Financial Intermediation at Any Scale For Quantitative Modelling (3/3)

Cours Bachelier

Charles-Albert Lehalle

Capital Fund Management, Paris and Imperial College, London

IHP, November 18, 2016 to December 4, 2016

Motivation

- Since the 2008-2009 crisis legislators’ and regulators’ viewpoint on financial markets changed,
- They target to monitor and limit the risk taken by the market participants,
- In one sentence: they want to ensure most participants plays a role of **intermediaries**, and nothing more.
- The notion of intermediation and the role of banks, investment banks, dealers, brokers, and now insurance companies and funds have evolved and continue to evolve;
- important concepts to understand this are: **microstructure** and **infrastructure**; they are linked to **liquidity**.
- These last 10 years, the field of Market Microstructure emerged. Related literature has grown...
- I am convinced **financial mathematics** can address quite efficiently core concepts, as partly an academic and partly a professional, I dedicated the last 12 years to understand these changes from a practical and a theoretical viewpoint.
- These sessions will be the occasion to share how, in my opinion, financial mathematics can **answer to new and important questions** raised by recent changes.

The Standard Pitch

Following the 2008 crisis, the financial system changed a lot:

- "Clients" (from inside or outside) have no more appetite for sophisticated products.

⇒ The system went **from a bespoke market to a mass market**.

Bespoke means to sell products that are very different: no economies of scale but high margins.

Mass market means a lot of similar products + optimized logistics.

- Regulators welcome this change because it can prevent an accumulation of risk in inventories (cf. optimized logistics).

⇒ The G20 of Pittsburgh (Sept. 2009) put the emphasis on **inventory control** (it is the root of improved clearing, segregated risk limits, etc).

⇒ Policy makers took profit of two existing regulations (Reg NMS in the US and MiFID in Europe) to push toward **electronification** of exchanges (i.e. improved traceability and less information asymmetry).

- Technology went into the game. Think about the kind of recent "innovations" (uber, booking.com, M-pesa, blockchain, etc): it is about **disintermediation**.

⇒ How do you desintermediate a system made of intermediates?

Historically, market **micro**structure stands for not reducing

- Sellers = Equity Shares and Bonds issuers
- Buyers = investors.

In practice, today, associated topics are

- Market impact, Fire sales and Flash Crashes
- Auction / Matching mechanisms (Limit Orderbooks, RFQ, conditional / fuzzy matching, etc)
- Optimal trading / Liquidation
- Market Making and High Frequency Trading
- Investment process while taking all this into account

CFM
INSIGHT.DATA.CLARITY.

My Viewpoint on Market Microstructure

I have been Global Head of Quantitative Research at Crédit Agricole Cheuvreux and CIB during years (including the crisis). I discuss a lot with regulators; previously inside the **working group on Financial Innovation** of the ESMA, now inside the **Scientific Committee** of the AMF. I am now in a large Hedge Fund.

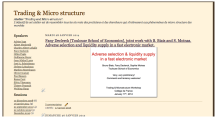

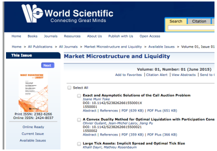

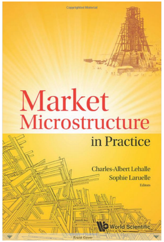

- From a **Financial Mathematics** perspective, it is nothing more than adding a variable to our models: the **Liquidity** .
- The interactions between liquidity and other (usual) variables is far from trivial.

Disclaimer: I express *my own opinion* and not the one of any of these institutions.

What Will We Do During 3 Lectures + 1 Seminar?

I will not go in the details of the models (except for few of them), because I target to give you enough information to include liquidity in the models you know better than me.

Hence, I will

- 18 Nov:
  - Start by the definition of **intermediation**
  - Focus on the two main **Liquidity** variables on financial market: inventories and flows
- 25 Nov:
  - Show you what **Liquidity** looks like when we can observe it
- 2 Dec:
  - Underline why **market making** (inventory keeping) and **optimal trading** (flow management) are core for the new role of market participants.
- 2 Dec [Seminar]:
  - Explain what practitioners are doing.

It is an on-going work

My own viewpoint on optimal trading:

- We have sophisticated (but tractable) methods to **optimize the strategy of one agent** (investment bank, trader, asset manager, etc) facing a “background noise” (stochastic control is now really mature),
- These methods are **used by practitioners** (already three books on this topic [Lehalle et al., 2013], [Cartea et al., 2015], [Guéant, 2016]),
- Differential games, and more specifically **mean field games** now propose very promising frameworks to replace most of the background noise by a mean field of explicitly modelled agents:
  - to provide robust results for practitioners [Cardaliaguet and Lehalle, 2016],
  - to obtain meaningful results for policy recommandations [Lachapelle et al., 2016].

Up to now most results on global modelling used a simplification of a reality. Now decisions are modelled and systematic, why not inject them into a global model?
It should enable you to produce very accurate models and draw powerful conclusions.

- Beyond optimal trading, these lectures should help you in introducing liquidity in any model of yours: **please ask question!**

Risks Transformation as The Primary Role of The Financial System

To understand the interactions between actors of financial markets, a first step is to understand the **role of the financial system**.

It takes its role at the root of capitalism:

- say you see a shoes shiner at Deli, India
- you pay $1 to have your shoes shined, and you ask to the guy
- "it seems you have around 30 customers each day, it let you with $30 every day, it is a good job."
- he answers: "not at all, I earn $1 a day... I do not own the brush, its owner loans its to me $29 a day. Since a brush costs $12 and I need my daily dollar to eat, I will never own one."

→ let's discuss about microcredit: loan him $12 during 2 days...

You have $30, you can ask to the guy some percents to cover the risk he will not have enough clients. If you are risk averse, you can even ask for the brush as collateral... A bank can "structures" the loan for you, it will take care of all the administrative aspects, it is a simple **risk transformation** (liquidity on you side, business of the shoes shiner side).

FIGURE 1: SAWTOOTH PATTERNS ON COCA-COLA, MCDONALD'S, IBM AND APPLE ON 19 JULY 2012

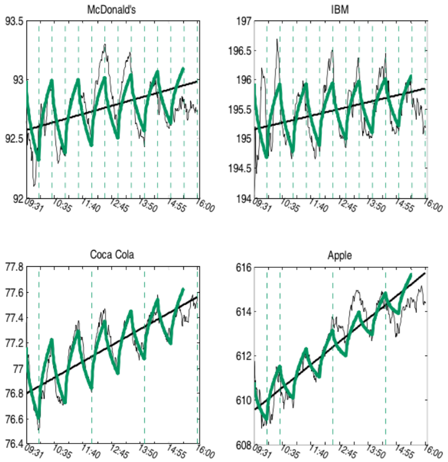

Even on liquid stocks and for vanilla options (close to maturity in this case), hedging can go wrong.

The 19th of July 2012, a trading algorithms bought and sold shares every 30 minutes without any views on its market impact [Lehalle et al., 2012].

For one visible mistake like this on liquid underlyings of vanilla products, how many bad sophisticated hedging processes on less liquid (even OTC) markets...
Anonymous continuous hedging of a remaining position outside of the bank does not mean all is going well.
Nevertheless we have solutions in recent literature: [Guéant and Pu, 2013], [Li and Almgren, 2014].

But nothing more generic, for instance the whole process of hedging books in presence of wrong way risk is not studied (as far as I know). One step in this direction is [Schied and Zhang, 2013].

My advices to an investment bank:

- **Net all your books** , maintain two opposite positions is costly and risky,
- If you can't it may be because you do not communicate enough internally (sometimes because of Chinese walls...), hence **be ready to hedge on the market** ,
- But **before try to match your small metaorders** : send them to an internal place and cross them as much as possible;
- You will have synchronization issues (at the level of these metaorders, no reason to be synchronized), ask to your traders to **implement facilitation-like market making** schemes inside the bank.
- The remaining quantity has to be sent to markets as smoothly as possible, but it does not mean you will have no impact. **Who is your counterpart in the market** should be an obsession: if you trade a one way risk, you will pay for this in the future...

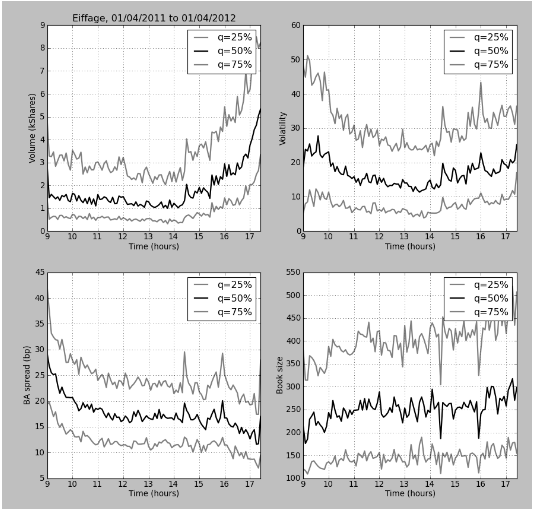

Intraday Seasonalities Essentials

- Volumes are U-shaped, log-volumes are close to Gaussian,
- Volatility are U-shaped too (less intense at the end than at the start of the day),
- Volatility is "more path dependent" than volumes,
- BA-spread is large at the start of the day, but finishes small because of market maker running to get rid of their inventory passively,
- "Volume on the Book" (i.e. $Q^A + Q^B)/2$) seasonality is the invert of the one of BA-spread. The more the spread is constraied by the tick, the more the seasonality is strong on the volume-ob-the-book.
- News implies peaks of volume / volatility,
- Activity on other markets has an influence.

Optimal trading is about optimizing a trading process.

- Once a decision to buy or sell has been taken (i.e. a **metaorder** has been issued),
- Market impact and trading costs can be **minimized**.
- The process of inserting and cancelling orders in orderbooks has to be automated.

Asset managers delegate to their **dealing desk** the process of buying or selling. They give to the trading desk specifications (speed, exposure, etc) to be fulfilled. For instance, some metaorders have to be executed fast, others more slowly.
The dealing desks are reporting to portfolio managers information about liquidity of traded instruments (to take liquidity into account during the allocation process). And give them "technical" advices, like effects of mergers, closed days, potentially cheaper instruments.

Specific participants (like high frequency traders of fast hedge funds) take decisions because of **liquidity signals**. In their case the buy / sell decision is intricated with their trading process.

Last but not least **Market Makers** use optimal trading to provide liquidity to other market participants.

- Batch (Fixing) auctions or Continuous auctions,
- price driven or order driven logic,
- bilateral or multilateral trading.

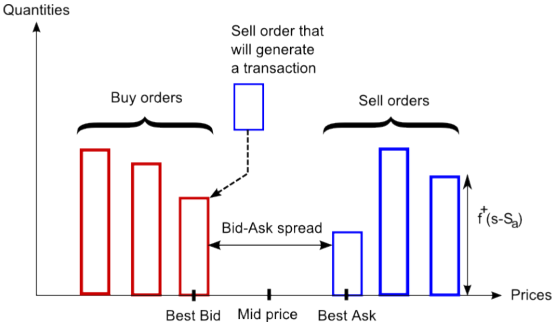

A Limit Order Book (LOB) hosts multilateral, order driven, continuous double auctions.

