with respect to μ_t in order to avoid infinite penalties. So we refine the definition of the set \mathcal{A} of admissible controls for the market maker as the set of random fields $(g_t(x))_{t \geq 0, x \in \mathbb{R}}$ such that $(g_t(x))_{t \geq 0}$ is progressively measurable for each $x \in \mathbb{R}$ fixed, and $\langle g, \mu_t \rangle = 1$ for each $t \geq 0$. Given an admissible control $g \in \mathcal{A}$, we define the function γ_t as the anti-derivative of the function $\gamma'(\alpha) = \langle g_t 1_{[0, \alpha]}, \mu_t \rangle$ satisfying $\gamma_t(0) = 0$. The objective of the market maker is to choose an admissible control in order to maximize his modified *objective function* defined as:

$$\int_0^\infty e^{-\beta t} \mathbb{E} [L_t \langle id, \beta_t \rangle + \langle -L_t \beta id + (id - \bar{\alpha}_t) \gamma_t' - \gamma_t - \epsilon f \circ g_t, \mu_t \rangle] dt \quad (5.1)$$

where the controlled dynamics of the state L_t of the system are given by:

$$dL_t = -\langle \gamma_t', \mu_t \rangle dt. \quad (5.2)$$

We used the functions γ_t and γ_t' for consistency with earlier discussions. See below for forms of the state dynamics and market maker objective function written exclusively in terms of the density g_t . The above quantities are well-defined because of the estimates proven for the measures μ_t , λ_t and β_t solving the SPDE.

5.2. Solving the Market Maker Control Problem. We denote by Y the adjoint variable of L so that the generalized (random) Hamiltonian of the problem can be defined as

$$H(t, L, g, Y) = -\langle \gamma_t', \mu_t \rangle Y + e^{-\beta t} (L \langle id, \beta_t \rangle + \langle -\beta L id + (id - \bar{\alpha}_t) \gamma_t' - \gamma_t - \epsilon f \circ g, \mu_t \rangle) \quad (5.3)$$

for any deterministic function g as long as we define γ and γ' by $\gamma'(\alpha) = \langle g_t 1_{[0, \alpha]}, \mu_t \rangle$ and $\gamma(\alpha) = \int_0^\alpha \gamma'(\tilde{\alpha}) d\tilde{\alpha}$. Note that in the above expression of the Hamiltonian, g is deterministic and does not depend upon t , but γ_t' which is the cumulative distribution function of the probability measure γ_t'' with density g with respect to μ_t is random and depends upon t . Clearly, so is γ_t . This Hamiltonian can be rewritten as

$$H(t, L, g, Y) = e^{-\beta t} L \langle id, \beta_t - \beta \mu_t \rangle + e^{-\beta t} \tilde{H}(g) \quad (5.4)$$

where the modified Hamiltonian \tilde{H} is defined by

$$\tilde{H}(g) = -\langle \gamma_t', \mu_t \rangle e^{\beta t} Y + \langle (id - \bar{\alpha}_t) \gamma_t' - \gamma_t - \epsilon f \circ g, \mu_t \rangle,$$

and the other terms do not depend upon the control. Since the stochastic maximum principle proven in appendix says that we can look for the optimal control by maximizing the Hamiltonian, we shall

maximize the modified Hamiltonian. For each choice of the admissible control $(g_t)_{t \geq 0}$, we consider the corresponding state process $(L_t)_{t \geq 0}$ given by (5.2) and the adjoint equation:

$$-dY_t = e^{-\beta t} \langle id, \beta_t - \beta \mu_t \rangle dt - Z_t dW_t. \quad (5.5)$$

Since the derivative of the Hamiltonian with respect to the state variable L does not depend upon the control or the state L , the adjoint process can be determined independently of the choice of the admissible control $(g_t)_{t \geq 0}$ and the associated state $(L_t)_{t \geq 0}$ solving (5.2). Given the explicit formulas we have derived for μ_t , we obtain:

Lemma 5.1. *The solution to the adjoint Backward Stochastic Differential Equation (5.5) is given by:*

$$Y_t = \frac{e^{-\beta t}}{\beta + \rho} \langle id, \beta_t - \beta \mu_t \rangle, \quad (5.6)$$

and it verifies the growth condition (A.6).

