Financial Intermediation at Any Scale For Quantitative Modelling (2/3)

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

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.

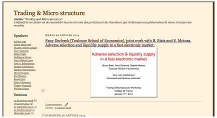

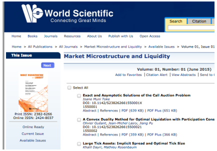

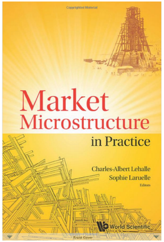

- From a **Financial Mathematics** perspective, it is nothing more than adding a variable to our models: the **Liquidity**.
- The interactions between liquidity and other (usual) variables is far from trivial.

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!**

Outline of the Sessions

1. The Financial System as a Network of Intermediaries

2. Stylized Facts on Liquidity

1. The Financial System as a Network of Intermediaries
   - Risks Transformation as The Primary Role of The Financial System
   - Making the Market: the Stakes of Liquidity Provision
   - The Market Impact of Large Orders
   - Quant Models For Common Practices
2. Stylized Facts on Liquidity
   - Seasonalities and Stationarity
   - Orderbook Dynamics

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

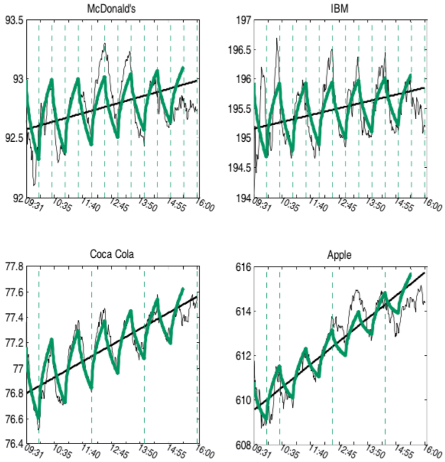

Even on liquid stocks and for vanilla options (close to maturity in this case), hedging can go wrong.

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...

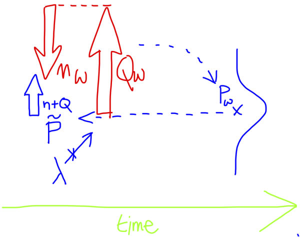

The market maker choose $\tilde{P}$ and λ to adjust her price to the flow

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.

in [Moro et al., 2009]

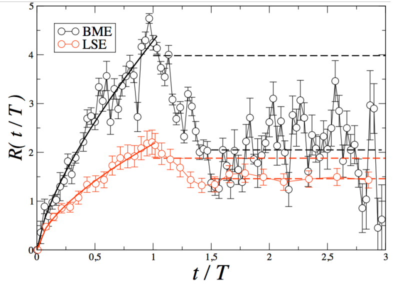

Market Impact takes place in different phases

- 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.

The Formula of the temporary market impact should be close to

$$\mathrm{MI} \propto \sigma \cdot \sqrt{\frac{\text{Traded volume}}{\text{Daily volume}}} \cdot T^{-0.2}$$

The term in duration is very difficult to estimate because you have a lot of conditioning everywhere:

On our database of 300,000 large orders [Bacry et al., 2015a]

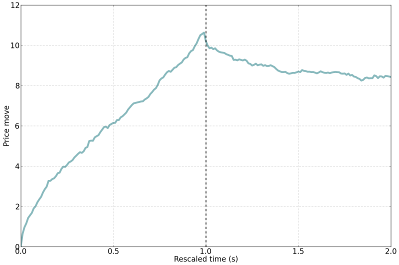

f

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.

The Formula of the temporary market impact should be close to

$$\mathrm{MI} \propto \sigma \cdot \sqrt{\frac{\text{Traded volume}}{\text{Daily volume}}} \cdot T^{-0.2}$$

The term in duration is very difficult to estimate because you have a lot of conditioning everywhere:

We had enough data to investigate long term impact, exploring the relationships between permanent impact and traded information.

Daily price moves

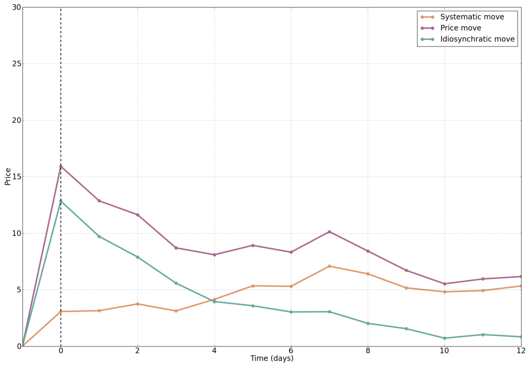

- If you plot the long term price moves (x-axis in days), you see an regular increase;
- But the same stock is traded today, tomorrow, the day after, etc.
- Once you remove the market impact of "future" trades (similarly to [Waelbroeck and Gomes, 2013]), you obtain a different shape.
- If you look each curve: the yellow one contains the CAPM β (the metaorders are trading market-wide moves), the green curve contains the idiosyncratic moves, this shape can be read as **the daily decay of metaorders impact**.

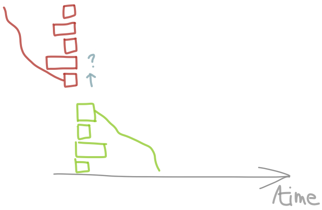

Adverse selection is the fact you obtained liquidity now and you could obtain it later and have a better price.

Adverse selection is the fact you obtained liquidity now and you could obtain it later and have a better price.

Price Impact and Adverse Selection

Price Impact for market orders implies Adverse selection for limit orders.

If the price in martingale after a price change, there is adverse selection ; the imbalance says you have a little less adverse selection than that since **once a full tick has been consumed, they are chances the discovered quantity is larger than average**.
When you owns a limit order in the book: **the more orders behind you, the more protected vs. adverse selection**.

2nd worst kept secret of HFT : if you have too few orders behind you, cancel your limit order.

“discovered quantity” means the quantity at second limit that is now a first limit

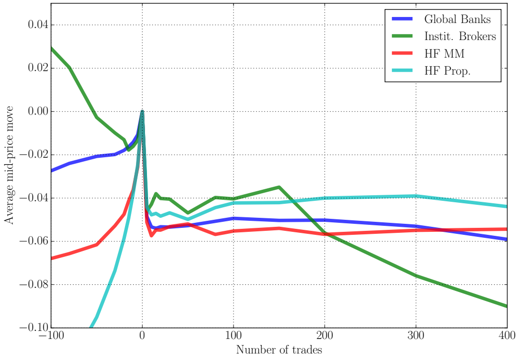

Price profiles: future and past of the mid-price (solid) or bid and asks (dotted), conditionally to an execution.

- Instutional Brokers (i.e. essentially "client flows": with a decision taken at a daily scale and large metaorders)
- HFT, split in HF market makers and HF proprietary traders
- Global banks, having a mix of client flows and proprietary trading flows.
- **Conditionally to the owner of the order**, the profile can be very different
- You can compare such graphs and make matrices [Brogaard et al., 2012]
- Nevertheless the big picture is dynamic...

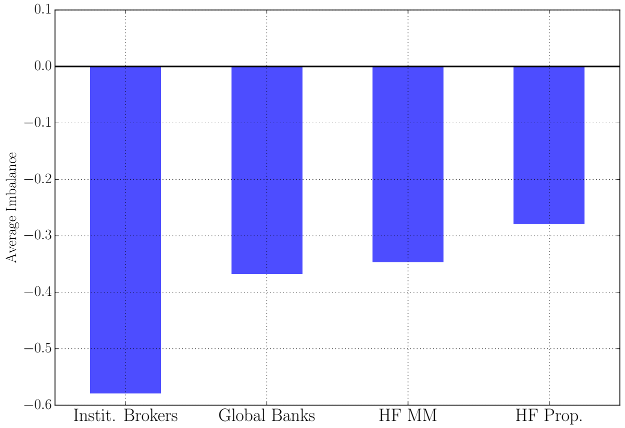

Imbalance profiles: State of the book conditionally to an execution, renormalized such as best opposite is 1, the green bar is your order size.

If someone trade at a given frequency $1/\delta t$ from 0, his price impact at $K\delta t$ will be (for an exponential kernel)

$$P(K\delta t) - P(0) = \sum_{k\leq K} \eta(1)\lambda e^{-k\delta t\,\lambda} \simeq \eta(1)(1 - e^{-K\delta t\,\lambda})/\delta t.$$

And for a power law

$$P(K\delta t) - P(0) = \eta(1)\left(1 - (1 + K\delta t)^{-(\gamma-1)}\right)/\delta t.$$

In both cases, if he stops trading at $K\delta t$, the price will **fully** revert according to an exponential (or a power lax).

