CFM
INSIGHT.DATA.CLARITY.

Machine Learning For Trading: Keeping Optimal Control Under Control

Charles-Albert Lehalle

Senior Research Advisor (Capital Fund Management, Paris); Visiting Researcher (Imperial College, London)

House of Finance Days, Université Paris Dauphine, April, 2018

Machine Learning For Trading

Recent successes of Machine Learning (ML, and AI) in signal processing (voice recognition, image classification, etc) and in sandboxes (Go game) pushe everyone to try it everywhere.

Financial markets are not spared from this overwhelming wave of attempts to apply Machine Learning to any possible data...
Of course we can think about **natural translations of existing applications**, like:

- Recommendations to manage relationships with customers (including robot-advisory);
- Nowcasting to harvest a lot of data describing the real economy to know as soon as possible its health.

Moreover, ML is known to be efficient when you **have access to a lot of data**. Intraday trading has **a very large amount of data** (one month of trading on one stock of the CAC 40 reads about 120,000 trades and 130 millions of orders). It is expected ML could find applications there.

The main specific aspects of trading: ① you automate the trading process; ② you will face a "**closed loop effect**" (i.e. via the market impact).

We review here some known results about intraday trading optimization and try to foresee some applications of machine learning in this area.

Machine Learning For Trading

Based On

This talk is based on few papers with co-authors:

① (Empirics + Theory) *Simulating and analyzing order book data: The queue-reactive model*, by W. Huang, C.-A. L and M. Rosenbaum [Huang et al., 2015]

② (Empirics + Theory) *Market Impacts and the Life Cycle of Investors Orders*, by E. Bacry, A. Iuga, M. Lasnier and C.-A. L [Bacry et al., 2015]

③ (Theory) *Limit Order Strategic Placement with Adverse Selection Risk and the Role of Latency*, by C.-A. L and O. Mounjid [L and Mounjid, 2016];

④ (Theory + empirics) ☞ *Optimal High Frequency Interactions with Orderbooks*, by O. Moundji, C.-A. L and M. Rosenbaum [L et al., 2018].

⑤ (Empirical) ☞ *The Behaviour of High-Frequency Traders Under Different Market Stress Scenarios*, by N. Megarbane, P. Saliba, C.-A. L and M. Rosenbaum [Megarbane et al., 2017];

⑥ (Theory + empirics) *Optimal split of orders across liquidity pools: a stochastic algorithm approach*, by G. Pagès, S. Laruelle and C.-A. L [Pagès et al., 2011]

Outline

Outline

What Do We Know About Intraday Dynamics?

Limit Orderbooks: The Queue Reactive Model

The **Queue Reactive Model** introduced by Weibing Huang during his PhD thesis [Huang et al., 2015] shows that

- The flows providing liquidity (i.e. limit orders) and consuming liquidity (i.e. cancel and market orders) and a queue of a limit orderbook can be modelled by Poisson processes
- There intensities are functions of the size of the considered queue and its nearest neighbours.

First Limit

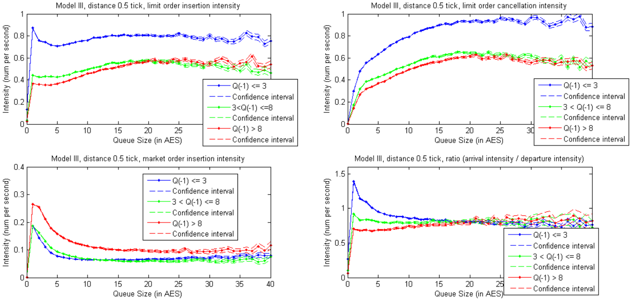

What Do We Know About Intraday Dynamics?

Limit Orderbooks: The Queue Reactive Model

The **Queue Reactive Model** introduced by Weibing Huang during his PhD thesis [Huang et al., 2015] shows that

- The flows providing liquidity (i.e. limit orders) and consuming liquidity (i.e. cancel and market orders) and a queue of a limit orderbook can be modelled by Poisson processes
- There intensities are functions of the size of the considered queue and its nearest neighbours.

Second Limit

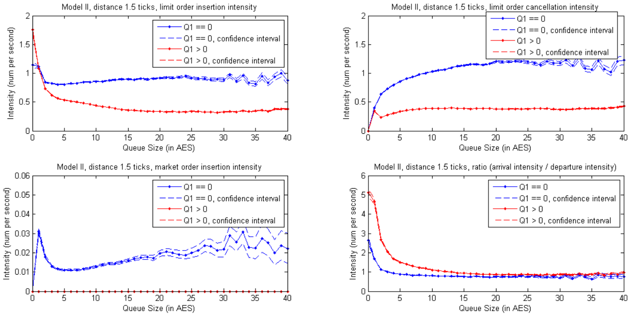

What Do We Know About Intraday Dynamics?