Proof. The exact form (5.6) of the solution can be guessed by going back to the explicit system of finitely many Ornstein-Uhlenbeck processes and taking the limit. However, for the proof, we show that the process $(Y_t)_{t \geq 0}$ given by (5.6) is the solution by direct inspection, computing the Itô differential of Y_t defined by (5.6) and using the fact that the random measures β_t and $\beta \mu_t$ also solve the SPDE (4.16). Using (4.26), we get:

$$\begin{aligned} dY_t &= -\beta Y_t dt - \rho Y_t dt + e^{-\beta t} \frac{\nu}{\beta + \rho} dW_t \\ &= -e^{-\beta t} \langle id, \beta_t - \beta \mu_t \rangle dt + e^{-\beta t} \frac{\nu}{\beta + \rho} dW_t \end{aligned}$$

and by the growth properties of the first moment, $\lim_{t \rightarrow \infty} Y_t = 0$. Moreover, $Z_t = e^{-\beta t} \nu / (\beta + \rho)$ clearly verifies the growth condition (A.6). \square

The form of the modified Hamiltonian justifies the introduction of the quantity:

$$\alpha_t^* = \bar{\alpha}_t + e^{\beta t} Y_t = \left\langle id, \lambda_t + \frac{\beta_t - \beta \mu_t}{\beta + \rho} \right\rangle \quad (5.7)$$

so that, if we compute the modified Hamiltonian along the path of the adjoint process we get:

$$\tilde{H}(g) = -\langle \alpha_t^* \gamma_t' - \gamma_t, \mu_t \rangle - \epsilon \langle f \circ g, \mu_t \rangle.$$

α_t^* can be viewed as the *shadow alpha* of the market maker.⁴ The first term in the definition of α_t^* is the average belief for alpha under the weighted client measure λ_t , while the second term takes into account mismatches in the time-horizons of the clients. For each $t \geq 0$, we define the profitability function m_t by:

$$\alpha \mapsto m_t(\alpha) = (\alpha - \alpha_t^*)[\mu_t([\alpha, \infty)) - 1_{(-\infty, 0]}(\alpha)]. \quad (5.8)$$

For each $\alpha \in \mathbb{R}$, $(m_t(\alpha))_{t \geq 0}$ is a progressively measurable stochastic process and for each fixed $t \geq 0$, the function $\alpha \mapsto m_t(\alpha)$ is almost surely continuous in α , except for a possible jump $m_t(0^+) - m_t(0^-)$ at $\alpha = 0$. m_t is bounded and vanishes at the infinities because of the integrability property (4.19) of the solutions of the SPDE. The profitability function quantifies the expected profit for an order placed at time t at the price level α . Indeed, the absolute value $|\alpha - \alpha_t^*|$ is equal to the *spread* the market maker expects to gain per filled order, and up to a possible change of sign, the term $\mu_t([\alpha, \infty)) - 1_{(-\infty, 0]}(\alpha)$ is equal to the proportion of clients that will fill the order. If the arrivals

⁴This terminology is justified by the fact that, when the market maker *does* have his own view on the market (that is, his own model for dp_t), we would obtain the same optimization problem replacing α_t^* by $\mathbb{E}[\int_t^\infty e^{-\beta(s-t)} dp_s | \mathcal{F}_t^0]$.

of the agents of our model occurred according to a Poisson process instead of simultaneously, this would be the *filling probability* of the order. The respective contributions of these two terms are commonly parsed by practitioners. Notice also that, in the degenerate case $\epsilon = 0$, the profitability function is the derivative of the Hamiltonian in the direction of the control.

We now identify the modified Hamiltonian in terms of the control g without involving the anti-derivatives γ'_t and γ_t .