Transient Impact and Decay

The concave increase of the impact with time and its reversion can be explained using propagator models.

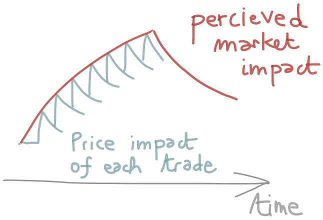

If someone trade at a given frequency $1/\delta t$ from 0, his price impact at $K\delta t$ will be (for an exponential kernel)

$$P(K\delta t) - P(0) = \sum_{k \leq K} \eta(1)\lambda e^{-k\delta t\,\lambda} \simeq \eta(1)(1 - e^{-K\delta t\,\lambda})/\delta t.$$

And for a power law

$$P(K\delta t) - P(0) = \eta(1)\left(1 - (1 + K\delta t)^{-(\gamma-1)}\right)/\delta t.$$

In both cases, if he stops trading at $K\delta t$, the price will **fully** revert according to an exponential (or a power lax).

Transient Impact and Decay

The concave increase of the impact with time and its reversion can be explained using propagator models.

From High Frequency to Low Frequencies

If someone trade at a given frequency $1/\delta t$ from 0, his price impact at $K\delta t$ will be (for an exponential kernel)

$$P(K\delta t) - P(0) = \sum_{k \leq K} \eta(1)\lambda e^{-k\delta t\,\lambda} \simeq \eta(1)(1 - e^{-K\delta t\,\lambda})/\delta t.$$

And for a power law

$$P(K\delta t) - P(0) = \eta(1)\left(1 - (1 + K\delta t)^{-(\gamma-1)}\right)/\delta t.$$

In both cases, if he stops trading at $K\delta t$, the price will **fully** revert according to an exponential (or a power lax).

Transient Impact and Decay

The concave increase of the impact with time and its reversion can be explained using propagator models.

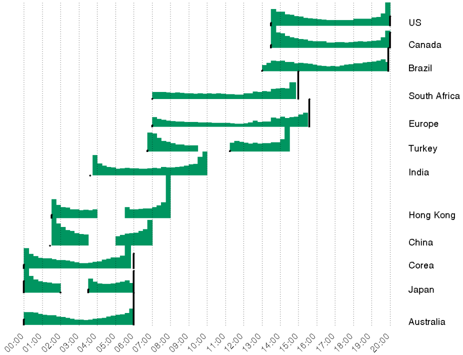

We have some "intraday patterns", the most famous being the U-shape of the traded volume. It come from the practices of participants, it is important to have them in mind.
We can see

- the effect of fixing auctions (more volume);
- the generic U shape;
- the influence of the opening of a market on another.

Reminder : if you want to compute correlations between volume in Europe, US and Asia, you need to **pay attention of simultaneity** (cf. [Hayashi and Yoshida, 2005] for returns).

Focus on Europe

The interesting aspect of European rhythms is they are affected by the opening of US markets. Moreover, 4 weeks per year the time difference change (the time change winter / summer does not take place the same week end).

Usual phases (in Europe):

- Open: uncertainty on prices and unwind of the overnight positions
- Macro economic news
- NY opens

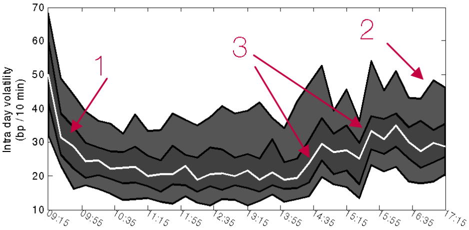

Volatility

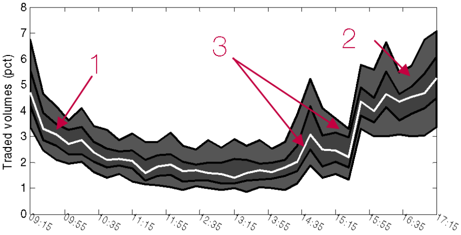

Volumes

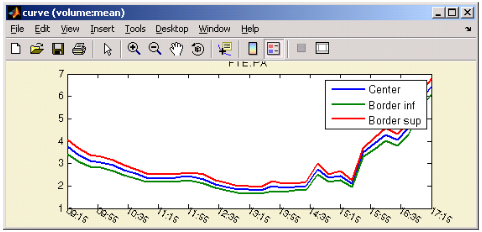

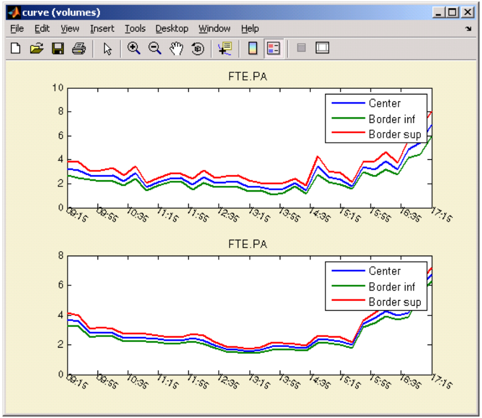

A large part of the variance comes from mixing Fridays with other days.
You can use auto correlations to obtain more robust estimators.

A "*simple*" model:

$$X_{n+1} = X_n + \sigma\sqrt{\delta t}\,\xi_n\ ,\ S_n = X_n + \varepsilon$$

Under these assumptions:

$$\sum_n (S_{n+1} - S_n)^2 = 2n\mathbb{E}(\varepsilon^2) + O(\sqrt{n})$$

Some possible ways to handle this

- Hawkes models [Bacry et al., 2015b],
- Uncertainty Zones Model [Robert and Rosenbaum, 2011],
- Clean statistics [Aït-Sahalia and Jacod, 2012],
- etc.

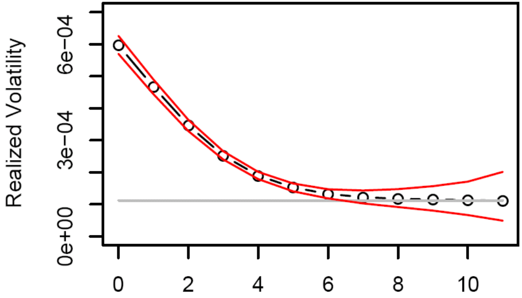

Courtesy of M. Rosenbaum

If only one microstructure effect should be kept, it is the Bid-Ask spread:

- Sell price $\neq$ Buy price
- Volume has an influence on the price
- Volatility estimations are not so simple

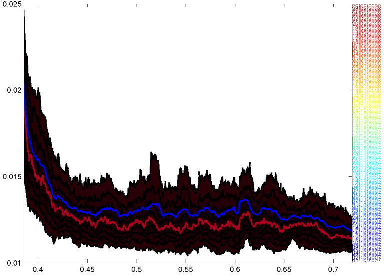

Spread-volatility relation' on three stocks

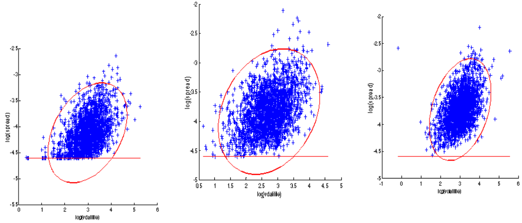

- What is the difference between these three stocks?

Spread-volatility relation' on three stocks

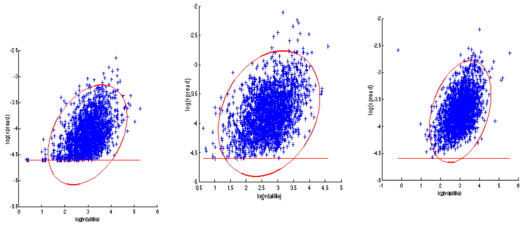

- What is the difference between these three stocks?
- The tick size does not constraint the bid-ask spread at the right, it does at the left.

Tick Size and Bid-Ask Spread

The tick size is the minimum increase between two consecutive prices. It is set in the **rulebook** of the exchange (or by the market maker). In the US it is regulated, in Europe it is not, hence in June 2009, European platforms competed to lower the tick as low as possible. It will be regulated within MiFID 2.

What should be the good value for the tick size? It is investigated in [Huang et al., 2015a]

S&P 500 July 2012

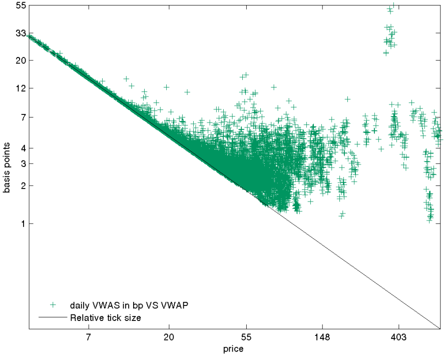

The two previous Sections can be used to explain the role of the tick size:

- The influence of the first, second and third queues on the flow distribution should be restated if the tick is smaller;
- Choosing the queue allow an agent to have a better control on the price improvement he obtained and the probability to be executed.
- Expressed in bp on a log scale, the average daily bid-ask spread as a function of the price of the instrument gives information about the efficiency of the tick [Lehalle et al., 2013].
- Tick size is (badly?) regulated in the US, and not regulated in Europe. MiFID 2 (Jan 2017) proposes to regulate it, but the nature of the “tick tables” is under discussion.