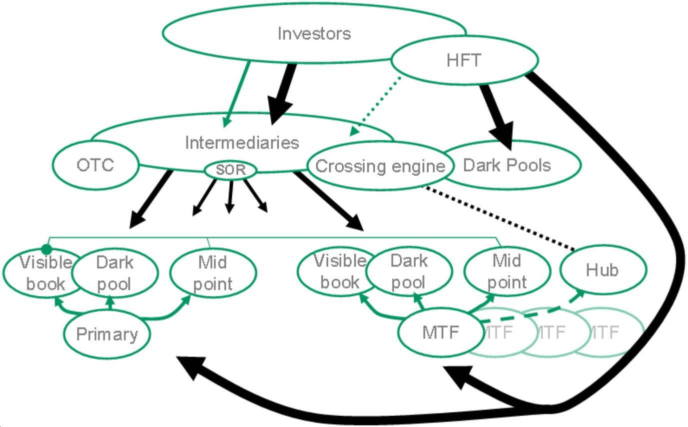

The trading now takes place on a distributed network of heterogenous trading platforms.

Fragmentation is the natural counterpart of **competition** : if you want the exchanges (i.e. trading venues) to compete.

Fragmentation is the natural counterpart of **competition** : if you want the exchanges (i.e. trading venues) to compete.

Have you ever seen a mailbox in France?

Fragmentation is the natural counterpart of **competition** : if you want the exchanges (i.e. trading venues) to compete.

Have you ever seen a mailbox in France?

This is competition (and fragmentation)

INSIGHT.DATA.CLARITY.

A typical fragmented stock

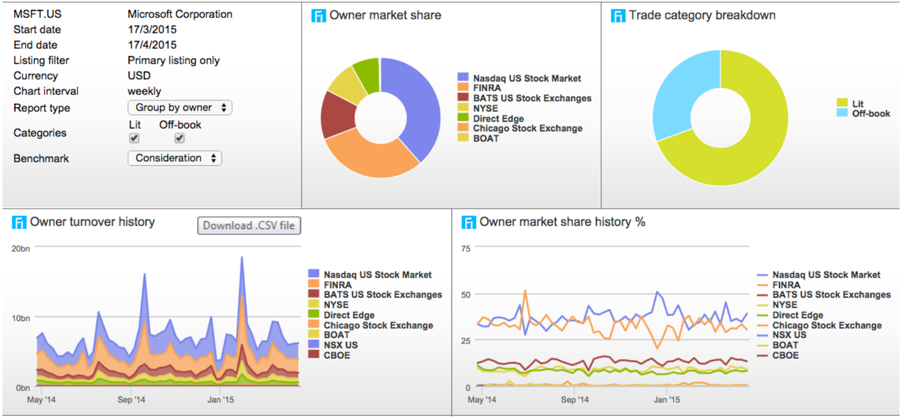

The Fragmentation of Microsoft the last 20 days (Source: Fidessa's fragulator)

A typical fragmented stock

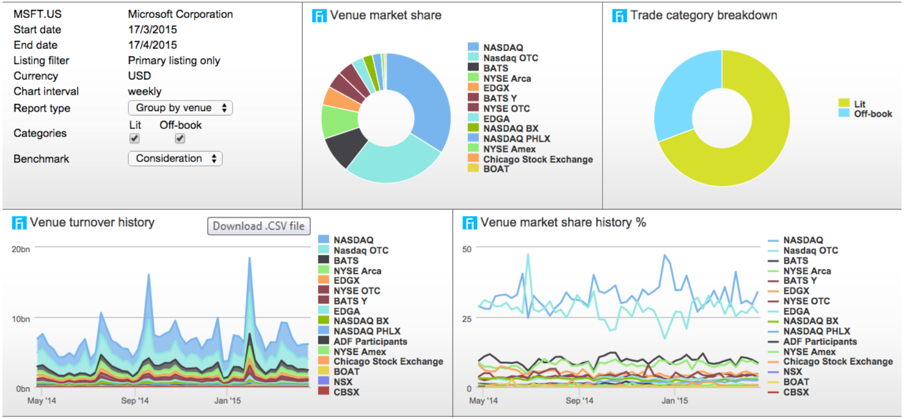

The Fragmentation of Microsoft the last 20 days (Source: Fidessa's fragulator) **more...**

Smart Routing of Limit Orders

- A limit order (i.e. buy –resp. sell– order at a price lower than the best ask –resp. best bid–). It has to be split too.
- There is structurally **uncertainty** on limit orders splitting, waiting on a bad queue generates **opportunity costs**.
- Say on venue n you seen a queue of size Q_n and you expect the **consuming flow** to follow a Poisson process with intensity Λ_n.

Goal **Queue Jumping**: Choose $(q_1, \cdot, q_N)$ so that the full quantity $Q^* = \sum_n q_n$ is *on average* executed as fast as possible.

- It means to focus on t_n such that $\int_0^{t_n} dN_t^n = q_n + Q_n$, which expectation reads $t_n \Lambda_n = q_n + Q_n$, where $\sum_n q_n = Q^*$.
- We do not need a Lagrangian here, it is enough to note *minimizing the minimum of all t_n implies $t^* = t_n$ for any n*. Hence $t^* = Q^* / \sum_n \Lambda_n + \sum_n Q_n / \sum_n \Lambda_n$. With the convenient notations $\rho_n := \Lambda_n / (\sum_\ell \Lambda_\ell / N)$ and $\bar{Q} := \sum_n Q_n / N$, we obtain

$$\boxed{q_n^* = \rho_n \frac{Q^*}{N} + (\rho_n \bar{Q} - Q_n)}.$$

When more than one trading destination are available (ECNs in the US, Multilateral Trading Facilities -MTF- in Europe):

- each of them provides a specific flow $\phi_t^{(i)}$,
- keeping ΔT constant over the trading destinations, each liquidity pool will be able to deliver a quantity D_i

Dark Pools are specific trading destinations because:

- they do not provide pre trade transparency about their limit order books
- you ask for V and you have $\min(V, D_i)$ back
- they allow "pegged" orders: you can specify δS rather than a limit price ("pinging" implies $\Delta T = 0$):

$$D_i = \int_{t=\tau}^{\tau+\Delta T} \phi_t^{(i)}(\delta S)\, dt$$

- The stationary solutions of the ODE: $\dot{x} = h(x)$ contains the extremal values of $F(x) = \int_0^x h(x)\,dx$
- A discretized version of the ODE is (γ is a step):

(1) $$x_{n+1} = x_n + \gamma_{n+1}\, h(x_n)$$

- A stochastic version of this being (ξ_n are i.i.d. realizations of a random variable, $h(X) = \mathbb{E}(H(X, \xi_1))$):

(2) $$X_{n+1} = X_n + \gamma_{n+1}\, H(X_n, \xi_{n+1})$$

- the stochastic algorithms theory is a set of results describing the relationship between these 3 formula and the nature of γ, H, h and ξ [Hirsch and Smith, 2005], [Kushner and Yin, 2003], [Doukhan, 1994]

Stochastic algorithms theory can be used when you only have a sequential access to a functional you need to minimize:

- to minimize a criteria $\mathbb{E}(F(X, \xi_1))$ of a state variable X
- if it is possible to compute:

$$H(X_n, \xi_{n+1}) := \frac{\partial F}{\partial X}(X_n, \xi_{n+1})$$

- the results of the stochastic algorithms theory (like the Robbins-Monro theorem [Pagès et al., 1990]) can be used to study the properties of the long term solutions of the recurrence equation:

$$X_{n+1} = X_n + \gamma_{n+1} H(X_n, \xi_{n+1})$$

- at high frequency, historical statistics are not so useful
- the limit price S and the quantity V are random variables,
- the executed quantity on dark pools has to be maximized (it is *market impact free*) and sometimes fees are different; this effect is modelled by a *discount factor* $\theta_i \in (0,1)$ (normalized with respect to a “reference” Lit pool)
- the quantity V is split into N parts (one for each DP): $r_i \times V$ is sent to the ith DP ($\sum_{i=1}^{N} r_i = 1$)

See *Optimal split of orders across liquidity pools: a stochastic algorithm approach* [Pagès et al., 2011] for more details.

The remaining quantity is to be executed on a reference Lit market, at price S.

The cost C of the whole executed order is given by

$$\begin{aligned} C &= S\sum_{i=1}^{N}\theta_i \min\left(r_i V, D_i\right) + S\left(V - \sum_{i=1}^{N}\min\left(r_i V, D_i\right)\right) \\ &= S\left(V - \sum_{i=1}^{N}\rho_i \min\left(r_i V, D_i\right)\right) \end{aligned}$$

where

$$\rho_i = 1 - \theta_i \in (0,1), i = 1,\ldots,N.$$

Minimizing the mean execution cost, *given the price S*, amounts to:

Maximization problem to solve

(3) $$\max\left\{\sum_{i=1}^{N}\rho_i\mathbb{E}\left(S\min\left(r_iV,D_i\right)\right),r\in\mathcal{P}_N\right\}$$

where $\mathcal{P}_N := \left\{r=(r_i)_{1\le i\le N}\in\mathbb{R}_+^N \mid \sum_{i=1}^N r_i=1\right\}$.

It is then convenient to *include the price S into both random variables V and D_i* by considering $\widetilde{V} := V\,S$ and $\widetilde{D}_i := D_iS$ instead of V and D_i. Assume that the distribution function of D/V is continuous on $\mathbb{R}_+$. Let $\varphi(r)=\rho\,\mathbb{E}\left(\min\left(rV,D\right)\right)$ be the mean execution function of a single dark pool ($\Phi=\sum_i\varphi_i(r_i)$), and assume that $V>0\ \mathbb{P}-a.s.$ and $\mathbb{P}(D>0)>0$

Skip Details

We consider the sequence $Y^n := \left(V^n, D_1^n, \ldots, D_N^n\right)_{n\geq 1}$.
We will take two types of stationarity assumptions on the sequence

$$(IID) \equiv \begin{cases} (i) & \text{the sequence } (Y^n)_{n\geq 1} \text{ is i.i.d. with distribution} \\ & \nu = \mathcal{L}(V, D_1, \ldots, D_N) \text{ on } \left(\mathbb{R}_+^{N+1}, \mathcal{B}(\mathbb{R}_+^{N+1})\right). \\ (ii) & V \in L^2(\mathbb{P}). \end{cases}$$

$$(ERG)_\nu \equiv \begin{cases} (i) & \text{the sequence } (Y^n)_{n\geq 1} \text{ is averaging } i.e. \\ & \mathbb{P}\text{-}a.s. \quad \frac{1}{n}\sum_{k=1}^{n} \delta_{(Y^k)} \overset{(\mathbb{R}_+^{N+1})}{\Longrightarrow} \nu = \mathcal{L}(V, D_1, \ldots, D_N), \\ (ii) & \sup_n \mathbb{E}(V^n)^4 < +\infty. \end{cases}$$

The Lagrangian Approach

We aim at solving the following maximization problem

$$\max_{r \in \mathcal{P}_N} \Phi(r) \tag{4}$$

The Lagrangian associated to the sole affine constraint is

$$L(r,\lambda) = \Phi(r) - \lambda \left(\sum_{i=1}^{N} r_i - 1 \right) \tag{5}$$

So,

$$\forall i \in \mathcal{I}_N, \quad \frac{\partial L}{\partial r_i} = \varphi_i'(r_i) - \lambda.$$

This suggests that any $r^* \in \arg\max_{\mathcal{P}_N} \Phi$ iff $\varphi_i'(r_i^*)$ is constant when i runs over $\mathcal{I}_N$ or equivalently if

$$\forall i \in \mathcal{I}_N, \quad \varphi_i'(r_i^*) = \frac{1}{N} \sum_{j=1}^{N} \varphi_j'(r_j^*). \tag{6}$$

To ensure that the candidate provided by the Lagragian approach is the true one, we need an additional assumption on φ to take into account the behaviour of Φ on the boundary of $\partial\mathcal{P}_N$.