Limit Orderbooks

This means that:

- Given you know the state of the liquidity offer (i.e. size of queues in the book)
- You have a good estimate of the distribution of the sequence of next events.
- This is more than just a prediction of the price in S seconds, it is a model of the dynamics (it predicts the next step, and can be iterated).
- Can this be used to **pilot a limit order**?
- In other terms: can market participants looking at orderbook state be **more efficient in providing liquidity**?
- Could Machine Learning help them?

This is an on-going research program.

What Do We Know About Intraday Dynamics?

Market Impact

in [Moro et al., 2009]

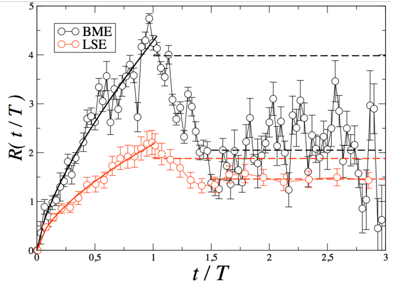

Market Impact takes place in different phases

- 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.

To be more than anecdotical, it is needed to make statistics, few papers document market impact at all scales: [Bershova and Rakhlin, 2013] (intraday impact), [Waelbroeck and Gomes, 2013] (daily impact of cash trades), [Brokmann et al., 2015] (daily impact of informed trades for a hedge fund), [Bacry et al., 2015] (intraday and daily impact of informed trades for a bank).

What Do We Know About Intraday Dynamics?

Market Impact

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

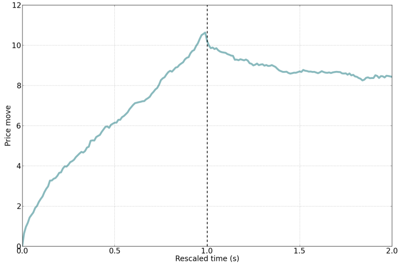

Market Impact takes place in different phases

- 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.

f

To be more than anecdotical, it is needed to make statistics, few papers document market impact at all scales: [Bershova and Rakhlin, 2013] (intraday impact), [Waelbroeck and Gomes, 2013] (daily impact of cash trades), [Brokmann et al., 2015] (daily impact of informed trades for a hedge fund), [Bacry et al., 2015] (intraday and daily impact of informed trades for a bank).

What Do We Know About Intraday Dynamics?

Market Impact : The Square Root Effect

Regression	Parameter	Coef. (log-log)	Coef. (L2)	Coef. (L1)
(R.1)				
	Daily participation	0.54	0.45	0.40
(R.2)				
	Daily participation	0.59	0.54	0.59
	Duration	−0.23	−0.35	−0.23
(R.3)				
	Daily participation	0.44	–	–
	Bid-ask spread	0.28	–	–
(R.4)				
	Daily participation	0.53	–	–
	Volatility	0.96	–	–
(R′.1)				
	Trading rate	0.43	0.33	0.43
(R′.2)				
	Trading rate	0.37	0.56	0.45
	Duration	0.15	0.24	0.23
(R′.3)				
	Trading rate	0.32	–	–
	Bid-ask spread	0.57	–	–
(R′.4)				
	Trading rate	0.32	–	–
	Volatility	0.88	–	–

Source: [Bacry et al., 2015]

The Formula 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:

- did you trading process reacted to market conditions?
- are you alone?
- etc.

We used different methods.

What Do We Know About Intraday Dynamics?

Market Impact : 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).

- The concave increase of the impact with time and its reversion can be explained using propagator models.
- But if you fit your "*price impact curves*" on data over a month, and you look at metaorders over the same month, the amplitude of the effects will not be compatible!

What Do We Know About Intraday Dynamics?

Market Impact : From High Frequency to Low Frequencies

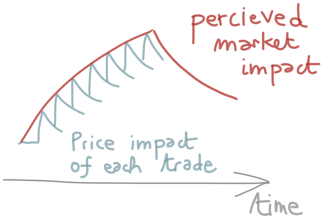

If someone trade at a given frequency $1/\delta t$ from 0, his price impact at $K\delta t$ will be (for an exponential kernel)

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

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).

- The concave increase of the impact with time and its reversion can be explained using propagator models.
- But if you fit your "*price impact curves*" on data over a month, and you look at metaorders over the same month, the amplitude of the effects will not be compatible!

What Do We Know About Intraday Dynamics?

Market Impact : From High Frequency to Low Frequencies

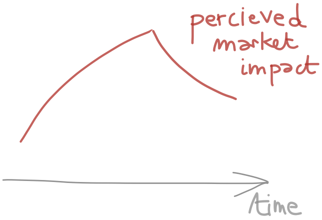

If someone trade at a given frequency $1/\delta t$ from 0, his price impact at $K\delta t$ will be (for an exponential kernel)

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

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).