Lemma 5.2. *For each $t \geq 0$, we have the identity*

$$\langle (id - \alpha_t^*)\gamma'_t - \gamma_t, \mu_t \rangle = \langle m_t, \gamma_t'' \rangle. \quad (5.9)$$

Proof. Successive integrations by parts and simplifications yield:

$$\begin{aligned} \langle \gamma_t'', (id - \alpha_t^*) (\mu_t([\cdot, \infty)) - 1_{(-\infty, 0]}(\cdot)) \rangle &= - \langle \gamma_t', (\mu_t([\cdot, \infty)) - 1_{(-\infty, 0]}(\cdot)) - (id - \alpha_t^*)\mu_t - \alpha_t^* \delta_0 \rangle \\ &= \langle \gamma_t, \mu_t - \delta_0 \rangle + \langle \gamma_t' (id - \alpha_t^*), \mu_t \rangle + \alpha_t^* \gamma_t'(0) \\ &= \langle (id - \alpha_t^*)\gamma_t' - \gamma_t, \mu_t \rangle \end{aligned}$$

where the first integration by parts is justified by the fact that γ' is bounded, the Dirac distribution has compact support and the simple facts:

$$\lim_{\alpha \rightarrow \infty} \mu_t([\alpha, \infty)) = 0, \quad \text{and} \quad \lim_{\alpha \rightarrow -\infty} \mu_t([\alpha, \infty)) - 1_{\alpha \leq 0} = - \lim_{\alpha \rightarrow -\infty} \mu_t((-\infty, \alpha]) = 0,$$

and $\alpha\mu(\alpha)$ vanishes at the infinities. The linear growth of γ justifies the second integration by parts. The last line uses the initial conditions of γ_t and γ_t' . \square

We can now state and prove the main result of this section:

Theorem 5.3. *The solution to the market maker's control problem is given by*

$$g_t(\alpha) = \frac{e^{m_t(\alpha)/\epsilon}}{\int e^{m_t(\alpha)\epsilon} d\mu_t(\alpha)} \quad (5.10)$$

The quantity in the right hand side of can be seen as the change of measure from the distribution of the clients alphas to the order book of the market maker.

Proof. Using Lemma 5.1 and equation (5.9) enables us to rewrite the modified Hamiltonian as:

$$\tilde{H}(g) = \langle m_t g, \mu_t \rangle - \epsilon \langle f \circ g, \mu_t \rangle \quad (5.11)$$

which we need to maximize for each fixed $t \geq 0$, over $g \in L^1(\mu_t)$ with $\langle g, \mu_t \rangle = 1$ and $g \geq 0$. Again, for each fixed t and almost surely, by convexity of f , \tilde{H} is concave in g and bounded from above, so that $\tilde{M} = \sup_{A_t} \tilde{H}(g)$ exists and is finite.

Let $(g_n)_{n \geq 1}$ be a maximizing sequence in A_t , i.e. a sequence of admissible g_n such that $H(g_n) \rightarrow \tilde{M}$. By Komlos's lemma (see for example [9]), there exists a subsequence $g_{\phi(n)}$ and an element $g_\infty \in L^1(\mu_t)$ such that $\frac{1}{n} \sum_{i=1}^n g_{\phi(\psi(i))} \rightarrow g_\infty$ μ_t -a.e. as $n \rightarrow \infty$ for any subsequence of the original subsequence. Clearly, $g_\infty \geq 0$ μ_t almost everywhere. To prove $\langle g_\infty, \mu_t \rangle = 1$ we show uniform integrability of $\{g_n\}_n$ using De la Vallée-Poussin's theorem (see [15]). Indeed, by definition of the maximizing sequence, $H(g_n) \geq \tilde{M} - 1$ for large enough n . Hence $\epsilon \langle f \circ g_n, \mu_t \rangle \leq -\tilde{M} + 1 + \|m_t\|_\infty$. f is non-negative, convex, increasing on $[1, \infty)$ and verifies $\lim_{x \rightarrow \infty} \frac{1}{x} f(x) = \infty$. This concludes the proof of the admissibility of g_∞ as a control. Finally, by concavity of the modified Hamiltonian, the supremum is attained at g_∞ .