Proposition 1

Assume that (V, D_i) satisfies upper assumptions for every $i \in \mathcal{I}_N$. Assume that the functions φ_i satisfy the following assumption

$$(\mathcal{C}) \equiv \min_{i\in\mathcal{I}_N} \varphi_i'(0) > \max_{i\in\mathcal{I}_N} \varphi_i'\left(\frac{1}{N-1}\right). \tag{7}$$

Then $\arg\max_{\mathcal{H}_N} \Phi = \arg\max_{\mathcal{P}_N} \Phi \subset int(\mathcal{P}_N)$ where

$$\arg\max_{\mathcal{P}_N} \Phi = \left\{r \in \mathcal{P}_N \mid \varphi_i'(r_i) = \varphi_1'(r_1), i = 1, \ldots, N\right\}.$$

Note that

$$\max \sum_{i=1}^{N} \rho_i \mathbb{E}(\min(r_i V, D_i)), \; \sum_i r_i = 1$$

$$\Leftarrow \quad \forall i : \; \mathbb{E}(\rho_i V \mathbf{1}_{r_i V < D_i}) = \lambda$$

$$\Leftrightarrow \quad \forall i : \; \mathbb{E}(\rho_i V \mathbf{1}_{r_i V < D_i}) = \frac{1}{N} \sum_j \mathbb{E}(\rho_j V \mathbf{1}_{r_j V < D_j})$$

i.e. “if each $f(i) = \lambda$ then each of them equals their average”.

Using the representation of the derivatives φ_i' yields that, if Assumption $(\mathcal{C})$ is satisfied, then

Characterization of the solution

$$r^* \in \arg\max_{\mathcal{P}_N} \Phi \Leftrightarrow \forall i \in \{1,\ldots,N\}\,, \mathbb{E}\left(V\left(\rho_i \mathbf{1}_{\{r_i^* V < D_i\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j^* V < D_j\}}\right)\right) = 0.$$

Using the representation of the derivatives φ_i' yields that, if Assumption $(\mathcal{C})$ is satisfied, then

Characterization of the solution

$$r^* \in \arg\max_{\mathcal{P}_N} \Phi \Leftrightarrow \forall i \in \{1, \ldots, N\}, \mathbb{E}\left(V\left(\rho_i \mathbf{1}_{\{r_i^* V < D_i\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j^* V < D_j\}}\right)\right) = 0.$$

Consequently, this leads to the following recursive zero search procedure

$$r_i^{n+1} = r_i^n + \gamma_{n+1} H_i(r^n, Y^{n+1}),\ r^0 \in \mathcal{P}_N,\ i \in \mathcal{I}_N, \tag{8}$$

where for $i \in \mathcal{I}_N$, every $r \in \mathcal{P}_N$, every $V > 0$ and every $D_1, \ldots, D_N \geq 0$,

$$H_i(r, (V, D_1, \ldots, V_N)) = V\left(\rho_i \mathbf{1}_{\{r_i V < D_i\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j V < D_j\}}\right)$$

Back to the Notation of Stochastic Algorithms

When we design a procedure using $H(r,\xi_1)$, we potentially converge to the extrema of $F(r) := \int h(r)$, where $h(r) := \mathbb{E}_{\xi_1} H(r,\xi_1)$.
Here $\xi_t := (V(t), D_1(t), \ldots, D_N(t))$, hence we can use the follwing stochastic procedure:

$$\begin{aligned} r_i(t+1) &= r_i(t) + \gamma(t) \cdot H_i(r(t),(V(t),D_1(t),\ldots,V_N(t))) \\ &= r_i(t) + \gamma(t) \cdot V(t) \cdot \left(\rho_i \mathbf{1}_{\{r_i(t)V(t)<D_i(t)\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j(t)V(t)<D_j(t)\}} \right) \end{aligned}$$

$\Rightarrow (r_1(\infty), \cdots, r_N(\infty))$ will be our solution, i.e. **the optimal split between Dark Pools**.

When we design a procedure using $H(r,\xi_1)$, we potentially converge to the extrema of $F(r) := \int h(r)$, where $h(r) := \mathbb{E}_{\xi_1} H(r,\xi_1)$.
Here $\xi_t := (V(t), D_1(t), \ldots, D_N(t))$, hence we can use the follwing stochastic procedure:

$$\begin{aligned} r_i(t+1) &= r_i(t) + \gamma(t)\cdot H_i(r(t), (V(t), D_1(t), \ldots, V_N(t))) \\ &= r_i(t) + \gamma(t)\cdot V(t)\cdot \left(\rho_i \mathbf{1}_{\{r_i(t)V(t)<D_i(t)\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j(t)V(t)<D_j(t)\}} \right) \end{aligned}$$

$\Rightarrow (r_1(\infty), \cdots, r_N(\infty))$ will be our solution, i.e. **the optimal split between Dark Pools**.

The underlying idea of the algorithm

is to reward the dark pools which outperform the mean of the N dark pools by increasing the allocated volume sent at the next step (and conversely).

In this algorithm, we took into account the constraint

$$\sum_{i=1}^{N} r_i = 1,$$

but not

$$r_i > 0, \ \forall 1 \le i \le N.$$

So the algorithm may exit from the simplex $\mathcal{P}_N$ stable. To overcome this problem, we have two possibilities

1. Use a Lyapunov function and a strong *mean-reverting* assumption out of $\mathcal{P}_N$: this solution is simpler from a mathematical point of view.
2. Force the coefficients r_i to stay in $\mathcal{P}_N$ by a truncation-projection procedure: this solution is more efficient for applications.

Theorem 1: Convergence

Assume that (V, D) satisfy upper assumptions, that $V \in L^2(\mathbb{P})$ and that Assumption $(\mathcal{C})$ holds. Let $\gamma := (\gamma_n)_{n\geq 1}$ be a step sequence satisfying the usual decreasing step assumption

$$\sum_{n\geq 1} \gamma_n = +\infty \quad \text{and} \quad \sum_{n\geq 1} \gamma_n^2 < +\infty.$$

Let $(V^n, D_1^n, \ldots, D_N^n)_{n\geq 1}$ be an i.d.d. sequence defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Then, there exists an $\mathrm{argmax}_{\mathcal{P}_N}\Phi$-valued random variable r^* such that

$$r^n \longrightarrow r^* \quad a.s.$$

Stochastic Algorithm for Optimal Trading (Conclusion)

We have seen

- how to improve a deterministic simple rule / optimization to a stochastic one.
- If you write properly the criterion to minimize / maximize, and if you can observe its derivative
- then **it is valuable to build a rigorous stochastic algorithm**
- to make the balance between exploration and exploitation.

In the case of Dark trading two other proposals:

- [Ganchev et al., 2010] uses cencored statistics to establish a robust confidence interval $[D_i^{\min}, D_i^{\max}]$ around each D_i, and then perform a determinstic optimization on $(D_1^{\min}, \ldots, D_N^{\min})$;
- [Agarwal et al., 2010] uses minimum regret and maintain a huge table of available *quantities* (and not *proportions*).

Moreover you can have a similar approach to choose the price to post in a limit orderbook (cf. *learning by trading* [Laruelle et al., 2013]).

Outline

1. The Financial System as a Network of Intermediaries
   - Risks Transformation as The Primary Role of The Financial System
2. Stylized Facts on Liquidity
3. Optimal Trading
   - Learning by Trading (in The Dark)
   - Trading Benchmarks
   - Optimal Trading Rate
   - Optimal Trading Against Permanent Impact: stylized facts
   - Optimal Control of Trading Robots
4. Conclusion
5. Closing The Loop
   - MFG of Controls
   - Kyle's Model

Trading Algorithms: Typical Features

Benchmark	Type of stock	Type of trade	Main feature
PoV	Medium to large market depth	(1) Long duration position	(1) Follows current market flow, (2) Very reactive, can be very aggressive, (3) More price opportunity driven if the range between the max percent and min percent is large
VWAP / TWAP	Any market depth	(1) Hedging order, (2) Long duration position, (3) Unwind tracking error (delta hedging of a fast evolving inventory)	(1) Follows the “usual” market flow, (2) Good if market moves with unexpected volumes in the same direction as the order (up for a buy order), (3) Can be passive
Implementation Shortfall (IS)	Medium liquidity depth	(1) Alpha extraction, (2) Hedge of a non-linear position (typically Gamma hedging), (3) Inventory-driven trade	(1) Will finish very fast if the price is good and enough liquidity is available, (2) Will “cut losses” if the price goes too far away
Liquidity Seeker	Poor a fragmented market depth	(1) Alpha extraction, (2) Opportunistic position mounting, (3) Already split / scheduled order	(1) Relative price oriented (from one liquidity pool to another, or from one security to another), (2) Capture liquidity everywhere, (3) Stealth (minimum information leakage using fragmentation)

Trading Algorithms: Typical Uses

Benchmark	Region of preference	Order characteristics	Market context	Type of hedged risk
PoV	Asia	Large order size (more than 10% of ADV: Average daily consolidated volume)	Possible negative news	Do not miss the rapid propagation of an unexpected news event (especially if I have the information)
VWAP / TWAP	Asia and Europe	Medium size (from 5 to 15% of ADV)	Any "unusual" volume is negligible	Do not miss the slow propagation of information in the market
Implementation Shortfall (IS)	Europe and US	Small size (0 to 6% of ADV)	Possible price opportunities	Do not miss an unexpected price move in the stock
Liquidity Seeker	US (Europe)	Any size	The stock is expected to "oscillate" around its "fair value"	Do not miss a liquidity burst or a relative price move on the stock

More on all this in the three "reference books" for practitioners:

- Market Microstructure in Practice [Lehalle et al., 2013]
- The Financial Mathematics of Market Liquidity [Guéant, 2016]
- Algorithmic and High-Frequency Trading [Cartea et al., 2015]

Optimal Trading Rate

The first papers [Almgren and Chriss, 2000], [Bertsimas and Lo, 1998], focussed on the optimal trading rate, or trading speed (i.e. how many shares to buy or sell every 5 minutes) for long metaorders.

- it does not deal with microscopic orderbook dynamics,
- it is a convenient way to take into account any information or constraint at this time scale.

It is very useful for asset managers, brokers, or hedgers. I.e. especially when the decision step is separated from the execution step.
Nevertheless it can be used for opportunistic trading too, when risk management at an intraday scale is important.