- The concave increase of the impact with time and its reversion can be explained using propagator models.
- But if you fit your "*price impact curves*" on data over a month, and you look at metaorders over the same month, the amplitude of the effects will not be compatible!

What Do We Know About Intraday Dynamics?

Trading Practices Take All This Into Account

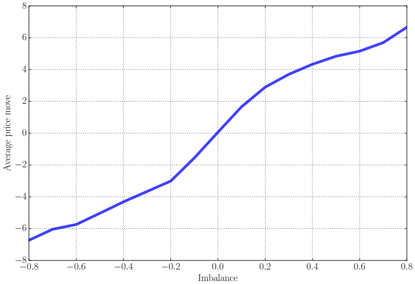

① The current imbalance predicts future price moves.

We just saw that the "market context" (i.e. expected news) could influence liquidity provision by market participants taking care of orderbooks (i.e. HFT).
① To see if they react to the state of the orderbook (and following the Queue Reactive Model), we can simply try to summarize the state of the book (i.e. queues sizes), by its **Imbalance**: $(Q^{ASK} - Q^{BID})/(Q^{ASK} + Q^{BID})$.

What Do We Know About Intraday Dynamics?

Trading Practices Take All This Into Account

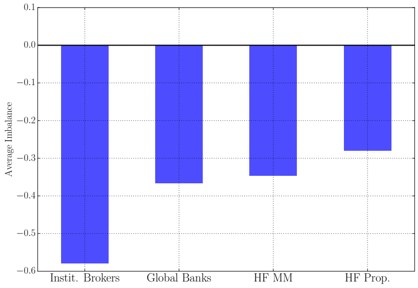

② The state of the imbalance given each type of participants traded with a limit order.

We just saw that the "market context" (i.e. expected news) could influence liquidity provision by market participants taking care of orderbooks (i.e. HFT).
① To see if they react to the state of the orderbook (and following the Queue Reactive Model), we can simply try to summarize the state of the book (i.e. queues sizes), by its **Imbalance**: $(Q^{ASK} - Q^{BID})/(Q^{ASK} + Q^{BID})$.

② We used a dataset of trades on NASDAQ-OMX (Nordic European Equity Markets), on which the identity of the buyer and a seller are know for each transaction, and synchronizing them to CFM's orderbook data. Thanks to this we can compute the average imbalance give each type of participant traded using a limit order.

What Do We Know About Intraday Dynamics?

Trading Practices Take All This Into Account

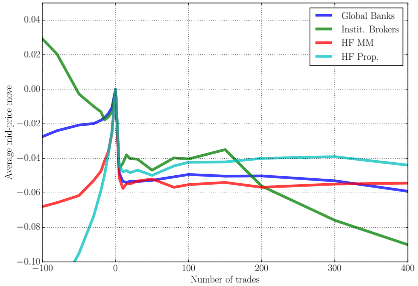

③ Price moves before and after a trade obtained via a limit order for each type of participant.

We just saw that the “market context” (i.e. expected news) could influence liquidity provision by market participants taking care of orderbooks (i.e. HFT).
① To see if they react to the state of the orderbook (and following the Queue Reactive Model), we can simply try to summarize the state of the book (i.e. queues sizes), by its **Imbalance**: $(Q^{ASK} - Q^{BID})/(Q^{ASK} + Q^{BID})$.

② We used a dataset of trades on NASDAQ-OMX (Nordic European Equity Markets), on which the identity of the buyer and a seller are know for each transaction, and synchronizing them to CFM's orderbook data. Thanks to this we can compute the average imbalance give each type of participant traded using a limit order.

③ **It is efficient**.

Outline

Some Theory: Optimal Control and Stochastic Algorithms

A Small Order: Optimal Trading Tactics Under Orderbook Dynamics

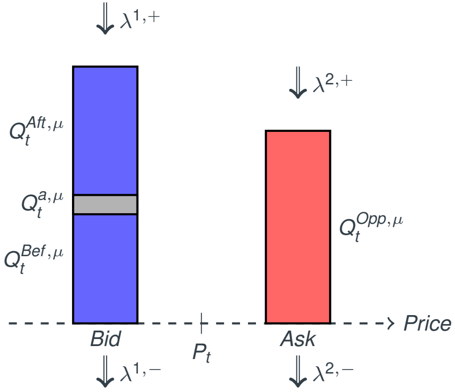

In [L et al., 2018] and [L and Mounjid, 2016], we design a procedure to control one limit order in an orderbook. Our model tracks the position of our limit order (of size Q^a) in the first queue. The flows adding and removing liquidity are similar to the ones of the QR Model (i.e. they are Poisson with intensities conditioned by the sizes of the queues).
The different transitions are:

- if no queue goes to zero, nothing special;
- if a queue goes to zero: a new queue is “discovered” on the same side and another queue is “inserted” on the opposite side. The sizes of these new queues are conditioned by the state of the orderbook.