We now characterize the maximal element we just constructed. We introduce a Lagrange multiplier $\eta \in \mathbb{R}$ to relax the problem to the set of $g \in L^1(\mu_t)$ such that $g \geq 0$, so that the Lagrangian reads:

$$\tilde{L}(g) = \langle (m_t - \eta)g - \epsilon f \circ g, \mu_t \rangle \quad (5.12)$$

Classical variational calculus yields the optimality condition

$$g = (f')^{-1}(\epsilon^{-1}(m_t - \eta)), \quad (5.13)$$

for some η . The existence and uniqueness of a Lagrange multiplier renormalizing g is obvious given the explicit formula for $(f')^{-1}$. We conclude that g is the desired optimum. \square

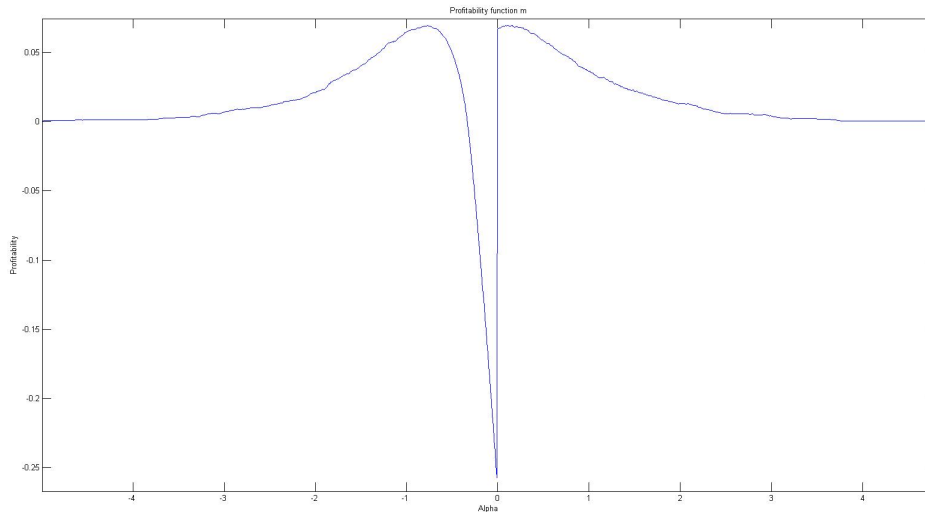


FIGURE 1. Graph of a simulated profitability curve m with $\alpha^* \approx -0.33$. The maximum profitability is only 5% of the volatility used to simulate prices. The average spread being comparable to the volatility (see [18]), this means that the market maker's margins are very thin in this simulation.

5.3. **Interpretation.** Figure 1 and Figure 2 illustrate what the market maker does:

- (1) He tries to stick to a shape not too far away from μ_t , his client alpha distribution. This is to avoid feedback effects and associated errors on the price estimation.
- (2) He also takes into account the profitability function m_t , which leads him “to make a big hole” in the center of the distribution μ_t .

The combined effects lead to the familiar “double hump” shape of the order book, as seen in [18]. Other consequences of the liquidity formula are

- (1) In the limit where one client is perfectly intelligent of the price ($\epsilon \rightarrow 0$), the market maker places a Dirac mass on the order book. In particular, he only trades on profitable sections of the book.
- (2) In the noise trade limit ($\epsilon \rightarrow \infty$), the market maker simply reproduces the client alpha distribution.
- (3) If the return distribution of the asset is mean-reverting (which is not the case for options with maturities, for example), then so will μ_t and hence γ_t and c_t . In the case of a European option, γ_t'' converges to a Dirac at the payoff at maturity.