I have to buy V^* share between time 0 and T ; I assume a regular temporal grid with a time step of δt (n goes from 1 to $N = [T/\delta t]$). My volume is splitted in N slices v_n such that $\sum_n v_n = V^*$ ([Almgren and Chriss, 2000]) The price follows a Brownian motion:

$$S_{n+1} = S_n + \alpha\,\delta t + \sigma_{n+1}\sqrt{\delta t}\,\xi_{n+1} \tag{9}$$

The additive, temporary only, market impact function is $\eta_n(v_n)$. Then the total cost is:

$$W = \sum_{n=1}^{N} v_n \cdot (S_n + \eta_n(v_n)) \tag{10}$$

Of course, far more sophisticated models have been studied: [Almgren and Lorenz, 2007], [Almgren, 2009], [Bouchard et al., 2011]

When the market impact is linear with respect to the participation rate v_n / V_n and proportional to the volatility (change of variable $x_n = \sum_{k=n}^{N} v_n$).

$$W = \underbrace{V^* S_0}_{\text{immediate cost in a "free" world}} + \underbrace{\sum_{n=1}^{N} x_n \sigma_n \xi_n}_{\text{market risk}} + \underbrace{\sum_{n=1}^{N} \eta \sigma_n \frac{v_n^2}{V_n}}_{\text{market impact}}$$

For a broker algo, we want to minimize a mean-variance criteria:

$$J_\lambda = \mathbb{E}(W | V_1, \ldots, V_N, \sigma_1, \ldots, \sigma_N) + \lambda \mathbb{V}(W | V_1, \ldots, V_N, \sigma_1, \ldots, \sigma_N) \quad (11)$$

We obtain a recurrence equation in x_n:

$$x_{n+1} = \left(1 + \frac{\sigma_{n-1}}{\sigma_n} \frac{V_n}{V_{n-1}} + \frac{\lambda}{\eta} \sigma_n V_n\right) x_n - \frac{\sigma_{n-1}}{\sigma_n} \frac{V_n}{V_{n-1}} x_{n-1}$$

A New Set of Equations to Account For Uncertainty on (V, σ) ▷

When the V_n and σ_n are i.i.d. random variables, and when V_n and ξ_n are correlated:

$$\mathbb{E}(W) = v^* S_0 + \sum_{n=1}^{N} \eta\, \mathbb{E}\left(\frac{\sigma_n}{V_n}\right) (x_n - x_{n+1})^2$$

$$\begin{aligned}\mathbb{V}(W) \quad = \quad & \sum_{n=1}^{N} x_n^2 \sigma_n^2 + \sum_{n=1}^{N} \eta\, x_n (x_n - x_{n+1}) \mathbb{E}(\sigma_n)\, \mathbb{E}\left(\frac{\xi_n}{V_n}\right) \\ & + \sum_{n=1}^{N} \eta^2\, \mathbb{V}\left(\frac{\sigma_n}{V_n}\right) (x_n - x_{n+1})^4\end{aligned}$$

It is easy to add a lot of other effects like: noise on “expected” market impact, auto-correlations, volume-volatility coupled-dynamic. No more closed-form formula.

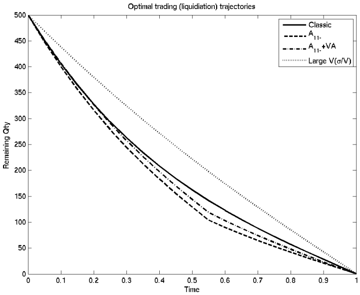

It introduced the idea of **optimal trading curves** $\rightarrow$ crucial for risk control.
A lot of effects can be easily added to the AC framework:

- seasonalities and predictions of V and σ can be plugged,
- arbitrage opportunities can be added, [Lehalle, 2013]

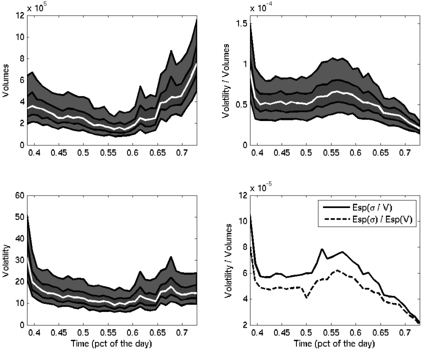

It introduced the idea of **optimal trading curves** $\rightarrow$ crucial for risk control.

A lot of effects can be easily added to the AC framework:

- seasonalities and predictions of V and σ can be plugged,
- arbitrage opportunities can be added, [Lehalle, 2013]
- $\mathbb{E}(W|V,\sigma) + \lambda\mathbb{V}(W|V,\sigma)$ can be replaced by $\mathbb{E}(W) + \lambda\mathbb{V}(W)$, to take uncertainty into account, [Lehalle, 2008],
- backtest parametric trading curves directly, etc.

Open questions (within this framework)

- the control is the trading rate,
- the choice of the criterion can be discussed (PoV, VWAP, TWAP, TC, etc.),
- the **variance** term has a strong influence,
- what can you choose ? (λ).

Why introduce a variance term in the cost function?

- with $\lambda = 0$ and an explicit (exponential) form of orderbook relaxation, Alfonsi and Schied elaborated in this direction ([Alfonsi et al., 2009, Gatheral et al., 2012]). A U-shaped trading curve generally stems from such choice (due to market impact relaxation and blindness after the last transaction).

$\Rightarrow$ Market impact decays implies a U-shaped trading rate.

- The choice of a variance term (instead of any *p*-variation) can be discussed, too [Labadie and Lehalle, 2014]. Especially if you think intraday price dynamics are *more mean reverting* than daily ones.

Practical Use of Automated Trading: Adding Constraints

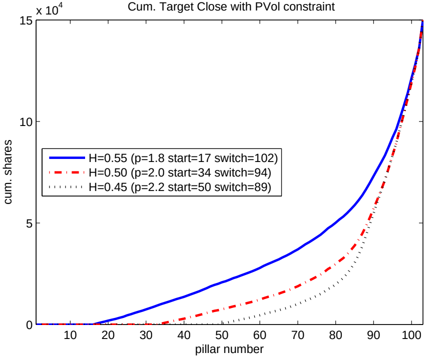

The Almgren-Chriss criterion is the **Implementation Shortfall** (i.e. $W(v)$), but other "trading styles" are possible, like **VWAP** (follow the usually traded volume [Cartea and Jaimungal, 2014]), **PoV** (follow the real-time market volume), and **Target Close**.

- The latter targets the closing fixing, trying to avoid a too large impact.
- It limits the volume in the fixing auction at q%,
- And does the remaining in an Almgren-Chriss way on $W(v) - S_T$

In practice users put some **participation constraints** to their trading flow (i.e. $v < \rho V$). For the Target Close it raises an interesting problem, especially if you use an estimate of the future volume or if you want to start you European trading for sure after the opening of US markets. With Mauricio L, we proposed a model and solved it for fractional Brownian motions [Labadie and Lehalle, 2014].

The usual (simplistic) example of (continuous time) optimal trading

1. Write the Markovian dynamics or the price P, the quantity to trade Q and the cash account X for a sell of Q_0 shares before $t = T$ (control is the trading speed r)

$$dQ = -\nu\, dt,\ dX = r\,(P - \kappa \cdot \nu)dt,\ dP = \mu\, dt + \sigma\, dW.$$

2. Write the cost function to minimize

$$V(t, p, q, x, \nu) = \mathbb{E}\left(X_T + Q_T(P_T - A \cdot Q_T) + \phi \int_{\tau=t}^{T} Q_\tau^2\, d\tau \middle| \mathcal{F}_t\right).$$

3. it gives the HJB and its terminal condition $V(T_f, \ldots) = x + q(p - Aq)$; (here $\mu = 0$)

$$0 = \partial_t V + \frac{\sigma^2}{2}\partial_P^2 V + \phi\, q^2 + min_\nu \left\{-\nu \partial_Q V\, dt + \nu(p - \kappa \cdot \nu)\partial_X V\right\}.$$

4. After the change of variable $V(t,p,q,x) = x + q\,p + v(t,q)$, you have

$$0 = \partial_t v + \phi\, q + \min_{\nu} \left\{ -\nu \partial_Q v - \kappa \nu^2 \right\}.$$

5. The optimal control is $\nu^* = -\partial_Q v/(2\kappa)$, and the PDE $0 = \partial_t v + \phi\, q + \kappa(\partial_Q v)^2/(4\kappa)$.
6. When the value function is quadratic: $v(t,q) = h_0(t) + q\,h_1(t) - q^2\,h_2(t)/2$, you can separate the PDE in three:

$$\left\{ \begin{array}{ccccc} 0 & = & h_0' & + & h_1^2 \\ 0 & = & h_1' & - & 2h_1 h_2 \\ 0 & = & -2\kappa h_2' & + & h_2^2 \end{array} \right.$$

Cartea and Jaimungal (with misc. co-authors) developed this framework for plenty versions: with a (slightly) different objective function (VWAP, PoV), with permanent market impact $\mu \to \mu + \nu$, with μ_t any (adapted) process, etc.

Optimal Trading Against Permanent Impact: stylized facts

On our database of 300,000 large orders

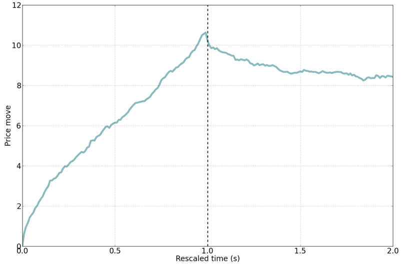

Market Impact takes place in different phases

- the **transient impact**, concave in time,
- reaches its maximum, the **temporary impact**, at the end of the metaorder,
- then it **decays**,
- up to a stationary level; the price moved by a **permanent** shift.

In [Bacry et al., 2015] we studied all the phases, using intraday and daily analysis (for the first time). We underlined the importance of some "normalization variables": the uncertainty on the price formation process , the capability of the orderbook to resist to volume pressure , and the duration of the metaorder.
Following [Waelbroeck and Gomes, 2013] and simultaneously with [Brokmann et al., 2014], we proposed an explanation of **permanent impact** .

Is it Really Possible to Minimize Permanent Impact?

We used an Hawkes-based toy model to show how the concavity of the market impact and the decay can form. Nevertheless an external parameter C is needed to control the permanent market impact. In such a framework the intensity of the Hawkes process can be seen as **an implicit inventory of the market makers**.

- Part of the price move while an asset manager is buying is due to its trading activity,
- But evidences on the permanent components could be explained by an **informational** effect: you buy because you anticipated the price will move. Buy or not: it will move in any case!

This effect seems to have been identified by Waelbroeck and Gomes on “cash trades”, by the CFM team on daily “deconvoluted trades”, and by us on the idiosyncratic component of price moves.

- The best way to take such permanent market impact into account could simply be to add a deterministic trend to price dynamics (i.e. “usually, when I decide to buy, the price will go up that way”)... Beside, it could justify the use of a variance term in the cost function.

This kind of analysis is linked with a potential understanding of the whole market dynamics and the way prices form.

CFM
INSIGHT.DATA.CLARITY.

We have seen

- Almgren and Chriss framework is **simple and flexible**. It can be seen as a way to include any statistical property of medium-term (5 to 30 min) dynamics in your trading style,
- It does not deal with **orderbook dynamics**. See [Guéant, 2016] for a lot of variations.
- It gives a **deterministic trading rate**.
- But it does not say how to choose the **risk aversion**. This parameter control the **urgency** of trading.
- With small changes to the criterion, Cartea and Jaimungal provide a way to obtain a **stochastic trading rate**. See [Cartea et al., 2015] for variations.
  Because the change affects the (not so well defined) risk aversion, it is not that a big deal.

⊕ Risk aversion is a way to deal with the **informational component of permanent market impact**. See [Lehalle et al., 2013] for a global overview.

⇒ 3 books for optimal trading...

Outline

In practice if you want to rely on previous frameworks (i.e. optimal trading rates), you need a way to realize it thanks to (optimal?) interactions with orderbook dynamics.

1. Use your favorite method to implement a solver and obtain (in real time) your **optimal trading rate** ,
2. Have statistics ready on a server with the needed parameters (market impact model, risk aversion, calendar with scheduled events, etc).

$\Rightarrow$ At any point in time, you know the **trading rate** $\nu_t\,\delta t$ you plan to obtain in the next δt seconds.

3. Give it to another logic, focussed on the short term and taking profit of orderbook dynamics. I will call this logic a **Trading Robot** , and a priori I will assume it should adopt an exploration-exploitation scheme.