Using the notation u for a state of the orderbook (including the controlled order), we can show that the process U_t is ergodic under reasonable conditions, and we can show the existence of a “price at infinity”:

$$g(u) = \mathbb{E}\left(P_\infty | \mathcal{F}_0, U_0 = u\right).$$

Some Theory: Optimal Control and Stochastic Algorithms

Orderbook Modelling For A Small Order: Comparing Empirics and Models

(a) Empirical Q^{Opp} after 20 events

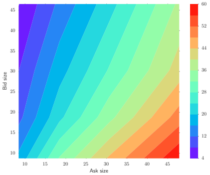

(b) Theoretical Q^{Opp} after 20 events

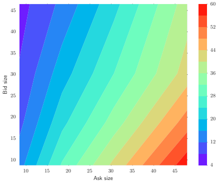

(c) Empirical Q^{Same} after 20 events

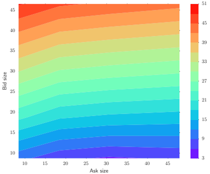

(d) Theoretical Q^{Same} after 20 events

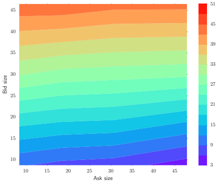

Some Theory: Optimal Control and Stochastic Algorithms

Definition of The Control Problem

The controls μ are taken from:

- Stay in the orderbook
- Cancel (and then reinsert at the top of the queue)
- Convert it in a market order.

You have two versions of the control problem: either the decision can be taken every Δ seconds, either it can be taken at any orderbook move.

Once the order is executed at time T^{μ}_{Exec} at price P, we value the strategy at

$$\sup_{\mu} \mathbb{E}\left[f \circ \mathbb{E}\left(P^{\mu}_{\infty} - P \,\middle|\, \mathcal{F}_{T^{\mu}_{Exec}} \right) - c\, q^{a}\, T^{\mu}_{Exec} \right].$$

Where c is a waiting cost, f can be any (Lipschitz) function, and $\mathbb{E}\left(P^{\mu}_{\infty} \,\middle|\, \mathcal{F}_t\right)$ is the price at infinity given the state of the orderbook at t

Some Theory: Optimal Control and Stochastic Algorithms

Definition of The Control Problem

The controls μ are taken from:

- Stay in the orderbook
- Cancel (and then reinsert at the top of the queue)
- Convert it in a market order.

You have two versions of the control problem: either the decision can be taken every Δ seconds, either it can be taken at any orderbook move.

Once the order is executed at time T^{μ}_{Exec} at price P, we value the strategy at

$$\sup_{\mu} \mathbb{E}\left[f \circ \mathbb{E}\left(P^{\mu}_{\infty} - P \,\middle|\, \mathcal{F}_{T^{\mu}_{Exec}} \right) - c\, q^a T^{\mu}_{Exec} \right].$$

Where c is a waiting cost, f can be any (Lipschitz) function, and $\mathbb{E}\left(P^{\mu}_{\infty} \,\middle|\, \mathcal{F}_t\right)$ is the price at infinity given the state of the orderbook at t

Dynamic Programming Equation (for the continuous time version)

Let $u = (q^{bef}, q^a, q^{aft}, q^{opp}, p, p^{exec})$ an initial state . The value function $V(t,u)$ satisfies:

$$\max \begin{pmatrix} g(.) - V(t,.) \\ \mathcal{A}V(t,.) - cq^a\mathbf{1} \\ V^{c-l}(t,.) - V(t,.) - cq^a\mathbf{1} \end{pmatrix} = 0, \text{ when } q^a > 0. \quad (1)$$

And $V(t,u) = u$ at execution and $V(T,u) = g(u)$ at T.

Using the DPP

We show how to make the numerics to solve (1), and we obtain results like

Estimate of the cost of latency

Let $V_T(0,u;\Delta_1)$ the optimal fast agent gain and $V_T(0,u;\Delta_2)$ the optimal slow agent gain.

$$|V_T(0,u;\Delta_1) - V_T(0,u;\Delta_2)| \le H_1 \left\lceil \frac{T}{\Delta_2} \right\rceil \left\lceil \frac{\Delta_2}{\Delta_1} \right\rceil e^{C_3 T} + H_2 \Delta_2 T,$$

where H_1, H_2 and C_3 are constants involving parameters of the problem.