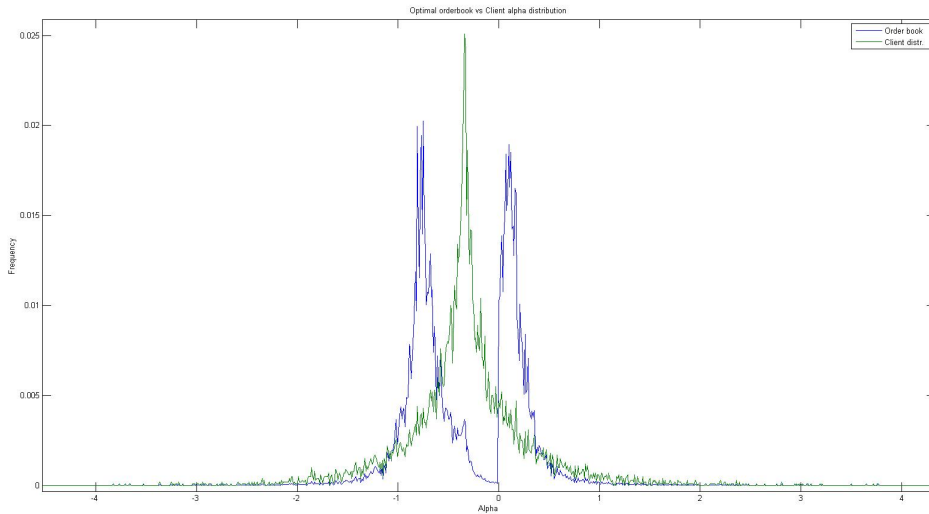


FIGURE 2. Bimodal distribution: Optimal order book γ'' for $\epsilon = 0.01$ using the simulated alpha distribution and the entropic penalizing factor. Unimodal distribution: Client alpha distribution. The profitability used is the same as in figure 1.

6. CONCLUSION

In this paper, we propose an equilibrium model for high frequency market making. We show that each client chooses his level of trade by an expected discount of future prices according to his beliefs. We call this the implied alpha of the client. This relationship can be used as a codebook to link trade and price dynamics. The different clients can be differentiated according to the time scales of their implied alphas and the variances of the errors they make estimating the price. This leads to an expression for the latter that the market maker can use to avoid having his own view on the market. He can then construct a profitability curve which dictates which sections of the order book are the most profitable. In solving the market maker optimization problem, we use a penalizing term to smooth the objective function and capture possible feedback effects. All these result in a tractable framework in which we can solve the market maker's control problem and identify the equilibrium order book dynamics.

APPENDIX A. A CONVENIENT FORM OF THE STOCHASTIC MAXIMUM PRINCIPLE

We present a form of the Pontryagin stochastic maximum principle tailored to the needs of the analysis of Section 5. The set-up is quite general, following loosely [17] Chapter 3. We assume that $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ is a filtered probability space, $(\mathcal{F}_t)_{t \geq 0}$ is generated by a Wiener process $(W_t)_{t \geq 0}$, and we let $A_t(\omega)$ be such that for all (t, ω) , A_t is a Borel convex subset of a Polish topological vector space E , adapted to $(\mathcal{F}_t)_{t \geq 0}$. The admissible control processes $(\alpha_t)_{t \geq 0}$ are the progressively measurable processes in E such that $\alpha_t \in A_t$ for all $t \geq 0$. We also assume that the dynamics of the controlled state X are governed by an Ordinary Differential Equation(ODE) with random

coefficients valued in \mathbb{R}^n :

$$dX_t = f_t(X_t, \alpha_t)dt \quad (\text{A.1})$$

where the coefficient $f : (\omega, t, X, a) \rightarrow f_t(X, a)$ is Lipschitz in X uniformly in all the other variables. Assume an estimate of the form

$$|X_t| \leq |X_0|e^{Ct} \quad (\text{A.2})$$

where C is a constant. We also assume that f is continuously differentiable (i.e. C^1) in X . The controller's objective function is given by

$$J(\alpha) = \mathbb{E} \left[\int_0^\tau j_t(X_t, \alpha_t)dt + g_\tau(X_\tau) \right] \quad (\text{A.3})$$

where τ is a stopping time adapted to $(\mathcal{F}_t)_{t \geq 0}$ and j and g are random real-valued concave functions which are C^1 in X . Furthermore, we assume that j , g and τ are such that $E[|g_\tau(X_\tau)|]$ and $\mathbb{E} \left[\int_0^\tau |j_t(X_t, \alpha_t)|dt \right]$ are finite and uniformly bounded from above. Next, we define the Hamiltonian