Typically I will use a stochastic algorithm to design my Trading Robot and thanks to backtests and production results I will have a clear idea of its “usual returns”. I.e. when I ask to Robot A to buy $\nu_t\,\delta t$ shares in δt on this kind of instrument, it obtain Q shares with a price improvement of ΔP (two random variables).

Optimal Trading Against Orderbook Dynamics

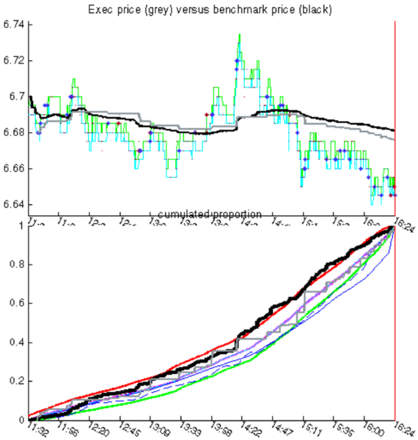

This idea of a decision process in two scales is close to reality:

1. the investor takes a decision according to its **views on the price / risk**, but he does not buy or sell himself;
2. he delegates to a **executing broker** or to a trading algorithm the trading process.

At the scale of the large order itself there is a similar split in two scales:

- a **scheduler** or a human trader takes care of a trading curve (close to the outcome of an Almgren-Chriss optimization),
- it uses **trading robots** for high frequency interactions with the orderbook.

Optimizing a Two-Scales Controller

The control

$$\nu = (\tau_i^\nu, \delta_i^\nu, \mathcal{E}_i^\nu)_i,$$

The dynamics

$$X^\nu(t) = \quad x_0 + \int_{s\leq t}(b\,dt + \sigma\,dW) + \sum_i \beta(\tau_i^\nu, \delta_i^\nu, \mathcal{E}_i^\nu)\mathbf{1}_{\tau_i^\nu \leq t},$$

The gain

$$\sum_{\tau_i^\nu < T} f(X(\tau_i^\nu + \delta_i^\nu), \mathcal{E}_i^\nu) + g(X^\nu(T)).$$

In [Bouchard et al., 2011] , we developed a model describing this two scales process:

- The Scheduler launches **trading robots**, known in probabilistic terms (the joint laws of their duration and their efficiency) at any time.
- It is an **impulse control problem** embedded in a continuous-time framework: the controls are the stopping times at which robots are launched and the quantities given to the robots. The scheduler has to wait the max between a given duration and the end of the robot work before launching the next one.

Optimizing a Two-Scales Controller

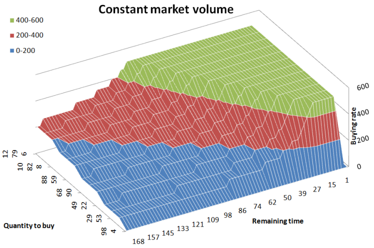

In [Bouchard et al., 2011] , we developed a model describing this two scales process:

- The Scheduler launches **trading robots**, known in probabilistic terms (the joint laws of their duration and their efficiency) at any time.
- It is an **impulse control problem** embedded in a continuous-time framework: the controls are the stopping times at which robots are launched and the quantities given to the robots. The scheduler has to wait the max between a given duration and the end of the robot work before launching the next one.

The outcome of this work has been a **better understanding of the discretization bias**. It is probably not really possible to go further with the trading rate as control.
Of course you can try to control directly (and solely) the interactions with orderbooks. We did it in several papers [Guéant et al., 2012, Guéant et al., 2013, Guéant and Lehalle, 2015] . I am not sure it is the best in terms of **risk control** (from an operational perspective at least).

Conclusion on Optimal Trading

We have seen optimal trading has two natural scales:

- At small scales (less than 1 minute on liquid instruments, but could be 1h on not liquid ones) you have to take into account the temporary variations of liquidity.

  An **exploration-exploitation scheme** has to be used. That for you need to

  1. Write your criterion,
  2. Derive its “mean field” gradient descent / ascent,
  3. Convert it in its stochastic counterpart (and be sure it converges).

  We did it for a Dark SOR.
- At longest scales (few hours), **the strategic interactions** between liquidity, urgency, and risk can be optimized too.

  The basic toolkit is the Almgren and Chriss one (it is a mean-variance scheme); it give birth to deterministic strategies. If you want to do something more subtle, the Cartea and Jaimungal toolkit is nice too, but a little more complex since it give birth to stochastic optimal strategies.

In practice you need to combine the two scales. The more explicitly, the better. Please do not write only one of the two framework properly, without paying attention to the other.

At least theoretically, you can try to optimize the all scales simultaneously. It is interesting to understand how they play with each others. But in terms of robustness I would recommand to **split a trading algorithm** in two parts: one strategic layer (slow and robust) and one tactical layer (fast and adaptive).

INSIGHT.DATA.CLARITY.

Takeaway Message 1: The Financial System

- The whole financial system is a partially **connected network of intermediaries**,
- They buy and sell risks, **netting** internally as exposure as possible,
- and using hedging (risk replication) techniques for the remaining, in the hope to **cross another intermediary's risk** on anonymous markets.
- The main activity of the financial system is hence to **make the market**, in an attempt to make margin while conveying a flow of transactions from some ultimate risk takers (like retail, corporate, asset managers, insurances) to others.
- Regulators have to **calibrate the amount of risk an intermediary can bear**: large enough to enable her to wait of next buyers or sellers (to balance her book), small enough to not allow directional risk.
- More continuous the trading, easier to balance the book and achieve very **low levels of risk for the whole system**.

Takeaway Message 2: Optimal Trading

- **Optimal trading** covers: execution of metaorders, opportunistic liquidity trading and market making.
- It makes the balance between **market impact** (trade slow) and the **inventory risk** (trade fast).
- At the smallest scale, I recommend the use of **forward methods**, like **statistical learning**. I have shown how **stochastic algorithms** could offer more accurate updates of variables of interest than reinforcement learning.
- At the largest scales, **stochastic control** is needed since the backward component is crucial (do not miss the target).
- You can try to optimize the two layers at once, as a practitioner I would not recommend this for robustness reasons (large exposure to small perturbations of HF data). Here the design of interactions between the slow layer (implementing backward strategies) and the fast one (focussed on tactical forward methods) in **an interesting algorithmic challenge**.

Takeaway Message 3: Stakes for financial maths

- introduce **liquidity** in existing models
- it means: instantaneous, transient, decay and permanent market impact
- or offer and demand (multiple agents)

- introduce **fragmentation**
- it means a choice between multiple "liquidity shapes"

- introduce **learning**
- dynamics of liquidity is sophisticated (cf the Queue Reactive model)
- it **predicts** short term price movement

Competition for liquidity is clearly an important topic.

INSIGHT.DATA.CLARITY.

Different Levels of Loops to Close

Two areas are not explored enough

- for **practitioners** : statistical learning; how to adapt online to regime switches (remember what we said about liquidity game vs. price game)? How to be robust to transitory phases? “Closing the loop” with learning is mixing exploration and exploitation.
- for **regulators** : game theory; what is the result of putting rational agents together? The more quants will read the 3 books, the more it will be needed to understand such interactions, and how changing “meta parameters” (ie rules) will modify the outcome of this game?

For game theory on financial market:

- few agents usually leads to principal - agent problems,
- a lot of agents usually leads to mean field games.

Moreover, game theory is a way to obtain **robust control** .

- the number of players needs to be "large enough"
- all players contribute to a "mean field" (i.e. a global variable: available shares, volatility, resource, etc)
- a function of this mean field (at least its mean, may be its standard deviation, etc) appear in this utility function of the players

$\rightarrow$ the name on the player cannot be used, but they can have a parameter (like a time horizon or risk aversion) of their own

The methodology is similar to the one to solve static Nash games:

- express the solution (for one agent) and find the solution as if the mean field was known
- you obtain a backward pde
- combine what you know about the mean field to find its forward pde

Liquidity is typically a mean field: the state of the inventory of participants influence their costs and can lead to fire sales [rené]. What practitioners call "velocity" of the liquidity (the flows) is a mean field too, it probably forms the prices along with market impact.

A continuum of agents trade optimally “à la Cartea-Jaimungal”. Main variables:

- public price P, exposed to permanent market impact. $\mu = \int_q \nu^a \, df(a)$ is the sum of the control of all agents
- the remaining qty of each agent Q^a (ν^a is the control)
- the wealth of each agent (public price is penalized by instantaneous impact)

Keep in mind ν^a is negative for a seller

$$dS_t = \alpha \mu_t \, dt + \sigma \, dW_t. \tag{12}$$

$$dQ^a_t = \nu^a_t \, dt, \tag{13}$$

since for a seller, $Q^a_0 > 0$ (the associated control ν^a will be mostly negative) and the wealth suffers from linear trading costs (or *temporary*, or *immediate market impact*):

$$dX^a_t = -\nu^a_t (S_t + \kappa \cdot \nu^a_t) \, dt. \tag{14}$$

The cost function of investor a selling from $t = 0$ and T is similar to the ones used in [Cartea et al., 2015]: the terminal inventory is penalized and a quadratic running cost is subtracted:

$$V^a_t := \sup_\nu \mathbb{E} \left(X^a_T + Q^a_T (S_T - A^a \cdot Q^a_T) - \phi^a \int_{s=t}^{T} (Q^a_s)^2 \, ds \Big| \mathcal{F}_t \right). \tag{15}$$

HJB For One Player

The Hamilton-Jacobi-Bellman associated to (15) is

$$0 = \partial_t V^a - \phi^a q^2 + \frac{1}{2}\sigma^2 \partial_S^2 V^a + \alpha\mu\partial_S V^a + \sup_\nu \{\nu \partial_Q V^a - \nu(s + \kappa\nu)\partial_X V^a\},\ V^a(T, x, s, q; \mu) = x + q(s - A^a q).$$

Following the Cartea and Jaimungal's approach, we will use the following ersatz:

$$V^a = x + qs + v^a(t, q; \mu). \tag{16}$$

Thus the HJB on v is

$$-\alpha\mu\, q = \partial_t v^a - \phi^a q^2 + \sup_\nu \left\{\nu \partial_Q v^a - \kappa \nu^2\right\},\ v^a(T, q; \mu) = -A^a q^2.$$

and the associated optimal feedback is

$$\nu^a(t, q) = \frac{\partial_Q v^a(t, q)}{2\kappa}. \tag{17}$$

The mean field of this framework is the distribution $m(t, dq, da)$ of the inventories Q_0^a and of their preferences (ϕ^a, A^a).
It is then straightforward to write the net trading flow μ at any time t

(18) $$\mu_t = \int_{(q,a)} \nu_t^a(q)\, m(t, dq, da) = \int_{q,a} \frac{\partial_Q v^a(t,q)}{2\kappa}\, m(t, dq, da).$$

v^a is an implicit function of μ, meaning we will have a fixed point problem to solve in μ. By the dynamics (13) of Q_t^a, we have

$$\partial_t m + \partial_q \left(m \frac{\partial_Q v^a}{2\kappa} \right) = 0 \text{ with initial condition } m_0 = m_0(dq, da).$$

(19) $$\begin{cases} -\alpha q \int_{(q,a)} \frac{\partial_Q v^a(t,q)}{2\kappa}\, m(t, dq, da) & = \partial_t v^a - \phi^a\, q^2 + \frac{(\partial_Q v^a)^2}{4\kappa} \\ \partial_t m + \partial_q \left(m \frac{\partial_Q v^a}{2\kappa} \right) & = 0 \end{cases}$$

