Mathematical Models to Study and Control the Price Formation Process

Charles-Albert Lehalle

U.F.R. DE MATHÉMATIQUES
UNIVERSITÉ PIERRE-ET-MARIE-CURIE (PARIS 6)

MÉMOIRE DE SYNTHÈSE

en vue de l'obtention de l'

HABILITATION À DIRIGER LES RECHERCHES

SPÉCIALITÉ : Mathématiques Appliquées

présenté et soutenu publiquement par

Charles-Albert LEHALLE

le 18 Décembre 2015

Titre :

MODÈLES MATHÉMATIQUES POUR L'ETUDE ET LE CONTRÔLE DU PROCESSUS DE FORMATION DES PRIX

Jury :

RAPPORTEURS

M. Robert	ALMGREN	New York University
M. Alvaro	CARTEA	University of Oxford
M. Nicolas	VAYATIS	ENS Cachan

EXAMINATEURS

M. Frédéric	ABERGEL	Ecole Centrale Supélec, Paris
M. Robert	AZENCOTT	Houston University
M. Bruno	BOUCHARD	Université Paris Dauphine
M. Nicole	EL KAROUI	Université Paris VI
M. Gilles	PAGÈS	Université Paris VI
M. Mathieu	ROSENBAUM	Université Paris VI

Modèles Mathématiques pour l'Etude et le Contrôle du Processus de Formation des Prix

Charles-Albert Lehalle

Imprimé le 18 décembre 2015

Table des matières

Introduction

Le processus de formation des prix est au coeur de tout modèles de mathématiques financières. D'abord approximé par un mouvement brownien (cf. [Karatzas & Shreve, 1998]), puis en y intégrant des sauts (voir [Shiryaev, 1999]) tant qu'il s'agissait d'échelles de temps longues, la prise en compte des volumes échangés est arrivée par la suite :

- Dans un cadre économétrique (voir par exemple [Tauchen & Pitts, 1983]) pour tenter de mieux expliquer la dynamique des prix
- dans le cadre d'études théoriques d'équilibres généraux où le "jeu d'enchère" est modélisé (comme dans [Kyle, 1985], [Ho & Stoll, 1983] ou [Glosten & Milgrom, 1985]).

La régulation poussant de plus en plus d'échanges vers des marchés électroniques, de très grosses bases de données sont aujourd'hui disponibles. Elles contiennent non seulement les transactions effectives, mais aussi les déclarations d'intérêt de tous les participants. Ceci à ouvert la porte à des études empiriques (comme [Lillo *et al.* , 2003]) qui fournissent des pistes intéressantes à de nouvelles familles de modèles (voir par exemple [Bacry & Muzy, 2013]).

Bien modéliser le processus de formation des prix permet de guider les régulateurs, qui tentent de favoriser les échanges sur des marchés électroniques (car il sont plus facilement traçables). Il permet aussi de mettre au point des techniques de *trading optimal*, qui, utilisées par les investisseurs et intermédiaires financiers, vont minimiser la perturbation des prix due à l'intensité des échanges. Nous verrons Section 2 que cette dégradation naturelle des prix (que l'on appelle communément le *market impact*) peut être modélisée, et la Section 1 montrera comment utiliser ces modèles pour choisir un *profil de trading* qui minimise cette dégradation, ce qui profite simultanément à l'utilisateur du profil et à l'efficience et la valeur des prix. L'apport des mathématiques appliquées à cette direction de recherche a par exemple été discutée lors d'un symposium de l'ESAIM à Seignosse en 2013 [Hoffmann *et al.* , 2013].

Les travaux présentés ici couvrent l'essentiel des thèmes abordés ces quinze dernières années par les mathématiques appliquées autour du processus de formation des prix :

- le trading optimal, qui considère un acteur qui doit acheter ou vendre une large quantité d'actions (cf. [Gökay *et al.* , 2011] pour un survey)
- la microstructure, qui se focalise sur le processus de formation des prix (i.e. la dynamique qui émerge de la confrontation des intérêts acheteurs et vendeurs [Lehalle *et al.* , 2013])
- l'apprentissage en ligne de paramètres sensés capturer des caractéristiques fugaces de la liquidité afin de les prendre en compte dans un processus de trading, ou de surveillance.

Ce mémoire se compose donc de trois parties, chacune détaillant un aspect de ces trois grands thèmes.

Chaque partie est composée d'une introduction en français, qui expose la problématique qui va être traitée en termes simples. Elle est suivie à chaque fois de sections en anglais, qui reprennent les éléments essentiels d'articles publiés. Ces sections ont en préambule un court propos introductif en anglais, qui positionne la problématique du ou des articles en question.

Liste des travaux

Voici la liste de mes travaux en lien avec les différentes parties de ce mémoire.

Trading optimal.

1. Bruno Bouchard, Ngoc-Minh Dang, and Charles-Albert Lehalle. Optimal control of trading algorithms : a general impulse control approach. *SIAM J. Financial Mathematics*, 2(1) :404–438, 2011.
2. O. Guéant, Charles-Albert Lehalle, and J. Fernandez-Tapia. Optimal Portfolio Liquidation with Limit Orders. *SIAM Journal on Financial Mathematics*, 13(1) :740–764, 2012.
3. Olivier Guéant and Charles-Albert Lehalle. General intensity shapes in optimal liquidation. *Mathematical Finance*, page n/a, October 2013.
4. Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Dealing with the inventory risk : a solution to the market making problem. *Mathematics and Financial Economics*, 4(7) :477–507, September 2013.
5. Mauricio Labadie and Charles-Albert Lehalle. Optimal starting times, stopping times and risk measures for algorithmic trading. *The Journal of Investment Strategies*, 3(2), March 2014.
6. Charles-Albert Lehalle. Market Microstructure knowledge needed to control an intra-day trading process. In Jean-Pierre Fouque and Joseph Langsam, editors, *Handbook on Systemic Risk*. Cambridge University Press, May 2013.
7. Marc Hoffmann, Mauricio Labadie, Charles-Albert Lehalle, Gilles Pages, Huyen Pham, and Mathieu Rosenbaum. Optimization and Statistical Methods for High Frequency Finance. In *ESAIM Proceedings*, Seignosse, France, 2013. SMAI.

Etude de la microstructure.

8. Frédéric Abergel, Jean-Philippe Bouchaud, Thierry Foucault, Charles-Albert Lehalle, and Mathieu Rosenbaum, editors. *Market Microstructure Confronting Many Viewpoints*. Wiley, 2012.
9. Emmanuel Bacry, Adrian Iuga, Matthieu Lasnier, and Charles-Albert Lehalle. Market Impacts and the Life Cycle of Investors Orders. *Social Science Research Network Working Paper Series*, December 2014.
10. Paul Besson and Charles-Albert Lehalle. The Deal/Book Split Analysis : A New Method to Disentangle the Contribution to Market and Limit Orders in Any Price Change. *Social Science Research Network Working Paper Series*, January 2014.
11. Weibing Huang, Charles-Albert Lehalle, and Mathieu Rosenbaum. Simulating and analyzing order book data : The queue-reactive model, December 2013.
12. Aimé Lachapelle, Jean-Michel Lasry, Charles-Albert Lehalle, and Pierre-Louis Lions. Efficiency of the Price Formation Process in Presence of High Frequency Participants : a Mean Field Game analysis, May 2013.
13. Charles-Albert Lehalle, Matthieu Lasnier, Paul Besson, Hamza Harti, Weibing Huang, Nicolas Joseph, and Lionel Massoulard. 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, August 2012.
14. Charles-Albert Lehalle, Olivier Guéant, and Julien Razafinimanana. High Frequency Simulations of an Order Book : a Two-Scales Approach. In F. Abergel, B. K. Chakrabarti, A. Chakraborti, and M. Mitra, editors, *Econophysics of Order-Driven Markets*, New Economic Windows. Springer, 2010.
15. Charles-Albert Lehalle. Le trading quantitatif : Une rationalisation de la négociation sur les marchés. *Les Cahiers de l'ILB*, 3, July 2011.

Algorithmes stochastiques pour le trading optimal.

16. Gilles Pagès, Sophie Laruelle, and Charles-Albert Lehalle. Optimal split of orders across liquidity pools : a stochastic algorithm approach. *SIAM Journal on Financial Mathematics*, 2 :1042–1076, 2011.
17. Sophie Laruelle, Charles-Albert Lehalle, and Gilles Pagès. Optimal posting price of limit orders : learning by trading. *Mathematics and Financial Economics*, 7(3) :359–403, June 2013.
18. Robert Azencott, Arjun Beri, Yutheeka Gadhyan, Nicolas Joseph, Charles-Albert Lehalle, and Matthew Rowley. Realtime market microstructure analysis : online Transaction Cost Analysis. *Quantitative Finance*, pages 0–19, March 2014.

1 Trading Optimal

Le "*Trading Optimal*" est un champ qui a débuté avec la publication de [Bertsimas & Lo, 1998] et [Almgren & Chriss, 2000]. Il s'agissait à l'époque de développer un point de vue sur l'achat ou la vente de grande quantité d'actions en peu de temps, de façon analogue à ce que Markowitz avait fait pour les portefeuilles. Markowitz [Markowitz, 1952] avait mis en place un cadre mathématique basé sur une optimisation moyenne-variance où la moyenne est tirée par les rendements attendus des actifs du portefeuille et la variance provient d'une estimation historique de leurs corrélations.

Dans le trading optimal, le rendement est remplacé par le "*Market Impact*" (dégradation des rendement due à la taille de l'achat ou la vente –cf. Section 2.3–) et les corrélations entre les composantes d'un portefeuille par l'autocorrélation du prix au cours de la journée. En effet la bonne analogie avec Markowitz consiste à considérer une allocation non pas au sein des composantes d'un portefeuille, mais tout au long de la journée. De façon un peu approximative pour l'instant, on peut écrire :

$$\underbrace{\left|\begin{array}{ll}\max & \mathbb{E}\left(\sum_k w_k r_k\right) - \lambda \mathbb{V}\left(\sum_k w_k r_k\right) \\ w_1, \dots, w_K & \\ \sum_k w_k = 1 & \end{array}\right.}_{\text{Portfolio Allocation}} \longrightarrow \underbrace{\left|\begin{array}{ll}\max & -\mathbb{E}\left(\sum_t \rho_t (S_t + \Delta S_t(\rho_t Q))\right) \\ \rho_1, \dots, \rho_T & -\lambda \mathbb{V}\left(\sum_t (S_t + \Delta S_t(\rho_t Q))\right) \\ \sum_t \rho_t = 1 & \end{array}\right.}_{\text{Mean-Variance Optimal Trading}},$$

où on noterait

- pour un portefeuille : r_k les rendements de la $k$ème ligne du portefeuille et w_k la proportion de l'investissement dans cette ligne ;
- pour le trading (i.e. l'achat de Q actions d'un instrument donné) : S_t le prix en t, $\Delta S(q)$ le market impact (ici additif) provoqué par un achat de q actions, et ρ_t la fraction de l'achat effectuée en t.

Ce *framework*, initialement proposé par Almgren et Chriss (là où Bertsimas et Lo se focalisaient surtout sur le market impact et l'aspect multi-lignes), est encore aujourd'hui utilisé par de nombreuses institutions financières : banques d'investissement, fonds, ou courtiers.

Extension du trading en moyenne-variance. Le recours encore très vivace à ce cadre de modélisation par l'industrie justifie les propositions universitaires qui le raffinent. Comme on va le voir dans la Section 1.1.1, il est en effet possible d'éclairer l'effet d'une minimisation moyenne-variance sur un portefeuille. Il s'agit alors de comprendre comment la combinaison du market impact et des correlations permet de trouver des *trajectoires de trading* (i.e. la distribution des $(\rho_t)_t$, qui trace un profil temporel, cf. Figure 5) jointes moins coûteuses que les rythmes individuels optimaux.

Par ailleurs, nous verrons en Section 2.3 que le market impact (noté un peu plus haut symboliquement $\Delta S(q)$) est une fonction de la volatilité intraday, du volume habituellement échangé sur le titre, voire de la fourchette offre-demande (i.e. *Bid-Ask spread* en anglais). Ces grandeurs ne sont jamais connues à l'avance, il s'agit donc a minima de variables aléatoires (obtenues grâce à des estimations). Il est donc tout à fait à propos de regarder (cf. Section 1.1.2) ce qui se passe lorsqu'on les considère comme telles dans l'expression de l'espérance et de la variance.

Les utilisateurs du trading optimal en moyenne-variance ont en outre parfois besoin de prendre en compte certains a priori sur la dynamique des prix. Cela peut par exemple se produire s'ils implémentent une stratégie de *retour à la moyenne* (typiquement utilisée en arbitrage statistique). Ou bien s'ils pensent que la dynamique des prix à court terme est sous diffusive (ce vers quoi certaines études empiriques tendent, cf. [Huang *et al.* , 2013] et sa bibliographie). La Section 1.1.3 ouvre des pistes dans ces directions.

En dernier lieu, les praticiens mettent souvent des bornes à la vitesse de leur trading. En effet, même s'ils utilisent des modèles de market impact qui rendent compte de la relation entre l'intensité de trading (i.e. la vitesse fois la quantité ; $\rho_T Q$ dans l'écriture précédente) et la dégradation du prix, ils préfèrent, dans le cadre d'un contrôle de risque rigoureux, s'astreindre à ne pas participer à plus de $q\%$ des volumes échangés (par exemple toutes les 5 minutes glissantes). Au lieu de borner les trajectoires de trading optimales par qV_t (où V_t seraient les volumes échangés de $t - 5\text{min}$ à t), les résultats de la Section 1.1.4 leur permettent d'obtenir des trajectoires optimales sous contrainte.

Recours au Contrôle Optimal Stochastique. Lorsqu'il a s'agit d'étendre le cadre moyenne-variance, les choix furent nombreux. Un peu comme si on voulait étendre la régression linéaire : on peut uniquement modifier la métrique (et récemment aboutir à des régressions Ridge ou LASSO [Wang *et al.* , 2007]), ou bien se concentrer sur l'orthogonalisation des variables explicatives (et par exemple proposer la régression *Partial Least Square*), ou même travailler sur des combinaisons non linéaires de ces variables (ce qui mène aux régressions du type *Support Vector Machine* [Vapnik, 1998]), etc.

En ce qui concerne le trading optimal, les chemins possibles sont :

- changer le critère à minimiser de telle sorte qu'il soit nécessaire de recourir au principe de programmation dynamique pour obtenir des résultats,
- ne pas se restreindre à un découpage du temps uniforme, mais le soumettre lui aussi à l'optimisation,
- modéliser plus finement le processus de trading (par exemple en prenant en compte les interactions avec le carnet d'ordres).

Comme la Section 1.1.3 le montre, il est possible de prendre en compte d'autres versions réalistes du critère en restant dans un cadre proche d'une moyenne-variance. En revanche la question de la discrétisation temporelle se pose clairement : rythmer le trading de tous les instruments que l'on achète ou vend par multiples de 5, 10 ou 15 minute n'a a priori pas de sens. Il faudrait en effet a minima s'adapter à la liquidité de chaque instrument. Qui plus est, on peut se demander s'il est opportun d'acheter ou vendre *continument* tout le long de la période de trading. Pourquoi ne pas faire de pause ?

On peut écrire encore un peu approximativement, en restant dans le cadre d'un contrôle à temps discret :

$$\underbrace{\left|\begin{array}{ll} \max & -\mathbb{E}\left(\sum_t \rho_t(S_t+\Delta S_t(\rho_t Q))\right) \\ \rho_1,\dots,\rho_T & -\lambda\mathbb{V}\left(\sum_t(S_t+\Delta S_t(\rho_t Q))\right) \\ \sum_t \rho_t = 1 & \end{array}\right.}_{\text{Mean-Variance Optimal Trading}} \longrightarrow \underbrace{\left|\begin{array}{lcl} \mathsf{EV}(t,x) & = & \max_{\rho(t,x)} \mathbb{E}\big(V(t,X_t=x,\rho(t,x)) \\ & & \qquad +\mathsf{EV}(t+1,X_{t+1}|x,\rho(t,x))\big) \\ \mathsf{EV}(T,x) & = & G(x) \end{array}\right.}_{\text{Stochastic Optimal Trading}},$$

où dans l'expression de droite, X est la variable d'état, ρ reste le contrôle, et $V(t,x,\rho)$ est une fonction de coût (qui doit inclure le market impact). G est un coût terminal en T, souvent le coût de liquidation de la position restante. La dynamique de X n'est pas bien spécifiée puisque on a simplement écrit $X_{t+1}|x,\rho$ pour "*la variable aléatoire* X_{t+1} *sachant que* $X_t = x$ *et que le contrôle* ρ *est appliqué entre* t *et* $t+1$".

On voit immédiatement que le premier intérêt de ce *framework* est d'introduire une dépendence en x du contrôle (i.e. en remplaçant ρ_t à gauche par $\rho(t,x)$ à droite).

Le modèle de la Section 1.2 propose un cadre pour le trading optimal avec pause. Pour se faire elle considère un cadre de contrôle impulsionnel qui consiste à déclencher à différents temps d'arrêt $\tau_1,\dots,\tau_N$ des *robots de trading*. Le $i$ème robot va acheter pendant une durée δ_i une quantité ε_i. Ces couples de paramètres sont eux soumis à l'optimisation. Chaque robot va provoquer un *market impact* qui va affecter son efficacité et celle des suivants (conformément à une dynamique markovienne arbitraire). Chacun de ces robots peut par ailleurs être modélisé de façon indépendante, par example conformément à ce que la Section 3.1 propose.

Cela étant fait, la puissance de la programmation dynamique peut être mise au service du trading optimal pour déterminer non pas un *rythme de trading optimal* (i.e. le nombre d'actions à acheter ou vendre à chaque instant), mais pour choisir *comment* acheter ou vendre. Pour cela il faut bien comprendre le processus d'achat ou de vente sur les marchés financier, qui est détaillé dans la Section 2 (pour plus de détail cf. [Lehalle *et al.* , 2013]). A ce stade disons simplement qu'il s'agit d'un processus *d'enchère double en temps continu* : les participants envoient des messages avec un sens (achat ou vente), prix et une quantité. Tous les messages participent aux enchères et permettent de fixer un prix walrasien : certains obtiennent des transactions (ceux qui offrent les prix les plus bas ou demande les prix les plus hauts), d'autres n'ont plus qu'à tenter leur chance à l'instant (infinitésimal) d'après. Sous l'hypothèse que l'on envoie des messages avec toujours la même quantité, le contrôle est alors uniquement constitué des différents prix proposés dans ce cadre. Le rythme de trading est une conséquence des prix choisis : si un acheteur choisit des prix très bas, le rythme sera lent ; si il choisit des prix très hauts, le rythme sera rapide. *Comment* acheter ou vendre revient alors à choisir les prix auxquels on envoie des messages.

La Section 1.3.1 expose un cadre de contrôle par les prix pour le *market making* (i.e. *tenue de marché* en français) optimal : le participant envoie simultanément des messages d'achat et de vente (à des prix différent bien entendu). Ce type de participant a un rôle très important : dans un processus d'enchère *peu liquide*, ce genre de comportement permet à des acheteurs et vendeurs mal synchronisés (au sens ou par exemple tous les vendeurs arrivent d'abord, puis tous les acheteurs quelques heures plus tard) les teneurs de marché achètent aux premiers et revendent aux seconds. Ils maintiennent ainsi les prix dans une *fourchette offre-demande* optimale, réalisant un équilibre entre le risque de porter la position (quelques heures) et le gain apporté par l'achat-vente.

Le même point de vue peut être adopté pour un acheteur ou un vendeur, qui serait ainsi comparable à un teneur de marché qui n'aurait pas la possibilité d'envoyer des messages d'achat et commencerait la journée avec un inventaire important. Quelle est le processus des prix optimaux à proposer tout au long de la journée ? Et le rythme de trading finalement obtenu est-il comparable à ce que préconise une stratégie de type moyenne-variance ? La Section 1.3.2 répond à cette question.

1.1 Mean-Variance Optimal Trading

The Almgren-Chriss framework allows an asset manager to schedule optimally the execution of a large trade, under a mean variance criterion taking into account the market impact and the volatility risk.

This framework is largely used in practice, justifying some ellaborations and fine tuning. The contributions presented here cover

- a portfolio version of this framework, implementing a compensation of the market impact on some components of the portfolio by slowing down their execution when the covariance reduces the market risk on a given market factor. The effect of this compensation is studied.
- the understanding of the difference between
  - plugging values of the average volatility and the average traded volume in the final reccurence equation
  - considering them as estimators (with their own expectation and variance)
- the modification of the criterion to take the mean reversion of intraday prices into account (following and somehow generalizing the viewpoint of [Gatheral *et al.* , 2012])
- taking into account risk bounds ; i.e. a maximum trading rate (preventing to impact too much the price formation process). This put some emphasis on the stopping time of an Implementation Shortfall as designed by Almgren and Chriss in [Almgren & Chriss, 2000]. It also put emphasis on the starting time of a Target Close, that had never been formalized before the work presented in Section 1.1.4, despite the associated criterion is widely used by practitioners.

This works does not ignore the interesting literature expanding the seminal work of Almgren-Chriss and Bertsimas-Lo, like the bridges with stochastic control established by Schied [Schied & Schöneborn, 2009], or the extension to a dynamic version of market impact (like in [Obizhaeva & Wang, 2005, Alfonsi & Schied, 2010, Graewe *et al.* , 2013]). It is worthwhile to note this last trend of research produces "U-shaped" trading curves, essentially because of the resiliency of the market model they use when the criterion is no more computed (i.e. at the last trade of the algorithm). The work presented here never goes into this direction for practical reasons : the authors wanted to not produce optimal strategies detrimental to the price formation process.

Recent work like [Alfonsi & Blanc, 2014] propose very interesting continuation of the Almgren-Chriss framework, using Hawkes processes to model the *background noise* in place of the "usual" Brownian ones. Introducing hedring and splitting behaviour in optimal trading is for sure a very realistic proposal.

1.1.1 Joint Mean-Variance Optimization (Balanced Portfolio)

This Section has been published as a part of *Rigorous Strategic Trading : Balanced Portfolio and Mean-Reversion* (2009 The Journal of Trading 4, 40–46) [Lehalle, 2009]

Once an asset manager or a proprietary trader decides to buy or sell a reasonably large portfolio with given weights, it is now well known that the micro structure scale has to be taken into account during the process of buying or selling the shares on the markets. The reference paper presenting the optimal way to build a multi-dimensional trading curve for such a portfolio ([Almgren & Chriss, 2000]) does not constrain the balance between the lines to be controlled during the trading process. This section deals with this constraint and analyses the results obtained.

The usual formalism is to take a discretized K-dimensional correlated Bachelier process for the diffusion of the "*fair price*" : $S_{n+1}^{(k)} = S_n^{(k)} + \xi_{n+1}^{(k)}$ (where ξ_n follows an i.i.d. K-dimensional Gaussian distribution with covariance matrix Σ_n supporting the natural filtration $(\mathcal{F}_n)_n$ used throughout this paper). The time discretization is taken as uniformly sampled (the time interval between any t_n and t_{n+1} is a constant equal to δt) ; but this hypothesis can be relaxed to any stopping time adapted to $(\mathcal{F}_n)_n$ (with an impact on the expression of the covariances Σ_n). The "*execution cost*" for the buy (or the sell) of v shares of the kth line of the portfolio during the nth time interval is $\eta_n^{(k)}(v)$. This execution cost is usually called "*temporary Market Impact*". In this paper, only temporary execution costs will be considered, mainly for notational simplifications : any permanent component can be added to execution costs without any structural change to the equations as long as it is linear.

Thus the paid price to buy v shares of the kth line of the portfolio between $n\,\delta t$ and $(n+1)\delta t$ is $v\,(S_n^{(k)} + \eta_n^{(k)}(v))$.

Finally, the wealth needed to buy a portfolio the kth component of which is a buy of $\varepsilon_k\,\pi_k v^*$ ($\varepsilon_k = +1$ for long lines and -1 for short ones) shares is :

$$W = \sum_{k=1}^{K}\sum_{n=1}^{N} \varepsilon_k\, v_k^{(n)} \left(S_n^{(k)} + \varepsilon_k\, \eta_n^{(k)}(v_n^{(k)}) \right)$$

where $\sum_{n=1}^{N} v_n^{(k)} = \pi_k v^*$ for any k. For notational convenience a long-only portfolio ($\varepsilon_k = +1$ for any k) will be considered. The optimal values for the $v_n^{(k)}$ are obtained through the minimization of a λ-combination of the mean and the variance of W :

$$J_\lambda = \mathbb{E}(W) + \lambda \mathbb{V}(W)$$

With the change of variable $\times(k,n) = \sum_{\ell=n}^{N} v_\ell^{(k)}$ (i.e. the remaining quantity of line k to buy at n) and a linear assumption for $\eta_n^{(k)}(v) = \eta_k \sigma_n^{(k)} v / V_n^{(k)}$ (where η_k is a parameter for the kth line, $\sigma_n^{(k)}$ and $V_n^{(k)}$ its volatility and the

usual traded volume at time n), we have ($\mathbf{x}_n$ is the vectorial notation for $(\times(1,n),\ldots,\times(K,n))$) :

$$J_\lambda = \sum_{n=1}^{N}\left(\lambda\,\mathbf{x}'_n\Sigma_n\mathbf{x}_n + \sum_{k=1}^{K}\eta_k\frac{\sigma_n^{(k)}}{V_n^{(k)}}(\times(k,n)-\times(k,n+1))^2\right)$$

The minimization of J_λ under terminal constraints $\times(k,1)=\pi_k v^*$ and $\times(k,N+1)=0$ for any k is solved by cancelling its gradient. This leads to the following K-dimensional recurrence equation ($\mathcal{D}_n$ is the $K\times K$ diagonal matrix the kth element of which is $\eta_k\,\sigma_n^{(k)}/V_n^{(k)}$) :

$$\mathbf{x}_{n+1} = (1+\lambda\,\mathcal{D}_n^{-1}\Sigma_n + \mathcal{D}_n^{-1}\mathcal{D}_{n-1})\,\mathbf{x}_n - \mathcal{D}_n^{-1}\mathcal{D}_{n-1}\,\mathbf{x}_{n-1} \tag{1}$$

that is a rewriting of the classical equation (31) of [Almgren & Chriss, 2000] under the notational simplifications of this paper.

Recurrence equation for a balanced portfolio. To project equation (9) on a balancing constraint it is sufficient to write $\pi_k\,\tilde{v}_n$ (where $\tilde{v}_n$ is a scalar) instead of $v_n^{(k)}$. This ensures control of the balance of the portfolio at a δt time scale, and translates the original $K\times N$ dimensional problem into a N dimensional one.

The corresponding recurrence equation is scalar :

$$\tilde{x}_{n+1} = \left(1+\lambda\frac{\pi'\Sigma_n\pi}{\mathcal{D}_n^\pi} + \frac{\mathcal{D}_{n-1}^\pi}{\mathcal{D}_n^\pi}\right)\tilde{x}_n - \frac{\mathcal{D}_{n-1}^\pi}{\mathcal{D}_n^\pi}\,\tilde{x}_{n-1} \tag{2}$$

where π is the vector $(\pi_1,\ldots,\pi_K)$ and $\mathcal{D}_n^\pi = \sum_{k=1}^K \eta_k\,\sigma_n^{(k)}\pi_k^2/V_n^{(k)}$ (i.e. a scalar). The terminal conditions for the recurrence are $\tilde{x}_0 = v^*$ and $\tilde{x}_{N+1}=0$.

Comparing the influence of execution costs. When execution cost coefficients (i.e. Market Impact) η_k are large compared to the risk aversion parameter λ, the recurrence equation (9) becomes K independent equations such as this one :

$$\times(k,n+1) = \left(1+\frac{\sigma_{n-1}^{(k)}}{\sigma_n^{(k)}}\,\frac{V_n^{(k)}}{V_{n-1}^{(k)}}\right)\times(k,n) - \frac{\sigma_{n-1}^{(k)}}{\sigma_n^{(k)}}\,\frac{V_n^{(k)}}{V_{n-1}^{(k)}}\;\times(k,n-1) \tag{3}$$

as the balanced recurrence equation (12) is quite similar replacing the coefficient $a_n^{(k)} = \sigma_{n-1}^{(k)}/\sigma_n^{(k)}\cdot V_n^{(k)}/V_{n-1}^{(k)}$ by its $\pi^2\eta$"*averaged version*" :

$$A_n = \left(\frac{\sum_{k=1}^K \pi_k^2\eta_k\,\sigma_{n-1}^{(k)}/V_{n-1}^{(k)}}{\sum_{k=1}^K\pi_k^2\eta_k}\right)\left(\frac{\sum_{k=1}^K \pi_k^2\eta_k\,\sigma_n^{(k)}/V_n^{(k)}}{\sum_{k=1}^K\pi_k^2\eta_k}\right)^{-1}$$

The first element to note is that the execution costs coefficients η_k no longer appear in the portfolio equation (12). The parameters that influence the optimal trading curves are ratios of usual volumes and usual volatility. As a consequence, the optimal trading curve will be the same for any asset that has the same intra-day volume and volatility shapes. If the volatility of an asset is regularly twice that of another asset, they will have the same optimal portfolio curves.