DAX July 2012

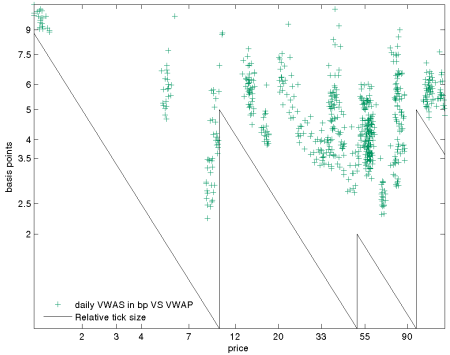

The two previous Sections can be used to explain the role of the tick size:

- The influence of the first, second and third queues on the flow distribution should be restated if the tick is smaller;
- Choosing the queue allow an agent to have a better control on the price improvement he obtained and the probability to be executed.

- Expressed in bp on a log scale, the average daily bid-ask spread as a function of the price of the instrument gives information about the efficiency of the tick [Lehalle et al., 2013].
- Tick size is (badly?) regulated in the US, and not regulated in Europe. MiFID 2 (Jan 2017) proposes to regulate it, but the nature of the “tick tables” is under discussion.

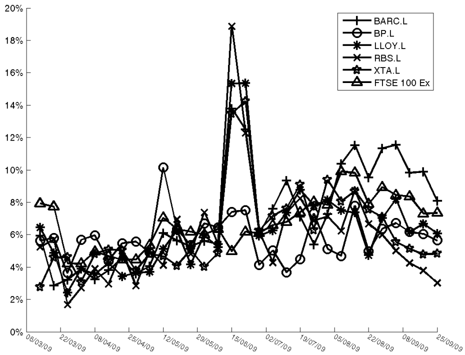

First of all, it is used in competition across trading venues.

- when a venue decreases its tick, it enable cheaper queue jumping in its orderbook;
- hence some traders will post there at +1 tick;
- finally the SORs will capture this price improvement.

Remark 1 (*The tick size*)

1. *Ideally, the information rate on an instrument should be able to generate a price change of one tick in few trades;*
2. *large tick stocks focus the trading strategies on queueing instead of splitting.*

Universal Relation Between Bid-Ask Spreads and Volatility

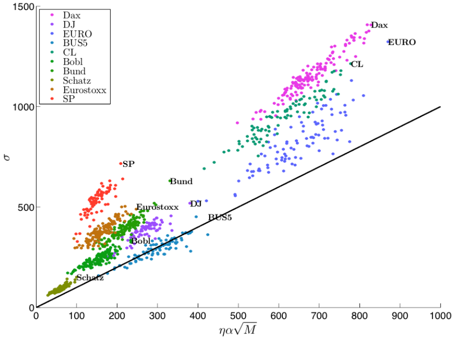

Taken From [Dayri and Rosenbaum, 2015]

- There is an universal relation between bid-ask spreads ψ and volatility σ,
  Economists told us the market makers had to be paid by the spread for the volatility risk they take:

  $$\psi \propto \sigma.$$

- On a stock by stock basis, the proportional factor seems to be close to the square root of the number of trades per day [Wyart et al., 2008]:

  $$\psi \propto \frac{\sigma}{\sqrt{N}}.$$

  The rational is the more trades per day, the easier to maintain an inventory.

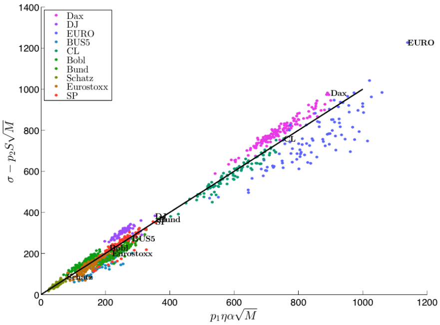

Taken From [Dayri and Rosenbaum, 2015]

But for large tick instruments, the relation breaks, a correction by $\eta = N^c/N^a/2$ has to be made (ψ is replaced by $\eta\delta$, where δ is the tick and where ϕ is an additional gain for market makers):

- There is an universal relation between bid-ask spreads ψ and volatility σ,
  Economists told us the market makers had to be paid by the spread for the volatility risk they take:

$$\psi \propto \sigma.$$

- On a stock by stock basis, the proportional factor seems to be close to the square root of the number of trades per day [Wyart et al., 2008]:

$$\psi \propto \frac{\sigma}{\sqrt{N}}.$$

  The rational is the more trades per day, the easier to maintain an inventory.

$$\boxed{\psi \simeq \eta\delta \propto \frac{\sigma}{\sqrt{N}} + \phi}.$$

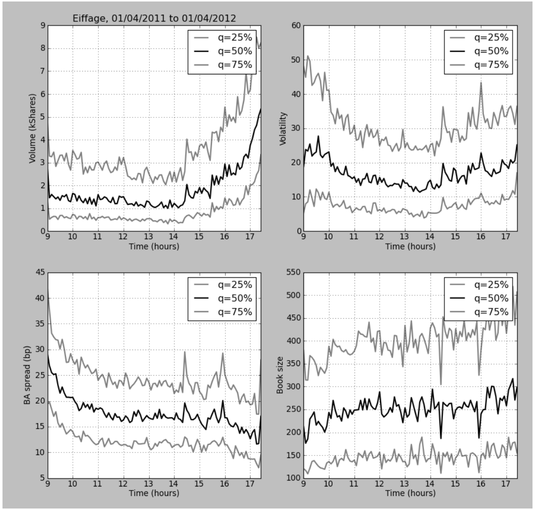

Intraday Seasonalities Essentials

- 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.

Orderbooks are the place where the matching of orders take place:

- an inserted order is compared to potential matching ones already in the engine,
- if it matches to one or more orders, it give birth to **transactions**,
- otherwise it is stored into the matching engine.

- **orderbook dynamics** is the study and the modelling of the insertions and cancellations of orders, conditionally to its "state". The more information in the "state", the more complex (and accurate?) model.
- There is now a large offer of models, from "zero-intelligence" ones (see [Smith et al., 2003]) to "game theoretic" ones (see [Lachapelle et al., 2016]), via empirics-driven PDE ones (see [Gareche et al., 2013]).

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

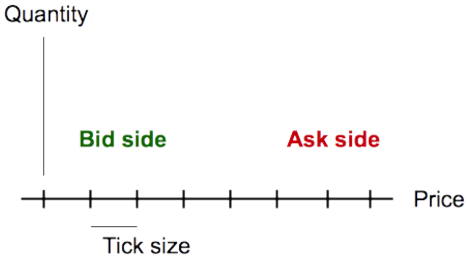

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

1. The initial state of the LOB is empty, then a trader send a first order. It is stored in the LOB.

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

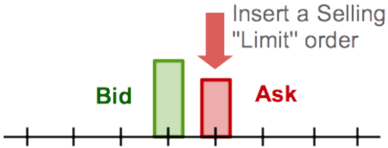

1. The initial state of the LOB is empty, then a trader send a first order. It is stored in the LOB.
2. Another trader send an order in the other direction at a non matching price.

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

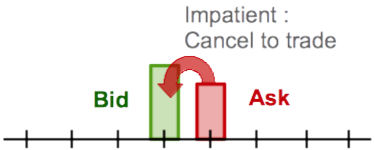

1. The initial state of the LOB is empty, then a trader send a first order. It is stored in the LOB.
2. Another trader send an order in the other direction at a non matching price.
3. The seller is impatient he cancels his resting order to insert it at a matching (less interesting price).

Empirics on Orderbooks Dynamics (continuous phase)

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

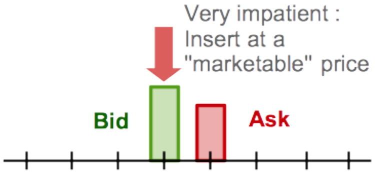

1. The initial state of the LOB is empty, then a trader send a first order. It is stored in the LOB.
2. Another trader send an order in the other direction at a non matching price.
3. The seller is impatient he cancels his resting order to insert it at a matching (less interesting price).
4. Or a very impatient trader can directly insert a “marketable” order.