We fit the model on data and we solve it numerically providing different qualitative results.

With the parameters: $\Delta = 1$ second, $T = 10\Delta$ $Q^{Disc} = 22$, $Q^{Ins} = 3$, $q = 1$, $c = 0$, and the tick is0.01.

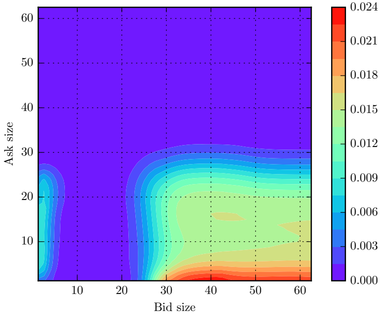

Difference between the value of a “join the bid” strategy and the value of the optimal one.

Using the DPP

We show how to make the numerics to solve (1), and we obtain results like

Estimate of the cost of latency

Let $V_T(0,u;\Delta_1)$ the optimal fast agent gain and $V_T(0,u;\Delta_2)$ the optimal slow agent gain.

$$|V_T(0,u;\Delta_1) - V_T(0,u;\Delta_2)| \leq H_1 \left\lceil \frac{T}{\Delta_2} \right\rceil \left\lceil \frac{\Delta_2}{\Delta_1} \right\rceil e^{C_3 T} + H_2 \Delta_2 T,$$

where H_1, H_2 and C_3 are constants involving parameters of the problem.

We fit the model on data and we solve it numerically providing different qualitative results.

With the parameters: $\Delta = 1$ second, $T = 10\Delta$, $\lambda^{Same,+} = \lambda^{Opp,+} = 0.06$, $\lambda^{Same,-} = \lambda^{Opp,-} = 0.12$, $Q^{Disc} = 5$, $Q^{Ins} = 2$, $q = 1$, $c = 0.0085$ and the tick is 0.01. Moreover $Q^{bef}(0) = 1$.

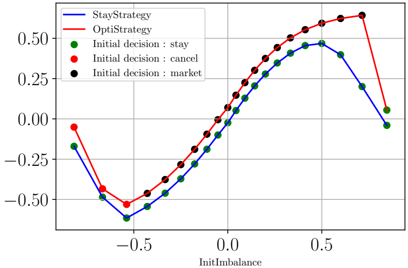

An extreme simulation to compare the "join the bid" strategy and the optimal one.

Some Theory: Optimal Control and Stochastic Algorithms

Stochastic Algorithms 101

- 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):

$$x_{n+1} = x_n + \gamma_{n+1}\, h(x_n) \tag{2}$$

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

$$X_{n+1} = X_n + \gamma_{n+1}\, H(X_n, \xi_{n+1}) \tag{3}$$

- 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]

Some Theory: Optimal Control and Stochastic Algorithms

Stochastic Algorithms: Machine Learning In Action

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:

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

Outline

Examples Of Statistical Learning In Trading

Exploration-Exploitation Of Dark-Pools: What is Dark Routing?

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

Examples Of Statistical Learning In Trading

Exploration-Exploitation Of Dark-Pools: What is Dark Routing?

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 $\mathtt{min}(V, D_i)$ back
- they allow "pegged" orders: you can specify δS rather than a limit price ("pinging" implies $\Delta T = 0$):

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

Examples Of Statistical Learning In Trading

Exploration-Exploitation Of Dark-Pools

- 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 ($\sum_{i=1}^{N} r_i = 1$)

Examples Of Statistical Learning In Trading

Dark Routing: Cost of the executed order (Details in [Pagès et al., 2011])

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

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

Examples Of Statistical Learning In Trading

Dark Pools: Mean Execution Cost

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

Maximization problem to solve

$$\max \left\{ \sum_{i=1}^{N} \rho_i \mathbb{E}\left(S \min\left(r_i V, D_i\right)\right), r \in \mathcal{P}_N \right\} \tag{4}$$

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 $\widetilde{D}_i := D_i S$ 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 $\mathbb{P}(D > 0) > 0$

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

Characterization of the solution

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

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

Characterization of the solution

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

Consequently, this leads to the following recursive zero search procedure

$$r_i^{n+1} = r_i^n + \gamma_{n+1} H_i(r^n, Y^{n+1}), \; r^0 \in \mathcal{P}_N, \; i \in \mathcal{I}_N, \tag{5}$$

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

The Stochastic Algorithm Version

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

Examples Of Statistical Learning In Trading

Dark Routing: 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**.

Examples Of Statistical Learning In Trading

Dark Routing: 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**.

The underlying idea of the algorithm

Do 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 punish the underperforming dark pools).

Examples Of Statistical Learning In Trading

Dark Routing: Stochastic Algorithms Provide Convergence Results

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

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

Outline

Good Practices, From [Megarbane et al., 2017]

Learning From HFT Around News

The data and some descriptive statistics.