Terminal conditions: $m(0, dq, da) = m_0(dq, da), \qquad v^a(T, q; \mu) = -A^a q^2.$

The System in The Case of Identical Preferences and Quad. V $\phi^a \equiv \phi$, $A^a \simeq A$

Set $E(t) = \mathbb{E}[Q_t] = \int_q q m(t, dq)$. Note that

(20) $$E'(t) = \int_q q\partial_t m(t, dq),\ E'(t) = -\int_q q\partial_q \left(m(t,q) \frac{\partial_Q v(t,q)}{2\kappa} \right) dq = \int_q \frac{\partial_Q v(t,q)}{2\kappa} m(t, dq).$$

When $v(t,q)$ can be expressed as a quadratic function of q: $v(t,q) = h_0(t) + q\, h_1(t) - q^2\, \frac{h_2(t)}{2}$,

(21) $$\mu(t) = \int_q \frac{\partial_Q v(t,q)}{2\kappa}\, dm(q) = \frac{h_1(t)}{2\kappa} - \frac{h_2(t)}{2\kappa} E(t).$$

and

$$E'(t) = \int_q m(t,q) \left(\frac{h_1(t)}{2\kappa} - \frac{h_2(t)}{2\kappa} q \right) dq = \frac{h_1(t)}{2\kappa} - \frac{h_2(t)}{2\kappa} E(t).$$

So we can supplement it with

(22) $$2\kappa E'(t) = h_1(t) - E(t) \cdot h_2(t).$$

We now collect all the equations. Recalling (21), we find:

$$\begin{cases} 4\kappa\phi = -2\kappa h_2'(t) + (h_2(t))^2, & \text{(23a)} \\ \alpha h_2(t)E(t) = 2\kappa h_1'(t) + h_1(t)\left(\alpha - h_2(t)\right), & \text{(23b)} \\ -\left(h_1(t)\right)^2 = 4\kappa h_0'(t), & \text{(23c)} \\ 2\kappa E'(t) = h_1(t) - h_2(t)E(t). & \text{(23d)} \end{cases}$$

with the boundary conditions

$$h_0(T) = h_1(T) = 0,\ h_2(T) = 2A,\ E(0) = E_0,$$

where $E_0 = \int_q q m_0(q)dq$ is the net initial inventory of market participants (i.e. the expectation of the initial density m).

To summarize, the equation satisfied by E is:

$$\left\{\begin{array}{ll} 0 = 2\kappa E''(t) + \alpha E'(t) - 2\phi E(t) & \text{for } t \in (0,T), \\ E(0) = E_0, \quad \kappa E'(T) + AE(T) = 0. & \end{array}\right. \tag{24}$$

Closed form for the net inventory dynamics $E(t)$

For any $\alpha \in \mathbb{R}$, the problem (24) has a unique solution E, given by

$$E(t) = E_0 a \left(\exp\{r_+ t\} - \exp\{r_- t\}\right) + E_0 \exp\{r_- t\}$$

where a is given by

$$a = \frac{(\alpha/4 + \kappa\theta - A)\exp\{-\theta T\}}{-\frac{\alpha}{2}\mathrm{sh}\{\theta T\} + 2\kappa\theta\mathrm{ch}\{\theta T\} + 2A\mathrm{sh}\{\theta T\}},$$

the denominator being positive and the constants $r_\alpha^\pm$ and θ being given by

$$r_\pm := -\frac{\alpha}{4\kappa} \pm \theta, \qquad \theta := \frac{1}{\kappa}\sqrt{\kappa\phi + \frac{\alpha^2}{16}}.$$

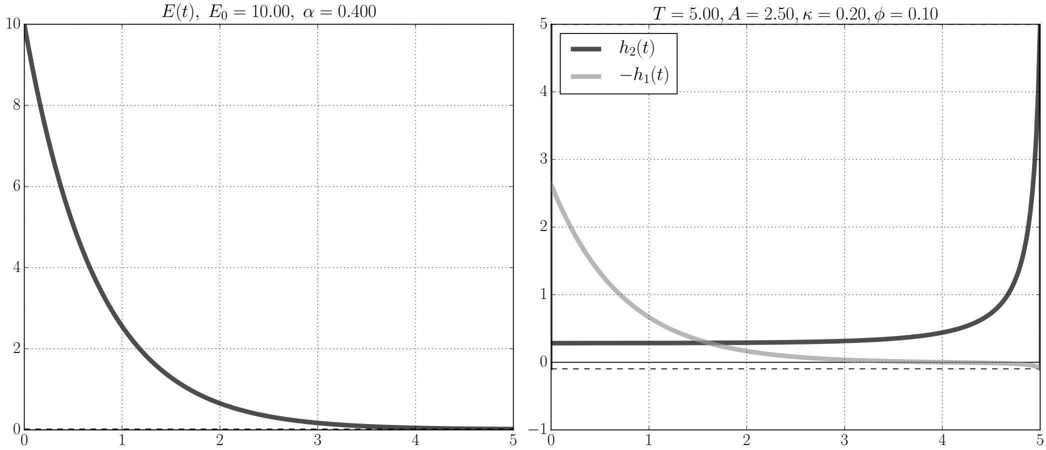

Dynamics of E (left) and $-h_1$ and h_2 (right) for a standard set of parameters: $\alpha = 0.4$, $\kappa = 0.2$, $\phi = 0.1$, $A = 2.5$, $T = 5$, $E_0 = 10$.

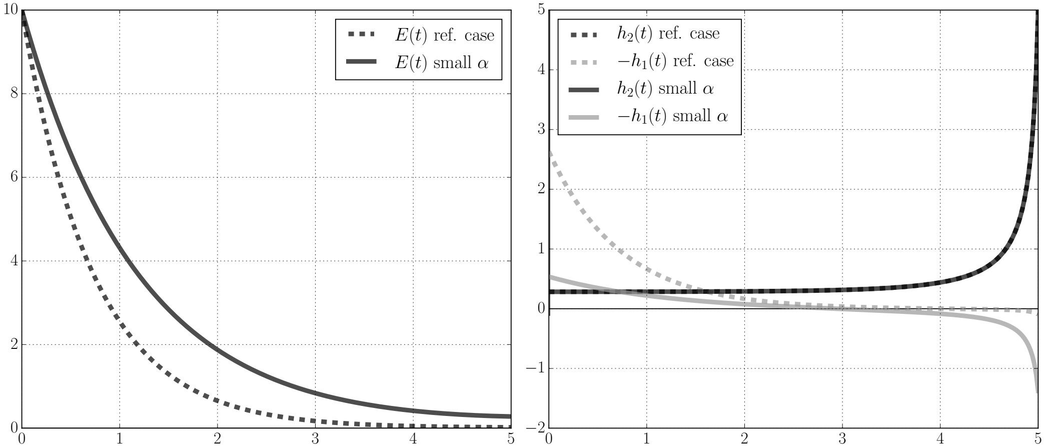

Comparison of the dynamics of E (left) and $-h_1$ and h_2 (right) between the "reference" parameters of Figure **??** and smaller α (i.e. $\alpha = 0.1$ instead of 0.4) such that $|h_1(0)|$ is smaller.

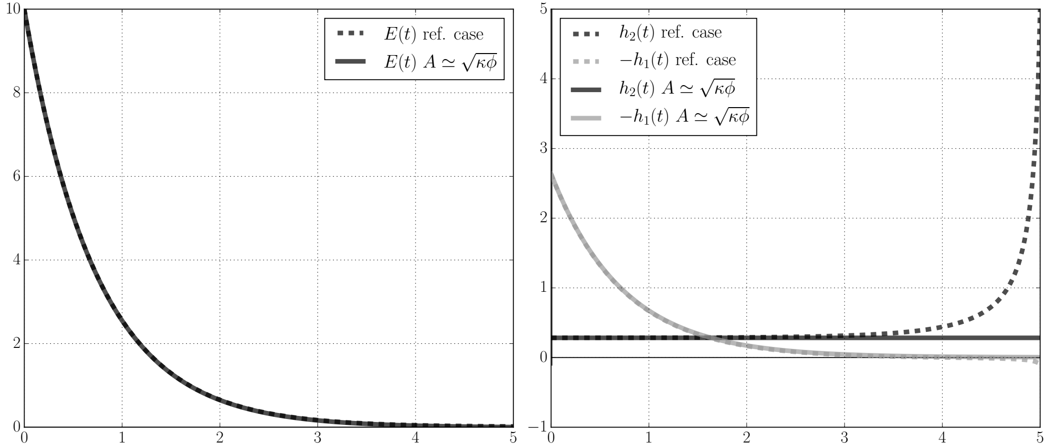

Comparison of the dynamics of E (left) and $-h_1$ and h_2 (right) between the "reference" parameters of Figure **??** and when $\sqrt{\kappa\phi} \simeq A$: in such a case h_2 is almost constant but E and h_1 are almost unchanged.

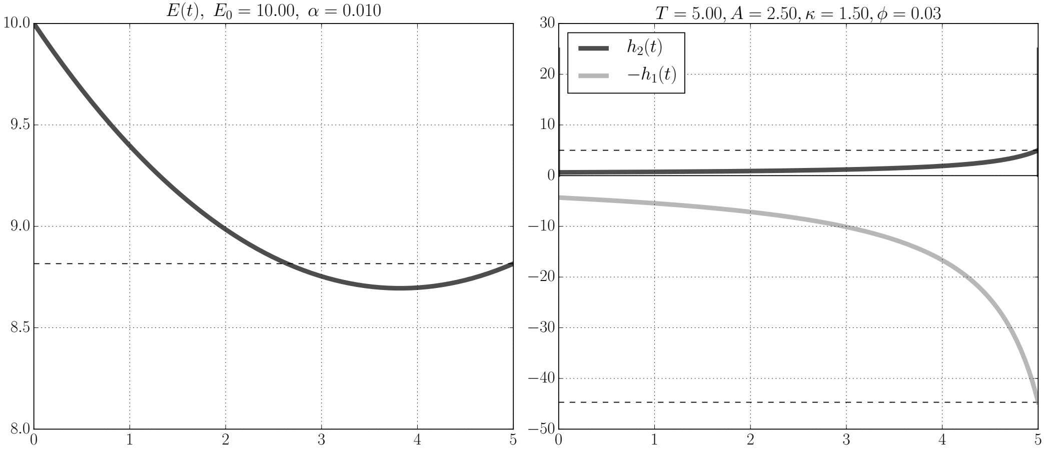

A specific case for which E is not monotonous: $\alpha = 0.01$, $\kappa = 1.5$, $\phi = 0.03$, $A = 2.5$, $T = 5$ and $E_0 = 10$.

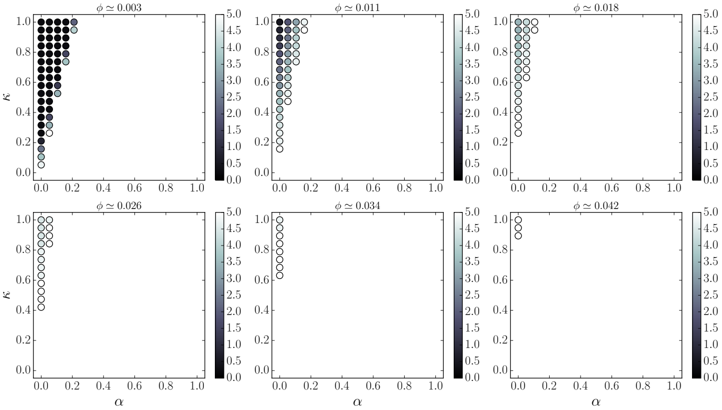

Numerical explorations of t^m for different values of ϕ (very small ϕ at the top left to small ϕ at the bottom right) on the $\alpha \times \kappa$ plane, when $T = 5$ and $A = 2.5$. The color circles codes the value of t^m: small values (dark color) when E changes its slope very early; large values (in light colors) when E changes its slope close to T.