$$\mathcal{H}_t(\omega, X, a, Y) = f_t(\omega, X, a)Y + j_t(\omega, X, a) \quad (\text{A.4})$$

for $Y \in \mathbb{R}^n$ and for each admissible control $(\alpha_t)_{t \geq 0}$, the adjoint equation:

$$-dY_t = \partial_X \mathcal{H}_t(X_t, \alpha_t, Y_t)dt - Z_t dW_t \quad (\text{A.5})$$

with $Y_\tau = \partial_X g(X_\tau)$ and Z_t such that $\mathbb{E} \int_0^\tau |Z_t|^2 dt < \infty$. Under these conditions we have the following result.

Theorem A.1 (Pontryagin's maximum principle). *Let $\hat{\alpha}$ be an admissible control and \hat{X} be the associated state variable. Suppose there exists a solution (Y, Z) to the adjoint equation (A.5) such that*

$$\mathbb{E} \left[\int_0^\tau e^{Ct} |Z_t|^2 dt \right] < \infty, \quad (\text{A.6})$$

and the Hamiltonian verifies for every $t \geq 0$

$$\mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t) = \max_{a \in A_t} \mathcal{H}_t(\hat{X}_t, a, Y_t), \quad a.s. \quad (\text{A.7})$$

and for every $t \geq 0$, almost surely, the function

$$(X, \alpha) \mapsto \mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t)$$

is concave, then $\hat{\alpha}$ is an optimal control, that is, $J(\alpha) \leq J(\hat{\alpha})$ for all admissible $\alpha = (\alpha_t)_{t \geq 0}$.

Proof. If $(X_t)_{t \geq 0}$ is the state associated to another admissible control $(\alpha_t)_{t \geq 0}$, we have the chain of relationships:

$$\begin{aligned} & \mathbb{E} \left[g_\tau(X_\tau) - g_\tau(\hat{X}_\tau) \right] \\ & \leq \mathbb{E} \left[(X_\tau - \hat{X}_\tau) \cdot Y_\tau \right] \end{aligned} \quad (\text{A.8})$$

$$= \mathbb{E} \left[\int_0^\tau \left(- (X_t - \hat{X}_t) \cdot \partial_X \mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t) + (f_t(X_t, \alpha_t) - f_t(\hat{X}_t, \hat{\alpha}_t)) \cdot Y_t \right) dt \right] \quad (\text{A.9})$$

$$\leq \mathbb{E} \left[\int_0^\tau \left(\mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t) - \mathcal{H}_t(X_t, \hat{\alpha}_t, Y_t) + (f_t(X_t, \alpha_t) - f_t(\hat{X}_t, \hat{\alpha}_t)) \cdot Y_t \right) dt \right] \quad (\text{A.10})$$

$$\begin{aligned} & = \mathbb{E} \left[\int_0^\tau \left(\mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t) - \mathcal{H}_t(X_t, \hat{\alpha}_t, Y_t) - \mathcal{H}_t(\hat{X}_t, \hat{\alpha}_t, Y_t) + \mathcal{H}_t(X_t, \alpha_t, Y_t) \right) dt \right] \\ & \quad - \mathbb{E} \left[\int_0^\tau \left(j_t(X_t, \alpha_t) - j_t(\hat{X}_t, \hat{\alpha}_t) \right) dt \right] \end{aligned} \quad (\text{A.11})$$

$$\leq -\mathbb{E} \left[\int_0^\tau j_t(X_t, \alpha_t) - j_t(\hat{X}_t, \hat{\alpha}_t) \right]. \quad (\text{A.12})$$