The limit orderbook (LOB) is a software receiving instructions from traders via members of the exchange (typically brokers). Orders contain a direction (side), a quantity, and (possibly) a price. I.e. buy 250 shares (of Vodafone) at 15.000. The price is a multiple of the **tick size**.

1. The initial state of the LOB is empty, then a trader send a first order. It is stored in the LOB.
2. Another trader send an order in the other direction at a non matching price.
3. The seller is impatient he cancels his resting order to insert it at a matching (less interesting price).
4. Or a very impatient trader can directly insert a "marketable" order.
5. If the buyer is really patient, he can cancel his order and wait at a more advantageous price.

Components of the Orderbook Dynamics

The **atomic events** impacting an orderbook are

- insertion,
- cancellation,
- transaction.

The available **information**

- the state of the orderbook itself:
  - queue size at each tick,
  - could be reduced to the queue size at each level (or limit),
  - or to the imbalance only (cf. the worst kept secret of HFTs)
- the past imbalances, trades, etc. (including the volatility)
- bids and asks on correlated instruments (Futures, ETF, Options, etc).

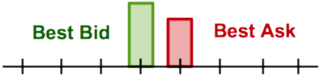

- To understand the dynamics of the queues, we will capture the inflows and outflows on the two sides of the book.

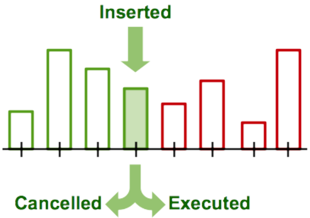

- To understand the dynamics of the queues, we will capture the inflows and outflows on the two sides of the book.
- Out first model will be the simplest one: each queue has its own dynamics. It is a little more than "zero-intelligence models", since the intensity of the three flows will be functions of the state (i.e. size) of the queue.

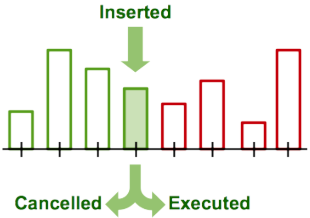

- To understand the dynamics of the queues, we will capture the inflows and outflows on the two sides of the book.
- Out first model will be the simplest one: each queue has its own dynamics. It is a little more than “zero-intelligence models”, since the intensity of the three flows will be functions of the state (i.e. size) of the queue.
- We renormalize the quantities in AES (“Average Event Size” of the queue), and work in “ticks” and not in “limits”.

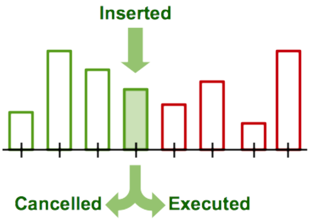
Inserted
Cancelled
Executed

First Limit

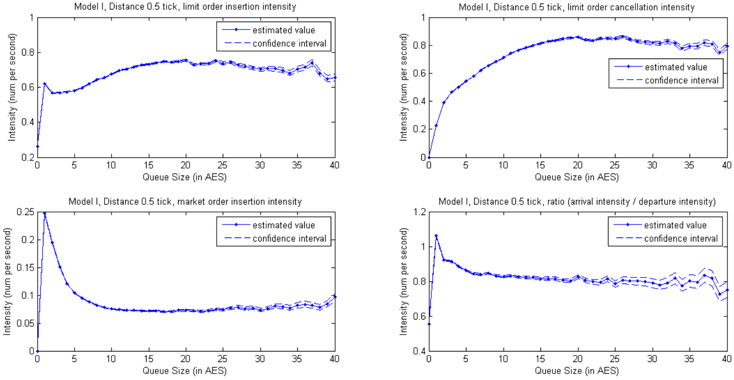
Model I, Distance 0.5 tick, limit order insertion intensity
estimated value
confidence interval
Intensity (num per second)
Queue Size (in AES)
Model I, Distance 0.5 tick, limit order cancellation intensity
estimated value
confidence interval
Intensity (num per second)
Queue Size (in AES)
Model I, Distance 0.5 tick, market order insertion intensity
estimated value
confidence interval
Intensity (num per second)
Queue Size (in AES)
Model I, Distance 0.5 tick, ratio (arrival intensity / departure intensity)
estimated value
confidence interval
Intensity (num per second)
Queue Size (in AES)

Second Limit

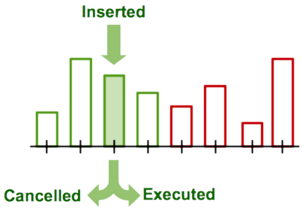

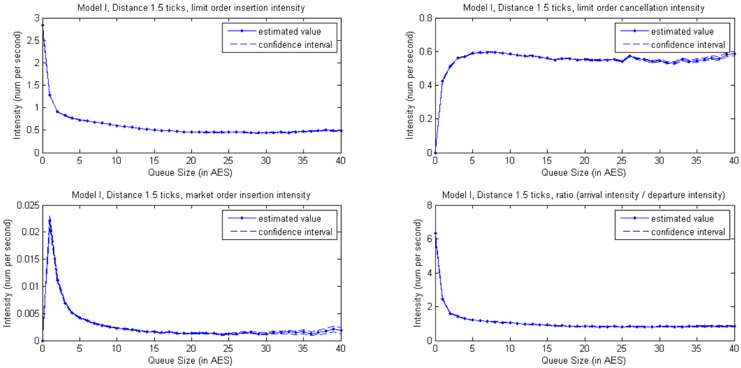

We can derive theoretically the invariant distribution of the quantity at a given distance of the mid price

- using the data,
- using our model fitted on the data,
- using a Poisson process fitted on the data for each queue.

  (do not forget once a queue is empty the mid price changes)

We can compare asymptotic distributions

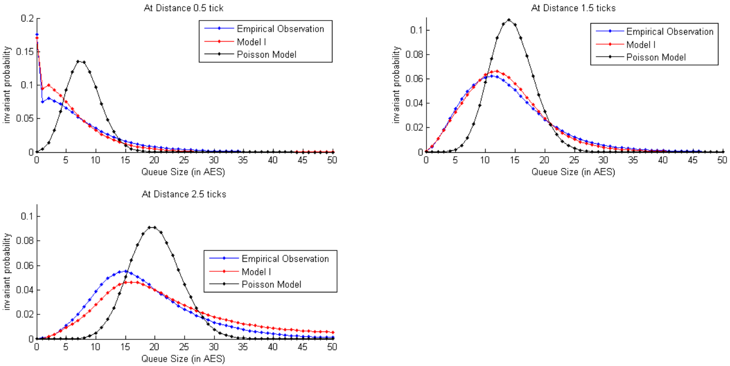

Events (insert, cancel, trade) are counted, modelled using three independent Point Process.

When the State is the Size of One Queue.

The intensity of Point Processes for cancels, inserts and trades are function of the size of the concerned queue.

- Asymptotic distributions of sizes are close to the real ones,
- But the probability of a limit order to be executed obtained by this model is not close to real ones.

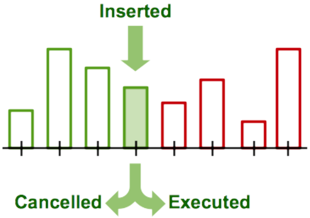

- Model II is a sophistication of Model I
- Flows at first queue are a function of its state (i.e. size)

Model II

- Model II is a sophistication of Model I
- Flows at first queue are a function of its state (i.e. size)
- and flows at the second queue are a function of its own state and the one of the first queue (same side).