The database is provided by the French regulator (AMF), all orders (and transactions) are labelled by the name of the owner, which allows us to identify HFTs. It covers the trades and orders on the most liquid French securities (36 of the CAC 40 stock), from November 2015 to July 2016 (approximately 40 millions of trades and 1.2 billions of orders to be processed).

❶ HFT are the main liquidity providers in the LOB

Presence in the LOB	Market share in (market depth)
At the best bid and offer	70.8 %
At the two best prices	77.3 %
At the three best prices	79.3 %

❷ And they are very diverse

	A/P ratio below 50%	A/P ratio over 50%
Part in nbe	60%	40%
Part in amount	45%	55%
Avg ratio (std)	25% (18%)	67% (10%)

❸ Everyone trades with everyone

Cons./Prov.	HFTs	non-HFTs	
HFTs	33.6%	31.2%	64.8%
non-HFTs	22.4%	12.8%	35.2%
	56.0%	44.0%	

But HFT are not providing that much liquidity to trades

Good Practices

Usual Intraday Behaviour of HFT

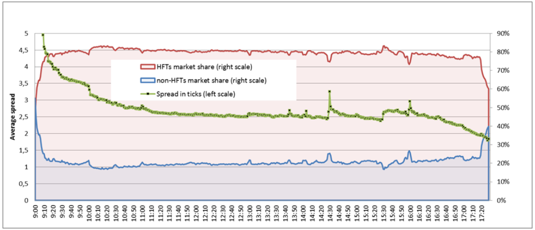

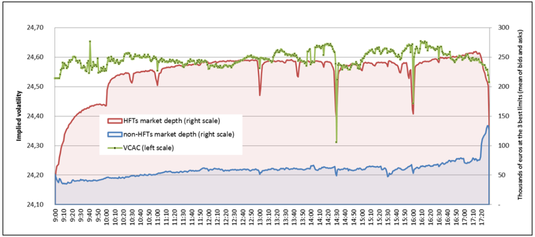

- TOP: pct of presence in the first 3 limits and the bid-ask spread,
- BOTTOM: amount in Euro on the first 3 limits and the implicit volatility.
- You can notice the macro news announcements (2:30pm and 4:00pm)

Good Practices

Does This Average Behaviour Changes When There Are News (1/2)

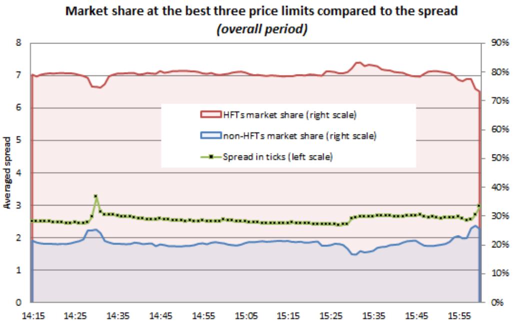

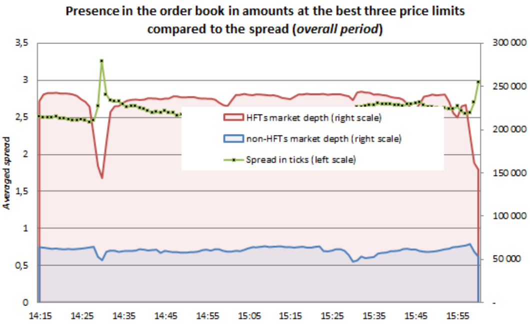

We selected the **10 most impacting News** around 2:30pm.

Left: market share (ie pct), Right: Size of the limit orders (in Euros).
The charts are different: first there is a scaling, second **the liquidity (in Euros) provided by HFT does not come back** after impacting news.

Good Practices

Does This Average Behaviour Changes When There Are News (2/2)

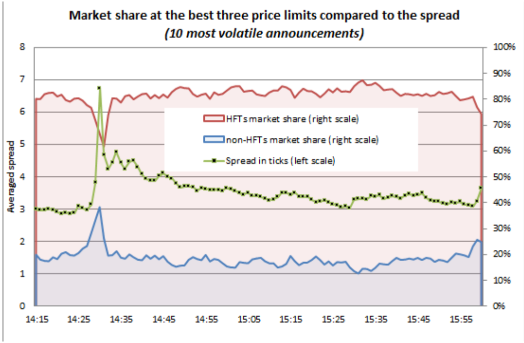

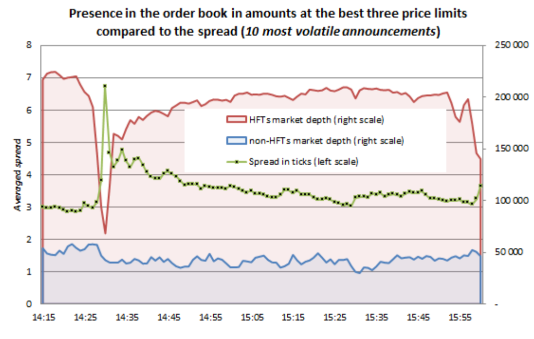

We selected the **10 most impacting News** around 2:30pm.
Left: market share (ie pct), Right: Size of the limit orders (in Euros).
The charts are different: first there is a scaling, second **the liquidity (in Euros) provided by HFT does not come back** after impacting news.