For the balanced portfolio, all optimal curves will be based on the same profile obtained from an averaged version of the execution costs. The difference between the two curves (the balanced one, to be used for instance by program trading algorithms, and the classic portfolio one) is due to the averaging : the values of the weights π_k and the scale parameter η_k of the execution costs can lead to very different A_n from the same values of $a_n^{(k)}$. A geometrical interpretation helps to understand the effect of the averaging : plotting $y_k = \sigma_{n-1}^{(k)}/V_{n-1}^{(k)}$ on the vertical axis and $x_k = \sigma_n^{(k)}/V_n^{(k)}$ on the horizontal one obtains $a_n^{(k)}$ as the slope to origin (Figure 1). A_n is the slope to origin of the barycenter of the $(x,y)_k$ associated with the weights $(\pi_k^2\eta_k)$. The wide spectrum of values that can be taken by A_n can be seen in Figure 1.

Figure 2 (left) compares the optimal curves found by the classical method (unbroken lines) and the one found for the balanced portfolio (dotted lines) on a numerical example. As expected, all the balanced curves have the same shape ; the difference between the classical curves and their balanced equivalent are plotted in Figure 2 (right).

Measuring the covariance effects. When execution costs do not dominate the market risk (i.e. η is not far larger than λ), we can study the stylized fact of two *identical* assets (in terms of usual volatility, market volume and execution costs) that have a correlation of ρ. For numerical illustrations all constants will be taken equal to one. The two remaning variables are the correlation ρ between the two lines of the portfolio and their weights π and $1-\pi$.

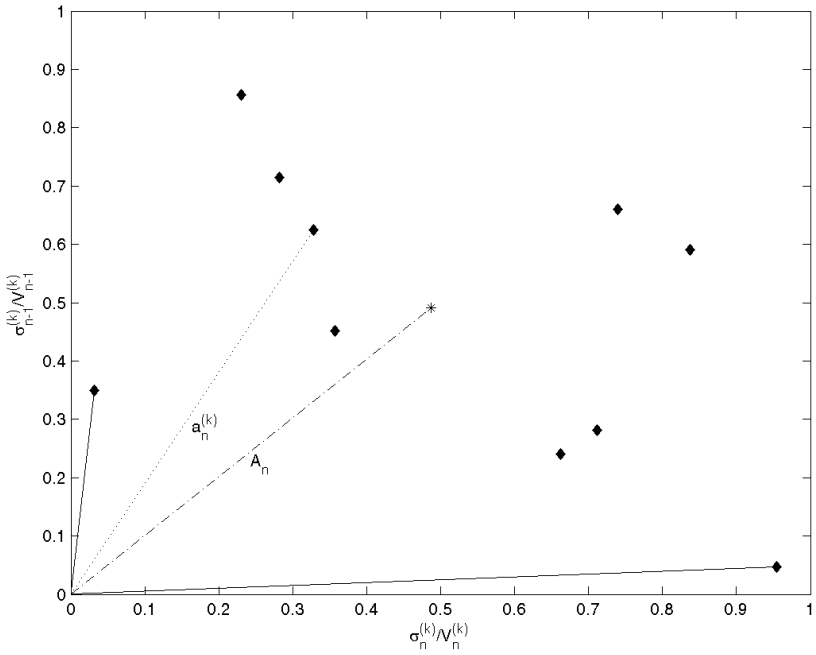

FIGURE 1 – Geometrical interpretation of the $\pi^2\eta$ averaging : for arbitrary weights and 10 lines, the barycenter to be used for the balanced portfolio is marked by a star, and the minimum and maximum slopes are the unbroken lines.

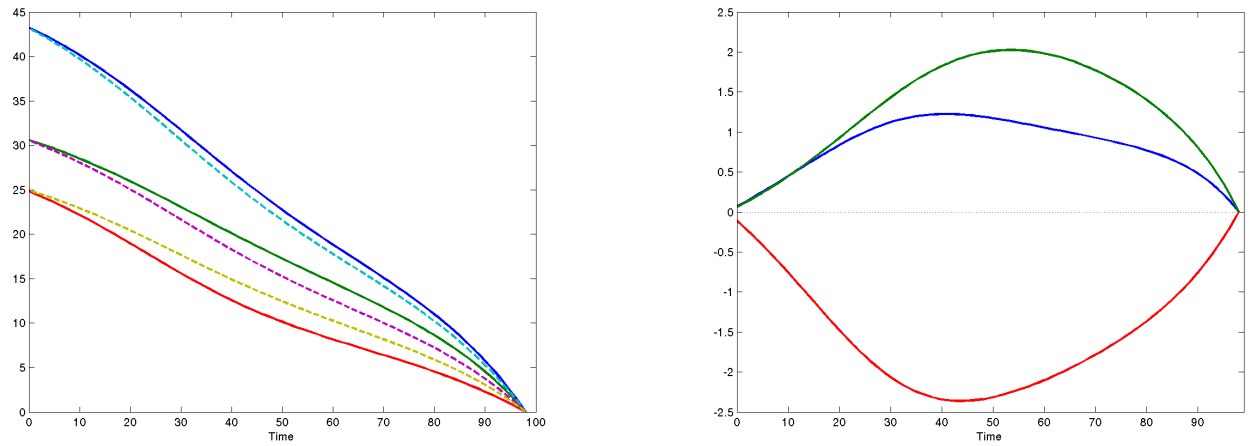

FIGURE 2 – Left : optimal liquidation curves for a classical portfolio (unbroken lines) and the balanced equivalent (dotted lines). Right : difference between optimal liquidation curves for a classical portfolio and its balanced equivalent.

Simple algebraic computations on the two recurrence equations lead to this matrix equivalent :

$$\begin{pmatrix} \times(1,n+1) \\ \times(2,n+1) \\ \times(1,n) \\ \times(2,n) \end{pmatrix} = \begin{bmatrix} 2+\lambda & \rho\lambda & -1 & 0 \\ \rho\lambda & 2+\lambda & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{pmatrix} \times(1,n) \\ \times(2,n) \\ \times(1,n-1) \\ \times(2,n-1) \end{pmatrix} \tag{4}$$

for the dynamic of the classical portfolio equation (9), and for the balanced one :

$$\begin{pmatrix} \tilde{x}_{n+1} \\ \tilde{x}_n \end{pmatrix} = \underbrace{\begin{bmatrix} 2+\lambda(1+w_\pi\rho) & -1 \\ 1 & 0 \end{bmatrix}}_{F} \begin{pmatrix} \tilde{x}_n \\ \tilde{x}_{n-1} \end{pmatrix} \tag{5}$$

where $w_\pi = 2\pi(1-\pi)/(\pi^2+(1-\pi)^2)$ is a weight depending only on the composition of the portfolio (cf Figure 3, left).

It is quite easy to solve explicitly those dynamics ; the terminal conditions (all the lines have to be fully executed) can be inferred to find $\tilde{x}_1 = -F^N_{1,2}/F^N_{1,1}v^*$ (where F is the matrix of equation (15)) and its vectorial equivalent $(\times(1,1),\times(2,1))'$.

To obtain illustrations, λ and v^* set at one and the weight $\pi = 25\%$ (dotted line on Figure 3, left).

As a qualitative analysis of recurrence equations (14) and (15), first note that they are the same when the assets are not correlated (i.e. $\rho = 0$), and that the eigenvalues of F increase with ρ : the more the assets are correlated and the more the algorithm "*fears*" market risk. As a consequence, a highly correlated portfolio should optimally be liquidated faster than a diversified one. Figures 3 (right) and 4 (left) show the optimal trajectories for each algorithm as functions of correlations, and Figure 4 (right) compares the two surfaces.

It can be seen that the surfaces are the same when $\rho = 0$ and the imbalance between the positive and the negative parts of the surface comes from the effect of w_π when π is not $1/2$.

Balanced portfolio vs. classical optimization. This analysis shows that trading algorithms dedicated to a balanced portfolio can be optimized within the Almgren-Chriss framework following the dynamics of equation (11). It can be extended to non-linear execution costs and take into account more trading effects (such as the variability of the main exogeneous variables : usual volumes as in [Lehalle, 2008] or usual volatilities).

The approach presented here uses an algebraic view of the dynamics through equations (14) and (15) that enables a computational efficiency to be reached even when no closed form solution exists. The execution cost analysis has shown that the geometrical interpretation of the execution costs (Figure 1) can be used to understand the way that the balanced portfolio takes individual costs into account. The correlation analysis has shown on the one hand that the imbalance between the lines (through the w_π coefficient) influence the averaged trading trajectory, and on the other hand that a *diversified portfolio can be traded slower that a concentrated one* (decreasing its overall market impact).

1.1.2 Pluging Estimators in the Almgren-Chriss Framework

This Section has been published as a part of *Market Microstructure knowledge needed to control an intra-day trading process* (2013 in : Handbook on Systemic Risk) [Lehalle, 2013]

A now widely used framework to control the overall costs of the liquidation of a portfolio has been proposed by Almgren and Chriss in the late nineties [Almgren & Chriss, 2000]. Applied to trade a single stock, this framework :

- cut the trading period into an arbitrary number of intervals N of a chosen duration δt,
- models the *fair price* moves thanks to a Gaussian random walk :

$$S_{n+1} = S_n + \sigma_{n+1}\sqrt{\delta t}\,\xi_{n+1} \tag{6}$$

- models the *temporary market impact* η_n inside each time bin using a power law of the trading rate (i.e. the ratio of the traded shares v_n by the trader over the market traded volume during the same period V_n) :

$$\eta(v_n) = a\,\psi_n + \kappa\,\sigma_n\sqrt{\delta t}\left(\frac{v_n}{V_n}\right)^\gamma \tag{7}$$

  where a, κ and γ are parameters, and ψ is the half bid-ask spread ;
- the *permanent market impact* is taken linear in the participation rate ;
- uses a mean-variance criterion and minimise it to obtain the optimal sequence of shares to buy (or sell) through time.

It is first important to notice that there is an implicit relationship between the time interval δt and the temporary market impact function : without changing η and simply by choosing a different slicing of time, the cost of trading is changed. It is in fact not possible to choose (a,κ,γ) and δt independently ; they have to be chosen accordingly to the decay of the market impact on the stock, provided that most of the impact is kept in a time bin of size δt. Not all the decay functions are compatible with this view (see [Gatheral & Schied, 2012] for details about available market impact

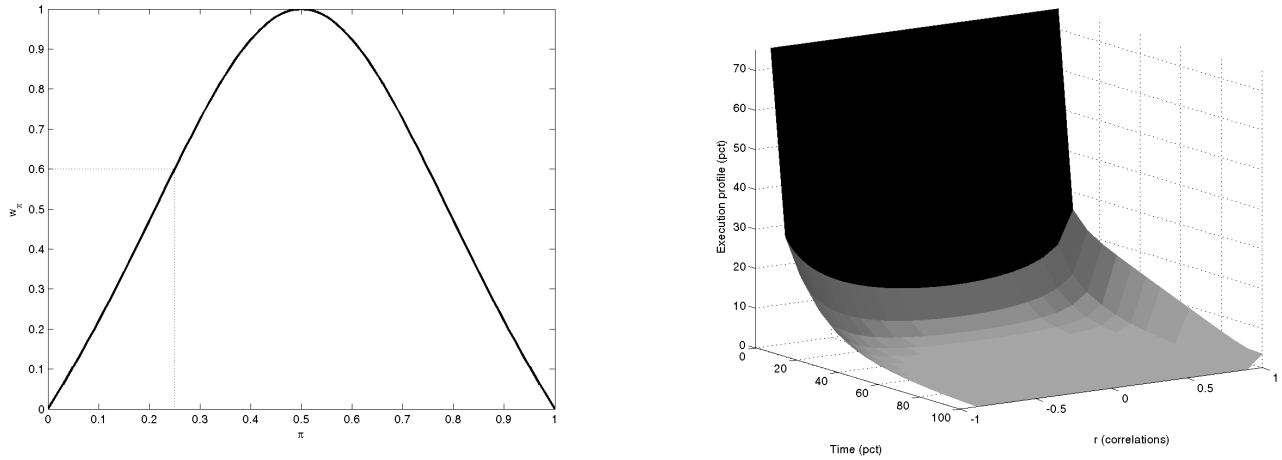

FIGURE 3 – Left : weighting to be used for the balanced portfolio dynamic. Right : optimal trajectory for one of the correlated assets as a function of time and correlation (classical portfolio case).

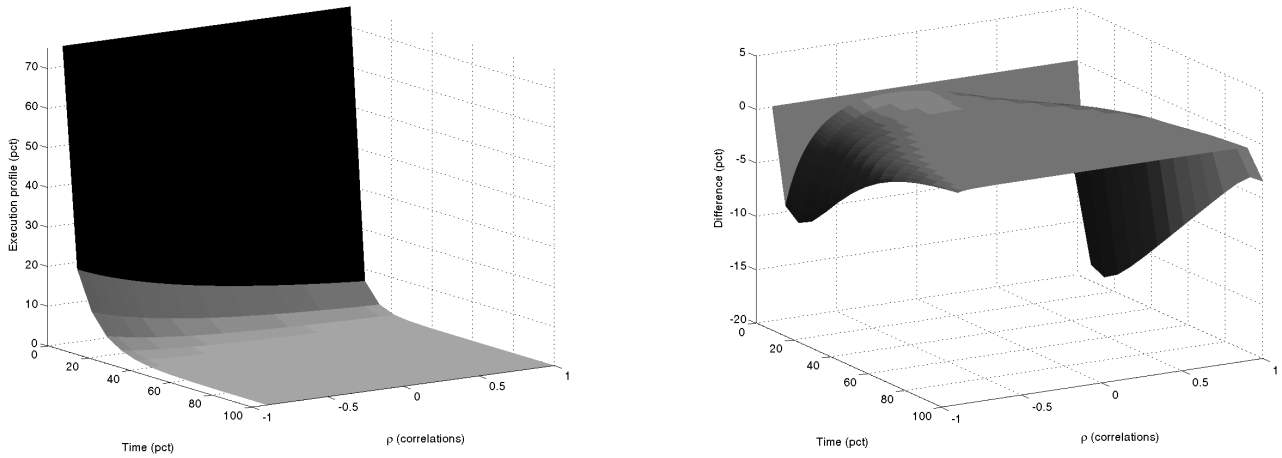

FIGURE 4 – Left : optimal trajectory for one of the correlated assets as a function of time and correlation (balanced portfolio case). Right : comparison of optimal trajectories associated to the classical case and the balanced one.

models and their interactions with trading). Up to now the terms in $\sqrt{\delta t}$ has been ignored. It will also be considered that the parameters (a, κ, γ) exist at this time scale.

It this not mandatory to see this framework as if it was based on structural model assumptions (i.e. that the market impact really have this shape, or that the price moves are really Brownian), but as a statistical one. With such a viewpoint, any practitioner can use the database of its past executed orders and perform an econometric study of its "trading costs" on any δt interval of time (see [Engle *et al.* , 2012] for an analysis of this kind on the whole duration of the order). If a given time scale succeeds into capturing the parameters of a model of trading cost with enough accuracy, then this model can be used to optimise trading. Formally, the result of such a statistical approach will be the same as the structural one, as it will be shown later in this paper, it is besides possible to go one step further, taking into account the statistical properties of the variables (and parameters) or interest.

Going back to the simple case of the liquidation of one stock without any permanent market impact, the value (which is a random variable) of a buy of v^* shares in N bins of size $v_1, v_2, \ldots, v_N$ is (see [Bouchard *et al.* , 2011] for a more sophisticated model and more generic utility functions ; since a stylized model is used here to obtain easier illustrations of phenomena of interest) is :

$$\begin{aligned} W(v_1, v_2, \ldots, v_N) &= \sum_{n=1}^{N} v_n(\quad S_n \quad + \quad \eta_n(v_n)) \\ &= S_0 v^* + \underbrace{\sum_{n=1}^{N} \sigma_n \xi_n x_n}_{\text{market move}} + \\ &\quad \underbrace{\sum_{n=1}^{N} a\, \psi_n (x_n - x_{n+1}) + \kappa\, \frac{\sigma_n}{V_n^{\gamma}} (x_n - x_{n+1})^{\gamma+1}}_{\text{market impact}} \end{aligned} \tag{8}$$

using the *remaining quantity to buy* : $x_n = \sum_{k \geq n} v_k$ instead of the instantaneous volumes v_n. To obtain as much closed form formula as possible, γ will be taken equal to 1 (i.e. linear market impact).

To add a practitioner-oriented flavor to our upcoming optimisation problems, just introduce a set of independent random variables $(A_n)_{1 \leq n \leq N}$ to model the *arbitrage opportunities* during time slices. It will reflect the anticipation that the trader will be able to buy shares at price $S_n - A_n$ during slice n rather than at price S_n.

Such an effect can be used to inject a statistical arbitrage approach into optimal trading or to take into account the anticipation of the opportunity to cross orders at mid price in Dark Pools or Broker Crossing Networks (meaning that the expected trading costs should be smaller during given time slices). Now the cost to buy v^* shares is :

$$W(\mathbf{v}) = S_0 v^* + \sum_{n=1}^{N} \sigma_n \xi_n x_n + \sum_{n=1}^{N} (a\, \psi_n - A_n) v_n + \kappa\, \frac{\sigma_n}{V_n}\, v_n^2 \tag{9}$$

Conditioned expectation optimisation. The expectation of this cost $\mathbb{E}(W | (V_n, \sigma_n, \psi_n)_{1 \leq n \leq N})$ given the market state writes :

$$C_0 = S_0 v^* + \sum_{n=1}^{N} (a\, \psi_n - \mathbb{E}A_n) v_n + \kappa\, \frac{\sigma_n}{V_n}\, v_n^2 \tag{10}$$

a simple optimisation under constraint (to ensure $\sum_{n=1}^{N} v_n = v^*$) gives :

$$v_n = w_n \left(v^* + \frac{1}{\kappa} \left(\left(\mathbb{E}A_n - \sum_{\ell=1}^{N} w_\ell \mathbb{E}A_\ell \right) - a \left(\psi_n - \sum_{\ell=1}^{N} w_\ell \psi_\ell \right) \right) \right) \tag{11}$$

where w_n are weights proportional to the inverse of the market impact factor :

$$w_n = \frac{V_n}{\sigma_n} \left(\sum_{\ell=1}^{N} \frac{V_\ell}{\sigma_\ell} \right)^{-1}$$

Simple effects can be deduced from this first stylized result :

1. without any arbitrage opportunity and without any bid-ask cost (i.e. $\mathbb{E}A_n = 0$ for any n and $a = 0$), the optimal trading rate is proportional to the inverse of the market impact coefficient : $v_n = w_n \cdot v^*$. Moreover, when the market impact has no intra-day seasonality, $w_n = 1/N$ implying that the optimal trading rate is linear.
2. following formula (11) it can be seen that : the largest the expected arbitrage gain (or the lower the spread cost) on a slice compared to the market-impact-weighted expected arbitrage gain (or spread cost) over the trading full interval, the largest quantity to trade during this slice. More quantitatively :

$$\frac{\partial v_n}{\partial \mathbb{E}A_n} = \frac{w_n}{2\kappa}(1 - w_n) > 0, \ \frac{\partial v_n}{\partial \psi_n} = -\frac{a}{2\kappa}(1 - w_n) w_n < 0$$

This result gives the *adequate weight* to apply to the expected arbitrage gain to translate it in an adequate trading rate to take profit on arbitrage opportunities on average. Just note that usually the expected arbitrage gains increase with market volatility, the w_n-weighting is consequently of interest to balance this effect optimally.

Conditioned mean-variance optimisation. Going back to a mean-variance optimisation of the cost to buy progressively v^* shares, the criteria to minimise (using a risk aversion parameter λ) writes :

$$\begin{aligned} C_\lambda &= \mathbb{E}(W|(V_n,\sigma_n,\psi_n)_{1\leq n\leq N}) + \lambda\mathbb{V}(W|(V_n,\sigma_n,\psi_n)_{1\leq n\leq N}) \\ &= S_0 v^* + \sum_{n=1}^{N}(a\psi_n - \mathbb{E}A_n)\mathbb{X}(n) + \left(\kappa\frac{\sigma_n}{V_n} + \lambda\mathbb{V}A_n\right)\mathbb{X}(n)^2 \end{aligned} \tag{12}$$

To minimise C_λ being only constrained by terminal conditions on x (i.e. $x_0 = v^*$ and $v_{N+1} = 0$) it is enough to cancel its derivatives with respect to any x_n, leading to a recurrence relation :

$$\begin{aligned} \left(\frac{\sigma_n}{V_n} + \frac{\lambda}{\kappa}\mathbb{V}A_n\right)x_{n+1} &= \frac{1}{2\kappa}(a(\psi_{n-1}-\psi_n) - (\mathbb{E}A_{n-1} - \mathbb{E}A_n)) \\ &\quad + \left(\frac{\lambda}{\kappa}\sigma_n^2 + \left(\frac{\sigma_n}{V_n} + \frac{\lambda}{\kappa}\mathbb{V}A_n + \frac{\sigma_{n-1}}{V_{n-1}} + \frac{\lambda}{\kappa}\mathbb{V}A_{n-1}\right)\right)x_n \\ &\quad - \left(\frac{\sigma_{n-1}}{V_{n-1}} + \frac{\lambda}{\kappa}\mathbb{V}A_{n-1}\right)x_{n-1} \end{aligned} \tag{13}$$

It shows that the variance of the arbitrage has a similar effect than the market impact (through a risk-aversion rescaling), and that the risk aversion parameter acts as a multiplicative factor to the market impact, meaning that within an arbitrage-free and spread-costs-free framework (i.e. $a = 0$ and $\mathbb{E}A_n = 0$ for all n), the market impact model by any constant b has no effect on the final result as far as λ is replaced by $b\lambda$.

Figure 5 compares optimal trajectories coming from different criteria and parameter values.

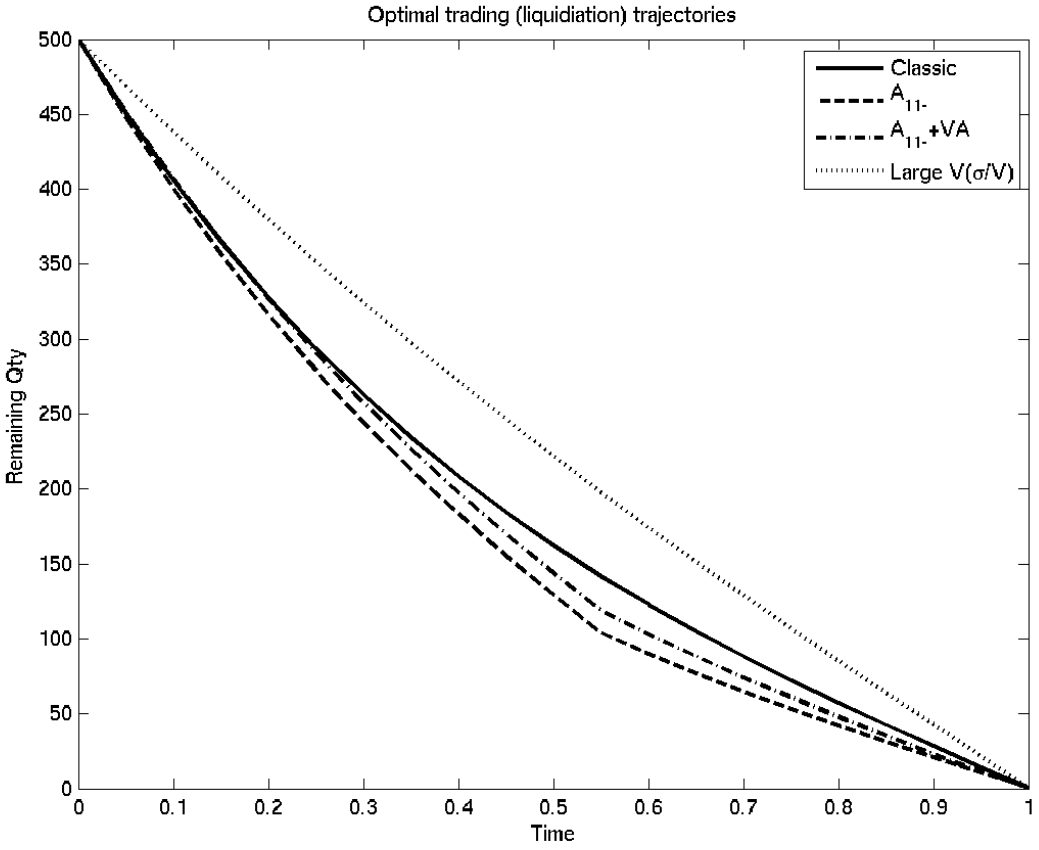

FIGURE 5 – Examples of optimal trading trajectories for mean-variance criteria : the classical result (Almgren-Chriss) in solid line, the dotted line is for high variance of the variable of interest (σ/V), the semi-dotted ones for an arbitrage opportunity (A_{11+} means : after the 11th period and $A_{11+} + VA$ means : adding expected variance to the arbitrage opportunity).

A statistical viewpoint. The two previous examples show how easy it is to include effects in this sliced mean-variance framework, the implicit assumptions are :

- on one time-slice, it is possible to capture the market impact (or *trading costs*) using model (7),
- the trader has a view on the traded volumes and market volatility at the same time-scale.

Practically, the two assumptions come from statistical modelling :

- the market impact parameters a, κ and γ are estimated on a large database of trades using a maximum likelihood or MSE methods ; the reality is consequently that the market model has the following shape :

$$\eta(v_n) = a\,\psi_n + \kappa\,\sigma_n\sqrt{\delta t}\left(\frac{v_n}{V_n}\right)^{\gamma} + \varepsilon \tag{14}$$

  where ε is an i.i.d. noise.
- moreover, the market volatility and traded volumes are estimated using historical data and market context assumptions (to take into account at least the scheduled news, like the impact of the expiry of derivative products on the volume of the cash market, see Figure 6 for typical estimates).