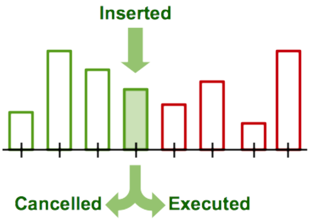
Inserted
Cancelled
Executed

First Limit

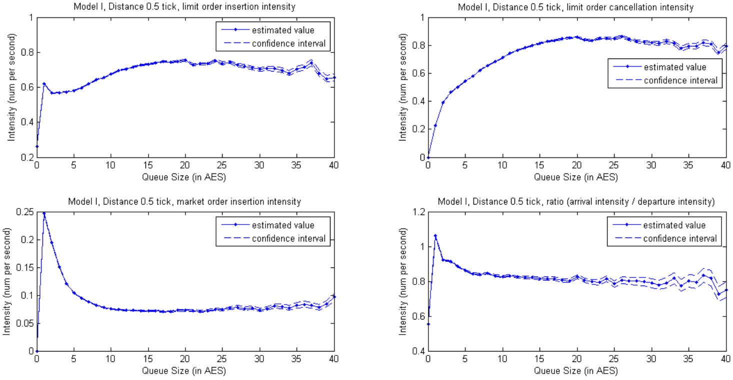
Model I, Distance 0.5 tick, limit order insertion intensity
Model I, Distance 0.5 tick, limit order cancellation intensity
Model I, Distance 0.5 tick, market order insertion intensity
Model I, Distance 0.5 tick, ratio (arrival intensity / departure intensity)
estimated value
confidence interval
Intensity (num per second)
Queue Size (in AES)

Second Limit

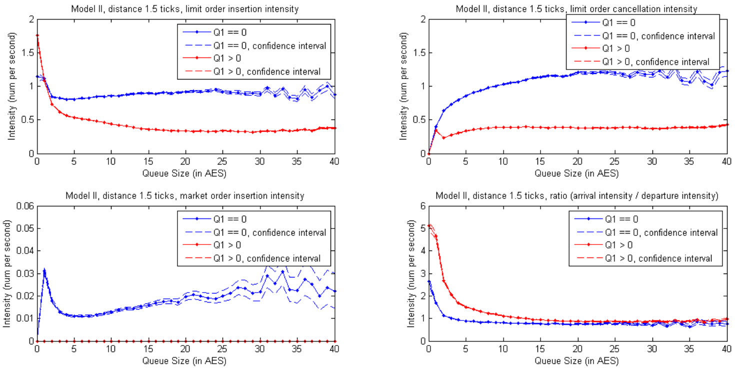

Adding the influence of the first queue on the second one

Adding as information the size of the first queue on the same side, we see the intensities on the second queue are really different when the curent queue is **protected** by a large first queue rather than are in fact close to a first limit.

- Model III is a sophistication of Model II
- Like for Model II: flows at the second queue are a function of its own state and the one of the first queue (same side).

- Model III is a sophistication of Model II
- Like for Model II: flows at the second queue are a function of its own state and the one of the first queue (same side).
- but flows of the first queue are influenced by its size and the best opposite queue size.


Inserted
Influences
Cancelled
Executed

First Limit

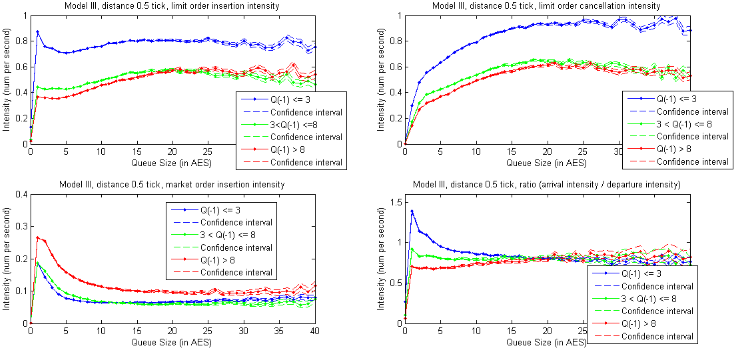
Model III, distance 0.5 tick, limit order insertion intensity
Model III, distance 0.5 tick, limit order cancellation intensity
Model III, distance 0.5 tick, market order insertion intensity
Model III, distance 0.5 tick, ratio (arrival intensity / departure intensity)
Intensity (num per second)
Queue Size (in AES)
Q(-1) <= 3
Confidence interval
3 < Q(-1) <= 8
Confidence interval
Q(-1) > 8
Confidence interval

Second Limit

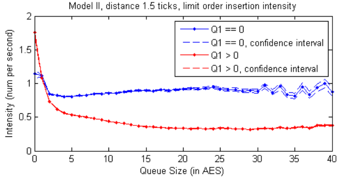

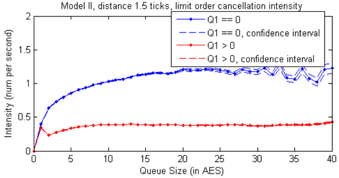

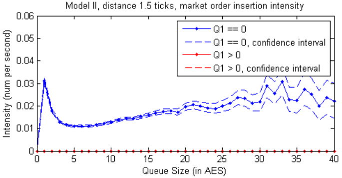

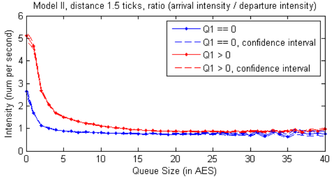

Influence of the opposite queue

Adding the size of the opposite queue in the state space increases the accuracy of the model. It is a way to take into account the **worst kept secret of High Frequency Trades** (i.e. influence of imbalance on next mid-price move).

The Queue Reactive Model is a sophistication of Model III

- Like for Model III: flows of the first queue are influenced by its size and the best opposite queue size.

Model IV: The Queue Reactive Model

The Queue Reactive Model is a sophistication of Model III

- Like for Model III: flows of the first queue are influenced by its size and the best opposite queue size.
- Like for Model II: flows at the second queue are a function of its own state and the one of the first queue (same side).

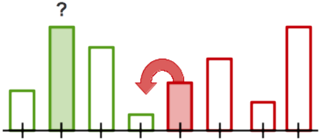

The Queue Reactive Model is a sophistication of Model III

- Like for Model III: flows of the first queue are influenced by its size and the best opposite queue size.
- Like for Model II: flows at the second queue are a function of its own state and the one of the first queue (same side).
- But we now model the behaviour of the third queue (previously it was Model I driven), when it become the second queue we keep it with proba $1-\theta$ or draw its size according to the observed distribution (with proba θ).

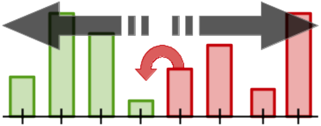

The Queue Reactive Model is a sophistication of Model III

- Like for Model III: flows of the first queue are influenced by its size and the best opposite queue size.
- Like for Model II: flows at the second queue are a function of its own state and the one of the first queue (same side).
- But we now model the behaviour of the third queue (previously it was Model I driven), when it become the second queue we keep it with proba $1-\theta$ or draw its size according to the observed distribution (with proba θ).
- And we had to add another effect: with a proba θ^{reinit}, the full orderbook is drawn from its stationary distribution when the price change.

Again we obtained asymptotic results.

We are more happy with execution probabilities, but we add one criterion: we measure the volatility and to compare it with the historical one.

We see θ controls the mean reversion of the price. Hence θ and θ^{reinit} have to be chosen to reproduce the mean reversion and the volatility of the modelled instrument (typically $\theta = 0.2$ and $\theta^{reinit} = 0.7$ for France Telecom / Orange).

Volatility and $\eta = N^c/2N^a$ with respect to θ

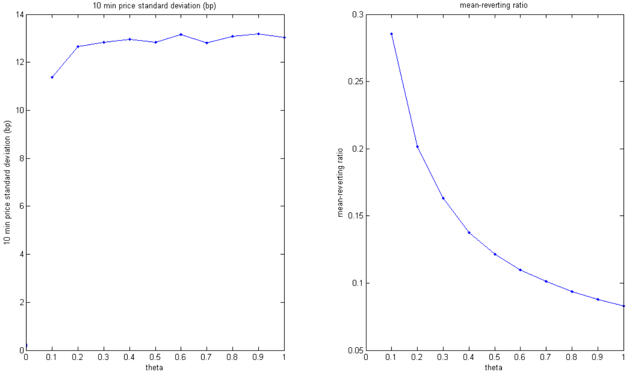

Adding exogenous information

Using the current sizes of the queues is not enough to recover realistic orderbook dynamics. Two more effects have to be added

- When a third limit is promoted to second limit[a], you cannot reuse its size. On average it is too large, impulsing an high mean-reversion to the price (and hence a very low volatility).
- Moreover after a price change you need to "reset" the sizes on the two sides of the book (to prevent a too high mean-reversion too).

$\Rightarrow$ memory of the recent past of events is needed, to prevent a too high level of mean reversion.

[a] The notions of "third" and "second" limits have to be properly defined. Here the stock is a large tick one.

Time to time (in around 10 to 20% of the time), everyone accept the new price.
Think about two previous macroscopic cases:

- The liquidation of Kerviel's inventory: prices changes because of liquidity consumption, no one accepts the new prices. $\Rightarrow$ **Large mean reversion**.
- The announce by President Sarkozy of the end of advertising on public TV channels. Prices change fast, and everyone accepts the new prices. $\Rightarrow$ **Large permanent price change**.

Conclusion on Orderbook Empirics

Using staged models of orderbook dynamics, we observed

- The sub-linear, increasing, cancellation rate.
- The decreasing limit order insertion rate for non-best limits.

- Agents are acting strategically in the orderbook and this has to be taken into account.
- A simple way is to unconditionally “reset” the orderbook with a given probability θ^{reinit}. In reality other (exogenous) factors are probably affecting this probability (i.e. the “acceptance” of a newly printed price).