(A.8) stems from the concavity of g and terminal condition of Y . (A.9) follows from the dynamics of Y and the relationship:

$$\mathbb{E} \left[\int_0^\tau (X_t - \hat{X}_t)^2 Z_t^2 dt \right] \leq 2X_0 \mathbb{E} \left[\int_0^\tau e^{Ct} Z_t^2 dt \right] < \infty \quad (\text{A.13})$$

which guarantees that the local martingale part is a martingale. (A.10) holds by the concavity of the Hamiltonian. (A.11) is just the definition of \mathcal{H} . Finally, (A.12) is a consequence of the fact that $\hat{\alpha}_t$ maximizes $\mathcal{H}_t(\hat{X}_t, \cdot, Y_t)$. \square

APPENDIX B. DERIVATION OF THE SPDE (4.16)

We give the main steps of the proof of Proposition 4.1. By independence of the common randomness W and the idiosyncratic randomness $(M^i, \alpha_0^i, \epsilon^i)_{i \geq 1}$, we can freeze the randomness of W and work after conditioning with respect to \mathcal{F}^W , the σ -algebra generated by W , without affecting the independence properties of the idiosyncratic random variables. By independence of the M^i , the α_0^i and ϵ^i , we have that, conditional on \mathcal{F}^W , the random variables

$$\alpha_t^i = \alpha_0^i e^{-\rho t} + \int_0^t e^{-\rho(t-s)} (\nu dW_s + \sigma dM_s^i) \quad (\text{B.1})$$

are iid and Gaussian. So if $f \in C^2$ is such that f and its two derivatives have at most polynomial growth, given the assumption (4.17) on the joint distribution m of all the couples (α_0^i, ϵ^i) , we can apply the law of large numbers and get:

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \epsilon^i f(\alpha_t^i) = \mathbb{E} \left[\epsilon^1 f(\alpha_t^1) \mid \mathcal{F}^W \right] \\ & = \int_{\mathbb{R}^3} \epsilon f \left(\alpha_0 e^{-\rho t} + \int_0^t e^{-\rho(t-s)} \nu dW_s + \sigma x \sqrt{\int_0^t e^{-2\rho(t-s)} ds} \right) \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} m(d\alpha_0, d\epsilon) dx \end{aligned}$$

The other results can be derived directly from this explicit representation.

REFERENCES

- [1] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.
- [2] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2000.
- [3] R. F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10(1):1–18, 2003.
- [4] Y. Amihud and H. Mendelson. Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2):223–249, 1986.
- [5] J. . Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets: The subtle nature of 'random' price changes. *Quantitative Finance*, 4(2):176–190, 2004.
- [6] J. . Bouchaud, M. Mzard, and M. Potters. Statistical properties of stock order books: Empirical results and models. *Quant. Finance*, 2(4):251–256, 2002.
- [7] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. *Operations research*, 58(3):549–563, 2010.
- [8] I. Csiszar. *Information-type measures of difference of probability distributions and indirect observations*, volume 2. 1967.
- [9] F. Delbaen and W. Schachermayer. *The Mathematics of Arbitrage*. Springer Verlag, 2010.
- [10] M. B. Garman. Market microstructure. *Journal of Financial Economics*, 3(3):257–275, 1976.
- [11] J. Hasbrouck. *Empirical market microstructure*. Oxford University Press, 2007.
- [12] T. G. Kurtz and J. Xiong. Particle representations for a class of nonlinear spdes. *Stochastic Processes and their Applications*, 83(1):103–126, 1999.
- [13] T. G. Kurtz and J. Xiong. Numerical solutions for a class of SPDEs with applications to filtering. *Stochastics in Finite and Infinite Dimension*, pages 233–258, 2000.
- [14] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
- [15] P.A. Meyer. *Probabilities and Potential*. North-Holland Pub., 1966.
- [16] A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *Preprint*, 2005.
- [17] H. Pham. *Continuous-time stochastic control and optimization with financial applications*. Springer Verlag, 2008.
- [18] M. Wyart, J. . Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between bid-ask spread, impact and volatility in order-driven markets. *Quantitative Finance*, 8(1):41–57, 2008.
- [19] I. Zovko and J. D. Farmer. The power of patience: A behavioral regularity in limit order placement. *Quantitative Finance*, 2(5):387–392, 2002.