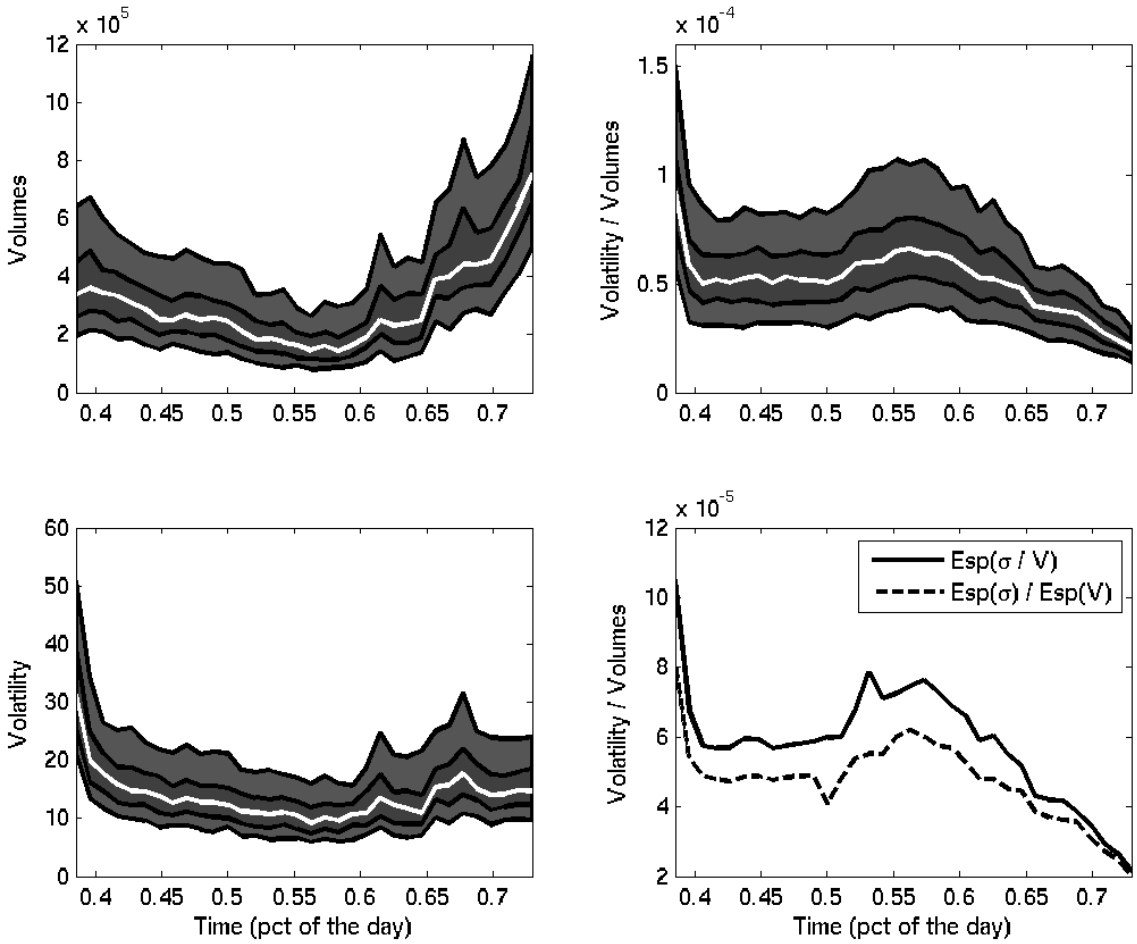

FIGURE 6 – Typical intra-day traded volume (top-left) and realized volatility (bottom-left) profiles (i.e. intra-day seasonalities on traded volumes and market volatility) with their quantiles of level 25% and 75%. X-axis is time. The top-right chart figures the quantiles of the ratio of interest σ/V. The bottom-right ones shows the difference between the expectation of the ratio (solid line) and the ratio of the expectations (dotted line).

Taking these statistical modelling steps into account in the classical mean-variance criteria of equation (12) changes it into its unconditioned version :

$$\begin{aligned}\tilde{C}_\lambda &= \mathbb{E}(W) + \lambda \mathbb{V}(W) \\ &= S_0 v^* + \sum_{n=1}^{N} (a\mathbb{E}\psi_n - \mathbb{E}A_n)\mathbb{X}(n) \\ &\quad + \left(\kappa\, \mathbb{E}\left(\frac{\sigma_n}{V_n}\right) + \lambda(a^2 \mathbb{V}\psi_n + \mathbb{V}A_n + \mathbb{V}\varepsilon)\right) \mathbb{X}(n)^2 \\ &\quad + \lambda\sigma_n^2 x_n^2 + \lambda\kappa^2 \mathbb{V}\left(\frac{\sigma_n}{V_n}\right) \mathbb{X}(n)^4\end{aligned} \tag{15}$$

The consequence of using this criteria rather than the conditioned one are clear :

- the simple plug-in of empirical averages of volumes and volatility in criteria (12) instead of the needed expectation of the overall trading costs leads to use $(\mathbb{E}\sigma_n)/(\mathbb{E}V_n)$ instead of $\mathbb{E}(\sigma_n/V_n)$, Figure 6 shows typical differences between the two quantities.
- if the uncertainty on the market impact is huge (i.e. the $\mathbb{V}\varepsilon$ term dominates all others), then the optimal trading strategy is to trade linearly, which is also the solution of a pure expectation-driven minimisation with no specific market behaviour linked with time.

Within this new statistical trading framework, the inaccuracy of the models and the variability of the market context are taken into account : the obtained optimal trajectories will no more have to follow sophisticated behaviours if the models are not realistic enough.

Moreover, it is not difficult to solve the optimisation program associated to this new criteria ; the new recurrence equation is a polynomial of degree 3. Figure 5 gives illustrations of obtained results.

A lot of other effects can be introduce in the framework, like auto-correlations on the volume-volatility couple. This statistical framework does not embed recent and valuable proposals like the decay of market impact [Gatheral & Schied, 2012] or a set of optimal stopping times to do not stick to an uniform and a priory sampled time [Bouchard *et al.* , 2011]. It is nevertheless simple enough so that most practitioners can use it to include their views of the market conditions and the efficiency of their interactions with the market at a given time scale ; it can be compared to the Markowitz approach for quantitative portfolio allocation [Markowitz, 1952].

1.1.3 Modifying the Framework to Take Into Account other Dynamics

This Section has been published as a part of *Optimal starting times, stopping times and risk measures for algorithmic trading* (2014 The Journal of Investment Strategies 3) [Labadie & Lehalle, 2014]

The p-variation model. Let $p > 1$ and $y = (y_1, \ldots, y_N) \in \mathbb{R}^N$ be a random vector of mean zero. We define the *p-variation* of y as

$$\mathbb{V}_p(y) := \sum_{n=1}^{N} \mathbb{E}\left[\, |\, y_n|^p \,\right].$$

The p-variation $\mathbb{V}_p(y)$ and the l_p-norm in $\mathbb{R}^N$ are related via

$$\|y\|_p = \mathbb{V}_p(y)^{1/p}.$$

Notice that if $y = (y_1, \ldots, y_N)$ is a time series of i.i.d. random variables then the 2-variation reduces to the variance, i.e.

$$\mathbb{V}_2(y) = \mathrm{Var}(y)\,.$$

Moreover, it is easy to show that the p-variation defines a metric on $\mathbb{R}^N$, and since all norms in $\mathbb{R}^N$ are equivalent there exist $0 < \beta_1 < \beta_2$ such that

$$\beta_1 \|y\|_p \leq \|y\|_2 \leq \beta_2 \|y\|_p.$$

Therefore, the variance (i.e. the 2-variation) and the p-variation are two equivalent metrics on $\mathbb{R}^N$, and in particular

$$\mathbb{V}_p(y) \sim \mathbb{V}_2(y)^{p/2}\,. \tag{16}$$

Now let us define the p-variation for a special family of functions of random variables. Let $y = (y_1, \ldots, y_N)$ a random vector of mean zero and consider the function $F : \mathbb{R}^N \to \mathbb{R}$ defined as

$$F(y) = \sum_{n=1}^{N} y_n.$$

We define the p-variation of F as

$$\mathbb{V}_p(F) = \sum_{n=1}^{N} \mathbb{E}\left[\, |\, y_n|^p \,\right].$$

Observe that if $y = (y_1, \dots, y_N)$ is a time series of mean zero then $\mathbb{V}_p(F)$ is the sample p-th moment of the time series y multiplied by N. Finally, for general functions F such that

$$F - \mathbb{E}(F) = \sum_{n=1}^{N} y_n$$

we define their p-variation as

$$\mathbb{V}_p(F) := \mathbb{V}_p(F - \mathbb{E}(F)) = \mathbb{V}_p(y).$$

It is worth o remark that if the random variables $y = (y_1, \dots, y_N)$ are i.i.d. of mean zero and variance 1 then the 2-variation and the variance of y coincide :

$$\mathbb{V}_2(F) = \mathrm{Var}(y)\,.$$

Optimal trading algorithms using p-variance as risk measure. Let $H \in (0,1)$ and assume that the price dynamics is self-similar, i.e.

$$S_{n+1} = S_n + \sigma_{n+1}\tau^H \varepsilon_{n+1}, \tag{17}$$

where $H \in (0,1)$ and $(\varepsilon_n)_{1\leq n\leq N}$ are i.i.d. random variables such that $\mathbb{E}[\varepsilon_n] = 0$. We will assume a power market impact of the form

$$h(v_n) = \kappa\sigma_n\tau^H \left(\frac{v_n}{V_n}\right)^{\gamma}. \tag{18}$$

In order to use the p-variation as a risk measure, we have to choose the right p. From (17) we see that if $H = 1/2$ we recover the classical Brownian motion, for which the variance is the most common choice for a risk measure. In this case we have $H = 1/2$ and $p = 2$, which implies that the risk measure is linear in time. This suggests that the correct choice of p is $p = 1/H$, since it is the only p that renders the risk p-variation as a risk measure linear in time.

We would like to remark that the idea of a risk measure that is linear in time was also introduced by [Gatheral & Schied, 2012], where the risk measure was the expectation of the time-average. The advantage of our approach is that we do not fix a priori the dynamics of the price process. Indeed, we first find empirically the right exponent of self-similarity H and then we choose the correct risk measure via $p = 1/H$.

In order to derive the recursive formula for a process following (17), we normalise the relative wealth as in the previous case of Brownian motion. Under this framework, the normalised relative wealth of a TC algorithm is

$$\tilde{W} = -\sum_{n=1}^{N-1} x_n\sigma_{n+1}\varepsilon_{n+1} + \sum_{n=1}^{N} \kappa\sigma_n \frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^{\gamma}}\,.$$

We will assume that the process (17) is normalised, i.e. $\mathbb{E}[|\varepsilon_n|^p] = 1$ for all n. In the case of Brownian motion ($H = 1/2$ and $p = 2$) this is equivalent to suppose that the increments $(\varepsilon_n)_{1\leq n\leq N}$ have variance 1. Under this framework, the average and p-variation of $\tilde{W}$ are

$$\mathbb{E}(\tilde{W}) = \sum_{n=1}^{N} \kappa\sigma_n \frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^{\gamma}}\,, \qquad \mathbb{V}_p(\tilde{W}) = \sum_{n=1}^{N-1} x_n^p\sigma_{n+1}^p\,.$$

Therefore, the corresponding p-functional is

$$\begin{aligned} J_p(x_1, \dots, x_N) &= \mathbb{E}(\tilde{W}) + \lambda\mathbb{V}_p(\tilde{W}) \\ &= \sum_{n=1}^{N} \kappa\sigma_n \frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^{\gamma}} + \lambda\sum_{n=1}^{N-1} x_n^p\sigma_{n+1}^p\,. \end{aligned}$$

The optimal trading curve is determined by solving

$$\frac{\partial J_p}{\partial x_n} = 0\,, \qquad n = 1, \dots, N\,,$$

i.e.

$$\kappa\sigma_n(\gamma+1)\frac{(x_n - x_{n-1})^{\gamma}}{V_n^{\gamma}} - \kappa\sigma_{n+1}(\gamma+1)\frac{(x_{n+1} - x_n)^{\gamma}}{V_{n+1}^{\gamma}} + p\lambda\mu_N(n)\sigma_{n+1}^p x_n^{p-1} = 0\,,$$

where $\mu_N(n) = 1$ if $n < N$ and $\mu_N(N) = 0$.

Returning to the variables v_n we get

$$\kappa\sigma_n(\gamma+1)\left(\frac{v_n}{V_n}\right)^{\gamma} - \kappa\sigma_{n+1}(\gamma+1)\left(\frac{v_{n+1}}{V_{n+1}}\right)^{\gamma} + p\lambda\mu_N(n)\sigma_{n+1}^p\left(\sum_{i=1}^{n} v_i\right)^{p-1} = 0\,.$$

We thus obtain the recursive nonlinear formula for the optimal TC trading curve for a self-similar process :

$$v_{n+1} = V_{n+1}\left[\frac{\sigma_n}{\sigma_{n+1}}\left(\frac{v_n}{V_n}\right)^{\gamma} + \frac{p\lambda\mu_N(n)}{\kappa(\gamma+1)}\sigma_{n+1}^{p-1}\left(\sum_{i=1}^{n} v_i\right)^{p-1}\right]^{1/\gamma} . \tag{19}$$

As in the Brownian case, if we were interested in the IS algorithm then an argument similar to the previous one would show that the optimal trading curve for IS satisfies

$$v_{n-1} = V_{n-1}\left[\frac{\sigma_n}{\sigma_{n-1}}\left(\frac{v_n}{V_n}\right)^{\gamma} + \frac{p\lambda\mu_1(n)}{\kappa(\gamma+1)}\frac{\sigma_n^p}{\sigma_{n-1}}\left(\sum_{i=n}^{N} v_i\right)^{p-1}\right]^{1/\gamma} , \tag{20}$$

where $\mu_N(n) = 1$ if $n > 1$ and $\mu_1(1) = 0$.

Examples of self-similar processes. Amongst the class of continuous stochastic processes that admit a discretisation of the form (17), we have three proceses in mind : Lévy processes, fractional Brownian motion and fractal processes (for more details we suggest [Bacry *et al.* , 2001], [Bouchaud & Potters, 2004][Embrechts, 2002], [Mandelbrot & Hudson, 2004] and [Mantegna & Stanley, 1994]) :

- **Truncated Lévy processes**. p-stable Lévy processes are the only self-similar processes satisfying (17) with $H = 1/p$ and with independent, stationary increments. If $p = 2$ we recover the classical Brownian motion. However, for such processes the p-th moment is infinite, and as such they cannot be used in our framework. Nevertheless, one can consider the so-called *truncated* Lévy distributions, which are Lévy within a bounded interval and exponential on the tails. This allows moments of any order, in particular the p-th moment, whilst within the bounded interval we keep the self-similarity given by (17).
- **Fractional Brownian motion**. The fractional Brownian motion is the only self-similar process with stationary, Gaussian increments. Its exponent of self-similarity H is called *Hurst exponent*). If $H = 1/2$ we recover the classical Brownian motion. The fractional brownian motion has moments of all orders, hence the p-variation is well-defined and we can apply our model. However, for $H \neq 1/2$ the increments are auto-correlated (positively if $H > 1/2$ and negatively if $H < 1/2$) and our model does not take into account the autocorrelations. Therefore, we can consider our model as an approximation when autocorrelations are weak with respect to the market impact and the p-variance.
- **Multifractal processes**. Multifractal processes are defined as follows. Given a stochastic process $X(t)$ its fluctuation is defined as

  $$\delta_l X(t) := X(t+l) - X(t) .$$

  For any $q > 0$ define

  $$m(q,l) := \mathbb{E}\left[|\delta_l X(t)|^q\right] .$$

  We say that $X(t)$ is multifractal of exponents $\zeta(q)$ if for any $q > 0$ there exists $K(q) > 0$ such that

  $$m(q,l) = K(q) l^{\zeta(q)} .$$

  In the case where $\zeta(q)$ is linear, i.e. $\zeta(q) = qH$ the process $X(t)$ is called monofractal. If $\zeta(q)$ is not linear then $X(t)$ is called multifractal. Notice that all self-similar processes are monofractal, in particular the fractional Brownian motion and Lévy processes. However, we will continue to use the term *self-similar*, even for monofractal process, since it is more common in the literature.

Numerical results. In Figure 7 we plotted three TC curves under the PVol constraint for three different self-similarity exponents H, which gives three different p's for the p-variation (recall $p = 1/H$).

H	p	start pillar	switch pillar
0.55	1.8	17	102
0.50	2.0	34	94
0.45	2.2	50	89

Our numerical example renders the following evidence, which has been found in all runs we have performed :

- If H increases then the starting pillar of the execution decreases, i.e. the execution starts earlier.
- If H increases then the pillar at with we switch from TC to PVol increases, i.e. the PVol constraint is saturated later.

Since starting the execution later and saturating the PVol constraint earlier is related to higher levels of aggressiveness, we can infer from Figure 7 that the level of aggressiveness of TC under the PVol constraint decreases as H increases. This finding is quite natural if we assume that the model is a fractional Brownian motion and H is the Hurst exponent :

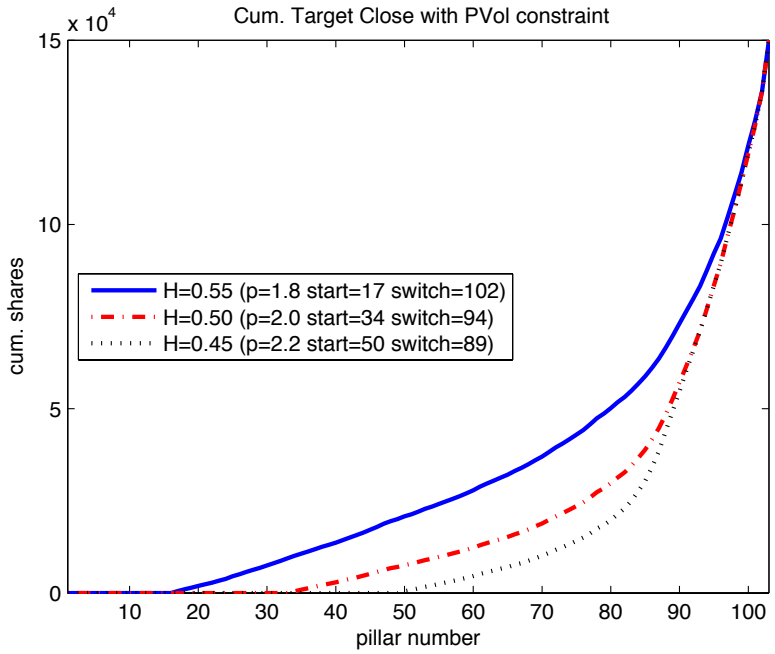

FIGURE 7 – Cumulative TC curves under PVol constraint for different values of H.

- For $H < 1/2$ the process has negative autocorrelations, i.e. it behaves as a mean-reverting process. Therefore, the market impact is reduced because prices go back to their level after an execution. In consequence, we can execute the order faster than in the case of a classical Brownian motion : we start the execution later and we go as fast as possible, and as such we saturate the constraint earlier.
- For $H > 1/2$ the process has positive autocorrelations, i.e. it has a trend. Therefore, the market impact is of paramount importance because if we execute too fast then prices will move in the wrong direction. In consequence, we start the execution earlier and we go as slow as possible, and as such we saturate the constraint later.

In the next section we will study in detail, in the TC algorithm without PVol constraint, the relation between the risk measures of p-variation type and both the starting time and the slope at the last pillar.

1.1.4 Mean-Variance Optimal Trading under Volume Constraints

This Section has been published as a part of *Optimal starting times, stopping times and risk measures for algorithmic trading* (2014 The Journal of Investment Strategies 3) [Labadie & Lehalle, 2014]

Inverting the optimal liquidation problem putting the emphasis on observables of the obtained trading process. We will show that the TC (Target Close) algorithm can be seen as a "*reverse IS*" (Implementation Shortfall) –see equation (23) and following for details–. In this framework, the starting time for a TC is the time-reverted equivalent of the ending time for an IS. However, for practitioners this distinction is even critical : shortening the trading duration of an IS to capitalise a profitable price opportunity can always be justified, but changing the "*expected optimal start time*" for a TC can be difficult to explain.

The paper also shows that the results obtained for the TC criterion can be applied to the IS criterion because TC and IS are both sides of the same coin. Indeed, on the one hand, the TC has a pre-determined stopping time, its benchmark is the price at the end of the execution and the starting time is unknown. On the other hand, IS has a pre-determined starting time, its benchmark is the price at the beginning of the execution and the stopping time is unknown. Therefore, there is no surprise that the recursive formula for IS turns out to be almost the same that for TC but with the time running backwards.

It is customary for practitioners to put constraints on the maximum participation rate of a trading algorithms (say 20% of the volume traded by the market). Therefore, it is of paramount importance to find a systematic way of computing the starting time of a TC under a percentage of volume (PVol) constraint. Such an "*optimal trading policy under PVol constraint*" is properly defined and solved in this paper. A numerical example with real data is provided, where the optimal trading curves and their corresponding optimal starting times are computed.

Solving the TC problem under constraints allows us to analyse the impact of the parameters of the optiomisation criterion on the observable variables of the trading process. It should be straightforward for quantitative traders to implement our results numerically, i.e. to choose the characteristics of the trading process they would like to target and then infer the proper value of the parameters of the criterion they need.

Derivation of the Target Close (TC) algorithm. As in [Almgren, 2003] and [Bouchaud, 2010], we will consider a power market impact function, i.e.

$$h(v_n) = \kappa\sigma_n\tau^{1/2}\left(\frac{v_n}{V_n}\right)^\gamma , \tag{21}$$

where v_n is the amount of shares executed at the n-th pillar (i.e. at time t_n), V_n is the (historical) volume at the n-th pillar, σ_n is the normalised volatility at the n-th pillar, and κ and γ are positive constants. Under this framework, the wealth process takes the form

$$W = \sum_{n=1}^N v_n S_n + \sum_{n=1}^N \kappa\sigma_n\tau^{1/2} v_n \left(\frac{v_n}{V_n}\right)^\gamma , \tag{22}$$

where N is the number of slices in the trading algorithm. The first term in the right-hand side of (22) is the cost of executing $v^* := \sum_{i=1}^N v_i$ shares; the second term models the market impact of the execution as a power law of the percentage of volume executed at each pillar $n = 1, \ldots, N$.

For a TC algorithm, the benchmark is the closing price. Therefore, the wealth process relative to this benchmark is

$$W^\sharp = W - S_N \sum_{n=1}^N v_n \,. \tag{23}$$

Let us assume again a Brownian motion model for the asset, i.e.

$$S_{n+1} = S_n + \sigma_{n+1}\tau^{1/2}\varepsilon_{n+1} \,.$$

Under the the change of variables

$$x_n := \sum_{i=1}^n v_i \qquad \Longleftrightarrow \qquad v_n = x_n - x_{n-1}$$

the relative wealth process takes the form

$$\begin{aligned} W^\sharp &= -\sum_{n=1}^N x_n\sigma_n\tau^{1/2}\varepsilon_n + \sum_{n=1}^N \kappa\sigma_n\tau^{1/2}\frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^\gamma} \\ &= \left(-\sum_{n=1}^{N-1} x_n\sigma_{n+1}\varepsilon_{n+1} + \sum_{n=1}^N \kappa\sigma_n\frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^\gamma}\right)\tau^{1/2} \,. \end{aligned}$$

Since the time-step $\tau^{1/2}$ is a constant multiplicative factor, we can consider a normalised relative wealth

$$\tilde{W} := \frac{W^\sharp}{\tau^{1/2}} \,.$$

We are not losing any generality with the normalisation because it is equivalent to use a normalised volatility $\tilde{\sigma}_n := \sigma_n\tau^{1/2}$. Under this new framework, the average and variance of $\tilde{W}$ are, respectively,

$$\mathbb{E}(\tilde{W}) = \sum_{n=1}^N \kappa\sigma_n\frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^\gamma} \,, \qquad \mathbb{V}(\tilde{W}) = \sum_{n=1}^{N-1} x_n^2\sigma_{n+1}^2 \,.$$

The corresponding mean-variance functional is thus

$$\begin{aligned} J_\lambda(x_1, \ldots, x_N) &= \mathbb{E}(\tilde{W}) + \lambda\mathbb{V}(\tilde{W}) \\ &= \sum_{n=1}^N \kappa\sigma_n\frac{(x_n - x_{n-1})^{\gamma+1}}{V_n^\gamma} + \lambda\sum_{n=1}^{N-1} x_n^2\sigma_n^2 \,. \end{aligned}$$

The optimal trading curve is determined by solving

$$\frac{\partial J_\lambda}{\partial x_n} = 0 \,, \qquad n = 1, \ldots, N \,,$$

i.e.

$$\kappa\sigma_n(\gamma+1)\frac{(x_n - x_{n-1})^\gamma}{V_n^\gamma} - \kappa\sigma_{n+1}(\gamma+1)\frac{(x_{n+1} - x_n)^\gamma}{V_{n+1}^\gamma} + 2\lambda\mu_N(n)\sigma_{n+1}^2 x_n = 0 \,,$$

where $\mu_N(n) = 1$ if $n < N$ and $\mu_N(N) = 0$.

Returning to the variables v_n we get

$$\kappa\sigma_n(\gamma+1)\left(\frac{v_n}{V_n}\right)^\gamma - \kappa\sigma_{n+1}(\gamma+1)\left(\frac{v_{n+1}}{V_{n+1}}\right)^\gamma + 2\lambda\mu_N(n)\sigma_{n+1}^2\left(\sum_{i=1}^n v_i\right) = 0\,. \tag{24}$$

Finally, we obtain an explicit, nonlinear recursive formula of the optimal trading curve for a TC algorithm :

$$v_{n+1} = V_{n+1}\left[\frac{\sigma_n}{\sigma_{n+1}}\left(\frac{v_n}{V_n}\right)^\gamma + \frac{2\lambda\mu_N(n)}{\kappa(\gamma+1)}\sigma_{n+1}\left(\sum_{i=1}^n v_i\right)\right]^{1/\gamma}. \tag{25}$$

Derivation of the Implementation Shortfall (IS) algorithm. For an IS algorithm, the starting time is given and we have to find the optimal stopping time for our execution. Since the benchmark is the price at the moment when the execution starts, the relative wealth of an IS algorithm is

$$W^\sharp = W - S_1\sum_{n=1}^N v_n\,.$$

Using the change of variables

$$x_n := \sum_{i=n}^N v_i \quad \Longleftrightarrow \quad v_n = x_n - x_{n+1}\,,$$

it can be shown that the relative wealth process is

$$\begin{aligned} W^\sharp &= \sum_{n=2}^N x_n\sigma_n\tau^{1/2}\varepsilon_n + \sum_{n=1}^N \kappa\sigma_n\tau^{1/2}\frac{(x_n-x_{n-1})^{\gamma+1}}{V_n^\gamma} \\ &= \left(\sum_{n=2}^N x_n\sigma_n\varepsilon_n + \sum_{n=1}^N \kappa\sigma_n\frac{(x_n-x_{n-1})^{\gamma+1}}{V_n^\gamma}\right)\tau^{1/2}\,. \end{aligned}$$

As in the TC case, we can consider a normalised relative wealth

$$\tilde{W} := \frac{W^\sharp}{\tau^{1/2}}\,,$$

whose average and variance are, respectively,

$$\mathbb{E}(\tilde{W}) = \sum_{n=1}^N \kappa\sigma_n\frac{(x_n-x_{n+1})^{\gamma+1}}{V_n^\gamma}\,, \qquad \mathbb{V}(\tilde{W}) = \sum_{n=2}^N x_n^2\sigma_n^2\,.$$

In consequence, the corresponding mean-variance functional is

$$\begin{aligned} J_\lambda(x_1,\ldots,x_N) &= \mathbb{E}(\tilde{W}) + \lambda\mathbb{V}(\tilde{W}) \\ &= \sum_{n=1}^N \kappa\sigma_n\frac{(x_n-x_{n+1})^{\gamma+1}}{V_n^\gamma} + \lambda\sum_{n=2}^N x_n^2\sigma_n^2\,. \end{aligned}$$

The optimal trading curve is determined by solving

$$\frac{\partial J_\lambda}{\partial x_n} = 0\,, \qquad n = 1,\ldots,N\,,$$

i.e.

$$\kappa\sigma_n(\gamma+1)\frac{(x_n-x_{n+1})^\gamma}{V_n^\gamma} - \kappa\sigma_{n-1}(\gamma+1)\frac{(x_{n-1}-x_n)^\gamma}{V_{n+1}^\gamma} + 2\lambda\mu_1(n)\sigma_n^2 x_n = 0\,,$$

where $\mu_1(n) = 1$ if $n > 1$ and $\mu_1(1) = 0$.

Returning to the variables v_n we get

$$\kappa\sigma_n(\gamma+1)\left(\frac{v_n}{V_n}\right)^\gamma - \kappa\sigma_{n-1}(\gamma+1)\left(\frac{v_{n-1}}{V_{n+1}}\right)^\gamma + 2\lambda\mu_1(n)\sigma_n^2\left(\sum_{i=n}^N v_i\right) = 0\,.$$

We thus obtain the recursive nonlinear formula for the optimal IS trading curve :

$$v_{n-1} = V_{n-1}\left[\frac{\sigma_n}{\sigma_{n-1}}\left(\frac{v_n}{V_n}\right)^\gamma + \frac{2\lambda\mu_1(n)}{\kappa(\gamma+1)}\frac{\sigma_n^2}{\sigma_{n-1}}\left(\sum_{i=n}^N v_i\right)\right]^{1/\gamma}. \tag{26}$$

Comparison between TC and IS. If the volatility σ_n is constant then (26) is exactly the same recursion as in (25), except for the time in IS *running backwards*. In the light of this *mirror* property, the analysis we will be performing for TC can be naturally extrapolated for IS, esp. adding a maximum participation rate constraint and computing the starting time.