- In the paper (with Huang and Rosenbaum) we do more: we use the Queue Ractive Model to simulate trading strategies.

We have seen participants act strategically

- when a seller enters into the “trading game”,
- he can choose to wait, sending a limit order,
- or to pay immediately a “price impact”.
- if he waits he can obtain a better price (the price impact of the opposite, liquidity consuming, order)

$\Rightarrow$ he needs to valuate the price of waiting in the queue.

We provided a theoretical study with numerical simulations [Lachapelle et al., 2016]. The closest existing study (in economics) is [Roşu, 2009]. Empirical studies investigated on high frequency data ([Gareche et al., 2013], [Huang et al., 2015b]).

Our Mean Field Game Model: One Queue

- Sellers only,
- one agent i arrives in "the game" at t according to a Poisson process N of intensity λ,
- it compares the value to wait in the queue ($y(x)$, where x is the size of the queue) to zero to choose to wait in the queue (when $u(x) > 0$) or not, its decision is δ^i
- the queue is consumed by a Poisson process $M^{\mu(x)}$ of intensity $\mu(x)$,
- in case of transaction, a "pro-rata" scheme is used ("equivalent" to infinitesimal possibility to modify orders): q/x of the order is part of it; can be relaxed for FIFO.

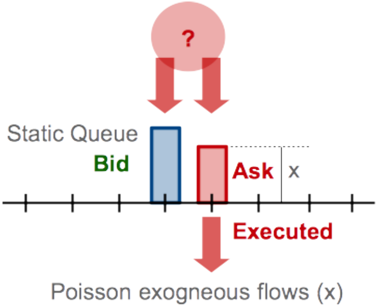

The **Mean Field** is the size of the queue (it is a **forward** process):

$$dx_t = q\left(dN_t^j \delta^j - dM_t^{\mu(x_t)}\right), \text{ remark: } j = N, \text{ I could have written } dN_t \delta^N.$$

The **Value function** the ith agent wants to minimize is driven by the following running cost

$$dJ(x_t) = \left[\frac{q}{x_t} P(x_t) + (1 - \frac{q}{x_t}) J(x_t - q)\right] dM_t^{\mu(x)} - cq\, dt.$$

$$u(x) := \mathbb{E} \int_{t_0}^{T} dJ(x_t),$$

and its control δ^i is to choose to be submitted to this cost function or to pay zero at t_0:

$$\mathcal{U}^i(x) := \max_{\delta^i \in \{0,1\}} \delta^i u(x).$$

The optimal decision δ^i is the solution of the **backward** associated dynamics.

The value function evolves following these four main possible events (we have price impact $P(x)$ and waiting cost proportional to $c\,q$; seen from *any agent*, the control δ^i is now replaced by the anonymous decision $\mathbf{1}_{u(x)>0}$):

$$\begin{aligned} u(x,t+dt) = \; & (1-\lambda \mathbf{1}_{u(x)>0} dt - \mu(x) dt) \cdot u(x) & \leftarrow \text{nothing happens} \\ & + \lambda \mathbf{1}_{\{u(x)>0\}} dt \cdot u(x+q) & \leftarrow \text{new entrance} \\ & + \mu(x) dt \cdot \Big(\frac{q}{x} P(x) + (1-\frac{q}{x}) u(x-q) \Big) & \leftarrow \text{transaction} \\ & - c\, q\, dt & \leftarrow \text{waiting cost} \end{aligned}$$

To solve it at the kth order, we will perform a Taylor expansion of $u(x+q)$ and $u(x-q)$ for small q at the kth order).

At the second order, we obtain:

$$0 = \frac{\mu(x)}{x}(P(x) - u) - c + (\lambda \mathbf{1}_{\{u>0\}} - \mu(x))u' + q\Big(\frac{1}{2}(\lambda \mathbf{1}_{\{u>0\}} - \mu(x))u'' + \frac{\mu(x)}{x}u'\Big),$$

Taylor expansion of the master equation

$$0 = \frac{\mu(x)}{x}(P(x) - u) - c + (\lambda \mathbf{1}_{\{u>0\}} - \mu(x))u'$$

Just keep the first order term

It corresponds to a (trivial) shared risk Mean Field Game monotone system with $N = 1$. The *mean field* aspect does not come from the continuum of agents (for every instant, the number of players is finite), but rather from the stochastic continuous structure of entries and exits of players.

At queue sizes x^* such that $x^* = \mu(x^*)P(x^*)/c$, u sign changes.
Moreover, for the specific case $\mu(x) = \mu_1 \mathbf{1}_{x<S} + \mu_2 \mathbf{1}_{x \geq S}$:

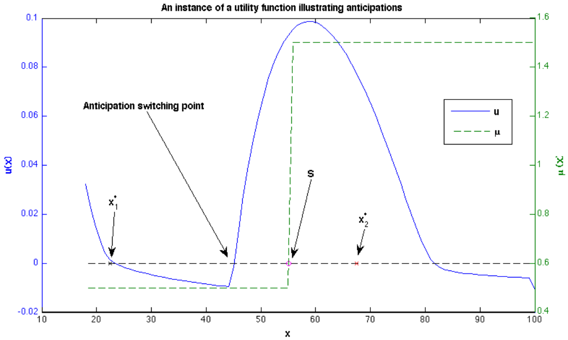

There is a point strictly before S where u switches from negative to positive. It means that participants anticipate service improvement.

The decision-taking process will follow this mechanism:

- consuming liquidity allows to obtain quantity immediately but at an impacted price, with respect to the liquidity available in the book,
- each time a market participant has to take a buy or sell decision, he tries to anticipate the "*long term*" value for him to be liquidity provider or liquidity consumer,
- each market participant can use a SOR (Smart Order Router [Foucault and Menkveld, 2008]) for this sophisticated valuation, otherwise he will just consume liquidity.

The decision-taking process will follow this mechanism:

- consuming liquidity allows to obtain quantity immediately but at an impacted price, with respect to the liquidity available in the book,
- each time a market participant has to take a buy or sell decision, he tries to anticipate the "*long term*" value for him to be liquidity provider or liquidity consumer,
- each market participant can use a SOR (Smart Order Router [Foucault and Menkveld, 2008]) for this sophisticated valuation, otherwise he will just consume liquidity.

The decision-taking process will follow this mechanism:

- consuming liquidity allows to obtain quantity immediately but at an impacted price, with respect to the liquidity available in the book,
- each time a market participant has to take a buy or sell decision, he tries to anticipate the "*long term*" value for him to be liquidity provider or liquidity consumer,
- each market participant can use a SOR (Smart Order Router [Foucault and Menkveld, 2008]) for this sophisticated valuation, otherwise he will just consume liquidity.

Model details

- Orders arrive at Poisson rate $\Lambda = \lambda + \lambda^-$
- strategic arrivals: λ, non-optimal: λ^- (can be read as “SOR” on “non-SOR” participants)
- $(Q_a, Q_b) :=$ number of orders on ask and bid sides
- Value functions: $u(Q^a, Q^b)$ for sellers and $v(Q^a, Q^b)$ for buyers
- Matching process for any quantity Q: Qq/Q^a
- Transaction price:

$$p^{\text{buy}}(Q^a) := P + \frac{\delta q}{Q^a - q}, \quad p^{\text{sell}}(Q^b) := P - \frac{\delta q}{Q^b - q}$$

where: q is the order size, δ is the market depth, P is the fair price

- cost to maintain inventory: c_a and c_b

- If $u(Q_t^a + q, Q_t^b) > p^{\text{sell}}(Q_t^b)$, it is more valuable to route the sell order to the ask queue $\rightarrow$ **Liquidity Consumer (LC) order**

- If $v(Q_t^a, Q_t^b + q) < p^{\text{buy}}(Q_t^a)$, it is more valuable to route the buy order to the bid queue $\rightarrow$ **Liquidity Provider (LP) order**

- Notations of the routing decisions:

$$R_{\text{buy}}^{\oplus}(v, Q_t^a, Q_t^b + q) := \mathbf{1}_{v(Q_t^a, Q_t^b+q) < p^{\text{buy}}(Q_t^a)}, \text{ LP buy order}$$
$$R_{\text{sell}}^{\oplus}(u, Q_t^a + q, Q_t^b) := \mathbf{1}_{u(Q_t^a+q, Q_t^b) > p^{\text{sell}}(Q_t^b)}, \text{ LP sell order}$$
$$R_{\text{buy}}^{\ominus}(Q_t^a, Q_t^b) := 1 - R_{\text{buy}}^{\oplus}(Q_t^a, Q_t^b), \text{ LC buy order}$$
$$R_{\text{sell}}^{\ominus}(Q_t^a, Q_t^b) := 1 - R_{\text{sell}}^{\oplus}(Q_t^a, Q_t^b), \text{ LC sell order is processed}$$