Conclusion on MFG of Controls For Liquidation

It is a proof of maturity of the use if stochastic control in financial math:

- Four years ago, it was difficult to think about a game theoretical version of the Almgren and Chriss optimal liquidation problem (schied and jaimungal).
- Our understanding of the problem itself improved (see Guéant and Cartead and Jaimungal books)
- and some extensions of MFG have been needed (see the paper).
- but we now know how to handle it (and in a specific case it is fully solved)

Solving game theoretical versions of what we know is important (instead of sophisticating it in a mean field game), because

- it is a way to obtain robust control
- it helps regulator to understand the system to adjust some meta parameters (κ is this example)

MFG is not the only way to answer to such questions.
Moreover learning should not be forgot (done in our paper): what does change when information is not complete?

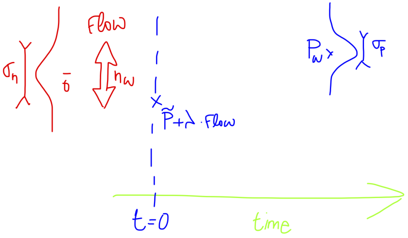

The framework

- An informed trader, knowing the future price
- Noise traders, knowing nothing
- A market makers, having only access to distributions (thanks to "backtests" / observations); she changes her price linearly according to the price pressure she observes: $f_P(q) = \tilde{P} + \lambda \cdot q$.

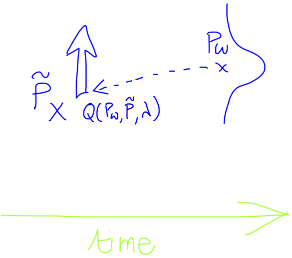

The informed trader chooses his quantity to maximize his expected profit

The framework

- An informed trader, knowing the future price
- Noise traders, knowing nothing
- A market makers, having only access to distributions (thanks to "backtests" / observations); she changes her price linearly according to the price pressure she observes: $f_P(q) = \tilde{P} + \lambda \cdot q$.
- The informed trader adjusts his participation to maximize its profit (given $\tilde{P}$ and λ),

Making the Market in Front of Adverse Selection: Kyle's 85 Model

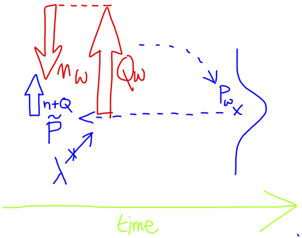

The market maker choose $\tilde{P}$ and λ to adjust her price to the flow

The framework

- An informed trader, knowing the future price
- Noise traders, knowing nothing
- A market makers, having only access to distributions (thanks to "backtests" / observations); she changes her price linearly according to the price pressure she observes: $f_P(q) = \tilde{P} + \lambda \cdot q$.
- The informed trader adjusts his participation to maximize its profit (given $\tilde{P}$ and λ),
- The market makers know the distribution of the informed price and set $\tilde{P}$ and λ so that her price is as close as possible to its expectation.

Following *Continuous Auctions and Insider Trading* – [Kyle, 1985]:

- Remember the market makers fear adverse selection.
- We have **informed traders**, they know the price will be p_ω after their trade, $p_\omega \sim \mathcal{N}(P^*, \sigma_p^2)$.
- Other traders, (i.e. *noise traders* for Kyle) trade for other reasons, their net direction is $n_\omega \sim \mathcal{N}(0, \sigma_n^2)$.
- The informed traders have to **choose a participation** $Q(p)$ (they know p) to maximize their profit,
- Knowing the market makers (MM) will react to the net perceived flow **linearly**: the *public price* will be

$$f_P(Q(p_\omega) + n_\omega) = \tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega).$$

- Moreover in their filtration, the MM should produce a price being the **best estimator of p_ω given** $Q(p_\omega) + n_\omega$.

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}(p_\omega | Q(p_\omega) + n_\omega)$.

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q | p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}\,(p_\omega | Q(p_\omega) + n_\omega)$.
- The solution is the linear regression of p on $Q(p) + n_\omega$:

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}\left(p_\omega | Q(p_\omega) + n_\omega\right)$.
- The solution is the linear regression of p on $Q(p) + n_\omega$:

$$\begin{cases} P^* & = & \tilde{P} + \lambda \mathbb{E}(Q(p_\omega) + n_\omega) \\ \lambda & = & \dfrac{\mathbb{C}\mathrm{ov}(p, Q(p) + n_\omega)}{\mathbb{V}(Q(p) + n_\omega)} \end{cases}$$

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}\,(p_\omega|Q(p_\omega) + n_\omega)$.
- The solution is the linear regression of p on $Q(p) + n_\omega$:

$$\begin{cases} P^* &= \tilde{P} + \lambda \mathbb{E}(Q(p_\omega) + n_\omega) \\ \lambda &= \dfrac{\mathbb{C}\mathrm{ov}(p, Q(p) + n_\omega)}{\mathbb{V}(Q(p) + n_\omega)} \end{cases} \Rightarrow \begin{cases} \tilde{P} &= P^* \\ \lambda &= \dfrac{\sigma_p^2/(2\lambda)}{\sigma_p^2/(2\lambda)^2 + \sigma_n^2} \end{cases}$$

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q + n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}\left(p_\omega | Q(p_\omega) + n_\omega\right)$.
- The solution is the linear regression of p on $Q(p) + n_\omega$:

$$\left\{ \begin{array}{ccc} P^* & = & \tilde{P} + \lambda \mathbb{E}(Q(p_\omega) + n_\omega) \\ \lambda & = & \dfrac{\mathbb{C}\mathrm{ov}(p, Q(p) + n_\omega)}{\mathbb{V}(Q(p) + n_\omega)} \end{array} \right. \Rightarrow \left\{ \begin{array}{ccc} \tilde{P} & = & P^* \\ \lambda & = & \dfrac{\sigma_p^2/(2\lambda)}{\sigma_p^2/(2\lambda)^2 + \sigma_n^2} \end{array} \right.$$

$\Rightarrow$ It can be solved with $\lambda = \sigma_p/(2\sigma_n)$.

- **Informed traders** maximize their expected price: $\arg\max_Q \mathbb{E}((p_\omega - f_p(Q+n_\omega))Q|p_\omega)$.

$\Rightarrow$ They decide to trade $Q(p_\omega) = (p_\omega - \tilde{P})/(2\lambda)$.

- The MMs have to choose $\tilde{P}$ and λ so that $\tilde{P} + \lambda \cdot (Q(p_\omega) + n_\omega) = \mathbb{E}(p_\omega | Q(p_\omega) + n_\omega)$.
- The solution is the linear regression of p on $Q(p) + n_\omega$:

$$\begin{cases} P^* & = & \tilde{P} + \lambda \mathbb{E}(Q(p_\omega) + n_\omega) \\ \lambda & = & \dfrac{\mathbb{C}\mathrm{ov}(p, Q(p) + n_\omega)}{\mathbb{V}(Q(p) + n_\omega)} \end{cases} \Rightarrow \begin{cases} \tilde{P} & = & P^* \\ \lambda & = & \dfrac{\sigma_p^2/(2\lambda)}{\sigma_p^2/(2\lambda)^2 + \sigma_n^2} \end{cases}$$

$\Rightarrow$ It can be solved with $\lambda = \sigma_p/(2\sigma_n)$.

- The more potential informational price move (i.e. large σ_p), the largest impact.
- The more non informative flow, the more difficult for the MM to identify information, hence the less she impacts the price.

On our database of 300,000 large orders

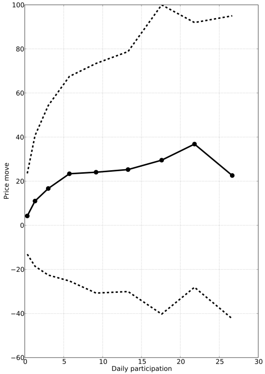

- This is renormalized market impact as a function of the *participation rate*, i.e. v/V, taken from [Bacry et al., 2015].
- It is not linear, it is **concave** .

Due to [Çetin and Danilova, 2015].

Notations and definitions.

- At $t = 0$, the insider trader knows the final "fundamental" value of the instrument. This value is drawn from a random variable V (the law of V is known by everyone). He tries to trade "optimally" knowing the market maker with react to the trading flow. His cumulated trading flow at t is X_t.
- Liquidity traders generate a (signed) cumulative demand of $\sigma_t B_t$ at t.
- The Market Maker (in reality N market makers competing "*à la Bertrand*") observes the net demand $Y_t := X_t + \sigma B_t$. She sets the price at $S_t := H(t, Y_t)$. She has a **CARA utility function** $U(x) = -\exp -\rho x$.

Note the terminal wealth of the insider is

$$W_1^X = \int_0^1 X_t dH(t, Y_t) + X_1(V - H(1, Y_1)) = \int_0^1 (V - H) dX. \tag{25}$$

On her side, the market maker wealth at t is

$$G_t := -\frac{1}{N} \int_0^1 Y_s dH(s, Y_s) + \mathbf{1}_{t=1} \frac{Y_1}{N} (H(1, Y_1) - V). \tag{26}$$

Say the infinitesimal increase of the market maker pricing rule H can be written using a controlled diffusion: $dS = Z_t dB^Y + \mu_t dt$, then the infinitesimal change in U is (using (26)): $dU(G_t) = U(G_t)\frac{\rho}{N} Y_t \left(\sigma_t dB^Y + (\mu_t + \frac{\rho}{2N} Y_t Z_t^2) dt\right)$. The dynamic programming principle states its deterministic part should be zero, hence

$$\mu_t = -\frac{\rho}{2N} Y Z^2. \tag{27}$$

And finally

$$dS = Z_t dB^Y - \frac{\rho}{2N} Y_t Z_t^2 dt. \tag{28}$$

Dynamics of the Pricing Rule

Because of the term YdS in the wealth of the market maker and because of her CARA(ρ) utility, the pricing rule necessarily follows equation(28): $dS = Z_t dB^Y - \frac{\rho}{2N} Y_t Z_t^2 dt$.

On his side the insider will act on his trading flow $dX := \alpha(t, Y, S, Z)dt$, the net demand will hence be a diffusion with a controlled drift:

$$dY = \sigma_t dB + \alpha dt. \tag{29}$$

Using α, the insider will hence maximize his utility; again the dynamic programming principle on $\Psi(t,y) := \sup_\alpha \mathbb{E}\left[\int_t^1 (V-H)\alpha ds \middle| Y_t = y\right]$ gives a useful information: the following HJB (subscripts stand for partial derivatives): $\Psi_t + \frac{\sigma^2}{2}\Psi_{yy} + \sup_\alpha\{\alpha(\Psi_y + (V-H))\} = 0$. In other terms: $\Psi_y = V - H$; $\Psi_t + \frac{\sigma^2}{2}\Psi_{yy} = 0$. This implies H must satisfy the heat equation too:

$$H_t + \frac{\sigma^2}{2} H_{yy} = 0. \tag{30}$$

On the one hand, writing Ito's formula on $S = H(t, Y)$ reads $dS = H_y\, dY + \underbrace{(H_t + \frac{\sigma^2}{2} H_{yy})}_{=0} dt = H_y \sigma_t dB + H_y \alpha dt.$

And on the other hand we have (28) because of the market maker optimal strategy. Equalling the deterministic parts reads

(31) $$H_y = \frac{Z}{\sigma}.$$

and the random parts reads $H_y \alpha = -\rho Y Z^2/(2N)$. Hence $(z/\sigma)\alpha^*(t, y, s, z) = -\rho/(2N)\; yz^2$, i.e.