If the volatility is not constant then there is a slight difference in the formula, namely a factor for $\sigma_{n+1} = \sigma_{n+1}^2/\sigma_{n+1}$ for TC and a factor σ_n^2/σ_{n-1} for IS. From the practitioner point of view, this difference is paramount. For TC it is the *forward* volatility σ_{n+1} which determines the weight of $\sum_{i=n}^{N} v_i$ i.e. the shares already executed. This comes from the fact that the ending time is fixed (it is the market close, in fact) and the number of shares to trade per pillar increases in time. Therefore, not only the trader has little room for changing her schedule, but also she has to anticipate the volatility one step ahead in order to avoid nasty surprises. For IS it is the *spot* volatility σ_n that counts. This is because the number of shares decreases in time, which gives more room to change the trading schedule the closer the trader is to the end of her execution, giving room for capitalising potential arbitrage opportunities.

Adding constraints : Percentage of Volume (PVol). The TC algorithm can have a constraint of , meaning that the size of each slice cannot exceed a fixed percentage of the current available volume. This algorithm is called Percentage of Volume (PVol). Under a PVol constraint, the trading slices v_n of the TC algorithm satisfy the constraint

$$v_n \leq qV_n\,, \qquad q \in (0,1)\,.$$

It is worth to notice that the PVol algorithm is not a solution of the Almgren-Chriss optimisation. Indeed, if it were then

$$\frac{v_n}{V_n} = q \quad \forall n = 1, \ldots, N\,,$$

and from (24) we would have that

$$\sum_{i=1}^{n} v_i = 0\,.$$

The two algorithms TC and PVol are mutually exclusive. This implies that a classical optimisation scheme of TC with the PVol constraint via Lagrange multipliers is not straightforward, to say the least. We thus have to find another way to obtain a solution of the TC algorithm under the PVol constraint.

From (25) we see that given v_n, the corresponding v_{n+1} depends on $\sum_{i=1}^{n} v_i$, i.e. the cumulative execution up to n, which implies that curve is in general increasing. Therefore, in order to satisfy the constraint of maximum percentage of volume (PVol), if the total volume to execute is large then the algorithm has to be divided into two patterns :

1. As long as the constraint of maximum participation rate (PVol) is not reached, we execute the slices according to the Almgren-Chriss recursive formula. This corresponds to the TC pattern.
2. As soon as the PVol constraint is attained, the algorithm executes the minimum between the TC curve and PVol curve.

Loosely speaking, we start with a TC algorithm, but once the slices are *saturated* we switch to a PVol algorithm until the end of the execution. However, it can happen that the algo switches back to TC if the PVol curve is bigger at a further pillar ; this situation is exceptional though, save for cases where the volume curve presents sharp peaks or gaps.

It is worth to mention that adding a PVol constraint to IS is the same as adding the constraint for TC and running the TC algorithm backwards.

Computing the optimal stopping time for TC. Let us describe in detail all the steps of our TC algorithm under PVol constraint. Let n_0 be the starting time and n_1 the switching time (i.e. when we change from TC to PVol).

1. According to the historical estimates of the available volume at the close auction, plus the desired participation rate, we define the execution at the pillar $n = 103$ (the close auction), denoted $v^\sharp$.
2. We compute the Almgren-Chriss algorithm for the residual shares $v^\flat = v^* - v^\sharp$ i.e. the shares to execute in continuous, outside of the close auction. We start with $n = 1$ and $n_1 = 102$ and launch the TC recursive argument (25). Since the algorithm is completely determined by $v_1 = \alpha$, it suffices to find the right α such that the cumulative shares at $n_1 = 102$ are equal to $v^\flat$.
3. We compare the trading curve of the previous step with the PVol curve. If the PVol constraint is satisfied then we are done. If not, we *saturate* pillar $n = 102$ with the PVol constraint and redefine the parameters : $v^\flat$ is now the shares to execute outside both pillars 103 and 102, i.e. $v^\flat = v^\flat - v_{102}$. We set $n_1 = 101$, i.e. $n_1 = n_1 - 1$ and repeat the procedure.
4. Eventually, we will obtain a TC curve starting at $n = 1$, switching to PVol at $n_1 \leq 102$ and satisfying the PVol constraint. Moreover, the algorithm finds the right number of shares $\alpha = \alpha_1$ to trade at pillar $n = 1$ such that the total execution from $n = 1$ to $n = 103$ is equal to v^*. Remark that the whole algorithm executes TC between $n = 1$ to n_1, PVol between n_1 and $n = 102$, and the desired participation at the close auction at $n = 103$.

5. In order to find the right starting time n_0, we choose a minimal trading size for each slice, that we denote α_{min}. For example, α_{min} can be either the average trading size or a percentage of the available liquidity. If $\alpha_1 < \alpha_{min}$ then we advance one pillar and we recompute the trading curve, which now starts at $n = 2$ and has a first slice of size α_2. We continue until we find the first pillar n_0 such that $\alpha = \alpha_{n_0} \geq \alpha_{min}$.

Therefore, n_1 is determined by the PVol constraint whilst n_0 is determined by the minimal trading size constraint α_{min}. Notice however that the optimal starting pillar n_0 is determined after n_1, which implies that n_0 depends not only on α_0 but also on the rest of the parameters like the PVol curve, the participation rate at the close auction and the market impact parameters.

Observe that there is a systematic way of computing the stopping pillar for an IS algorithm : it corresponds to the *backwards or symmetrical image* of the starting time we computed for the TC algorithm.

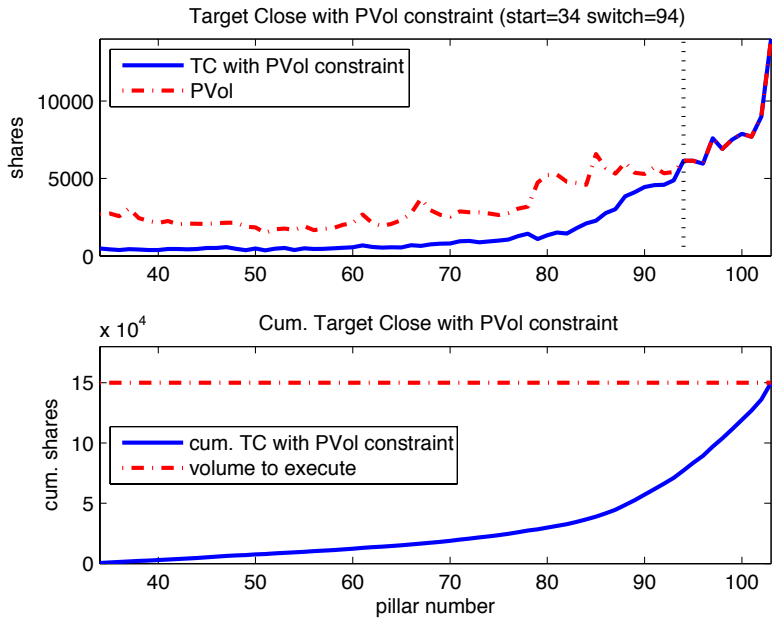

FIGURE 8 – TC (above) and cumulative TC curves under PVol constraint. The starting and switching pillars are resp. $n_0 = 34$ and $n_1 = 94$.

Numerical results. In the first plot of Figure 8 we have the TC curve (solid line) under PVol constraint vs the PVol curve (broken line) of stock AIRP.PA (Air Liquide). In the second plot we have the cumulative execution of the TC curve (solid line) under PVol constraint vs the volume to execute v^* (broken line). The parameters we used are $v^* = 150,000$ shares, $\alpha_{min} = 500$ shares and a maximum participation of 20% in both the continuous trading period and the close auction. The historical volatility and volume curves, the market impact parameters κ and γ and the risk-aversion coefficient λ were provided by the Quantitative Research at Cheuvreux - Crédit Agricole.

In Figure 8 we can also observe that the algorithm finds the optimal starting time at pillar $n = 34$ (beginning of the horizontal axis), at which it starts to execute the order following the TC algorithm based on the Almgren-Chriss optimisation. At pillar $n = 94$ (vertical line) the algorithm switches to PVol in order to satisfy the constraint. Notice that during the whole execution, the PVol constraint has been satisfied, and that the number of traded shares per pillar is increasing in time. Finally, in the second plot we can see that the TC algorithm under PVol constraint successfully executed the whole order.

1.2 Going Further : Impulse Control for Optimal Liquidation by Robots

Stochastic control is the next expected step after mean variance optimization along discretized trajectories. The work presented here is with Pham [Kharroubi & Pham, 2010] one of the first proposals in this direction.

The unique added values of this work is to answer to two practical questions :

- how to choose trading time during an optimal trading scheme ? Using impulse control is the only way to answer. The answers we provided have not really be improved yet.
- how to render the realistic aspect of trading ? In real world algorithmic traders rely on two decision layers :
  - one being "strategic", planning the trading at an horizon from few hours to few days
  - one being "tactic", very close to orderbook dynamics and liquidity seeking.

  The framework presented here allows to model this useful trading architecture.

Moreover, one should note the "tactical layer", modeled here by a "trading robot", is modelled in the work presented in this section by a "black box". It is perfectly suited to be modelled by a stochastic algorithm like the ones presented in section 3.1 or 1.3.

This remark sheds light on a possible articulation of the work presented here :

- the section on microstructure, and in particular Section 2.1 allows to understand orderbooks and liquidity dynamics, and to isolate some "invariants" (in the sense that they are Markovian, and if possible with transitions independent of time)
- The section on statistical learning allows to build "optimal" updating rules (i.e. learning robots) at a very small time scale, relying on the upper invariants.
- last but not least, this section arrange the optimal use of such learning robots along an optimal trajectory (in the sens of stochastic control).

It is worthwhile to note one of the unique aspect of the work presented thereafter is the capability offered by mathematical tools to interact with the price formation process, in a way practitioners can use them and hence improve the efficiency and the risk control of their practices.

This Section has been published as a part of *Optimal control of trading algorithms : a general impulse control approach* (2011 SIAM J. Financial Mathematics 2, 404–438) [Bouchard *et al.* , 2011]

1.2.1 Problem formulation

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting a d-dimensional Brownian motion W, $d \geq 1$. Let $\mathbb{F} := (\mathcal{F}_t)_{t\geq 0}$ denote the right-continuous complete filtration generated by W, and let $T > 0$ be a finite time horizon.

Control policies We first describe how the trading algorithms are controlled, precise dynamics will be imposed in Section 1.2.1 below. Note that different algorithms can be viewed as a single parameterized one. In what follow, we therefore consider that we have only one algorithm.

A control policy of the trading algorithm is described by a non-decreasing sequence of stopping times $(\tau_i)_{i\geq 1}$ and a sequence of $[\underline{\delta}, \infty) \times E$-valued random variables $(\delta_i, \mathcal{E}_i)_{i\geq 1}$. The stopping times τ_i describe the times at which an order is given to the algorithm, $\mathcal{E}_i$ is the value of the parameters with which the algorithm is run and δ_i the length of the period (duration of the trading for this "robot") during which it is run with the value $\mathcal{E}_i$. The set E is a compact subset of $\mathbb{R}^d$, which represents the possible values of the parameters, the quantity

$$0 < \underline{\delta} < T$$

denotes the minimum length of the time period during which the algorithm can be run. To be consistent we impose that

$$\tau_i + \delta_i \leq \tau_{i+1} \text{ and } (\delta_i > 0 \Rightarrow \tau_i + \delta_i \leq T) \ , \ i \geq 1 \ . \tag{27}$$

The first condition expresses the fact that a new order can not be given before the end of the time period associated to the previous order. The second one means that an order should be given only if it ends before the final time horizon T.

Remark 1 1. The minimal duration constraint, $\delta_i \geq \underline{\delta}$ with $\underline{\delta} > 0$, has been justified from a practical point of view in the introduction. From the mathematical point of view, the problem described in Sections 1.2.1 and 1.2.1 below would not make sense without this condition, if no additional cost related to launching the algorithm with new parameters is introduced. Indeed, for $\underline{\delta} = 0$, and without additional costs, the controller could, at the limit, control the parameters continuously, and this would actually certainly be optimal. The controller will then act, at the limit, as a trader acting continuously on the market.

2. Models with $\underline{\delta} = 0$ and with an additional cost (paid each time the algorithm is launched with new parameters) could be discussed by following the lines of this paper. Such a cost is actually already embedded in the general dynamics of Section 1.2.1, up to an additional assumption on the function β, see example below. This would however require to justify this cost, and to evaluate it in practice. Moreover, certain bounds like the one stated in Remark 2 below

would not be true anymore and other conditions would be required. For sake of simplicity, we therefore stick to the case $\underline{\delta} > 0$ which corresponds more, from our point of view, to practical situations.

3. In the absence of a cost penalizing frequent changes in the parameters, it may be optimal to choose $\delta_i = \underline{\delta}$ most of the time (depending whether $\underline{\delta}$ is small or not). However, other values of δ_i may also be optimal. Our algorithm provides a way to select the maximal optimal value, the one for which the trader changes the parameters the least often (which is desirable in practice).

As usual the value of the parameters and the size of the durations can not be chosen in some anticipative way, i.e. we impose that

$$(\delta_i, \mathcal{E}_i) \text{ is } \mathcal{F}_{\tau_i}\text{-measurable }, \ i \geq 1 \ . \tag{28}$$

At time $t \in [\tau_i, \tau_i + \delta_i)$, the value of the parameter of the trading algorithm is denoted by ν_t. For $t \in A((\tau_i, \delta_i)_{i\geq 1})$, defined as

$$A((\tau_i, \delta_i)_{i\geq 1}) := \mathbb{R}_+ \setminus \left(\bigcup_{i\geq 1} [\tau_i, \tau_i + \delta_i) \right) \ ,$$

we set $\nu_t = \varpi$, where $\varpi \in \mathbb{R}^d \setminus E$ can be viewed as a cemetery point, recall that E is compact.

It follows that the value of the parameters of the trading algorithm ν can be written as

$$\nu_t = \varpi \mathbf{1}_{t \in A((\tau_i, \delta_i)_{i\geq 1})} + \sum_{i\geq 1} \mathcal{E}_i \mathbf{1}_{t \in [\tau_i, \tau_i + \delta_i)} \ , \ t \in [0, T] \ , \tag{29}$$

where $\nu_t = \varpi$ means that the algorithm is not running at time t.

In the following, we denote by $\mathcal{S}$ the set of adapted processes ν that can be written in the form (29) for some sequence of stopping times $(\tau_i)_{i\geq 1}$ and of $[\underline{\delta}, \infty) \times E$-valued random variables $(\delta_i, \mathcal{E}_i)_{i\geq 1}$ satisfying (27) and (28).

For ease of notations, we shall now write

$$(\tau_i^\nu, \delta_i^\nu, \mathcal{E}_i^\nu)_{i\geq 1} \ \text{ the sequence associated to } \nu \in \mathcal{S},$$

and define, for all stopping times ϑ_1 and ϑ_2 satisfying $\vartheta_1 \leq \vartheta_2$ $\mathbb{P}$-a.s., the set of indices corresponding to orders whose execution ends between ϑ_1 and ϑ_2 :

$$\mathbb{I}^\nu_{\vartheta_1, \vartheta_2} := \{ i \geq 1 : \vartheta_1 < \tau_i^\nu + \delta_i^\nu \leq \vartheta_2 \} \ .$$

Remark 2 *Note that the constraint $\delta_i^\nu \geq \underline{\delta}$ for all $i \geq 1$ and $\nu \in \mathcal{S}$ implies that*

$$\text{card}(\mathbb{I}^\nu_{0,T}) \leq \text{card}\left(\{ \tau_i^\nu \leq T, \ i \geq 1 \} \right) \leq T/\underline{\delta} \ .$$

For ease of notations, we also set

$$\bar{E} := E \cup \{\varpi\} \subset \mathbb{R}^d \ ,$$

and introduce the process $\Delta^\nu_\cdot$

$$\Delta_t^\nu := \sum_{i\geq 1} [\tau_i^\nu + \delta_i^\nu - t]^+ \mathbf{1}_{t \geq \tau_i^\nu} \ , \ t \in [0, T] \ .$$

The quantity Δ_t^ν denotes the remaining duration during which no new order can be passed to the algorithm. When $\Delta_t^\nu > 0$, the algorithm is running with a value of the parameters ν_t. When $\Delta_t^\nu = 0$, the algorithm is not running anymore, $\nu_t = \varpi$, and a new order can be passed.

The following picture sums up the dynamics of the control.

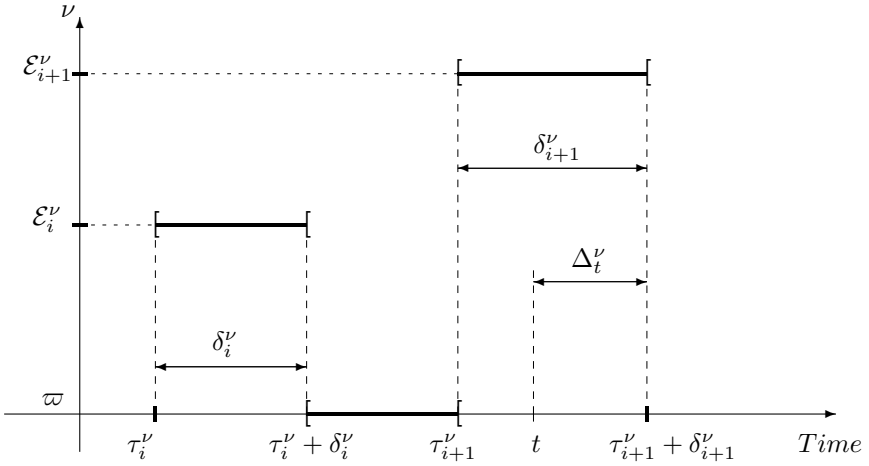

Output of the trading algorithm. Given some initial data $(t,x) \in [0,T] \times \mathbb{R}^d$, the output of the trading algorithm associated to some control policy $\nu \in \mathcal{S}$ is defined as the strong solution $X^\nu_{t,x}$ on $[0,T]$ of the stochastic differential equation

$$X^\nu_{t,x}(s) = x + \mathbf{1}_{s \geq t} \left(\int_t^s b(X^\nu_{t,x}(r), \nu_r) dr + \int_t^s a(X^\nu_{t,x}(r), \nu_r) dW_r + \sum_{i \geq 1} \beta(X^\nu_{t,x}(\tau_i^\nu -), \mathcal{E}_i^\nu, \delta_i^\nu) \mathbf{1}_{t < \tau_i^\nu \leq s} \right) , \tag{30}$$

where $\beta : \mathbb{R}^d \times \bar{E} \times [\underline{\delta}, T] \to \mathbb{R}^d$ and $(b,a) : \mathbb{R}^d \times \bar{E} \to \mathbb{R}^d \times \mathbb{M}^d$ are continuous functions such that there exists $K > 0$ for which, for all $x, x' \in \mathbb{R}^d$, $e, e' \in \bar{E}$, $\delta, \delta' \in [\underline{\delta}, T]$,

$$\left\{ \begin{array}{rcll} |\psi(x,e,\delta) - \psi(x',e,\delta)|) & \leq & K|x-x'| & \\ |\psi(x,e,\delta)| & \leq & K(1+|x|) & \text{for } \psi = b, a, \beta \, . \\ |\psi(x,e,\delta) - \psi(x,e',\delta')| & \leq & K(1+|x|)(|e-e'| + |\delta - \delta'|) & \end{array} \right. \tag{31}$$

We do not differentiate here between the components that correspond to real outputs of the algorithm (cumulated gains, cumulated volumes executed by the algorithm, etc...) and others that simply describe the evolution of financial data or market factors (prices of the traded assets, global traded volumes on the markets, volatilities, etc...).

The jumps on the dynamics are introduced to model the changes in the initial conditions on the variables of interest for the trading algorithm when it is launched (e.g. volume to be executed between τ_i^ν and $\tau_i^\nu + \delta_i^\nu$).

Moreover, there is no loss of generality in assuming that X, W and ν have the same dimension d. One can always reduce more general situations to this case by playing with the coefficients b, a, β and with the choice of E. Time dependent coefficients can similarly be considered by putting the first line of a and β equal to 0, and the first component of b equal to 1, so that the first component of X actually coincides with the time parameter.

We refer to Section 1.2.2 for the description of simple examples of application which illustrate the flexibility of the above model.

Gain function. The aim of the controller is to maximize the expected value of the gain functional

$$\nu \in \mathcal{S} \mapsto \Pi_{t,x}(\nu) := g(X^\nu_{t,x}(T)) + \sum_{i \in \mathbb{I}^\nu_{t,T}} f(X^\nu_{t,x}(\tau_i^\nu + \delta_i^\nu -), \mathcal{E}_i^\nu) \, ,$$

with the usual convention $\sum_\emptyset = 0$, among the set

$$\mathcal{S}_{t,\delta,e} := \left\{ \begin{array}{ll} \big\{ \nu \in \mathcal{S} \ : \ \nu_s = e \text{ for } s \in [t, t+\delta) \text{ and } \Delta^\nu_{t+\delta} = 0 \big\} & \text{if } e \neq \varpi \text{ and } \delta > 0 \\ \{ \nu \in \mathcal{S} \ : \ \nu_t = \varpi \} & \text{otherwise} \end{array} \right.$$

where $(\delta, e) \in \mathbb{R}_+ \times \bar{E}$ denotes the initial state of the remaining duration and value of the parameters.

In this section (cf. [Bouchard *et al.* , 2011] for all the details), we provide a rigorous formulation which follows ideas introduced in [Bouchard & Touzi, 2011]. Namely, we only provide a weak formulation in terms of test functions. The main advantage of this approach is that it does not require any regularity on the value function V itself, but only some lower-semicontinuity of the objective function $J(\cdot;\nu)$, see below. We refer to [Bouchard & Touzi, 2011] for a general discussion.

Lemma 1 (Weak Dynamic Programming Principle) *Fix $(t,x,\delta,e) \in D$ and let $\{\vartheta^\nu, \ \nu \in \mathcal{S}^a_{t,\delta,e}\}$ be a family of $[t,T]$-valued stopping times independent of $\mathcal{F}_t$. Then, we have*

$$V(t,x,\delta,e) \leq \sup_{\nu \in \mathcal{S}^a_{t,\delta,e}} \mathbb{E} \left[[V^*, g](\vartheta^\nu, X^\nu_{t,x}(\vartheta^\nu), \Delta^\nu_{\vartheta^\nu}, \nu_{\vartheta^\nu}) + \sum_{i \in \mathbb{I}^\nu_{t,\vartheta^\nu}} f(X^\nu_{t,x}(\tau_i^\nu + \delta_i^\nu -), \mathcal{E}_i^\nu) \right] , \tag{32}$$

where $[V^*, g](s,\cdot) := V^*(s,\cdot)\mathbf{1}_{s<T} + g\mathbf{1}_{s=T}$, and

$$\sup_{\nu \in \mathcal{S}^a_{t,\delta,e}} \mathbb{E} \left[\varphi(\vartheta^\nu, X^\nu_{t,x}(\vartheta^\nu), \Delta^\nu_{\vartheta^\nu}, \nu_{\vartheta^\nu}) + \sum_{i \in \mathbb{I}^\nu_{t,\vartheta^\nu}} f(X^\nu_{t,x}(\tau_i^\nu + \delta_i^\nu -), \mathcal{E}_i^\nu) \right] \leq V(t,x,\delta,e) \tag{33}$$

for all upper semi-continuous function φ such that $V \geq \varphi$ on $\bar{D}$.

As in [Bouchard & Touzi, 2011], the proof of the above result relies on some lower-semicontinuity property of the function J.

1.2.2 An Illustration

A Simple Example. Consider the case where the aim of the controller is to buy a number Q_0 of one stock S between 0 and $T > 0$. The dynamics of (S, V) is given as in the previous example.

Here $\mathcal{E}_i^\nu$ stands for the intensity at which the stocks are bought, i.e. the algorithm buys a number $\mathcal{E}_i^\nu dt = \nu_t \mathbf{1}_{\nu_t \neq \varpi} dt$ of stocks on $[t, t+dt]$, $t \in [\tau_i^\nu, \tau_i^\nu + \delta_i^\nu)$. The dynamics of the remaining number of stocks to be bought before T is thus given by :

$$Q_t^\nu = Q_0 - \int_0^t q(\nu_s) ds$$

where q is now defined as $q(e) = e\mathbf{1}_{e \neq \varpi}$. It follows that the cumulated wealth's dynamic is

$$Y_t^\nu = 0 + \int_0^t \tilde{S}_r q(\nu_r) dr = 0 + \int_0^t (S_r + \eta(\nu_r, S_r, V_r)) q(\nu_r) dr \ ,$$

where η is a given market impact function.

If the number Q_0 of shares is not liquidated at time T, the remaining part Q_T^ν is instantaneously bought on the market at the price $S_T + c(Q_T^\nu, S_T, V_T)$, for some Lipschitz continuous function c.

The total cost after the final transaction is thus given by :

$$Y_T^\nu + (S_T + c(Q_T^\nu, S_T, V_T)) (Q_T^\nu)^+ \ .$$

The aim of the controller is to minimize the expectation of the quantity

$$\ell(Y_T^\nu + (S_T + c(Q_T^\nu, S_T, V_T)) (Q_T^\nu)^+)$$

for some convex function ℓ with polynomial growth.

Numerics. We consider the following set of parameters. The trading period corresponds to a period of 3 hours. The price process is assumed to follow a Black and Scholes dynamics with zero drift $S_t = S_0 e^{-\frac{1}{2}\sigma^2 t + \sigma W_t}$, where $S_0 := 13$ and the annualized volatility is 25%. Adding a drift would only change the optimal strategy in an obvious manner, depending on its sign. We assume a deterministic evolution of the instantaneous volume traded on the market $(V_t)_{t \leq T}$ as given below. It corresponds to an intensity in minutes. The impact function η is given by $\eta(e, v) = 0.4(e/v)^{1.1}$. This coincides with plausible calibrated data. We take $\underline{\delta} = 5$ minutes. For this numerical test, we restrict to values of the duration in the set $5, 10, \ldots, 60$ minutes. The different values of the buying rate are 50, 100, 150, ..., 500. It correspond to numbers of bought stocks per minute.

The final cost is given by $c(Q, v) = \eta(Q/(0.417v))$, which implicitly means that the trader has 25 seconds to finalize the operation, i.e. he must buy Q in 0.417 minutes at a rate $Q/0.417$.

We consider two different types of functions ℓ : either ℓ is the identity, $\ell(r) = r$, or ℓ is of exponential type, $\ell(r) = e^{10^{-5}r} \wedge 100$. The value 100 corresponds to more than four times the cost evaluated with the exponential function $r \mapsto e^{10^{-5}r}$ of the operation which consists in buying 15000 stocks with a constant rate, assuming that the market volume takes the minimal value corresponding to the U-shaped path defined below, and for a constant stock price equal to 26. This is an extreme scenario. However, truncating the exponential function is needed in order to ensure that the value function is finite, since a log normal distribution does not admit exponential moments.

We first consider the case where ℓ is linear. In this case, the controller is risk neutral so that he has no incentive to buy the stock quickly because of a risk of increase of the price (recall that here the price is a martingale). In Figure 9 and Figure 10, we compare the case where the market volume is constant $V_t = 50000$, on the left, to the case where the market volume is strongly U-shaped : $V_t = 50000(1.1 - 0.9\sin(\pi t/T))$ with $T = 180$ minutes. This volume is also given per minutes. Both figures provide the optimal buying rate in terms of the remaining time $T - t$ and the remaining quantity to buy Q_t, for $S_t = S_0$. A typical path has to be read from north-west to south-east, since Q decreases as time goes buy. As expected, when the path of the market volume is U-shaped the optimal rate strongly decreases in the middle of the period, when the market volume is low and the impact on the price of the stock is high. This is compensated by a higher rate at the beginning of the period.

In Figure 13 and Figure 14, we provide the maximal value of the optimal duration δ. Recall that the existence of multiple optima is possible since there is no additional cost related to the launching of a new slice, see Remark 1. In most cases, the maximal value is strictly above the minimal threshold of 5 minutes. This support our choice of considering the duration as a control, even in the absence of a cost associated to a change of parameters : when a duration of, e.g., 15 minutes is optimal, the trader can let the algorithm run for this time period without having to launch it again every 5 minutes, and thus take care of his other positions.

In the case where the market volume is strongly U-shaped, it is smaller at the beginning of the trading period, in comparison to the constant volume case. This is due to the fact that the algorithm knows that the market volume is going to decrease strongly (since it is deterministic) and that he will need to reduce the buying rate. It is small near the terminal time because of the constraint $t + \delta \leq T$. When looking at the picture backward in time, i.e. as $T - t$ increases, we see that the maximal value first increases and then drops down very quickly. The first phenomena is due

to the fact that the buying rate is essentially kept constant at its maximal value near the terminal time. Then, this rate decreases as the algorithm has more time to buy the shares. The period during which the algorithm reduces the buying rate naturally coincides with a lower duration.