The **2D mean field** is the size of the two queues (Q_t^a, Q_t^b); it evolves according to the **forward** dynamics (j, k, ℓ are strategic ask providing, strategic ask consuming, and blind ask consuming agents):

$$dQ_t^a = (dN_t^{\lambda_{\text{sell}}(j)} \underbrace{R_{\text{sell}}^{\oplus}(j)}_{\delta^j} - (dN_t^{\lambda_{\text{buy}}(k)} \underbrace{R_{\text{buy}}^{\ominus}(k)}_{\delta^k} + dN_t^{\lambda^-(\ell)}))\, q,$$

and for the cost function at the ask:

$$dJ^u(Q^a, Q^b) = \left[\frac{q}{Q^a} p^{\text{buy}}(Q^a) + \left(1 - \frac{q}{Q^a}\right) J^u(Q^a - q, Q^b)\right] (dN_t^{\lambda_{\text{buy}}(k)} \underbrace{R_{\text{buy}}^{\ominus}(k)}_{\delta^k} + dN_t^{\lambda^-(\ell)}) - c_a q\, dt.$$

Again, with T large enough, $u(\mathcal{Q}^a, \mathcal{Q}^b) = \mathbb{E} \int_{t=0}^T J(Q_t^a, Q_t^b)\, dt$ given $Q_0^a = \mathcal{Q}^a$, $Q_0^b = \mathcal{Q}^b$, and

$$\mathcal{U}(\mathcal{Q}^a, \mathcal{Q}^b) := \max_{\delta^i \in \{0,1\}} \delta^i u(\mathcal{Q}^a, \mathcal{Q}^b) + (1 - \delta^i) p^{\text{sell}}(Q^b).$$

The control δ^i is thus the result of a **backward** process.

Equation of a Seller Utility Function

$$
\begin{aligned}
u(Q_t^a, Q_t^b) \quad = \quad & (1 - \lambda_{\text{buy}} dt - \lambda_{\text{sell}} dt - 2\lambda^- dt)\, u(Q_t^a, Q_t^b) && \leftarrow \text{nothing} \\
& + (\lambda_{\text{sell}} R_{\text{sell}}^{\ominus}(u, Q_t^a + q, Q_t^b) + \lambda^-) dt\, u(Q_t^a, Q_t^b - q) && \leftarrow \text{sell order, LC} \\
& + \lambda_{\text{sell}} R_{\text{sell}}^{\oplus}(u, Q_t^a + q, Q_t^b) dt\, u(Q_t^a + q, Q_t^b) && \leftarrow \text{sell order, LP} \\
& + (\lambda_{\text{buy}} R_{\text{buy}}^{\ominus}(v, Q_t^a, Q_t^b + q) + \lambda^-) dt \cdot [&& \leftarrow \text{buy order, LC} \\
& \qquad \underbrace{\frac{q}{Q_t^a} p^{\text{buy}}(Q_t^a)}_{\text{trade part (ask)}} \; \underbrace{+ (1 - \frac{q}{Q_t^a})\, u(Q_t^a - q, Q_t^b)]}_{\text{removing (ask)}} \\
& + \lambda_{\text{buy}} R_{\text{buy}}^{\oplus}(v, Q_t^a, Q_t^b + q) dt\, u(Q_t^a, Q_t^b + q) && \leftarrow \text{buy order, LP} \\
& - c_a q\, dt. && \leftarrow \text{waiting cost}
\end{aligned}
$$

And symmetrical dynamics for a buyer utility function.

Why this Model is Important?

It provides a way to compute the utility function of a trader in the book:

- For the owner of a limit order, just write the equation a non stationarized way, you obtain the dynamics of the utility function. Add a terminal condition (the limit order is executed) and you can write an optimal strategy for a limit order.
- For the owner of a market order, should it be better to wait inside the book? just read the stationarized utility function.
- and you can elaborate...

- Symmetric case: $c_a = c_b = c$
- Simpler notations: $x := Q_a$ and $y := Q_b$

First Order Equations

$$0 = [(\lambda R_b^{\ominus} + \lambda^-)\frac{1}{x}(p^b(x) - u) - c] + [\lambda R_s^{\oplus} - \lambda R_b^{\ominus} - \lambda^-] \cdot (\partial_x u + \partial_y u)$$
$$0 = [(\lambda R_s^{\ominus} + \lambda^-)\frac{1}{y}(p^s(y) - v) + c] + [\lambda R_s^{\oplus} - \lambda R_b^{\ominus} - \lambda^-] \cdot (\partial_x u + \partial_y u)$$

General form: shared risk MFG where $m := (x, y) \in \mathbb{R}^2$

$$0 = \beta_a(u, v, x, y) + \alpha(u, v, x, y)(\partial_x u + \partial_y u)$$
$$0 = \beta_b(u, v, x, y) + \alpha(u, v, x, y)(\partial_x v + \partial_y v)$$

Four mixes of LC and/or LP agents:

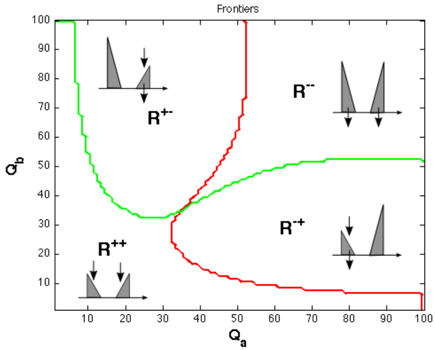

Sellers and buyers are Liquidity Providers

$$R^{++} = \{(x,y), R_s^{\oplus}(x,y) = R_b^{\oplus}(x,y) = 1\},$$

Sellers and buyers are Liquidity Consumers

$$R^{--} = \{(x,y), R_s^{\ominus}(x,y) = R_b^{\ominus}(x,y) = 1\},$$

Sellers provide liquidity and buyers consume it

$$R^{+-} = \{(x,y), R_s^{\oplus}(x,y) = R_b^{\ominus}(x,y) = 1\},$$

Sellers consume liquidity and buyers provide it

$$R^{-+} = \{(x,y), R_s^{\ominus}(x,y) = R_b^{\oplus}(x,y) = 1\}.$$

Lemma

$$\forall x, y,\ R^{\oplus}_{\text{sell}}(u, x, y) = R^{\oplus}_{\text{buy}}(2P - v, y, x)$$

And as a direct consequence, if there is a unique solution (u, v) to the previous system, then

$$\forall x, y,\quad u(x, y) + P = P - v(y, x).$$

- The characteristic lines of the solutions to the system have the form $y = x + k$
- Thanks to the lemma: we solve the equations along the characteristics, only on the region $y \leq x$ (same reasoning on $y \geq x$).
- Think of small bid and ask queues on a given $y = x + k$, $k > 0$. First both buyers and sellers are LP.

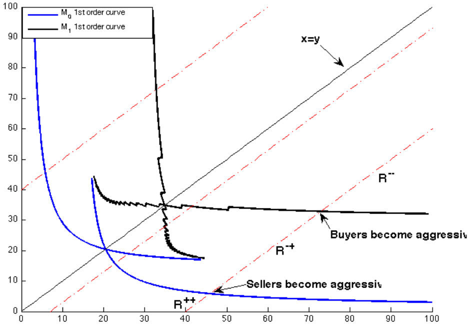

The curves analytically computed at order 1.