Good Practices

Presence in The Book Around 4:00p.m. Announcements

- We only consider news related to the U.S economy (Bloomberg news): 140 days with announcements, vs. 51 without announcements. Data are restricted between 3:40pm and 4:50pm and we consider 1min bins.
- We create 3 dummy variables: *B* (for *Before*), *D* (for *During*) and *A* (for *After*) 4:00pm.
- The empirical volatility during each 1min bin is renormalized by the avg volatility of the day.
- **Methodology:** Do a model using days without announcements only, work on the residuals of this model and try to explain these residuals on announcement days.

Explaining the pct of HFT liquidity in the book

Variable	Coef.	Std. err.	t	$P > \|t\|$	95% Conf. Int.
Const.	0.7866	0.003	302.564	0	[0.781, 0.792]
Const.	0.011	0.002	6.302	0	[0.008, 0.015]
σ_{norm}	−0.0045	0.002	−2.664	0.008	[-0.008, -0.001]
B	−0.0520	0.004	−14.404	0	[-0.059, -0.045]
D	−0.1507	0.005	−28.941	0	[-0.161, -0.141]
A	−0.0283	0.004	−7.797	0	[-0.035, -0.021]

Good Practices

HFT Agressive/Passive Ratio Around 4:00p.m. Announcements

- We only consider news related to the U.S economy (Bloomberg news): 140 days with announcements, vs. 51 without announcements. Data are restricted between 3:40pm and 4:50pm and we consider 1min bins.
- We create 3 dummy variables: *B* (for *Before*), *D* (for *During*) and *A* (for *After*) 4:00pm.
- The empirical volatility during each 1min bin is renormalized by the avg volatility of the day.
- **Methodology:** Do a model using days without announcements only, work on the residuals of this model and try to explain these residuals on announcement days.

Explaining HFT Agressive/Passive Ratio

Variables	Coef.	Std. err.	t	$P > \|t\|$	95% Conf. Int.
Const.	0.5340	0.002	228.198	0	[0.529, 0.539]
σ_{norm}	0.0111	0.002	5.023	0	[0.007, 0.015]
D	0.0169	0.007	2.494	0.013	[0.004, 0.03]
Const.	0.0113	0.001	9.029	0	[0.009, 0.014]
σ_{norm}	-0.0053	0.001	-4.475	0	[-0.008, -0.003]
B	0.0184	0.003	7.116	0	[0.013, 0.023]
D	0.0268	0.004	7.237	0	[0.02, 0.034]

Good Practices

HFT Market Share on Trades Around 4:00p.m. Announcements

- We only consider news related to the U.S economy (Bloomberg news): 140 days with announcements, vs. 51 without announcements. Data are restricted between 3:40pm and 4:50pm and we consider 1min bins.
- We create 3 dummy variables: *B* (for *Before*), *D* (for *During*) and *A* (for *After*) 4:00pm.
- The empirical volatility during each 1min bin is renormalized by the avg volatility of the day.
- **Methodology:** Do a model using days without announcements only, work on the residuals of this model and try to explain these residuals on announcement days.

Explaining HFT market share on trades

Variables	Coef.	Std. err.	t	$P > \|t\|$	95% Conf. Int.
Const.	0.5557	0.003	208.543	0	[0.551, 0.561]
σ_{norm}	0.0473	0.003	18.740	0	[0.042, 0.052]
Const.	0.0097	0.001	16.799	0	[0.009, 0.011]
B	−0.0346	0.003	−10.228	0	[-0.041, -0.028]
D	−0.0469	0.005	−9.869	0	[-0.057, -0.038]

Good Practices

Summary of HFT Behaviour Around News

All these regressions point out in a quantitative way that **the behaviour of HFTs around announcements** cannot be read as a simple reaction to associated variations of volatility.
Around a scheduled announcement, on top of usual reactions to volatility, HFTs:

- provide 15% less liquidity,
- are slightly more aggressive,
- trade less.

On the contrary, **when no announcement is planned**, their attitude towards an increase of volatility goes in the opposite direction (trading more). We thus identify a “change of regime" in the presence of scheduled news.