We next consider the case where ℓ is of exponential type. The optimal buying rates are reported in Figures 11-12. Because the controller is now risk adverse, he has an incentive to buy the stocks more quickly in order to avoid an increase of the price. This can be seen by comparing Figures 11-12 with Figures 9-10. However, we do not observe significant changes in the maximal optimal duration.

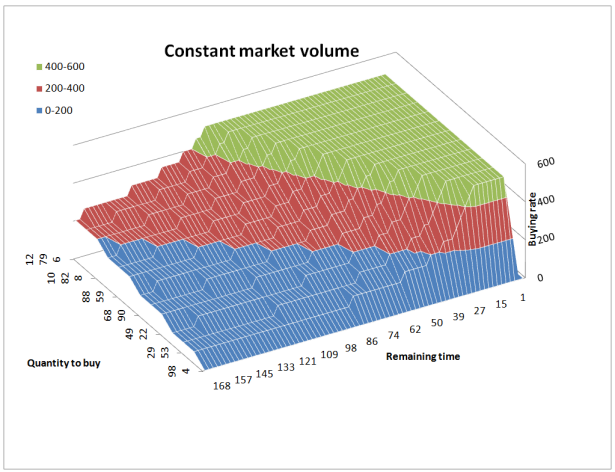

FIGURE 9 – Buying rate - flat volume - Linear cost

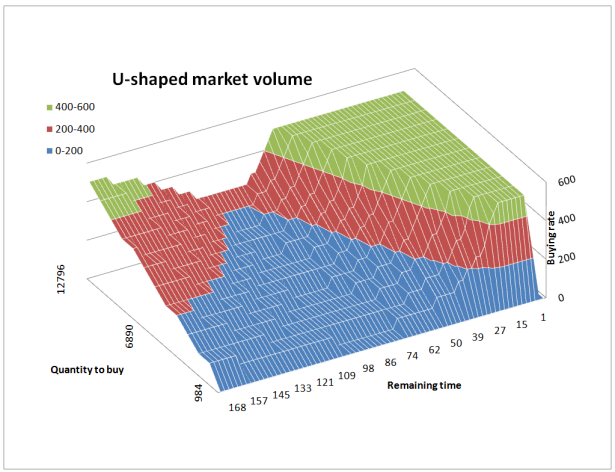

FIGURE 10 – Buying rate - U-shaped volume - Linear cost

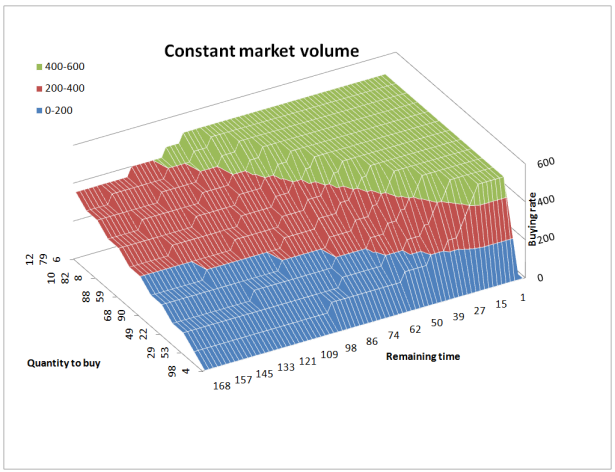

FIGURE 11 – Buying rate - flat volume - Exponential cost

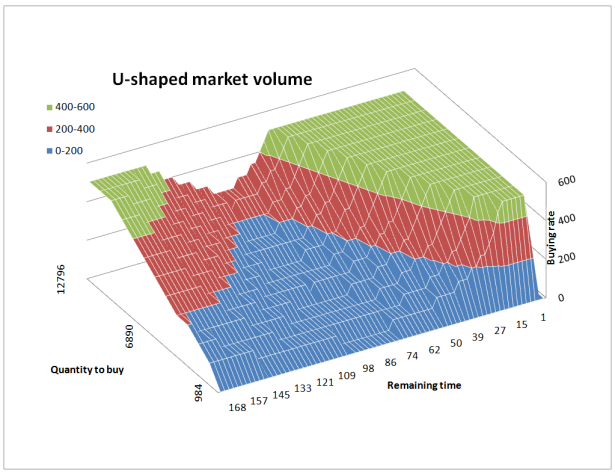

FIGURE 12 – Buying rate - U-shaped volume - Exponential cost

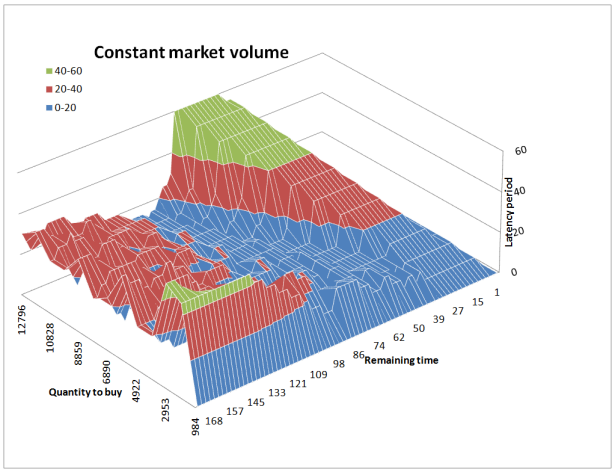

FIGURE 13 – Latency period - flat volume - Linear cost

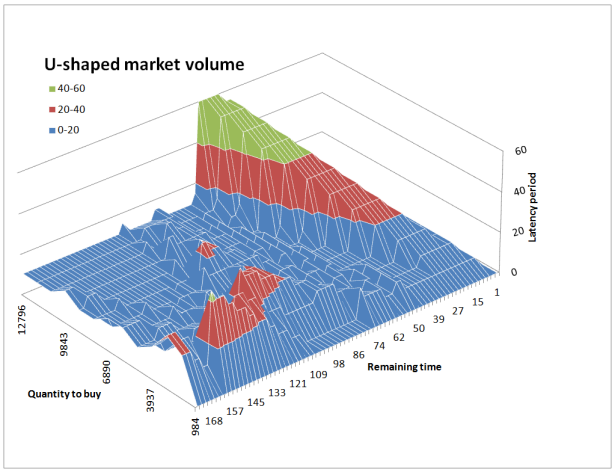

FIGURE 14 – Latency period - U-shaped volume - Linear cost

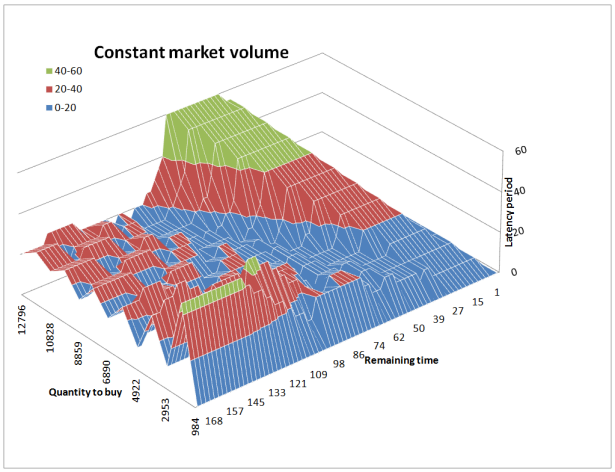

FIGURE 15 – Latency period - flat volume - Exponential cost

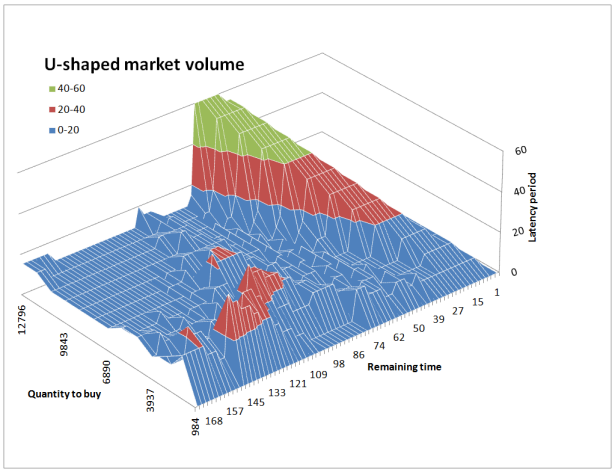

FIGURE 16 – Latency period - U-shaped volume - Exponential cost

1.3 Controlling the Trading Process by Prices

This section belong to the most covered viewpoint on optimal trading by applied mathematics : the use of stochastic control to deduce optimal trading strategies from a model on which the dynamic programming principle applies.

In these kind of article, all is modelled a way suited to obtain easy to write PDEs : time is continuous, point processes intensities have an exponential dependence to controls, etc.

The two works presented here share the same underlying logic :

- controlling the trading using "contribution prices", like a market maker would do,
- a model of the probability to be "hit" at a given distance δ of the midpoint (or other any reference price), follows a point process with an exponential intensity $A \exp -k\delta$, following [Avellaneda & Stoikov, 2008]
- the capability to include a trend to the underlying process, but no dependence between the trend and a potential asymmetry of the intensities of incoming selling and buying orders.

This work has to be put in perspective with other optimal trading frameworks. Like [Bayraktar & Ludkovski, 2012] with a less generic utility function than ours. [Kharroubi *et al.* , 2013, Guilbaud & Pham, 2011] cover a lot of different aspects of optimal trading. A recent and very interesting stream of research targets to exploit exogenous information ; [Cartea *et al.* , 2011] is a typical reference. Optimal trading of options is not covered here but a very good survey Soner [Gökay *et al.* , 2011] gives good references, just cite [Li & Almgren, 2014, Guéant & Pu, 2013, Stoikov & Saglam, 2009, El Aoud & Abergel, 2014] as entry points, the latter being very interesting in the sense it mixes options trading and hedging on the underlying as a coupled optimal market making problem.

1.3.1 Optimal Market Making

This Section has been published as a part of *Dealing with the inventory risk : a solution to the market making problem* (2013 Mathematics and Financial Economics 4, 477–507) [Guéant *et al.* , 2013]

Setup of the model. Let us fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. We assume that all random variables and stochastic processes are defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$.

We consider a high-frequency market maker operating on a single stock We suppose that the mid-price of this stock or more generally a reference price of the stock moves as an arithmetic Brownian motion

$$dS_t = \sigma dW_t$$

The market maker under consideration will continuously propose bid and ask prices denoted respectively S_t^b and S_t^a and will hence buy and sell shares according to the rate of arrival of market orders at the quoted prices. His inventory q, that is the (signed) quantity of shares he holds, is given by

$$q_t = N_t^b - N_t^a$$

where N^b and N^a are the point processes (independent of $(W_t)_t$) giving the number of shares the market maker respectively bought and sold (we assume that transactions are of constant size, scaled to 1). Arrival rates obviously depend on the prices S_t^b and S_t^a quoted by the market maker and we assume, in accordance with the model proposed by Avellaneda and Stoikov [Avellaneda & Stoikov, 2008] , that intensities λ^b and λ^a associated respectively to N^b and N^a depend on the difference between the quoted prices and the reference price (*i.e.* $\delta_t^b = S_t - S_t^b$ and $\delta_t^a = S_t^a - S_t$) and are of the following form :

$$\lambda^b(\delta^b) = Ae^{-k\delta^b} = A\exp(-k(s - s^b)) \qquad \lambda^a(\delta^a) = Ae^{-k\delta^a} = A\exp(-k(s^a - s))$$

where A and k are positive constants that characterize the liquidity of the stock. In particular, this specification means – for positive δ^b and δ^a – that the closer to the reference price an order is posted, the faster it will be executed.

As a consequence of his trades, the market maker has an amount of cash evolving according to the following dynamics :

$$dX_t = (S_t + \delta_t^a)dN_t^a - (S_t - \delta_t^b)dN_t^b$$

To this original setting introduced by Avellaneda and Stoikov ([Avellaneda & Stoikov, 2008], itself following partially Ho and Stoll [Ho & Stoll, 1981]), we add a bound Q to the inventory that a market maker is authorized to have. In other words, we assume that a market maker with inventory Q ($Q > 0$ depending in practice on risk limits) will never set a bid quote and symmetrically that a market maker with inventory $-Q$, that is a short position of Q shares in the stock under consideration, will never set an ask quote. This realistic restriction may be read as a risk limit and allows to solve rigorously the problem.

Now, coming to the objective function, the market maker has a time horizon T and his goal is to optimize the expected utility of his P&L at time T. We will focus on CARA utility functions and we suppose that the market maker optimizes :

$$\sup_{(\delta_t^a)_t, (\delta_t^b)_t \in \mathcal{A}} \mathbb{E}\left[-\exp\left(-\gamma(X_T + q_T S_T)\right)\right]$$

where $\mathcal{A}$ is the set of predictable processes bounded from below, γ is the absolute risk aversion coefficient characterizing the market maker, X_T is the amount of cash at time T and $q_T S_T$ is the evaluation of the (signed) remaining quantity of shares in the inventory at time T (liquidation at the reference price S_T

Characterization of the optimal quotes. The optimization problem set up in the preceding section can be solved using the classical tools of stochastic optimal control. The first step of our reasoning is therefore to introduce the Hamilton-Jacobi-Bellman (HJB) equation associated to the problem. More exactly, we introduce a system of Hamilton-Jacobi-Bellman partial differential equations which mainly consists of the following equation (for $|q| < Q$, other cases can be found in [Guéant *et al.* , 2013]) indexed by $q \in \{-Q, \dots, Q\}$ for $(t, s, x) \in [0, T] \times \mathbb{R}^2$:

$$\begin{aligned}
&\partial_t u(t,x,q,s) + \frac{1}{2}\sigma^2 \partial^2_{ss} u(t,x,q,s) \\
&+ \sup_{\delta^b} \lambda^b(\delta^b) \left[u(t, x - s + \delta^b, q+1, s) - u(t,x,q,s)\right] \\
&+ \sup_{\delta^a} \lambda^a(\delta^a) \left[u(t, x + s + \delta^a, q-1, s) - u(t,x,q,s)\right] = 0.
\end{aligned}$$

To solve these equations we will use a change of variables based on two different ideas. First, the choice of a CARA utility function allows to factor out the Mark-to-Market value of the portfolio $(x + qs)$. Then, the exponential decay for the intensity functions λ^b and λ^a allows to reduce the Hamilton-Jacobi-Bellman (HJB) equations associated to our control problem to a linear system of ordinary differential equations :

Proposition 1 (Change of variables for (HJB)) *Let us consider a family $(v_q)_{|q|\leq Q}$ of positive functions solution of :*

$$\forall q \in \{-Q+1, \dots, Q-1\}, \quad \dot{v}_q(t) = \alpha q^2 v_q(t) - \eta \left(v_{q-1}(t) + v_{q+1}(t)\right)$$

$$\text{and on the borders : } \dot{v}_Q(t) = \alpha Q^2 v_Q(t) - \eta v_{Q-1}(t),\ \dot{v}_{-Q}(t) = \alpha Q^2 v_{-Q}(t) - \eta v_{-Q+1}(t)$$

with $\forall q \in \{-Q, \dots, Q\}, v_q(T) = 1$, where $\alpha = \frac{k}{2}\gamma\sigma^2$ and $\eta = A(1 + \frac{\gamma}{k})^{-(1+\frac{k}{\gamma})}$.

Then, $u(t,x,q,s) = -\exp(-\gamma(x+qs))v_q(t)^{-\frac{\gamma}{k}}$ is solution of (HJB).

Then, the following proposition proves that there exists such a family of positive functions :

Proposition 2 (Solution of the ordinary differential equations) *Let us introduce the matrix M defined by :*

$$M = \begin{pmatrix}
\alpha Q^2 & -\eta & 0 & \cdots & \cdots & \cdots & 0 \\
-\eta & \alpha(Q-1)^2 & -\eta & 0 & \ddots & \ddots & \vdots \\
0 & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
\vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
\vdots & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\
\vdots & \ddots & \ddots & 0 & -\eta & \alpha(Q-1)^2 & -\eta \\
0 & \cdots & \cdots & \cdots & 0 & -\eta & \alpha Q^2
\end{pmatrix}$$

where $\alpha = \frac{k}{2}\gamma\sigma^2$ and $\eta = A(1 + \frac{\gamma}{k})^{-(1+\frac{k}{\gamma})}$.

Let us define

$$\begin{aligned}
v(t) &= (v_{-Q}(t), v_{-Q+1}(t), \dots, v_0(t), \dots, v_{Q-1}(t), v_Q(t))' \\
&= \exp(-M(T-t)) \times (1, \dots, 1)'
\end{aligned}$$

Then, $(v_q)_{|q|\leq Q}$ is a family of positive functions solution of the equations of Proposition 1.

Using the above change of variables and a verification approach, we are now able to solve the stochastic control problem, that is to find the value function of the problem and the optimal quotes :

Theorem 1 (Solution of the control problem) *Let consider $(v_q)_{|q|\leq Q}$ as in Proposition 2.*

Then $u(t,x,q,s) = -\exp(-\gamma(x+qs))v_q(t)^{-\frac{\gamma}{k}}$ is the value function of the control problem.

Moreover, the optimal quotes are given by :

$$s - s^{b*}(t,q,s) = \delta^{b*}(t,q) = \frac{1}{k}\ln\left(\frac{v_q(t)}{v_{q+1}(t)}\right) + \frac{1}{\gamma}\ln\left(1 + \frac{\gamma}{k}\right), \quad q \neq Q$$

$$s^{a*}(t,q,s) - s = \delta^{a*}(t,q) = \frac{1}{k}\ln\left(\frac{v_q(t)}{v_{q-1}(t)}\right) + \frac{1}{\gamma}\ln\left(1 + \frac{\gamma}{k}\right), \quad q \neq -Q$$

and the resulting bid-ask spread quoted by the market maker is given by :

$$\psi^*(t,q) = -\frac{1}{k}\ln\left(\frac{v_{q+1}(t)v_{q-1}(t)}{v_q(t)^2}\right) + \frac{2}{\gamma}\ln\left(1 + \frac{\gamma}{k}\right), \quad |q| \neq Q$$

1.3.2 Optimal Trading with Limit Orders

This Section has been published as a part of *Optimal Portfolio Liquidation with Limit Orders* (2012 SIAM Journal on Financial Mathematics 13, 740–764) [Guéant *et al.* , 2012]

Setup of the model. Let us fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. We assume that all random variables and stochastic processes are defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$.

We consider a trader who has to liquidate a portfolio containing a large quantity q_0 of a given stock. We suppose that the reference price of the stock (which can be considered the mid-price or the best bid quote for example) moves as a brownian motion with a drift :

$$dS_t = \mu dt + \sigma dW_t$$

The trader under consideration will continuously propose an ask quote denoted $S^a_t = S_t + \delta^a_t$ and will hence sell shares according to the rate of arrival of aggressive orders at the prices he quotes.

His inventory q, that is the quantity he holds, is given by $q_t = q_0 - N^a_t$ where N^a is the jump process counting the number of shares he sold . We assume that jumps are of unitary size and it is important to notice that 1 share may be understood as 1 bunch of shares, each bunch being of the same size . Arrival rates obviously depend on the price S^a_t quoted by the trader and we assume that intensity λ^a associated to N^a is of the following form :

$$\lambda^a(\delta^a) = A \exp(-k\delta^a) = A \exp(-k(s^a - s))$$

This means that the lower the order price, the faster it will be executed.

As a consequence of his trades, the trader has an amount of cash whose dynamics is given by :

$$dX_t = (S_t + \delta^a_t) dN^a_t$$

The trader has a time horizon T to liquidate the portfolio and his goal is to optimize the expected utility of his P&L at time T. We will focus on CARA utility functions and we suppose that the trader optimizes :

$$\sup_{(\delta^a_t)_t \in \mathcal{A}} \mathbb{E}\left[-\exp\left(-\gamma(X_T + q_T(S_T - b))\right)\right]$$

where $\mathcal{A}$ is the set of predictable processes on [0,T], bounded from below, where γ is the absolute risk aversion characterizing the trader, where X_T is the amount of cash at time T, where q_T is the remaining quantity of shares in the inventory at time T and where b is a cost (per share) one has to incur to liquidate the remaining quantity at time T

Optimal quotes. The optimization problem set up in the preceding section can be solved using classical Bellman tools. To this purpose, we introduce the Hamilton-Jacobi-Bellman equation associated to the optimization problem, where u is an unknown function that is going to be the value function of the control problem :

$$\begin{aligned}\text{(HJB)} \qquad & \partial_t u(t,x,q,s) + \mu \partial_s u(t,x,q,s) + \frac{1}{2}\sigma^2 \partial^2_{ss} u(t,x,q,s) \\ & + \sup_{\delta^a} \lambda^a(\delta^a)\left[u(t, x+s+\delta^a, q-1, s) - u(t,x,q,s)\right] = 0\end{aligned}$$

with the final condition :

$$u(T,x,q,s) = -\exp\left(-\gamma(x + q(s-b))\right)$$

and the boundary condition :

$$u(t,x,0,s) = -\exp\left(-\gamma x\right)$$

To solve the Hamilton-Jacobi-Belmann equation, we will use a change of variables that transforms the PDEs in a system of linear ODEs.

Proposition 3 (A system of linear ODEs) *Let us consider a family of functions $(w_q)_{q\in\mathbb{N}}$ solution of the linear system of ODEs $(\mathcal{S})$ that follows :*

$$\forall q \in \mathbb{N}, \dot{w}_q(t) = (\alpha q^2 - \beta q) w_q(t) - \eta w_{q-1}(t)$$

with $w_q(T) = e^{-kqb}$ and $w_0 = 1$, where $\alpha = \frac{k}{2}\gamma\sigma^2$, $\beta = k\mu$ and $\eta = A(1+\frac{\gamma}{k})^{-(1+\frac{k}{\gamma})}$.

Then $u(t,x,q,s) = -\exp(-\gamma(x+qs)) w_q(t)^{-\frac{\gamma}{k}}$ is solution of (HJB).

The change of variables used in Proposition 2.1 is based on two different ideas. First, the choice of a CARA utility function allows to factor out the Mark-to-Market value of the portfolio $(x + qs)$. Then, the exponential decay for the intensity allows to introduce $w_q(t)$ and to end up with a linear system of ordinary differential equations.

Now, using this system of ODEs, we can find the optimal quotes through a verification theorem :

Theorem 2 (Verification theorem and optimal quotes) *Let us consider the solution w of the system $(\mathcal{S})$ of Proposition 2.1. Then, $u(t,x,q,s) = -\exp(-\gamma(x+qs)) w_q(t)^{-\frac{\gamma}{k}}$ is the value function of the control problem and the optimal ask quote can be expressed as :*

$$\delta^{a*}(t,q) = \frac{1}{k}\ln\left(\frac{w_q(t)}{w_{q-1}(t)}\right) + \frac{1}{\gamma}\ln\left(1+\frac{\gamma}{k}\right)$$

Simple case : optimal price. We now derive tractable formulae for w and for the optimal quotes :

Proposition 4 (The no-drift/no-volatility case) *Assume that $\sigma = 0$ and that there is no drift ($\mu = 0$).*
Let us define :

$$w_q(t) = \sum_{j=0}^{q} \frac{\eta^j}{j!} e^{-kb(q-j)} (T-t)^j$$

Then w defines a solution of the system $(\mathcal{S})$ and the optimal quote is :

$$\delta^{a*}(t,q) = -b + \frac{1}{k} \ln\left(1 + \frac{\frac{\eta^q}{q!}(T-t)^q}{\sum_{j=0}^{q-1} \frac{\eta^j}{j!} e^{-kb(q-j)}(T-t)^j}\right) + \frac{1}{\gamma} \ln\left(1+\frac{\gamma}{k}\right)$$

In this no-drift/no-volatility case, the optimal quote still is an increasing function of A and a decreasing function of γ . If the above closed-form formula does not shed any particular light on the dependence on k, it highlights the role played by the liquidation cost b. Differentiating the above formula with respect to b, we indeed get a negative sign and therefore that optimal quote is a decreasing function of b. Since b is the cost to pay for each share remaining at time T, an increase in b gives an incentive to speed up execution and hence to lower the quotes.

We also see that the optimal quote is bounded from below by $-b + \frac{1}{\gamma}\ln\left(1+\frac{\gamma}{k}\right)$. Since execution is guaranteed at price $s-b$ at time T, it is in particular natural in the absence of price risk, that quotes never go below $-b$.

Now, if one wants to remove risk aversion with respect to both price risk and non-execution risk, one can consider the limit of the above solution when γ tends to 0.

One then obtains :

$$\delta^{a*}(t,q) = -b + \frac{1}{k} \ln\left(1 + \frac{\frac{A^q}{e^q q!}(T-t)^q}{\sum_{j=0}^{q-1} \frac{A^j}{e^j j!} e^{-kb(q-j)}(T-t)^j}\right) + \frac{1}{k}$$

and this is the result of Bayraktar and Ludkovski [Bayraktar & Ludkovski, 2012] in the case $b = 0$, because they do not consider any liquidation cost. In particular, the optimal quote of [Bayraktar & Ludkovski, 2012] does not converge to a limit value as T tends to $+\infty$, but rather increases with no upper bound. This is an important difference between the risk-neutral case and our risk-adverse framework.

Simple case : optimal trading curve. Let us now consider the limiting case $b \to +\infty$. Sending b to infinity corresponds to a situation in which a very high incentive is given to the trader for complete liquidation before time T. If we look at the Almgren-Chriss-like literature on optimal execution, the authors are often assuming that $q_T = 0$. Hence, if one writes the value functions associated to most liquidity-consuming optimal strategies, it turns out that they are equal to $-\infty$ at the time horizon T except when the inventory is equal to nought (hence $b = +\infty$, in our framework). However, here, due to the uncertainty on execution, we cannot write a well-defined control problem when b is equal to $+\infty$. Rather, we are interested in the limiting behavior when $b \to +\infty$, *i.e.* when the incentive to liquidate before time T is large.

By analogy with the initial literature on optimal liquidation [Almgren & Chriss, 2000], we can also have some limiting results on the trading curve.

Hereafter we denote $w_{b,q}(t)$ the solution of the system $(\mathcal{S})$ for a given liquidation cost b, $\delta_b^{a*}(t,q)$ the associated optimal quote and $q_{b,t}$ the resulting process modeling the number of stocks in the portfolio.

More results can be obtained in the no-volatility case :

Proposition 5 (no-volatility case, $b \to +\infty$) *Assume that $\sigma = 0$ and consider first the case $\mu \neq 0$. We have :*

$$\lim_{b\to+\infty} w_{b,q}(t) = \frac{\eta^q}{q!}\left(\frac{e^{\beta(T-t)}-1}{\beta}\right)^q$$

The limit of the optimal quote is :

$$\delta_\infty^{a*}(t,q) = \lim_{b\to+\infty} \delta_b^{a*}(t,q) = \frac{1}{k}\ln\left(\frac{A}{1+\frac{\gamma}{k}}\frac{1}{q}\frac{e^{\beta(T-t)}-1}{\beta}\right)$$

The limit of the associated trading curve is $V(t) = q_0\left(\frac{1-e^{-\beta(T-t)}}{1-e^{-\beta T}}\right)^{1+\frac{\gamma}{k}}$.

Now, in the no-volatility/no-drift case ($\sigma = \mu = 0$), similar results can be obtained, either directly or sending μ to 0 in the above formulae :

$$\lim_{b\to+\infty} w_{b,q}(t) = \frac{\eta^q}{q!}(T-t)^q$$

The limit of the optimal quote is given by :

$$\delta_\infty^{a*}(t,q) = \lim_{b\to+\infty} \delta_b^{a*}(t,q) = \frac{1}{k}\ln\left(\frac{A}{1+\frac{\gamma}{k}}\frac{1}{q}(T-t)\right)$$

The limit of the associated trading curve is $V(t) = q_0\left(1-\frac{t}{T}\right)^{1+\frac{\gamma}{k}}$.