- As $(x, y = x + k)$ grow, sellers turn to be LC first, while buyers remain LP (boundary between R^{++} and R^{-+}). Looking at the equations, we can get the parametric curve of points $M_0 = (x_0, l(x_0))$ where the switch happens.
  (The coeff multiplying the derivatives switches sign + equality in the routing decision criterium)
- We can do the same kind of reasoning for the second switching curve M_1, when buyers turn to be LC (boundary between R^{-+} and R^{--}).
  (Here no sign switch. M_0 is used for ODE resolution + antisymmetry argument + equality in the routing decision criterium)

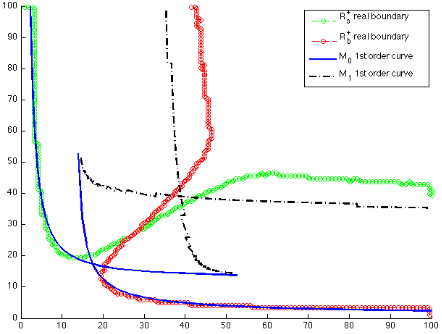

Computations corroborate the first order expansion ... and show second order terms effects

- As $(x, y = x + k)$ grow, sellers turn to be LC first, while buyers remain LP (boundary between R^{++} and R^{-+}). Looking at the equations, we can get the parametric curve of points $M_0 = (x_0, l(x_0))$ where the switch happens.
  (The coeff multiplying the derivatives switches sign + equality in the routing decision criterium)
- We can do the same kind of reasoning for the second switching curve M_1, when buyers turn to be LC (boundary between R^{-+} and R^{--}).
  (Here no sign switch. M_0 is used for ODE resolution + antisymmetry argument + equality in the routing decision criterium)

General form

$$0 = \beta_a(u,v,x,y) + \alpha(u,v,x,y)(\partial_x u + \partial_y u)$$
$$+ q\Big(\rho(v,x,y))\partial_x u + \xi_1(u,v,x,y)\partial_{xx} u + \xi_2(u,v,x,y)\partial_{yy} u\Big),$$
$$0 = \beta_b(u,v,x,y) + \alpha(u,v,x,y)(\partial_x v + \partial_y v)$$
$$+ q\Big(\rho(2P - u,y,x))\partial_y v + \xi_1(u,v,x,y)\partial_{xx} v + \xi_2(u,v,x,y)\partial_{yy} v\Big),$$

- where:

$\rho = \frac{1}{x}(\lambda R_b^{\ominus} + \lambda^-)$,
$\xi_1 = (\lambda(R_s^{\oplus} + R_b^{\ominus}) + \lambda^-)/2$, and $\xi_2 = (\lambda(R_s^{\ominus} + R_b^{\oplus}) + \lambda^-)/2$.

Now, we use this toolbox to model various markets...

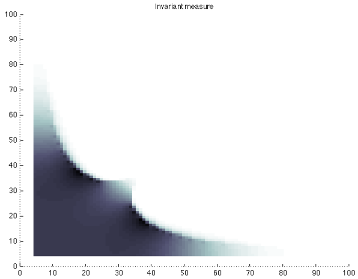

In the paper, we use this framework to study:

- the equilibrium with one type of agents: **Stable liquidity imbalance states** are possible.
- When another type of agent is added (faster): **The imbalanced states are fewer**, and the bid-ask spread (i.e. average cost for liquidity consumers) decreases. But its decrease is **in favour of the faster traders**.
- The results are compatible with empirical studies ([Gareche et al., 2013], [Huang et al., 2015b]).

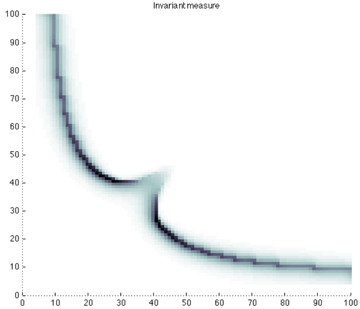

In the paper, we use this framework to study:

- the equilibrium with one type of agents: **Stable liquidity imbalance states** are possible.
- When another type of agent is added (faster): **The imbalanced states are fewer**, and the bid-ask spread (i.e. average cost for liquidity consumers) decreases. But its decrease is **in favour of the faster traders**.
- The results are compatible with empirical studies ([Gareche et al., 2013], [Huang et al., 2015b]).

Conclusion on MFG Modelling

Mean Field Games seems to be an adequate framework to model the controlled dynamics of the orderbooks

- It needed to identify carefully the intensives of agents,
- for simplicity, we modelled pro-rata rules, but it could be extended,
- we obtain results that are in line with empirical observations.
- in the paper, we derive results about the effect of mixing time scales in the same orderbooks.
- in another paper [Lehalle et al., 2010] we attempt to introduce MFG at a largest time scale (i.e. once a participant traded, he needs to unwind this position). In this other model we introduced the idea of "latent orderbook" (i.e. our mean field: the aggregation of the views of all market participants).
- There is another paper using MFG at an intermediate time scale to design kind of "robust" optimal trading for liquidation [Jaimungal and Nourian, 2015].

The Liquidity Game

- It is possible to write properly the value of one limit order in a book (and to obtain its stationarized value);
- In the MFG model, liquidity imbalance can be stable, it is **in contradiction to the worst kept secret of HFTs**;
- But if we mix trading speeds, no imbalance is stable anymore.

$\Rightarrow$ It is like orderbook dynamics alternate between two states

- a **liquidity game**, during which traders compete for liquidity;
- when one of the queue is very small, every one run to be the last to capture the remaning quantity; i.e. the worst kept secret of HFT is reinforced this way.
- just after the best queue has been fully consumed, a **price game** takes place: do the traders accept the new price or let the current sizes of the queues drive mean reversion?

Traders look at “exogenous” information to accept the new price or not: indices, Futures, news, long lasting consuming pressure.

CFM
INSIGHT.DATA.CLARITY.

Thank You For Your Attention

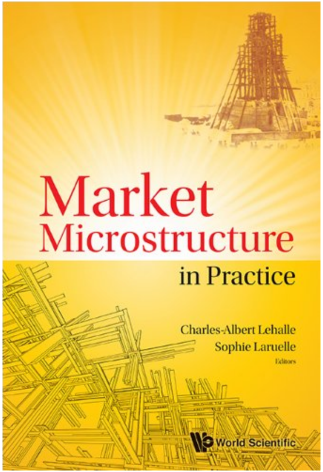

References I

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.

Aït-Sahalia, Y. and Jacod, J. (2012).
Analyzing the Spectrum of Asset Returns: Jump and Volatility Components in High Frequency Data.
Journal of Economic Literature, 50(4):1007–50.

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. (2015a).
Market Impacts and the Life Cycle of Investors Orders.
Market Microstructure and Liquidity, 1(2).

Bacry, E., Mastromatteo, I., and Muzy, J.-F. (2015b).
Hawkes Processes in Finance.
Market Microstructure and Liquidity, 01(01):1550005.

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.

Brogaard, J., Hendershott, T. J., and Riordan, R. (2012).
High Frequency Trading and Price Discovery.
Technical report, Berkeley University.

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., Donnelly, R. F., and Jaimungal, S. (2013).
Algorithmic Trading with Model Uncertainty.
Social Science Research Network Working Paper Series.

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

References III

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.

Dayri, K. and Rosenbaum, M. (2015).
Large Tick Assets: Implicit Spread and Optimal Tick Size.
Market Microstructure and Liquidity, 01(01):1550003.

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

Foucault, T. and Menkveld, A. J. (2008).
Competition for Order Flow and Smart Order Routing Systems.
The Journal of Finance, 63(1):119–158.

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

Gareche, A., Disdier, G., Kockelkoren, J., and Bouchaud, J.-P. (2013).
A Fokker-Planck description for the queue dynamics of large tick stocks.
Phys. Rev. E, 88(3):32809–32820.

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.

Hayashi, T. and Yoshida, N. (2005).
On Covariance Estimation of Non-synchronously Observed Diffusion Processes.
Bernoulli, 11(2):359–379.

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

Huang, W., Lehalle, C.-A., and Rosenbaum, M. (2015a).
How to predict the consequences of a tick value change? Evidence from the Tokyo Stock Exchange pilot program.

Huang, W., Lehalle, C.-A., and Rosenbaum, M. (2015b).
Simulating and analyzing order book data: The queue-reactive model.
Journal of the American Statistical Association, 10(509).

Jaimungal, S. and Nourian, M. (2015).
Mean-Field Game Strategies for a Major-Minor Agent Optimal Execution Problem.
Social Science Research Network Working Paper Series.

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.

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., Guéant, O., and Razafinimanana, J. (2010).
High Frequency Simulations of an Order Book: a Two-Scales Approach.
In Abergel, F., Chakrabarti, B. K., Chakraborti, A., and Mitra, M., editors, *Econophysics of Order-Driven Markets*, New Economic Windows. Springer.

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

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.

Moro, E., Vicente, J., Moyano, L. G., Gerig, A., Farmer, D. J., Vaglica, G., Lillo, F., and Mantegna, R. N. (2009).
Market impact and trading profile of large trading orders in stock markets.

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.

Robert, C. Y. and Rosenbaum, M. (2011).
A New Approach for the Dynamics of Ultra-High-Frequency Data: The Model with Uncertainty Zones.
Journal of Financial Econometrics, 9(2):344–366.

Roşu, I. (2009).
A Dynamic Model of the Limit Order Book.
Review of Financial Studies, 22(11):4601–4641.

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

Smith, E., Farmer, D. J., Gillemot, L., and Krishnamurthy, S. (2003).
Statistical Theory of the Continuous Double Auction.
Quantitative Finance, 3(6):481–514.

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.

Wyart, M., Bouchaud, J.-P., Kockelkoren, J., Potters, M., and Vettorazzo, M. (2008).
Relation between Bid-Ask Spread, Impact and Volatility in Double Auction Markets.
Technical Report 1.