For the purpose of this talk: It seems that HFT take exogenous information into account to avoid to focus too much on high frequency data.

Good Practices

Mixing Learning and Control: More References

Other papers to have ideas on ML for trading:

- **Mean Field Game theory** [Cardaliaguet and L, 2016]: to take into account the market impact of other traders: If you know you are not the only one to trade, you can improve your strategy.
  The difficult point is to **learn what others are doing**. It can be shown that in this case such a learning is efficient (in the sense you learn the optimal strategy, as if you knew others' inventories).
- **Using a directly a price signal** [L and Neuman, 2017]: if you have not a model, but a signal, it is possible (but difficult) to include it your optimal strategy.
- **Automated monitoring** of hundreds of trading algorithms [Azencott et al., 2014], using real-time (adaptive) modelling of performances, you can identify the potential causes of bad functioning of algos and provide online decision support to traders.

Good Practices

Mixing Learning and Control (1/2)

Predictive Models

- You can use ML to model future price changes or the future state of the offer and demand of liquidity,
- In such a case you will have an approximation error ϵ,
- What is the influence of not using ϵ in closed loop control? By construction $\mathbb{E}(\epsilon) = 0$ but if your cost function is not linear in the predictor, you may have some problems...
- Moreover how can you **model interactions with others**?

Direct Control

- If you automate a system, you end up with using the dynamic programming principle:

$$V_t(x) = \min_c \sum_{X(t+1)} \mathbb{P}(x \xrightarrow{c} X(t+1)) \cdot \left\{ V_{t+1}(X(t+1)) - \mathrm{cost}(x \xrightarrow{c} X(t+1)) \right\}.$$

  This means you will implement a backward reasoning (at $t = 0$ you will use all your prediction from $t = T$ to now). As a consequence a small modelling error will have a huge influence on your strategy: do you want to drive few hours of trading with few seconds ahead predictions?

Good Practices

Mixing Learning and Control (2/2)

"Robust" Control

- It is probably better to implement two layers of control [Bouchard et al., 2011]

② One slow "**risk control**" layer to bound your liquidity and market risks, it can include few **slowly evolving meta-parameters**, or **fast reaction** to exogenous stimuli (like HFT reaction to news).

③ Inside such bounds, you can take fast decision exploiting stationarity of liquidity dynamics thanks to **statistical learning**.

Other techniques

- **Monitoring** is very important,
- If you are a **bilateral market maker**, you can use **recommandation-like methods** to guess what kind of client to accept, reject or call.
- In any case, more tasks are automated each month on trading floor, one of the biggest challenge is to **establish an efficient human-machine-interface**: when to give back the control to humans? what information to give them so that they can take enlightened decisions?

Thank You For Your Attention

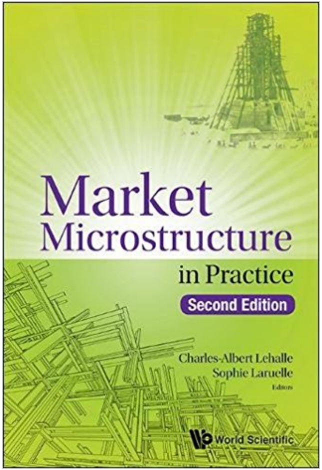

To submit papers:
Market Microstructure and Liquidity .

More on market microstructure:
Market Microstructure in Practice by C.-A. L and Sophie Laruelle (World Scientific Publisher, 1s ed. 2013, 2nd ed. 2018).

References I

Azencott, R., Beri, A., Gadhyan, Y., Joseph, N., Lehalle, C.-A., and Rowley, M. (2014).
Realtime market microstructure analysis: online Transaction Cost Analysis.
Quantitative Finance, pages 0–19.

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).

Bershova, N. and Rakhlin, D. (2013).
The Non-Linear Market Impact of Large Trades: Evidence from Buy-Side Order Flow.
Social Science Research Network Working Paper Series.

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. (2015).
Slow decay of impact in equity markets.
Market Microstructure and Liquidity, 1(2).

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

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

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

References II

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

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.

Lehalle, C.-A. and Mounjid, O. (2016).
Limit Order Strategic Placement with Adverse Selection Risk and the Role of Latency.

Lehalle, C.-A., Mounjid, O., and Rosenbaum, M. (2018).
Optimal liquidity-based trading tactics.

Lehalle, C.-A. and Neuman, E. (2017).
Incorporating Signals into Optimal Trading.

Megarbane, N., Saliba, P., Lehalle, C.-A., and Rosenbaum, M. (2017).
The Behaviour of High-Frequency Traders Under Different Market Stress Scenarios.
Technical report, SSRN.

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.

References III

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.

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.