2 Etude de la Microstructure

La section précédente a détaillé plusieurs cadres pour l'optimisation du passage de *gros ordres* sur un marché à enchères doubles. On y a vu apparaître des quantités comme la volatilité intraday, les volumes habituellement échangés, le Bid-Ask spread, le *mid-price*, ou la probabilité d'obtenir une transaction lorsqu'on est à une distance fixée de ce mid-price. La Section 1.1.2 a notamment détaillé le cas où on veut prendre en compte la variabilité des estimateurs de certaines de ces grandeurs. En pratique il est indispensable de bien comprendre pourquoi ces grandeurs sont particulièrement importantes, et comment les estimer au mieux.

L'étude de la microstructure tente de répondre à ces questions. Elle va même plus loin dans le sens où ce champ de recherche passe parfois par l'étude théorique des interactions entre peu de participants, afin de comprende en quoi des équilibres long terme peuvent émerger. Le modèle de market impact de Kyle [Kyle, 1985] est typiquement de cette sorte : il n'utilise pas de données mais postule quelques interactions simples entre un *market maker* (i.e. un *teneur de marché* en français) et un investisseur qui désire acheter un grand nombre d'actions. Il en déduit l'existence d'un *market impact* bien superieur à ce que les économistes avaient en tête à l'époque

D'autres approches plus empiriques complètent ces points de vue théorique : à l'aide de bases de données des ordres passés par un gestionnaire de portefeuille et en les confrontant aux données publiques de toutes les transactions atomiques survenue sur le ou les marchés concernés, on peut essayer de dégager des modèles statistiques. Le modèle de market impact utilisé Section 1.1 a été construit de cette façon.

Dynamique des carnets d'ordres. L'augmentation des capacité d'acquisition, de stockage, et de traitement (que l'on qualifie souvent de *big data*) permet une étude de plus en plus fine des comportements joints pendant les jeux d'enchères qui se déroulent en temps continu sur les marchés électroniques. [Lehalle *et al.* , 2013] détaille l'évolution récente de ces jeux d'enchères, et explique pourquoi la crise financière de 2008 et les évolutions règlementaires amènent de plus en plus de participants vers des marchés électroniques.

Comme cela a été détaillé dans l'introduction de la section précédente (cf. [Lehalle, 2013] pour plus de détails), les participants aux jeux d'enchère envoient des messages à un serveur central avec leurs intentions. Ces messages détaillent le sens (achat ou vente), le prix et la quantité des intérêt du participant. Tout nouveau message est comparé à la liste de ceux déjà stockés dans la mémoire du serveur d'appariement (i.e. *Matching Engine* en anglais) :

Si un message entrant à la vente (respectivement à l'achat) a un prix inférieur (resp. supérieur) à celui d'un message d'achat (resp. de vente) déjà présent en mémoire :

Alors une transaction est générée entre les deux parties (i.e. l'acheteur et le vendeur) ;

Sinon le message entrant rejoint le stock des messages mémorisés par le serveur. Ce stock s'appelle le *Carnet d'Ordres* (ou *Carnet d'Ordres Limites* : LOB, *Limit Order Book* en anglais) , en référence aux petits carnets de papier sur lesquels les courtiers notaient leurs achats et leurs ventes pendant le déroulement de la criée, afin de procéder à leur dépouillement en fin de journée (cf. [Muniesa, 2003]).

Ces carnets d'ordres sont organisés par *limite de prix* ; croissantes pour les ventes et décroissantes pour les achats. On s'attend à ce que la dynamiques d'un LOB (acronyme anglais de *Limit Order Book*) soit soumise à quelques effets simples :

(*i*) Les "*premières limites*" du LOB sont manifestement les plus informatives puisque les acheteurs proposant les plus hauts prix et les vendeurs proposant les plus hauts prix sont les plus susceptibles de participer aux prochaines transactions.

(*ii*) Le déséquilibre entre les deux premières limites (achat et vente) conditionne fortement le *signe* de la prochaine transaction (i.e. sera-t-elle initiée par un nouveau message d'achat ou de vente ?).

(*iii*) Les autres limites, plus profondes, auront quant à elles une influence sur la dynamique future des transaction et des messages lorsque le prix "les atteindra". Ce sont elles qui permettront ou non au prix de continuer à diffuser. Ou au contraire qui le ramèneront à ses niveaux précédents.

La Section 2.1 apporte un éclairage empirique et théorique à cette dynamique a priori sophistiquée. Elle documente aussi les réponses possibles à des interrogations du type : "*sachant l'état des limites, quelle est la probabilité d'obtenir une transaction à horizon T si on envoie un message d'achat au $k$ème meilleur prix ?*".

Dégageant ainsi une modélisation markovienne du LOB, cette section montre comment il est nécessaire d'ajouter une composante de saut au prix (paramètrée par θ dans ce modèle dit *Queue Reactive*) pour obtenir un modèle réaliste à toutes les échelles.

La nécessité d'ajouter ainsi une composante exogène à la dynamique des carnets d'ordres pour compenser la composante "*retour à la moyenne*" d'un modèle markovien de carnets peut être interprété de plusieurs façons. Du point de vue du choix d'un modèle, il faut noter qu'il est par exemple possible remplacer cette composante exogène θ par la distribution conditionnelle d'une limite de prix *sachant que la limite précédente vient de se vider*, comme dans [Cont & De Larrard, 2013].

Dans les deux cas émerge l'idée que certains participants décident de tenir un changement de prix pour aquis ou non suivant un critère extérieur à la simple dynamique markovienne des carnets (ce qui est conforme au fait stylisé (*iii*) énuméré précédemment). Cela ne devrait pas surprendre le lecteur de la Section 1 sur le *trading optimal* puisque qu'on y détaille comment la fonction d'utilité d'un participant (prenant en compte les prix courants, leur historique,

et surtout son propre inventaire, qui n'est pas observable par le reste des joueurs) le pousse à participer plus ou moins au jeu d'enchères (et donc à accepter plus ou moins le prix).

Chaque participant implémente une optimisation à la volée, qui prend en compte non seulement les observables microscopiques passés, mais aussi ses propres informations et surtout son inventaire et ses contraintes individuelles.

Afin de bien comprendre le rôle de ces stratégies individuelles, la Section 2.2 propose une modélisation de la dynamique des deux premières limites du carnet d'ordres (celle des vendeurs et celle des acheteurs), lorsque chacun met en œuvre une stratégie de trading optimal. Elle recours à un modèle de Jeux à Champ Moyen (MFG : *Mean Field Game* en anglais).

Modélisation du Market Impact. Au delà d'une meilleure compréhension de la dynamique des carnets, qui peut nourrir les modèles utilisés dans les Sections 1.3 et 3.2, une des composantes primordiales de la microstructure reste le *Market Impact*. Il s'agit du mouvement de prix dû à un pression durable à l'achat ou à la vente d'un *gros ordre*. De cas chroniques démontrent l'existence de ce phénomène (comme la liquidation de l'inventaire de Jérôme Kerviel par la Société Générale,ou bien les mouvements de prix du 19 Juillet 2012 sur les marchés américains [Lehalle *et al.* , 2012]). Des bases de données de plusieurs centaines de milliers de *métaordres* (i.e. un gros ordre qui va être découpé en plusieurs transactions, souvent suivant une stratégie de trading optimal) permettent de mettre à jour le même effet avec moins d'intensité pour des tailles nettement moins importantes que l'inventaire de Kerviel.

La Section 2.3 explicite et propose des modèles pour toutes les phases du market impact :

- le market impact *temporaire* : le mouvement de prix entre le début et la fin d'un métaordre ;
- le market impact *transient* : la dynamique temporelle du prix entre le début et la fin d'un métaordre ;
- la *relaxation* du market impact : le mouvement de prix après la fin du métaordre ;
- l'*effet permanent* suite à l'impact : ce qui reste dans le prix plusieurs jours après la fin du métaordres.

Cette section propose aussi un modèle jouet basé sur les processus de Hawkes.

English introduction to the section

Orderbooks modelling. Orderbook modelling is at the crossing of a lot of different research strands :

Approach Z – Zero intelligence models (like [Smith *et al.* , 2003]), that are unconditioned models using data to fit marginal distributions on their first moments only. This has been the first attempts of modelling orderbooks dynamics a statistical way, and that for they are of importance (for a quite exhaustive review see [Chakraborti *et al.* , 2011]).

Approach F – Flow driven (or event driven) dynamics, with papers focussing on the link between the price moves and the quantities at first limits. Typical papers of this kind are [Cont *et al.* , 2008] and [Cont & De Larrard, 2013]. their major take away is to provide asymptotics of the price diffusive behaviour. One of their very interesting feature is they replace the marked process of price moves by the stopping times generated by the crossing of the vertical and horizontal axis by a two dimensional Markovian process (a two dimensional Brownian motion at the limit). Another interesting paper using a similar angle is [Gareche *et al.* , 2013], where authors extend the previous dynamics in a PDE framework and comfort it with real data.

Approach E – Economic models, like [Roşu, 2009] or [Jaisson, 2014] are more theoretic and adopt a global equilibrium viewpoint.

Approach D – more descriptive than microfunded empirical studies exist too. They span a broad spectrum of models from the "uncertainty bands" one ([Robert & Rosenbaum, 2011], focussing on the tick size and the local time of Brownian motions), to more liquidity driven ones (like Pham [Fodra & Pham, 2013]). The Hawkes models (see [Bacry *et al.* , 2015]) are part of this family of descriptive models, with a success due to their self exciting or inhibitory natural properties.

The works presented here is at the crossing of these different approaches. First of all the Queue Reactive model starts with simple statistics on the distribution of different flows increasing or decreasing the queues of liquidity providers waiting for a transaction : insertion of new limit orders, cancellations and trades (this flow being the intensity of liquidity consumers, accepting to pay the bid-ask spread in exchange of the immediacy and the certainty of a transaction). Going further than the first limits only (unlike most models of Z and F), it considers these intensities for any limit. It shows these point processes have to be conditioned by the state of, not only the limit on the same side (ask for sellers and bid for buyers), but also of the best opposite one. At this stage it shows that for the considered stocks, the three first limits are enough to understand liquidity formation.

Going one step forward, it shows this Markovian process of liquidity is not enough to reproduce the price formation. At this point it joins the models of F, and use a straightforward way to catch up with the understanding provided by models E : since the price generated by the Markovian process is not diffusive enough, let's add a component (parametrized by θ) modelling the "acceptance by market participants" of a new price. It reproduces the diffusive property of [Cont & De Larrard, 2013] at the limit, in which the size of any "discovered limit" (after the consumption of 100% of the 1st limit) is drawn according to an iid random variable F. Without the effect provided by θ in the QR model, the size of any "discovered limit" is too large to enable diffusion.

The need of "more memory" than the one provided by a Markovian process of liquidity to obtain a diffusive limit can be obtain different ways : Z and F do not care that much about the size of limits after the first ones, D starts from

a diffusive process and adjusts it to emulate mean reversion at small scales. The assumption underlying the E approach is tho distinguish the behaviours at small time scales (where agents competes for liquidity) from the behaviours at large time scales (where they compete for prices), stems from *agents utility functions*. To understand this viewpoint it is enough to come back to our Section 1 : agents are implementing stochastic control (explicitly if they read our work, implicitly otherwise). In Section 1, it is assumed, via a *implicit mean field assumption*, the actions of other agents sum up to a background noise. On the path from small scales to large ones, it seems this assumption no more holds.

Section 2.2 explores a continuation of this mean field approximation and proposes a mean field game (MFG) model for orderbook dynamics.

Thanks to this MFG view, one can question and study the importance of taking into account the optimal behaviour of all market participants, without any mean field approximation, but inside a mean field game. The work of Section 2.2 focus on orderbook dynamics and obtain the stady-state behaviour of strategic agents trading in a pro-rata orderbook (see [Field & Large, 2008] for details of this kind of matching mechanisms). The result on pro-rata books can be extended to first in first out (FIFO) ones under the assumption (like in [Roşu, 2009]) that participant can cancel and re enter their orders at any time.

More than obtaining the stationary joint distribution of the size of the two first queues, the study deciphers the effect of having participants with different speeds in the same market. Paying attention that all participants have the same *flow* on average (i.e. size of their orders times the intensity of the Poisson process governing the decision rate), we have shown how liquidity *stabilizes* with the diversity of participants speeds. The more diverse speed, the more stable states around the $Q_{\text{bid}} = Q_{\text{ask}}$ diagonal and the smaller the effective bid-ask spread. We could also show that, under our MFG model, this apparent improvement of the liquidity benefits more the faster agents than to the others.

Market Impact. Another important component of market microstructure is the market impact arising from large *metaorders*. This topic appear at largest scale than the one of orderbook dynamics. It appears when a large trader sells or buys a lot of shares ; the trading pressure on the auction mechanisms pushes the price a detrimental way (i.e. up for a buy order and down for a sell order).

We will not go in the details of impact of atomic trades, see [Lehalle & Dang, 2010] for bibliographic references and an empirical study of the part price formation process due to market orders, and ref :deal/book for a decomposition of prices moves between quotes tmoves and trade impact.

A good understanding of market impact and of its components allows traders to minimize it, and therefore is at the advantage of all market participants, since it maintains the price closer to its *faire value* or *latent value*. The link between this latent value and the impact is of importance : does the trading pressure from t_0 to t_f drives the price to a value that would be nevertheless attained at t_f, because the price move comes from an informational effect rather than a mechanical one on orderbooks ? or does the price moves only when participants trade ? Being a little more formal, we could follow [Lehalle *et al.* , 2010] and design a theoretical model with two main components

- market participants' views on the price, aggregated in a latent or theoretical orderbook L_t,
- the real orderbooks B_t,

and condition the arrival of events on B_t to the distance between B_t and L_t. It can be seen as a continuation of the previous discussion on mixing the Markovian process of liquidity of B_t by the acceptance or not of the price by market participants because of the distance between B and L (i.e. participants "accept" a price move when it drives B to L, and let it follow the upper Markov process when B is already to L).

The work presented in Section 2.3 is more data-driven that a theoretical model. Nevertheless there is an attempt to explain our observations using a Hawkes process (coming from approach D). This study is the only one describing market impact at all scales :

- first the intraday formation of the impact (transient phase),
- then after a peak (the temporary impact, that can be approximated by a square root of the daily participation rate), its decay,
- and after few days, once the impact of correlated trades is removed (in an attempt to reproduce the methodology of [Gomes & Waelbroeck, 2013] and [Brokmann *et al.* , 2014]), the remaning price move can be explained by investors anticipations (given the assumption investors of our database expect to see the systematic moves of the price in a CAPM like model, see [Fama & French, 2003] about the CAPM).

The added value of this work is to give a consistent explanation of metaorders market impact at any scale, at the daily scale it provides a separation between information-driven move and trading-pressure one. At this stage it is not possible to make the difference between pure information-driven price moves (like when a news is simultaneously disclosed to all market participants), and the coordinated market impact of participants that are not in the database, but are progressively trading this information, hence building the price moves by their own impact.

May be studies on market impact at small scales during the trading of large orders (like [Lehalle & Dang, 2010]), could provide insight to answer to this *informational puzzle*. Full databases like the bitcoin one used in [Donier & Bonart, 2014] could also provide answers.

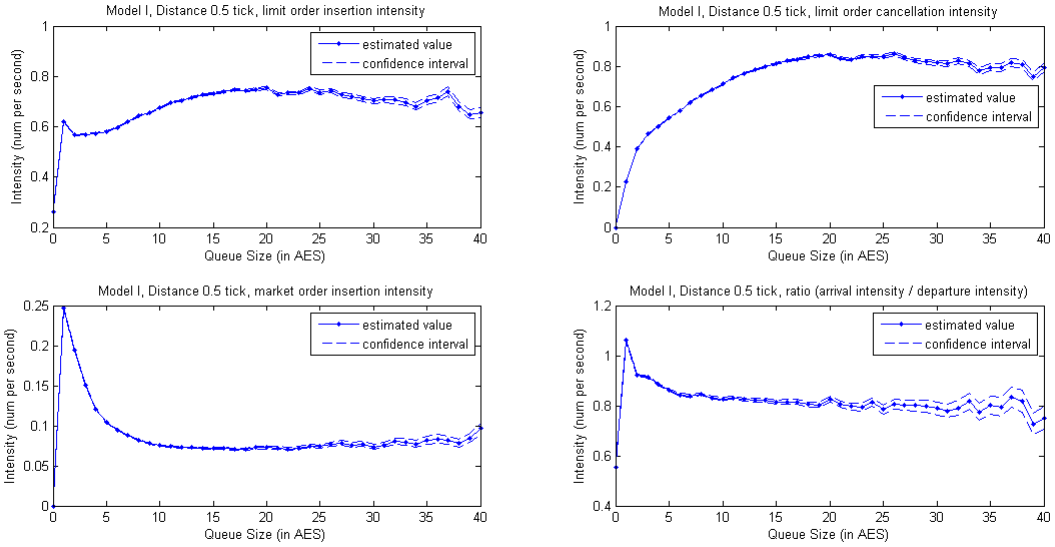

FIGURE 17 – First Limit Intensities, France Telecom

2.1 Limit Orderbook Dynamics : the Queue Reactive Model

This Section has been published as a part of *Simulating and analyzing order book data : The queue-reactive model* (accepted by UASA : Journal of the American Statistical Association) [Huang *et al.* , 2013]

2.1.1 A Collection of Models

Model I : Collection of independent queues. In this model, we assume independence between the queues at different limits. Thus, the order flow intensities at different limits are functions of their own queue size. To have independence between the different limits, we also suppose that market orders sent to Q_i consume directly the quantities available at Q_i, without consuming firstly the liquidity which may be present at limits closer to p_{ref}. This assumption is reasonable for large tick assets for which the intensity of the market order flow is almost fully concentrated on the first limits to the left and to the right of the reference price (since the spread is essentially equal to one tick). Note that consequently, we replace the market buy/sell flows N^M_{buy} and N^M_{sell} by market order flows sent at each limit N^M_i. Hence, the principle of "price priority" that states that market orders should be matched with the current best bid/ask offer is neglected in this model.

The flows $N^L_i(t)$, $N^C_i(t)$ and $N^M_i(t)$ are modeled as independent point processes (with respect to the index i). Furthermore, the associated jump sizes are equal to one at each jump time. Their intensities λ^L_i,λ^C_i and λ^M_i are assumed to be functions of the corresponding queue size Q_i. The values of these intensities when $Q_i = n$ are denoted by $\lambda^L_i(n)$, $\lambda^C_i(n)$ and $\lambda^M_i(n)$. Under the above assumptions, the LOB becomes a queuing system of $2K$ independent queues, and each queue is a birth and death process.
In Figures 17, 18, 19, based on data collected from Jan 2010 to March 2012, we present the estimated intensities and the arrival/departure ratios $\rho_i(n)$, defined as :

$$\rho_i(n) = \frac{\lambda^L_i(n)}{(\lambda^C_i(n+1) + \lambda^M_i(n+1))}$$

Data at Q_i and Q_{-i} are aggregated together (by combining simply the two collected samples) in order to estimate the intensities at distance $i - 0.5$ ticks to the reference price. Confidence intervals in these figures are computed using central limit approximations detailed in the appendix.

Model II : Two sets of dependent queues. One important element not considered in Model I is the notion of "best offer" limit. Institutional traders and brokers tend to place most of their limit orders in these best limits (defined as the nearest non-empty bid/ask queue to the reference price), while many market makers, arbitragers and other high frequency traders stand also in queues behind these best limits. This suggests that the dynamics at $Q_{\pm 2}$ may not only depend on their own queue siz, but also on whether they are the current best offer limits, that is, whether $Q_{\pm 1} = 0$. Our empirical studies of intensity functions at $Q_{\pm 2}$ under different state of $(Q_{\pm 1}, Q_{\pm 2})$ show also a significant pattern change in the estimated intensity functions at $Q_{\pm 2}$ when $Q_{\pm 1}$ becomes zero, while intensity functions at $Q_{\pm 1}$ stay almost the same for different values of $Q_{\pm 2}$.

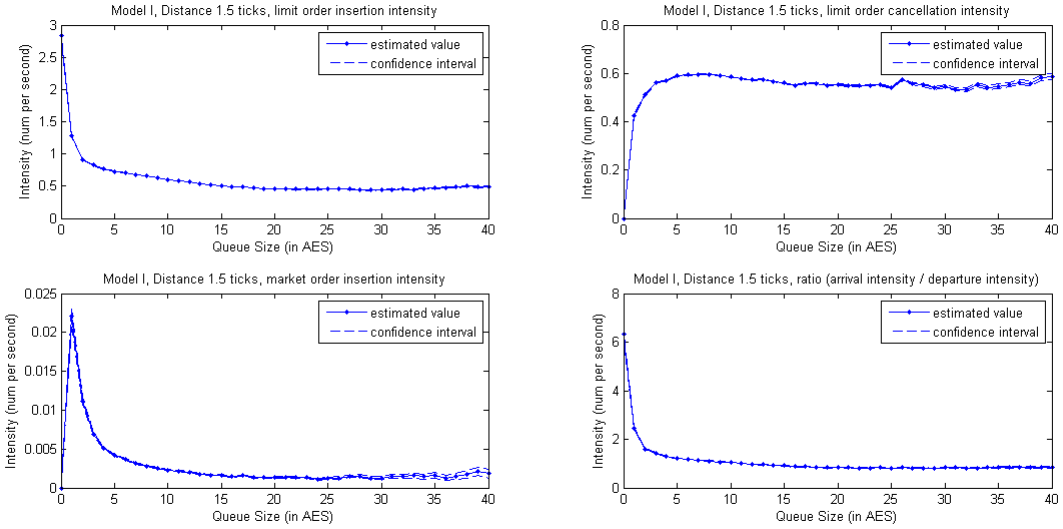

FIGURE 18 – Second Limit Intensities, France Telecom

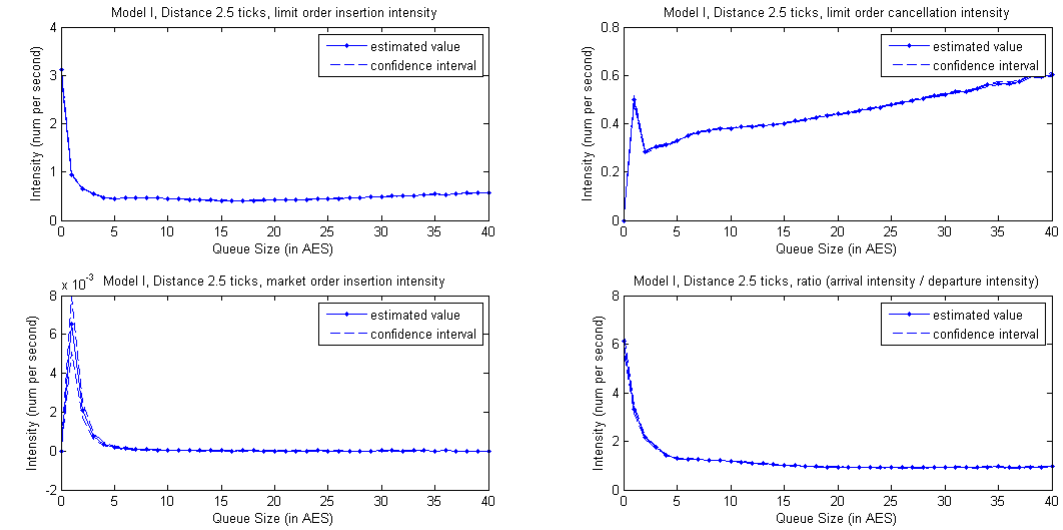

FIGURE 19 – Third Limit Intensities, France Telecom

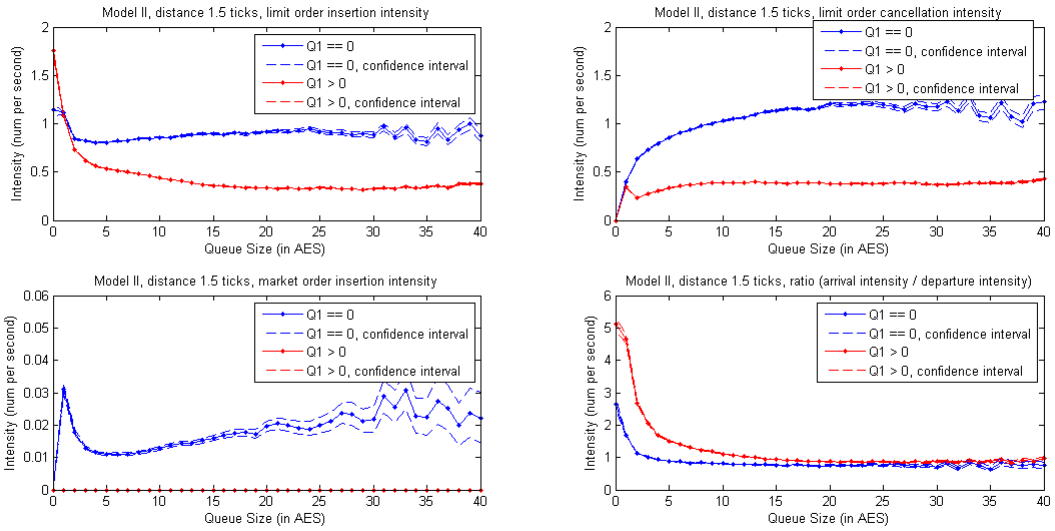

FIGURE 20 – Q_2 Intensities as a function of $M(Q_1)$ and its own size, France Telecom

We thus propose to use the following intensity functions for the queue $Q_{\pm 2}$: λ_2^L and λ_2^C (resp. λ_{-2}^L and λ_{-2}^C) are functions of $1_{Q_1>0}$ and Q_2 (resp. $1_{Q_{-1}>0}$ and Q_{-2}). Intensities at $Q_i, i \neq \pm 2$ remain functions of Q_i only. For large tick assets, the probability that $Q_{\pm 3}, i \geq 3$ is the best ask/bid offer limit being negligible, it is thus reasonable to assume that market orders are only sent to $Q_{\pm 1}, Q_{\pm 2}$. This enables us to keep the independence property between $Q_{\pm 3}$ and $(Q_{\pm 1}, Q_{\pm 2})$. When Q_1 (resp. Q_{-1}) > 0, the market order consumption intensity Q_{buy}^M (resp. Q_{sell}^M) is a function of Q_1 (resp. Q_{-1}). When Q_1 (resp. Q_{-1}) $= 0$, the market order consumption intensity Q_{buy}^M (resp. Q_{sell}^M) is a function of Q_2 (resp. Q_{-2}).

Compared to Model I, Model II includes the price priority principle into the market order arrival processes and a regime switching for the dynamics at $Q_{\pm 2}$, depending on whether $Q_{\pm 2}$ is the best limit. This model neglects, however, the interactions between the bid side $(Q_i, i < 0)$ and the ask side $(Q_i, i > 0)$, which will be studied in Section 2.1.1. The $2K$ dimensional Markov process can be split into two identical (in law) 2-dimensional Markov processes $((Q_{-2}, Q_{-1})$ and $(Q_1, Q_2))$ and $(2K - 4)$ independent queues. Thus the problem is reduced to the study of the 2-dimensional continuous time Markov jump process (Q_1, Q_2).
The dynamics at Q_1 actually depend only on its size, the estimated values of λ_1 are thus very close to those estimated in Model I (they are not exactly the same because the recording process is different) and will not be shown here. The intensity functions at Q_2 are given in Figure 20.

An agent based interpretation. We develop here a simple agent based trading model in order to explain these two curves. Suppose we have two kinds of traders in the market, noise traders and arbitragers. Noise traders make trading decisions regardless to the state of the LOB, while arbitragers adjust carefully their trading behaviors according to what they see in the LOB. So, we assume that the noise traders send their orders according to a time-homogeneous Poisson process whereas arbitragers insert orders when they think they can make profit. Consequently, the insertion rate of limit orders of arbitragers decreases to zero when the corresponding queue size increases to infinity, as the priority value of a newly inserted limit order is a decreasing function of the corresponding queue size. The observed limit order insertion intensity function is then the sum of a decreasing function (due to the arbitragers) and a constant function (due to the noise traders).

Model II belongs to a special class of Markov processes, called Quasi Birth and Death processes (QBD). Their asymptotic behaviors can be studied by the matrix geometric method. Definitions for QBD processes and explanations of the matrix geometric method can be found in the appendix.