(32) $$\alpha^* = -\frac{\rho\sigma}{2N} yz.$$

Insider's Optimal Strategy

Because of Ito on the pricing rule seen as a function of (t, Y) and because of the shape of the pricing rule fixed by the market maker, we necessarily have equations (31) $H_y = Z/\sigma$ and (32) $\alpha^* = -\rho\sigma \cdot yz/(2N)$. It implies equation (33) for the net demand dynamics: $dY = \sigma_t dB - \frac{\rho\sigma^2}{2N} YH_y\, dt$.

Terminal Condition

Because of the shape of the insider's wealth $\int_0^1 (V - H)dX$ and because the expectation of dY is dX, we necessarily have $H(1, Y_1) = V$.

Remark: Obtaining the terminal condition is tricky.

Replacing α and z by their values in (29) allows to write

$$\begin{aligned} dY &= \sigma_t dB + \alpha dt \\ &= \sigma_t dB - \frac{\rho\sigma}{2N} YZ\, dt, \text{ using (32)} \\ &= \sigma_t dB - \frac{\rho\sigma^2}{2N} YH_y\, dt, \text{ using (31).} \end{aligned} \tag{33}$$

Putting the main equations (30)-(33) side to side, we now have

$$\begin{cases} H_t + \dfrac{\sigma^2}{2} H_{yy} = 0, \text{ under the terminal condition } H(1, Y_1) = V \\ dY = \sigma_t dB - \dfrac{\rho\sigma^2}{2N} YH_y\, dt \end{cases} \tag{34}$$

The iterations to solve the closed loop are as follow (use Schauder's fixed point theorem, i.e. a strict compactness inclusion argument)

- **Initialization.**
  - V is the final value distribution, it is revealed to the insider trader at $t = 0$.
    Typically $V \sim \mathcal{N}(0, \sigma_V^2)$.
  - The market maker makes a guess for $Y_1^{(1)}$, the sum of noise traders' flows and the insider's one.
    It can be a Gaussian too, or something more "protective" (in the sense a distribution with fat tails).
  - She solves (a backward way) the heat equation

$$H_t^{(1)} + \frac{\sigma^2}{2} H_{yy}^{(1)} = 0$$

    with the terminal condition $H^{(1)}(1, Y_1^{(1)}) = V$. She obtain a "first guess" pricing rule to apply.
- **Steps to Iterate.**
  - The insider trader now sees $H^{(n)}(t, Y^{(n)})$ in real-time, and reacts by setting his flow so that

$$dY^{(n+1)} = \sigma_n d\beta - \frac{\rho\sigma^2}{2N} Y_t^{(n+1)} H_y^{(n)} dt.$$

  - At the end of the day the market maker discovers $Y^{(n+1)}$ and know V in distribution, hence she sets a new pricing rule solving again the heat equation, but with a different terminal condition, now

$$H^{(n+1)}(1, Y^{(n+1)}) = V.$$

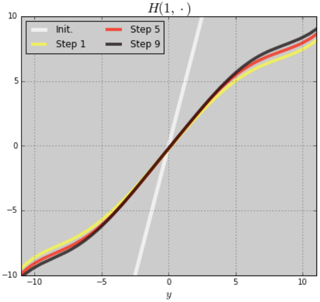

I implemented the numerics, here is the "market impact" (i.e. pricing rule) function (and its evolution during the tracking of the fixed point). It is **concave**, accordingly with the data.

This master system of equations can be understood as

- The market maker(s) tunes her pricing rule H such that it follows the heat equation and it exactly fits the law of the fundamental prices V at terminal time $t = 1$;
- The insider mean reverts the net demand (his demand plus the noisy one) flow (i.e. net demand) through time using a mean reversion strength proportional to $\partial_y H(t, y)|_{(t,Y_t)}$: when the slope of the pricing rule is intense, it mean reverts more than when it is not.

The complexity of this system is the market maker needs to know in advance the insider's strategy to have an access to Y_1 to satisfy her terminal condition. That for a "fixed point" theorem will be needed, the good one is Schauder's one, and that for it will be needed Y has a smooth transition density. This property will need to be stable via the application of the system (34).

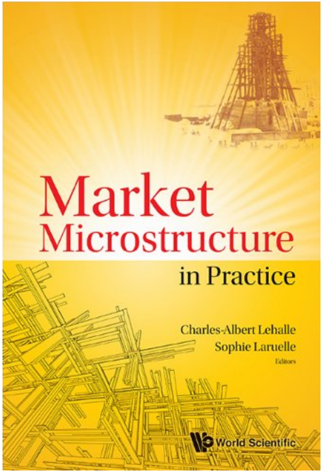
Market
Microstructure
in Practice
Charles-Albert Lehalle
Sophie Laruelle
Editors
World Scientific


Market Microstructure
and Liquidity (MML)
www.worldscientific.com/mml
Call for Papers
MARKET MICROSTRUCTURE and LIQUIDITY
Managing Editors
Associate Editors
Industry Advisory Board
Managing Editors
Frédéric Abergel
Jean-Philippe Bouchaud
Joel Hasbrouck
Charles-Albert Lehalle
Mathieu Rosenbaum
World Scientific

References I

Agarwal, A., Bartlett, P. L., and Dama, M. (2010).
Optimal Allocation Strategies for the Dark Pool Problem.
In Teh, Y. W. and Titterington, M., editors, *Proceedings of The Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 9, pages 9–16.

Alfonsi, A., Fruth, A., and Schied, A. (2009).
Optimal execution strategies in limit order books with general shape functions.
Quantitative Finance, 10(2):143–157.

Almgren, R. (2009).
Optimal Trading in a Dynamic Market.
Technical Report 2.

Almgren, R. and Lorenz, J. (2007).
Algorithmic Trading III: Precision, Control, Execution.
Institutional Investor Journals.

Almgren, R. F. and Chriss, N. (2000).
Optimal execution of portfolio transactions.
Journal of Risk, 3(2):5–39.

Bacry, E., Iuga, A., Lasnier, M., and Lehalle, C.-A. (2015).
Market Impacts and the Life Cycle of Investors Orders.
Market Microstructure and Liquidity, 1(2).

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

Bouchard, B., Dang, N.-M., and Lehalle, C.-A. (2011).
Optimal control of trading algorithms: a general impulse control approach.
SIAM J. Financial Mathematics, 2(1):404–438.

Brokmann, X., Serie, E., Kockelkoren, J., and Bouchaud, J. P. (2014).
Slow decay of impact in equity markets.

Cardaliaguet, P. and Lehalle, C.-A. (2016).
Mean Field Game of Controls and An Application To Trade Crowding.

Cartea, A. and Jaimungal, S. (2014).
A Closed-Form Execution Strategy to Target VWAP.
Social Science Research Network Working Paper Series.

Cartea, A., Jaimungal, S., and Penalva, J. (2015).
Algorithmic and High-Frequency Trading (Mathematics, Finance and Risk).
Cambridge University Press, 1 edition.

Çetin, U. and Danilova, A. (2015).
Markovian Nash equilibrium in financial markets with asymmetric information and related forward-backward systems.

Doukhan, P. (1994).
Mixing: Properties and Examples (Lecture Notes in Statistics).
Springer, 1 edition.

Ganchev, K., Nevmyvaka, Y., Kearns, M., and Vaughan, J. W. (2010).
Censored exploration and the dark pool problem.
Commun. ACM, 53(5):99–107.

INSIGHT.DATA.CLARITY.

References III

Gatheral, J., Schied, A., and Slynko, A. (2012).
Transient Linear Price Impact and Fredholm Integral Equations.
Mathematical Finance, 22:445–474.

Guéant, O. (2016).
The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making.
Chapman and Hall/CRC.

Guéant, O. and Lehalle, C.-A. (2015).
General intensity shapes in optimal liquidation.
Mathematical Finance, 25(3):457–495.

Guéant, O., Lehalle, C.-A., and Fernandez-Tapia, J. (2012).
Optimal Portfolio Liquidation with Limit Orders.
SIAM Journal on Financial Mathematics, 13(1):740–764.

Guéant, O., Lehalle, C.-A., and Fernandez-Tapia, J. (2013).
Dealing with the inventory risk: a solution to the market making problem.
Mathematics and Financial Economics, 4(7):477–507.

Guéant, O. and Pu, J. (2013).
Option pricing and hedging with execution costs and market impact.

Hirsch, M. W. and Smith, H. (2005).
Monotone dynamical systems, volume 2, pages 239–357.

Kushner, H. J. and Yin, G. G. (2003).
Stochastic Approximation and Recursive Algorithms and Applications (Stochastic Modelling and Applied Probability) (v. 35).
Springer, 2nd edition.

Kyle, A. P. (1985).
Continuous Auctions and Insider Trading.
Econometrica, 53(6):1315–1335.

Labadie, M. and Lehalle, C.-A. (2014).
Optimal starting times, stopping times and risk measures for algorithmic trading.
The Journal of Investment Strategies, 3(2).

Lachapelle, A., Lasry, J.-M., Lehalle, C.-A., and Lions, P.-L. (2016).
Efficiency of the Price Formation Process in Presence of High Frequency Participants: a Mean Field Game analysis.
Mathematics and Financial Economics, 10(3):223–262.

Laruelle, S., Lehalle, C.-A., and Pagès, G. (2013).
Optimal posting price of limit orders: learning by trading.
Mathematics and Financial Economics, 7(3):359–403.

Lehalle, C.-A. (2008).
Rigorous optimisation of intra day trading.
Wilmott Magazine.

Lehalle, C.-A. (2013).
Market Microstructure knowledge needed to control an intra-day trading process.
In Fouque, J.-P. and Langsam, J., editors, *Handbook on Systemic Risk*. Cambridge University Press.

Lehalle, C.-A., Laruelle, S., Burgot, R., Pelin, S., and Lasnier, M. (2013).
Market Microstructure in Practice.
World Scientific publishing.

References V

Lehalle, C.-A., Lasnier, M., Bessson, P., Harti, H., Huang, W., Joseph, N., and Massoulard, L. (2012).
What does the saw-tooth pattern on US markets on 19 July 2012 tell us about the price formation process.
Technical report, Crédit Agricole Cheuvreux Quant Note.

Li, T. M. and Almgren, R. (2014).
Option Hedging with Smooth Market Impact.
Technical report.

Pagès, G., Lapeyre, B., and Sab, K. (1990).
Sequences with low discrepancy. Generalization and application to Robbins-Monro algorithm.
Statistics, 21(2):251–272.

Pagès, G., Laruelle, S., and Lehalle, C.-A. (2011).
Optimal split of orders across liquidity pools: a stochastic algorithm approach.
SIAM Journal on Financial Mathematics, 2:1042–1076.

Schied, A. and Zhang, T. (2013).
A state-constrained differential game arising in optimal portfolio liquidation.

Waelbroeck, H. and Gomes, C. (2013).
Is Market Impact a Measure of the Information Value of Trades? Market Response to Liquidity vs. Informed Trades.
Social Science Research Network Working Paper Series.