Model III : modeling bid-ask dependence. The bid and ask side are treated separately in Model II : the dynamics at Q_1 and Q_{-1} are supposed to be independent. In Model III, we remove this assumption and study interactions between the bid queues and the ask queues. First we define the function $L_{m,l}(x)$:

$$\begin{aligned} L_{m,l}(x) &= 0, \text{if, } x <= m \\ L_{m,l}(x) &= 1, \text{if, } m < x <= l \\ L_{m,l}(x) &= 2, \text{if, } x > l \end{aligned} \tag{34}$$

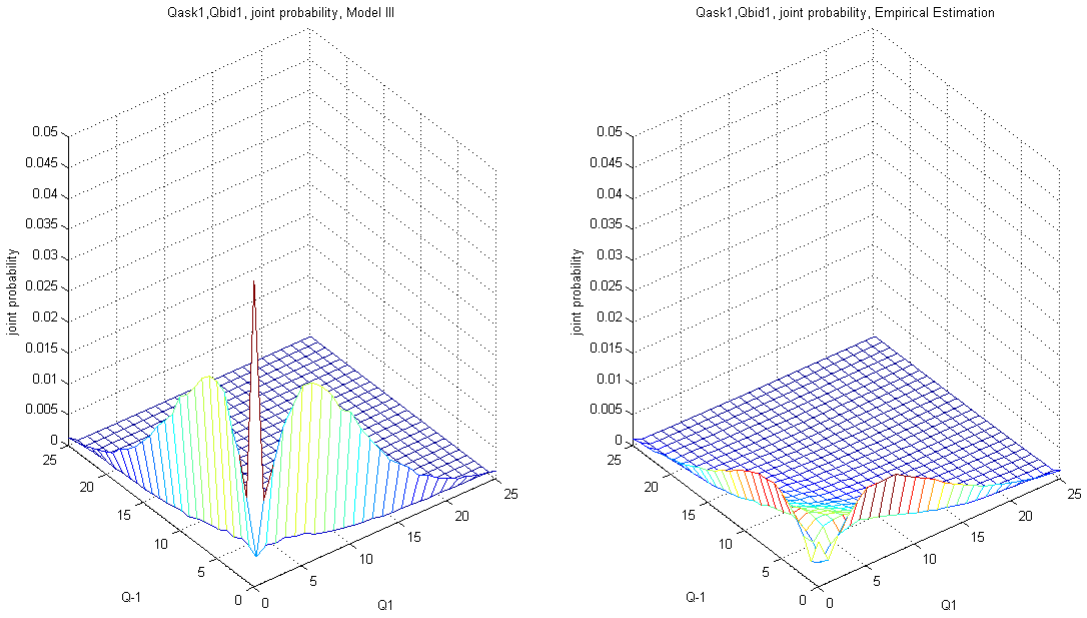

FIGURE 21 – Model III : Joint Distribution of Q_{-1}, Q_1, France Telecom

This function splits the queue size into three different ranges : small ($[0, m]$), usual ($(m, l]$) and large ($(l, +\infty)$). We set m the lower 33% quantile value of $Q_{\pm 1}$ and l the higher 33% quantile value of $Q_{\pm 1}$. We assume market participants at $Q_{\pm 1}$ adjust their behaviors not only according to the size of the target queue, but also to whether the opposite queue size is small, usual or large. The functions λ_1^L and λ_1^C (resp. λ_{-1}^L and λ_{-1}^C) are modeled as functions of Q_1 and $L_{m,l}(Q_{-1})$. As in Model II, we suppose that market orders consume quantities at the best limits and are sent only to $Q_{\pm 1}$ and $Q_{\pm 2}$. When $Q_1 > 0$ (resp. $Q_{-1} > 0$), the market order consumption intensity Q_{buy}^M (resp. Q_{sell}^M) is assumed to be function of Q_1 and $L(Q_{-1})$ (resp.Q_{-1} and $L_{m,l}(Q_1)$). Price priority and regime switching at $Q_{\pm 2}$ are kept in Model III : $\lambda_{\pm 2}^L$, $\lambda_{\pm 2}^C$ are, as in Model II, functions of $1_{Q_{\pm 1}>0}$ and $Q_{\pm 2}$, and when $Q_1 = 0$ (resp. $Q_{-1} = 0$), the market order consumption intensity Q_{buy}^M (resp. Q_{sell}^M) is modeled as a function of Q_2 (resp. Q_{-2}).

In Model III, the $2K$ dimensional Markov process can be split into a 4-dimensional Markov process and $(2K-4)$ independent queues. Thus the problem is reduced to the study of the 4-dimensional continuous time Markov jump process $(Q_{-2}, Q_{-1}, Q_1, Q_2)$. One important feature of Model III is that the queues $Q_{\pm 2}$ have no influence on the dynamics at $Q_{\pm 1}$. Therefore, we only need to study the 3-dimensional process (Q_{-1}, Q_1, Q_2) (or even the two dimensional problem Q_{-1}, Q_1 if one is only interested by the dynamics at Q_{-1} and Q_1).

Of course, other choices are possible for the intensity functions at $Q_{\pm 1}$. For example, one can consider them as functions of the first level bid/ask imbalance, defined as $\frac{Q_1 - Q_{-1}}{Q_1 + Q_{-1}}$, or simply as functions of the spread size.

- Limit order insertion : we notice that the limit order insertion rate is a decreasing function of its opposite limit's size. The curves have similar shapes but different asymptotic values. In particular, when the opposite queue is small (blue curve), it is significantly larger than the others. When quantities available at Q_{-1} are small, the implicit fair price is probably close to p_{-1}. Market participants are thus happy to insert orders at Q_1, where they either get cheap executions (with respect to the far price), or gain priority value by staying at the queue.
- Limit order cancellation : the cancellation rates for different ranges of Q_{-1} are similar in their forms but have different asymptotic values. This rate, being an indicator of market participants' patience, is not surprisingly a decreasing function of liquidity level at the opposite side.
- Market order consumption : we see that when the liquidity available at Q_{-1} is abundant, more market orders are sent to Q_1. Indeed, in that case, the fair price is closer to the side of Q_1 so that market order are not too costly.

Asymptotic behaviors in Model III. To obtain the invariant distribution of the order book in Model III, we use the Monte-Carlo method. We consider 10^5 simulations of periods of 10^4 seconds. The theoretical and empirical joint distributions of Q_{-1} and Q_1 are shown in Figure 21. As for Model II, the shape of the empirical distribution is well approximated by the theoretical one. However, the error in estimating the boundary probability at $Q_{\pm 1} = 0$ leads to some overestimation in the theoretical values for this example.

2.1.2 The Queue Reactive Market Model : Going Further than Orderbooks-driven Dynamics

In the previous paragraphs, we present three different models. Model I assumes that the trading activities at different limits are independent. The only factor that influences traders intention to send limit/market orders, or to cancel existing orders is the queue size at which these orders will be sent. This assumption is very convenient in order to focus on the features of the limits at different distances to the reference price. Furthermore, invariant measures can be computed explicitly, and they approximate very well the empirical LOB distributions. Model II enlarges the information set used by Model I by including the notion of "best", that is, whether the studied limit is the current best offering limit. Regarding $Q_{\pm 2}$, it means that market participants adjust their trading behaviors according to whether $Q_{\pm 1}$ is empty. Empirical results show significantly different behaviors of market participants between this two situations. Moreover, an agent based explanation for this difference is provided. Also, we design a method which allows for numerical computation of the LOB distribution. Model III adds dependence between the bid and ask queues. In this setting, our empirical study enables us to understand the impact of the shape of one side of the order book on the trading activity on the other side.

Two large tick stocks (France Telecom and Alcatel-Lucent) are considered. We have presented in this section the empirical results of the stock France Telecom. The comments made for the intensity functions of this stock under Model I, Model II and Model III apply also to the stock Alcatel-Lucent (in the appendix). Two particularly interesting empirical facts have been discovered for these large tick assets :

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

The first of these facts can be viewed as a consequence of the priority value. The second one is probably related to the existence of arbitragers at those secondary limits.

Asymptotic studies are carried for all these three models. Their results strongly suggest the following important conclusion :

- For large tick assets, the empirical LOB distribution is actually an asymptotic equilibrium created by market participants as they act differently in different states of the LOB.

The Queue Reactive Model. A key aspect in the price dynamics is the shape of the LOB right after a price change, which is closely linked to the mean reversion ratio of the asset and thus to its volatility. We have seen in the preceding section that the price fluctuations generated by LOB dynamics alone are usually not sufficient to reproduce the empirical volatility level, as large queues behind the two best limits create barriers that prevent further price movement in the same direction and thus push the price back to its previous level.

To overcome this difficulty, we assume that the price process is not only driven by its order book dynamics but also by some exogenous movements. Hence the "realized" volatility is considered as the combined effect of the "mechanical" volatility due to temporary fluctuations of the LOB and the "informational" volatility due to exogenous market movements. To that purpose, we suppose that with probability θ^{reinit}, the state of the LOB is redrawn from its invariant distribution around the new reference price when it changes. The parameter θ^{reinit} can be understood as the percentage of price changes due to exogenous price movements. Indeed, in this case, market participants readjust very quickly their order flows around the new reference price, as if a new state of the LOB were drawn from its invariant distribution. A similar approach has been used by Cont *et al.* in [Cont & De Larrard, 2013] , where they fix the value θ^{reinit} to 1 and use the empirically estimated queue distributions immediately after a price move instead of the invariant distribution at the corresponding queue to reinitialize the state of the LOB.

This new element related to exogenous information is added to the purely order book driven model to build the "Queue reactive market model". We calibrate the LOB reinitialization probability θ^{reinit} and the reference price move probability θ using the 10 minutes standard deviation of the price returns (the volatility) and the mean reversion ratio η. In Figure 22, we show the surface of the 10 min volatility and η for different values of these two parameters.

2.1.3 Market impact profile

Tactic 1 and Tactic 2 have shown very different performances when being coupled with the same scheduling strategy. We now study in detail their market impact profiles during the order placement period. Recall that the execution algorithm has two parameters : the period length T and the total quantity to execute n. In the following experiments, T will still be set to 10 minutes, and we vary the value of n from 1 to 60 AES. We denote by $\mathrm{MI}^i(t,n)$ the market impact of Tactic i with total quantity n at the moment t, defined by : $\mathrm{MI}^i(t,n) = \mathrm{Exp}[\frac{S_t - S_0}{S_0}|$. We launch 1000 simulations for each value of n, t between $1-60$ AES and $1-600$ seconds.

Impact profiles of Tactic 1 and Tactic 2 are given in Figure 23. In agreement with the famous square-root law, see [Gatheral, 2010a] , the market impact curve is shown here to be concave both in time and total execution volume. The price relaxation (for a buying order, it means that the price falls on average to a lower level after the trader stops his

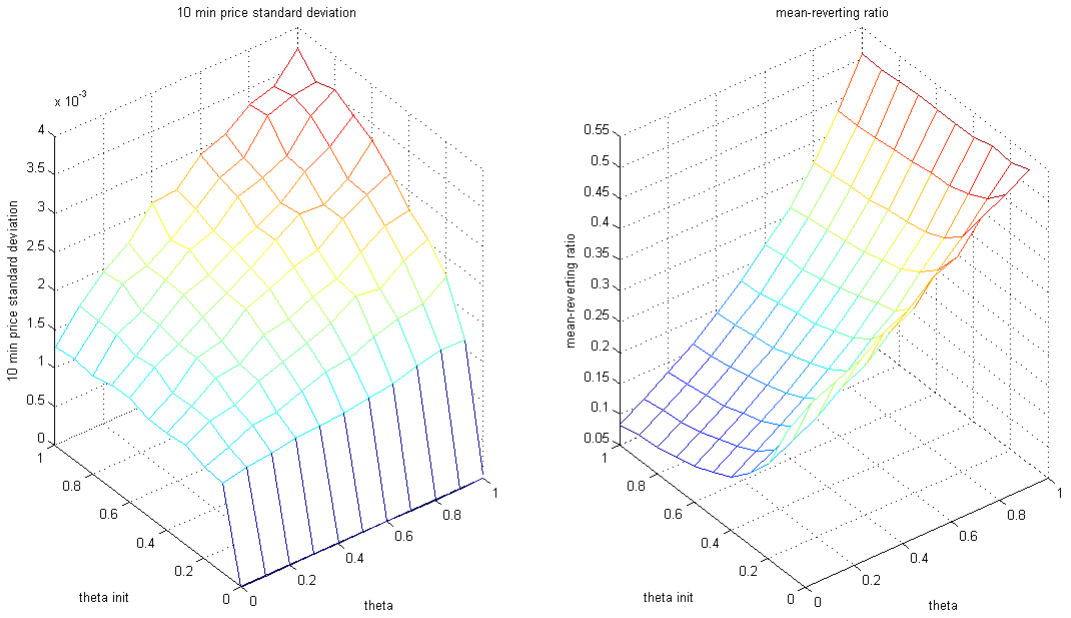

FIGURE 22 – 10 min volatility and Mean-reversion ratio, France Telecom

trading activity, see [Bacry & Muzy, 2013]) cannot be observed in our framework. Actually, under the Markov setting, the price dynamics will continue its usual path after the terminal time T. By staying passive in the best limit until full execution, Tactic 2 causes larger impact than Tactic 1, especially when dealing with large quantities. To conclude, Tactic 2 performs better when dealing with small quantities. However, when one needs to trade large quantities, Tactic 1 becomes a more suitable choice as now the cost of market impact outweighs the benefit of passive execution.

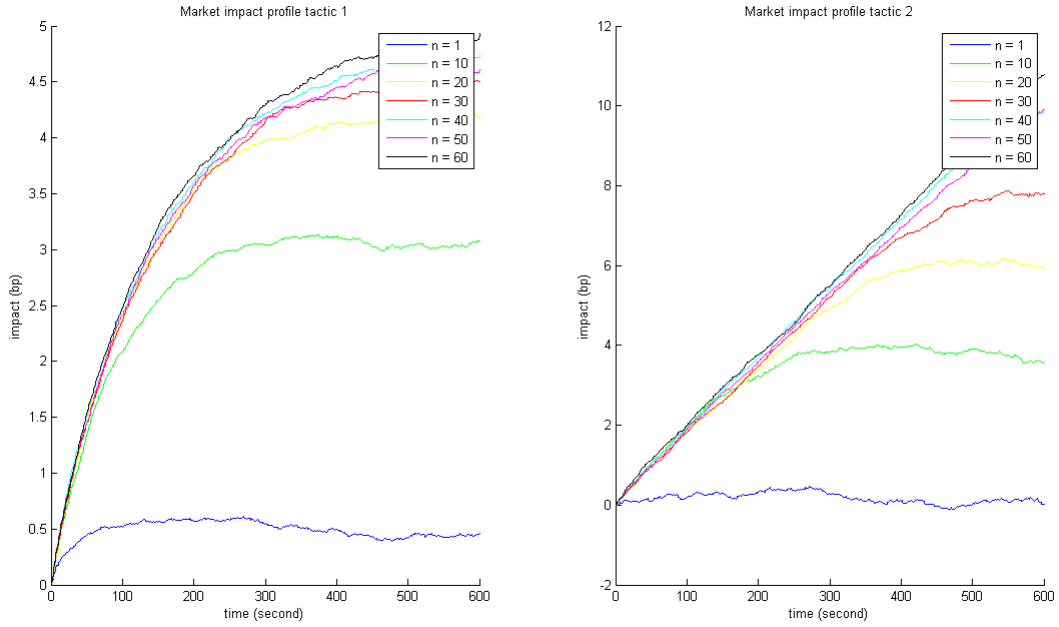

FIGURE 23 – Market Impact Profile

2.2 A Controlled View on the Limit Orderbook dynamics using a MFG approach

This Section has been published as a part of *Efficiency of the Price Formation Process in Presence of High Frequency Participants : a Mean Field Game analysis* (in review process at MAFE : Mathematics and Financial Economics) [Lachapelle *et al.* , 2013]

2.2.1 A simple single-queue model with anticipations

The purpose of introducing first a single queue model is didactic and does not aim at directly providing insights on order book modeling. However we believe this single queue is the occasion to introduce some key concepts, such as endogenous strategic entries of agents that anticipate the future. Consequently, sellers entering the system are also called players since we locate the modeling approach in the game theoretic framework (agents perform actions optimizing their respective pay-off).
In particular, when new sellers arrive, they look at the queue size and decide whether to enter the queue or not (action), after considering their expected pay-off (value function assessment).

With this simplified model we introduce anticipatory behaviors in a very stylized one-sided order book, where patient sellers arrive at exogenous Poisson rate and where the arrival rate of impatient buyers increases as soon as the queue size increases. We will finally use it to provide insights on the modeling of distinct execution protocols, namely process sharing and First In First Out protocol.

The model. The arrival rate of players is continuous and stochastic. In this simplified model, it is exogenous.
- As usual, they arrive following a Poisson process with intensity λ.
- Impatient buyers arrive at rate $\mu(x) \geq 0$, a given increasing function of x ; i.e. the more patient sellers in the queue, the higher arrival rate of impatient buyers.
- The unit size of an order in the queue is q. The queuing discipline is a process sharing one (with no priority), i.e. individual service in a queue of size x is worth q/x. In terms of trading rules, one may think about a pro-rata one [Field & Large, 2008].
- The pay-off gained by a player per unit of order is a nonnegative decreasing function of the queue size : $P(x)$. Typical cases are $P(x) := p > 0$ and $P(x) = 1/x$. On the other hand, there is a cost c of waiting in the queue.

Now, as usual in game theory, there is a value function u for any player. The value function depends upon the queue size x. It is the expected Profit & Loss (P&L) of a player entering the queue. Note that we assume that agents are risk neutral and that their reservation utility is set to 0, which means that an agent decides to enter the queue as soon as the value function is positive : $u(x) > 0$.

The value function dynamic comes from an infinitesimal expression of events impacting it :
- a newcomer enters the queue as soon as $u(x) > 0$ (remember u is the "*expected value received if you enter the queue*").
- the queue is consumed by a Poisson process of intensity $\mu(x)$. Each time an order already waiting in the queue is partially executed (according to a prorata rule) : its owner will sell q/x shares for a price $P(x)$. The new expected value for a participant waiting in the queue in this case is thus $q/x \cdot P(x) + (1 - q/x) \cdot u(x - q)$ (i.e. the first part of the expression comes from the sell of q/x shares and the second one from the expected value of the queue that is now of size $x - q$).
- in all other cases, the expected value does not change.
- the waiting cost is proportional to q (the size of the orders) ; it decreases the expected value of u by $c\,q\,dt$, where dt is the time unit.

More formally (with the notations $dN_t^{\mu(x)}$ for the queue-consuming point process and $dN_t^{\lambda \mathbf{1}_{\{u(x)>0\}}}$ for the queue-increasing one), the dynamics of the value function $u(x)$ is driven by :
- a stochastic differential equation on the size of the queue x_t :

$$dx_t = \left(dN_t^{\lambda \mathbf{1}_{\{u(x)>0\}}} - dN_t^{\mu(x)} \right) q,$$

- and a running cost function $J(x)$:

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

The additive waiting costs are compatible with the very short time scale having a sense for orderbook dynamics. It can be noted here that another cost function $\mathcal{J}$ could be defined here as :

$$d\mathcal{J}(x_t) = \left[\omega(q, x_t) P(x_t) + (1 - \omega(q, x_t)) \mathcal{J}(x_t - q) \right] dN_t^{\mu(x)} - cq\,dt,$$

where $\omega(q, x_t)$ is a random variable taking value 1 with a probability q/x and 0 otherwise. In such a case, instead of a prorata rule, we will have a trading rule for which an order is fully executed with a probability q/x, or not at all. This case covers the trading model of [Roşu, 2009], in which the orderbook matching rule is FIFO (First

In, First Out), but any agent can modify and reinsert his order at any time. In such a case the probability for one specific agent to be first in the queue (and thus be fully filled), is q/x.
Since $\mathbb{E}dJ = \mathbb{E}d\mathcal{J}$, the emerging dynamics are the same.

With such a formalism, the value function can be defined as $u(X) = \mathbb{E}\int_{t=0}^{T} J(x_t)\,dt$ given $x_0 = X$, with T "large enough" at the intraday time sale.

Equilibrium : the value function equation. In this paragraph we introduce the equation verified by the value function at the equilibrium. Below we detail the equilibrium equation for each probability event.

$$\begin{aligned} u(x) = \; & (1 - \lambda \mathbf{1}_{\{u(x)>0\}} dt - \mu(x)dt) \cdot u(x) && \leftarrow \text{ nothing happens} \\ & + \lambda \mathbf{1}_{\{u(x)>0\}} dt \cdot u(x+q) && \leftarrow \text{ new queue entrance} \\ & + \mu(x)dt \cdot \Big(\frac{q}{x}P(x) + (1 - \frac{q}{x})u(x-q)\Big) && \leftarrow \text{ service} \\ & - cq\,dt && \leftarrow \text{ waiting cost} \end{aligned} \tag{35}$$

We can perform a Taylor expansion for small q in the discrete equation above. In this way we derive the following differential equation :

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

where the second order term is the last one (blue term).

First order analysis. Before approximating numerically the solution to (35), we propose to get some insights on the shape of the solution by doing a first order analysis. More precisely, the solution to the queuing system described above is characterized by the sign of the value function u. Consequently we are interested in finding potential sign switching points of u.
The core modeling ingredient is the value of the Poisson arrival rate λ relative to $\mu(x)$.
For the first order analysis we look at the first order equation :

$$0 = \frac{\mu(x)}{x}\Big(P(x) - u(x)\Big) - c + \Big(\lambda \mathbf{1}_{\{u(x)>0\}} - \mu(x)\Big)u'(x). \tag{36}$$

Remark 1 *Let us remark that equation (36) corresponds to a trivial shared risk Mean Field Game monotone system with $N = 1$, as described in the previous section. Note that in the framework of this model, the mean field aspect does not come from the continuum of agents (for every instant, the number of players is finite), but rather to the stochastic continuous structure of entries and exits of players.*

Now we look at the case where the stylized limit order book presented here has an *infinite resiliency*, meaning once the orderbook is partially consumed by a marketable order, the remaining liquidity rearranges itself to fill the gap. Moreover, we will consider the non degenerated case where sellers arrive at rate λ, larger than the exogenous consuming rate $\mu(x)$, for all x.

An example with anticipatory behavior. Assume the arrival rate of buyers has the specificity to take two values :
- a low value μ_1 below a certain queue size threshold S,
- a higher value μ_2 ($\mu_2 > \mu_1$), above the threshold S.

As a function depending upon the queue size variable x, it reads :

$$\mu(x) = \mu_1 \mathbf{1}_{x<S} + \mu_2 \mathbf{1}_{x \geq S},\ 0 \leq \mu_1 < \mu_2.$$

Here there are at least two points where u changes sign :

$$x_1^* = \mu_1 P(x_1^*)/c \text{ and } x_2^* = \mu_2 P(x_2^*)/c. \tag{37}$$

Figure 24 shows the plot of the solution (numerical approximation of the solution to equation (35)) for a certain set of parameters (for P constant). We can see that the first switching point is close to the first order approximation x_1^*, while the last sign switch significantly deviates from the first order approximation x_2^*. It means that higher order terms have a non-negligible effect.
But most importantly, we observe that there is another sign switch strictly below the threshold S. The existence of such a switching point means that players anticipate the improved service before the threshold is reached. Indeed, their value function becomes positive meaning that players enter the queue strictly before the improved service starts. This is why we talk about an *anticipation switching point*. Consequently, we can conclude that at the equilibrium, the strategical players adopt an anticipatory behavior.

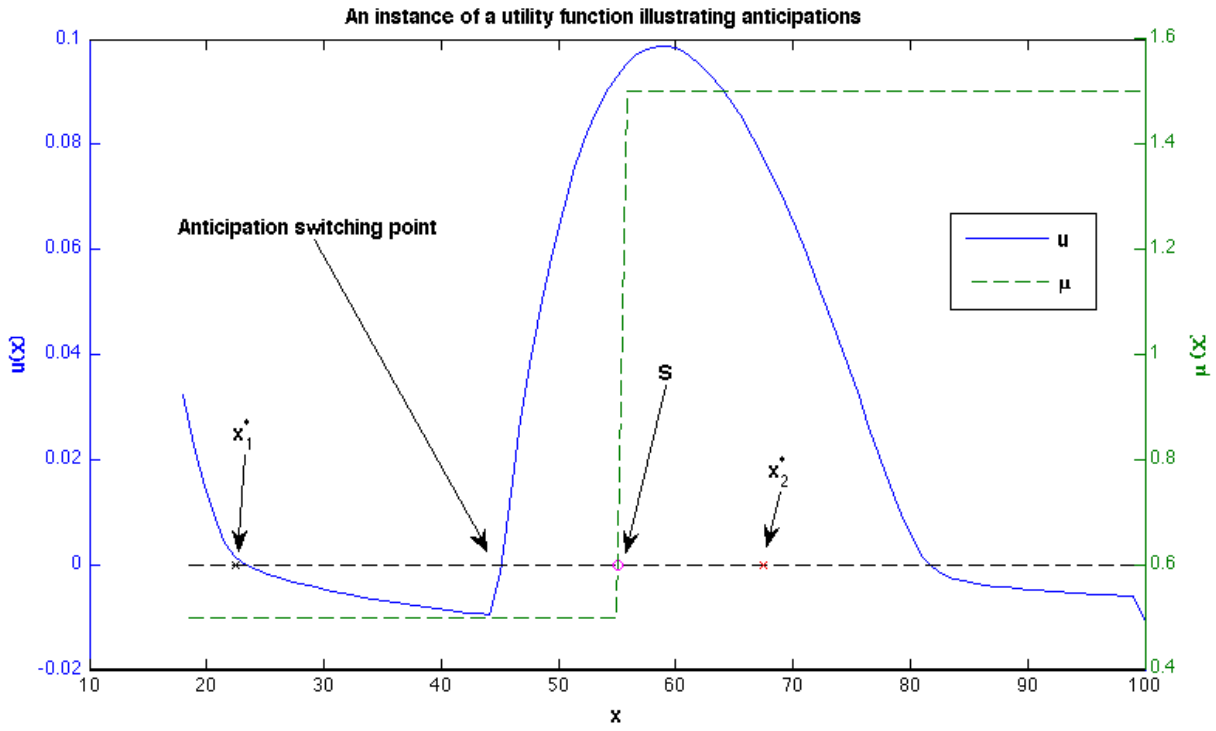

FIGURE 24 – Here we notice that there is a point strictly before S where u switches from negative to positive. It means that players anticipate improved service μ_2 and therefore are newly interested in entering the queue.

First In First Out model. Finally we want to show that our approach allows to model distinct execution processes, and how the resulting equilibrium equations are impacted.
To do so, we consider the First In First Out (FIFO) protocol. This is the only change we make in the model. To consider such a priority protocol, we have to introduce a new variable z denoting the position of a trader in the queue of size x. Consequently the problem becomes bi-dimensional.
The equation becomes :

$$\begin{aligned} u(z,x) = \; & (1-\lambda \mathbf{1}_{u(x,x)>0}dt - \mu(x)dt)\cdot u(z,x) && \leftarrow \text{ nothing happens} \\ & + \lambda \mathbf{1}_{\{u(x,x)>0\}}dt \cdot u(z,x+q) && \leftarrow \text{ new queue entrance} \\ & + \mu(x)dt \cdot u(z-q,x-q) && \leftarrow \text{ execution of the first order} \\ & - cqdt && \leftarrow \text{ waiting cost,} \end{aligned} \tag{38}$$

in the domain $q < z < x$, and the boundary condition for $z = q$ is :

$$\begin{aligned} u(q,x) = \; & (1-\lambda \mathbf{1}_{u(x,x)>0}dt - \mu(x)dt)\cdot u(q,x) && \leftarrow \text{ nothing happens} \\ & + \lambda \mathbf{1}_{\{u(x,x)>0\}}dt \cdot u(q,x+q) && \leftarrow \text{ new queue entrance} \\ & + \mu(x)dt \cdot P(x) && \leftarrow \text{ execution of the first order} \\ & - cqdt && \leftarrow \text{ waiting cost.} \end{aligned} \tag{39}$$

System (38-39) can be easily solved numerically.

2.2.2 The Mean Field Game Orderbook Model

The matching mechanisms of order books. One of the roles of financial markets is to *form prices* according to the balance between offer and demand. In modern markets, this mechanism takes place inside *electronic order books where multilateral trading takes place*. They implement the following dynamic :

1. buyers and sellers can send electronic messages to a "*matching engine*". These messages, called *orders*, contain a side ("buy" or "sell"), a limit price and a quantity.
2. The matching engine contains a list of all pending orders it received in its memory. When it receives a new buy (respectively sell) order, it looks if pending sell (resp. buy) orders at a lower (resp. higher) price are available.
   - If it is the case, it generates transactions between the owner of the incoming order and the owners of the compatible opposite orders, and removes the corresponding quantities in its list of pending orders ;
   - if the incoming order has a remaining quantity, it is inserted in the list of pending orders.

The list of pending orders waiting in the matching engine is called its "*limit order book*" (LOB).

During the matching process, it is possible that the quantity of an incoming order does not match exactly the quantity made available at a compatible price (i.e. lower prices for a buy order and higher prices for a sell order) by opposite orders in the order book. To handle such cases, matching engines need to implement a priority mechanism. The most used (see [Mendelson & Amihud, 1991] for more details) are :

- *time priority* : the "oldest" pending orders in the order book are matched first ;
- *size priority* : the largest pending orders are matched first in case of competition between resting orders at the same price ;
- *pro rata* : pending orders are matched for a fraction of their quantity proportionally to their relative size to the one of the whole queue (see [Field & Large, 2008]).

Each trading platform discloses its matching mechanism in detail to market participants in a *rulebook* (like [Euronext, 2006]).

Rendering different trading styles in an order book model. To understand the features of our MFG model, we will first study its dynamics in a market with homogenous participants. Since we are in a MFG framework, it will render a continuum of agents, at this stage they share the same macroscopic parameters :

- the same messaging intensity λ,
- the same size of orders they send q,
- the same waiting cost c.

Beside, we enrich the model with one more feature : the use of SOR (Smart Order Router). A Smart Order Router (see [Foucault & Menkveld, 2008] for an efficiency study or [Lehalle *et al.* , 2013] for a generic presentation) is a device containing a software dedicated to "smartly route" orders. In our model, only SOR users will be able to act strategically instead of being blindly impatient.

It can be considered that agents not using a SOR have an infinite waiting cost. Since institutional investors take decisions independently of the current state of the orderbook, it is realistic to consider that a fraction of them will not take time to implement sophisticated microscopic strategies on some of their orders.

The proportion of market participants using a SOR (i.e. not infinitely impatient market participants) will be parametrized thanks to a specific flow of intensity λ^-.

	Instit. Investors	HFT
Order size	large	small
Speed	normal	fast
SOR	often used	always used

TABLE 1 – Qualitative modeling of Institutional Investors and HFT.

In a second stage we will mix heterogenous agents, with different behaviours summarized in Table 1 :

1. *Institutional investors*, trading large quantities not using systematically a SOR ;
2. *HFT (High Frequency Traders)*, faster than the former participants, using smaller orders, more patient (in the sense that they bare a lower cost per share waiting in a queue), and all of them using a SOR.

Transaction price. The *market price* will be centered on a constant P. The *market depth* is δ, meaning that no transaction will take place at a price lower than $P-\delta$ or higher than $P+\delta$. The (time varying) size of the bid queue (waiting buy orders) is Q_t^b and the size of the ask one (waiting sell orders) is Q_t^a.

When a market (buying) order hits the ask queue, the transaction price is p^{buy} and when the bid queue is lifted by a market (selling) order, the transaction price is p^{sell}. The price takes into account instantaneous queue size adjustments depending upon the order size q.

$$p_q^{\text{buy}}(Q_t^a) := P + \frac{\delta q}{Q_t^a - q}, \quad p_q^{\text{sell}}(Q_t^b) := P - \frac{\delta q}{Q_t^b - q} \tag{40}$$

Qualitatively, it means that the market impact is linear. Boundary conditions, to be introduced later, impose $Q_t^a, Q_t^b > q$, so that there is no definition problem of the transaction prices

Value functions. The value function for a trader submitting a buy order in the bid queue is $v(Q_t^a, Q_t^b)$ and the one of a sell order in the ask queue is $u(Q_t^a, Q_t^b)$. In the model agents have risk-neutral preferences, thus the utility functions coincide with price expectations.

Orders arrival rates. We distinguish between SOR and non-SOR orders. The proportion of these two types of orders is exogenous, and set as an input of the model.

Buy and sell SOR orders arrive according to two Poisson processes with intensity λ_{buy} and λ_{sell}. Several cases can be considered :

1. Homogeneous Poisson processes :
$$\lambda_{\text{buy}} = \lambda_{\text{sell}} = \lambda. \tag{41}$$
2. Heterogeneous (in space) Poisson processes
$$\lambda_{\text{buy}} = \lambda f(Q_t^b),\ \lambda_{\text{sell}} = \lambda f(Q_t^a),$$
where $f(x)$ is a decreasing function. Typical instances are $f(x) = 1/x$, $f(x) = \mathbf{1}_{x \leq \bar{Q}}$ likewise.

However, we will focus in this paper on the homogenous case.
Let us remark that the previous rates could be endogenized and set as the result of an optimization problem involving the utility functions, consequently depending upon the queue sizes $Q_t^\bullet$.

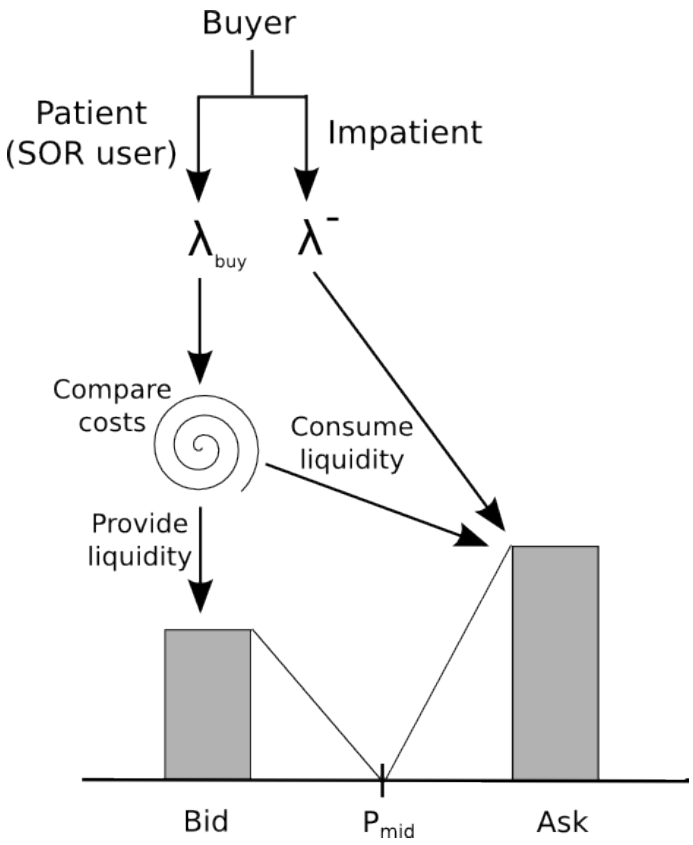

FIGURE 25 – Idealized diagram of the decision tree of agents in the model.

Non-SOR orders (i.e. belonging to very impatient investors or traders) are always liquidity remover, with arriving rate $2\lambda^-$ (equally distributed between buyers and sellers).

Market participants decision processes. When a buy (resp. sell) order arrives, its owner has to make a *routing decision* (see Figure 25 for an idealized diagram of this process) :

– if $v(Q_t^a, Q_t^b + q) < p^{\text{buy}}(Q_t^a)$ (resp. $u(Q_t^a + q, Q_t^b) > p^{\text{sell}}(Q_t^b)$) it is more valuable to route the order to the bid (resp. ask) queue (i.e. sending a limit order). In such a case the order will be a *Liquidity Provider (LP)*. We define symmetrically *Liquidity Consumer (LC)* orders. This decision is formalized in the model by setting the variable $R^{\oplus}_{\text{buy}}(v, Q_t^a, Q_t^b + q)$ to 1 when $v(Q_t^a, Q_t^b + q) < p^{\text{buy}}(Q_t^a)$, and to zero otherwise :
$$\begin{aligned} R^{\oplus}_{\text{buy}}(v, Q_t^a, Q_t^b + q) &:= \mathbf{1}_{v(Q_t^a, Q_t^b+q) < p^{\text{buy}}(Q_t^a)}, \text{ LP buy order} \\ R^{\oplus}_{\text{sell}}(u, Q_t^a + q, Q_t^b) &:= \mathbf{1}_{u(Q_t^a+q, Q_t^b) > p^{\text{sell}}(Q_t^b)}, \text{ LP sell order.} \end{aligned} \tag{42}$$

– otherwise the order goes Liquidity Consumerly to the ask (resp. bid) queue to obtain a trade. It will be a liquidity remover in this case :
$$\begin{aligned} R^{\ominus}_{\text{buy}}(Q_t^a, Q_t^b) &:= 1 - R^{\oplus}_{\text{buy}}(Q_t^a, Q_t^b), \text{ LC buy order} \\ R^{\ominus}_{\text{sell}}(Q_t^a, Q_t^b) &:= 1 - R^{\oplus}_{\text{sell}}(Q_t^a, Q_t^b), \text{ LC sell order.} \end{aligned}$$
The price of such a transaction is p^{buy} (resp. p^{sell}) as defined by equality (40). Note that we omit the dependence on u, v when it is unnecessary for the understanding of the equations.

We impose the following boundary conditions :

$$\begin{aligned} \text{Min liquidity condition} &: R^{\oplus}_{\text{buy}}(r, Q^b_t) = 1,\ R^{\oplus}_{\text{sell}}(Q^a_t, r) = 1,\ \forall r \le q, \\ \text{Technical condition} &: R^{\oplus}_{\text{buy}}(Q^a_t, r) = 1,\ R^{\oplus}_{\text{sell}}(r, Q^b_t) = 1,\ \forall r < q. \end{aligned} \tag{43}$$

In particular, conditions (43) ensure that $(Q^a_0, Q^b_0) \ge (q, q) \Rightarrow (Q^a_t, Q^b_t) \ge (q, q),\ \forall t > 0$.

SDE formalism. The dynamics associated with this matching mechanism can be written :

– for the size of the ask queue Q^a_t :

$$dQ^a_t = \left(dN^{\lambda_{\text{sell}} R^{\oplus}_{\text{sell}}} - (dN^{\lambda_{\text{buy}} R^{\ominus}_{\text{buy}}} + dN^{\lambda^-}) \right) q,$$

– and for the cost function at the ask :

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

Again, with T large enough, $u(\mathcal{Q}^a, \mathcal{Q}^b) = \mathbb{E} \int_{t=0}^{T} J(Q^a_t, Q^b_t)\, dt$ given $Q^a_0 = \mathcal{Q}^q, Q^b_0 = \mathcal{Q}^b$.

Matching process. The matching process is close to a pro-rata one [Field & Large, 2008] : in case of a liquidity consuming buy order of size Q to be matched, all market participants having a quantity q resting in the ask queue will obtain a transaction for a fraction $Q \cdot q / Q^a_t$ of its order at price $p^{\text{buy}}(Q^a_t)$, the remaining quantity staying in the orderbook.
At a first glance one may think that this matching process will induce intricate terms in the equations, but in fact it will not since we only consider utilities by units of good transactions.

– The orderbook shape is assumed to be linear (in the price), meaning that if a newcomer decide to provide liquidity to the market, her order will be split proportionally to the liquidity already present in the book : the orderbook will remain linear in price with a higher slope.
– Hence when a Liquidity Consumer order occurs, it will partially fill all Liquidity Provider orders according to a proportional rule.

2.2.3 Introducing the equations of the Price Formation Process dynamics

We characterize an equilibrium via recursive equations of the expected value of future payoffs (value functions).

$$\begin{aligned} & u(Q^a_t, Q^b_t) = && \\ & (1 - \lambda_{\text{buy}} dt - \lambda_{\text{sell}} dt - 2\lambda^- dt)\ u(Q^a_t, Q^b_t) && \leftarrow \text{nothing} \\ & + (\lambda_{\text{sell}} R^{\ominus}_{\text{sell}}(u, Q^a_t + q, Q^b_t) + \lambda^-) dt\ u(Q^a_t, Q^b_t - q) && \leftarrow \text{sell order, LC} \\ & + \ \lambda_{\text{sell}} R^{\oplus}_{\text{sell}}(u, Q^a_t + q, Q^b_t) dt\ u(Q^a_t + q, Q^b_t) && \leftarrow \text{sell order, LP} \\ & + (\lambda_{\text{buy}} R^{\ominus}_{\text{buy}}(v, Q^a_t, Q^b_t + q) + \lambda^-) dt \cdot [&& \leftarrow \text{buy order, LC} \\ & \qquad \underbrace{\frac{q}{Q^a_t} p^{\text{buy}}(Q^a_t)}_{\text{trade part (ask)}} \quad + \underbrace{(1 - \frac{q}{Q^a_t})\, u(Q^a_t - q, Q^b_t)}_{\text{removing (ask)}}] && \\ & + \ \lambda_{\text{buy}}\ R^{\oplus}_{\text{buy}}(v, Q^a_t, Q^b_t + q) dt\ u(Q^a_t, Q^b_t + q) && \leftarrow \text{buy order, LP} \\ & - \ c_a q\, dt. && \leftarrow \text{cost to maintain inventory} \end{aligned} \tag{44}$$

Symmetrically, we have :

$$\begin{aligned} & v(Q^a_t, Q^b_t) = && \\ & (1 - \lambda_{\text{buy}} dt - \lambda_{\text{sell}} dt - 2\lambda^- dt)\ v(Q^a_t, Q^b_t) && \leftarrow \text{nothing} \\ & + (\lambda_{\text{buy}} R^{\ominus}_{\text{buy}}(v, Q^a_t, Q^b_t + q) + \lambda^-) dt\ v(Q^a_t - q, Q^b_t) && \leftarrow \text{buy order, LC} \\ & + \ \lambda_{\text{buy}} R^{\oplus}_{\text{buy}}(v, Q^a_t, Q^b_t + q) dt\ v(Q^a_t, Q^b_t + q) && \leftarrow \text{buy order, LP} \\ & + (\lambda_{\text{sell}} R^{\ominus}_{\text{sell}}(u, Q^a_t + q, Q^b_t) + \lambda^-) dt \cdot [&& \leftarrow \text{sell order, LC} \\ & \qquad \underbrace{\frac{q}{Q^b_t} p^{\text{sell}}(Q^b_t)}_{\text{trade part (bid)}} \quad + \underbrace{(1 - \frac{q}{Q^b_t})\, v(Q^a_t, Q^b_t - q)}_{\text{removing (bid)}}] && \\ & + \ \lambda_{\text{sell}}\ R^{\oplus}_{\text{sell}}(u, Q^a_t + q, Q^b_t) dt\ v(Q^a_t + q, Q^b_t) && \leftarrow \text{sell order, LP} \\ & - \ c_b q\, dt. && \leftarrow \text{cost to maintain inventory} \end{aligned} \tag{45}$$

Remind that R_{buy} and R_{sell} are functionals of Q^a and Q^b and also implicitly depends on u and v. Of course the previous principles hold for $Q^a_t, Q^b_t > q$, which is always the case thanks to conditions (43). In the equations above, c_a and c_b are positive constants modeling the cost to maintain inventory per unit, that is the cost of never being processed once waiting in the queue.

Symmetric case. In the case where $\lambda_{\text{sell}} = \lambda_{\text{buy}} = \lambda$, and $c_a = c_b = c$, we have the following results.
For the sake of simplicity we will often use new notations for the queue size variables : x and y stand for Q_a and Q_b.

Lemma 2

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

This simple symmetry result is useful to get a necessary condition for the solution.

Proposition 6 *If system (44)-(45) has a unique solution (u,v), then*

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

That is, u and v are antisymmetric up to the constant P.

Take Equation (44) then perform the change of variable $w(y,x) = 2P - u(x,y)$, then apply the previous Lemma, switch the roles of x and y and multiply by -1. Then you get equation (45), hence the conclusion.

Continuous approximation. In this paragraph we formally derive differential equations corresponding to the PFP dynamic discrete equations (44-45) as presented in the previous section. Hopefully, this will lead us to get easily some qualitative insights on the solutions u and v.
To do so, we write the Taylor expansion of order 1 at the point (x,y) in system (44-45). After a quick computation, we get the following system of Partial Differential Equations (PDEs). Note that for the sake of simplicity we shorten the notations as follows : *sell* becomes s, *buy* becomes b, Q^a becomes x and Q^b becomes y.

$$\begin{aligned}
\text{(Ask)}\quad 0 &= [(\lambda_b R^{\ominus}_b + \lambda^-)\frac{1}{x}(p^b(x) - u) - c_a] \\
&\quad + [\lambda_s R^{\oplus}_s - \lambda_b R^{\ominus}_b - \lambda^-] \cdot \partial_x u + [\lambda_b R^{\oplus}_b - \lambda_s R^{\ominus}_s - \lambda^-] \cdot \partial_y u, \\
\text{(Bid)}\quad 0 &= [(\lambda_s R^{\ominus}_s + \lambda^-)\frac{1}{y}(p^s(y) - v) + c_b] \\
&\quad + [\lambda_s R^{\oplus}_s - \lambda_b R^{\ominus}_b - \lambda^-] \cdot \partial_x v + [\lambda_b R^{\oplus}_b - \lambda_s R^{\ominus}_s - \lambda^-] \cdot \partial_y v.
\end{aligned}$$

Recall that u, v, R_b, R_s are estimated at (x,y) and R_b depends upon v, resp. R_s depends upon u. Consequently, R_b and R_s are the coupling terms in the PDE system (Ask)-(Bid).
The system has to be understood locally in the four regions

$$\begin{aligned}
R^{++} &= \{(x,y), R^{\oplus}_s(x,y) = R^{\oplus}_b(x,y) = 1\}, R^{--} = \{(x,y), R^{\ominus}_s(x,y) = R^{\ominus}_b(x,y) = 1\}, \\
R^{+-} &= \{(x,y), R^{\oplus}_s(x,y) = R^{\ominus}_b(x,y) = 1\}, R^{-+} = \{(x,y), R^{\ominus}_s(x,y) = R^{\oplus}_b(x,y) = 1\}.
\end{aligned}$$

Now we can write the general form of the first order system of coupled PDEs.

$$0 = \gamma_a(u,v,x,y) + \alpha(u,v,x,y)\partial_x u + \beta(u,v,x,y)\partial_y u \tag{46}$$
$$0 = \gamma_b(u,v,x,y) + \alpha(u,v,x,y)\partial_x v + \beta(u,v,x,y)\partial_y v, \tag{47}$$

where γ_a, γ_b, α, β have some good symmetry properties to be described later on.

The MFG framework. The model is of course a Mean Field Game. As mentioned in section 2.2, there are continuous entries and exits of players (modeled with Poisson processes). Therefore the basis assumptions are fulfilled : continuum of atomized and anonymous players.

Second order terms. We kept only the first order terms in the equations. The second order terms to be added to the equations are :

$$\text{In (Ask)}\ \frac{q^2}{2}\left[\frac{2}{x}(\lambda_b R^{\ominus}_b + \lambda^-)\partial_x u + \lambda^- \Delta u + (\lambda_s R^{\oplus}_s + \lambda_b R^{\ominus}_b)\partial_{xx} u + (\lambda_s R^{\ominus}_s + \lambda_b R^{\oplus}_b)\partial_{yy} u\right],$$

$$\text{In (Bid)}\ \frac{q^2}{2}\left[\frac{2}{y}(\lambda_s R^{\ominus}_s + \lambda^-)\partial_y v + \lambda^- \Delta v + (\lambda_s R^{\oplus}_s + \lambda_b R^{\ominus}_b)\partial_{xx} v + (\lambda_s R^{\ominus}_s + \lambda_b R^{\oplus}_b)\partial_{yy} v\right].$$

2.2.4 Equilibrium analysis

Change of variables. From now on we focus on the symmetric case where $\lambda_s = \lambda_b = \lambda$ and $c_a = c_b = c$. First it is convenient to notice that in this important case, we have the following property :

$$\alpha = \beta = [\lambda(R_s^{\oplus}(u,x,y) - R_b^{\ominus}(v,x,y)) - \lambda^-].$$

We will see later that this property allows to solve the problem thanks to the characteristics method.
There is a very welcome change of variables that we will use throughout this section. We define

$$\tilde{u} = (u-P)/q \quad \text{and} \quad \tilde{v} = (v-P)/q. \tag{48}$$

Then the (Ask)-(Bid) system reads

$$\begin{aligned} 0 &= [(\lambda\tilde{R}_b^{\ominus} + \lambda^-)\frac{1}{x}(\frac{\delta}{x-q} - \tilde{u}) - \frac{c}{q}] + [\lambda\tilde{R}_s^{\oplus} - \lambda\tilde{R}_b^{\ominus} - \lambda^-] \cdot (\partial_x\tilde{u} + \partial_y\tilde{u}), \\ 0 &= [(\lambda\tilde{R}_s^{\ominus} + \lambda^-)\frac{1}{y}(\frac{-\delta}{y-q} - \tilde{v}) + \frac{c}{q}] + [\lambda\tilde{R}_s^{\oplus} - \lambda\tilde{R}_b^{\ominus} - \lambda^-] \cdot (\partial_x\tilde{v} + \partial_y\tilde{v}). \end{aligned} \tag{49}$$

Proposition 7 *Assume that system (49) admits a unique solution $(\tilde{u}, \tilde{v})$, then it is antisymmetric, that is :*

$$\forall (x,y),\ \tilde{v}(x,y) = -\tilde{u}(y,x).$$

The general form of the system (49) is as follows :

$$0 = \gamma(\tilde{u},\tilde{v},x,y) + \alpha(\tilde{u},\tilde{v},x,y)(\partial_x\tilde{u} + \partial_y\tilde{u}) \tag{50}$$

$$0 = -\gamma(\tilde{v},\tilde{u},y,x) + \alpha(\tilde{u},\tilde{v},x,y)(\partial_x\tilde{v} + \partial_y\tilde{v}). \tag{51}$$

First Order Analysis. Here we explore formally some aspects of the first order approximation to the solution.
The key point of the analysis is that in the two equations of system (49), the derivative terms are the same, so that we conclude that the characteristics satisfy

$$\dot{x} = \dot{y} = \alpha \Rightarrow x = y + k.$$

Note that the reasoning of this paragraph holds on the region below the diagonal, but can be trivially extended to the whole domain by symmetry arguments.
We heuristically suppose that for a given k, and along the characteristic line $y = x-k$, there is a first point $M_0 = (x_0, y_0)$ where the sellers become Liquidity Consumer, that is M_0 is a point at the boundary of the regions R^{++} and R^{-+}. Then there is a second point $M_1 = (x_1, y_1)$, with $x_1 \geq x_0$ and $y_1 \geq y_0$ where the buyers become Liquidity Consumer, that is M_1 is a point at the boundary of the regions R^{-+} and R^{--}.
First recall that :

$$\begin{aligned} R^{++} &\text{ is defined by } R_s^{\oplus} = 1 \text{ and } R_b^{\oplus} = 1, \\ R^{-+} &\text{ is defined by } R_s^{\oplus} = 0 \text{ and } R_b^{\oplus} = 1, \\ R^{--} &\text{ is defined by } R_s^{\oplus} = 0 \text{ and } R_b^{\oplus} = 0. \end{aligned}$$

We can write the differential equations on the three regions mentioned above :

$$\begin{aligned} (A_{R^{++}}) \quad 0 &= \Big[\frac{\lambda^-}{x}(\frac{\delta}{x-q} - \tilde{u}) - \frac{c}{q}\Big] + [\lambda - \lambda^-] \cdot (\partial_x\tilde{u} + \partial_y\tilde{u}), \\ (B_{R^{++}}) \quad 0 &= \Big[\frac{\lambda^-}{y}(\frac{-\delta}{y-q} - \tilde{v}) + \frac{c}{q}\Big] + [\lambda - \lambda^-] \cdot (\partial_x\tilde{v} + \partial_y\tilde{v}), \\ (A_{R^{-+}}) \quad 0 &= \Big[\frac{\lambda^-}{x}(\frac{\delta}{x-q} - \tilde{u}) - \frac{c}{q}\Big] + [-\lambda^-] \cdot (\partial_x\tilde{u} + \partial_y\tilde{u}), \\ (B_{R^{-+}}) \quad 0 &= \Big[\frac{\lambda+\lambda^-}{y}(\frac{-\delta}{y-q} - \tilde{v}) + \frac{c}{q}\Big] + [-\lambda^-] \cdot (\partial_x\tilde{v} + \partial_y\tilde{v}), \\ (A_{R^{--}}) \quad 0 &= \Big[\frac{\lambda+\lambda^-}{x}(\frac{\delta}{x-q} - \tilde{u}) - \frac{c}{q}\Big] + [-\lambda - \lambda^-] \cdot (\partial_x\tilde{u} + \partial_y\tilde{u}), \\ (B_{R^{--}}) \quad 0 &= \Big[\frac{\lambda+\lambda^-}{y}(\frac{-\delta}{y-q} - \tilde{v}) + \frac{c}{q}\Big] + [-\lambda - \lambda^-] \cdot (\partial_x\tilde{v} + \partial_y\tilde{v}). \end{aligned} \tag{52}$$

The equations are relatively simple in each region. The tricky point is, as always, to stick together the solutions of each region. First we compute the boundaries of the regions.

First order boundaries. Let's note M_0 the first order boundary between R^{++} and R^{-+} and M_1 between R^{-+} and R^{--}.

Proposition 8 (First order boundary between R^{++} and R^{-+}) *The diagonal point of the boundary M_0 is the point*

$$(x_0^*, x_0^*) = (q + \sqrt{q^2 + 8/\eta})/2 \tag{53}$$

and the boundary M_0 is given by the set of points (x_0, y_0) verifying :

$$(x_0, y_0) = \left(x_0, l(x_0) := q + \left(\eta x_0 - \frac{1}{x_0 - q}\right)^{-1}\right),\ \forall x_0 \geq x_0^*, \tag{54}$$

where $\eta := c/(\delta q \lambda^-)$.

Proposition 9 (First order boundary between R^{-+} and R^{--}) *The boundary M_1 is defined by the set of points $(y_1 + x_0 - l(x_0), y_1),\ \forall x_0 \geq x_0^*$, where*

$$y_1 \text{ verifies } f_{x_0 - l(x_0)}(y_1) = \frac{\delta}{y_1 + x_0 - l(x_0) - q}. \tag{55}$$

Figure 26 exhibits an instance of the first order curves. We observe that near the diagonal, there is a region where several solutions could happen. The first order analysis thus shows the global form of the shape of the solution (since it is based on the curves $M_0,\ M_1$), and that considering higher order terms is necessary to understand what happens in the region near the diagonal.

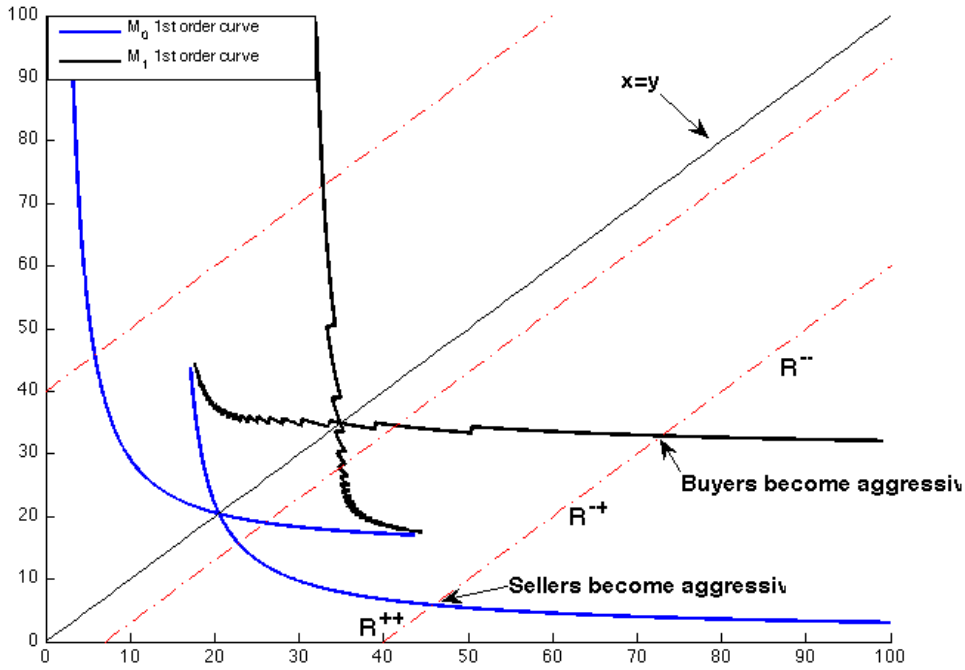

FIGURE 26 – First order decision curves

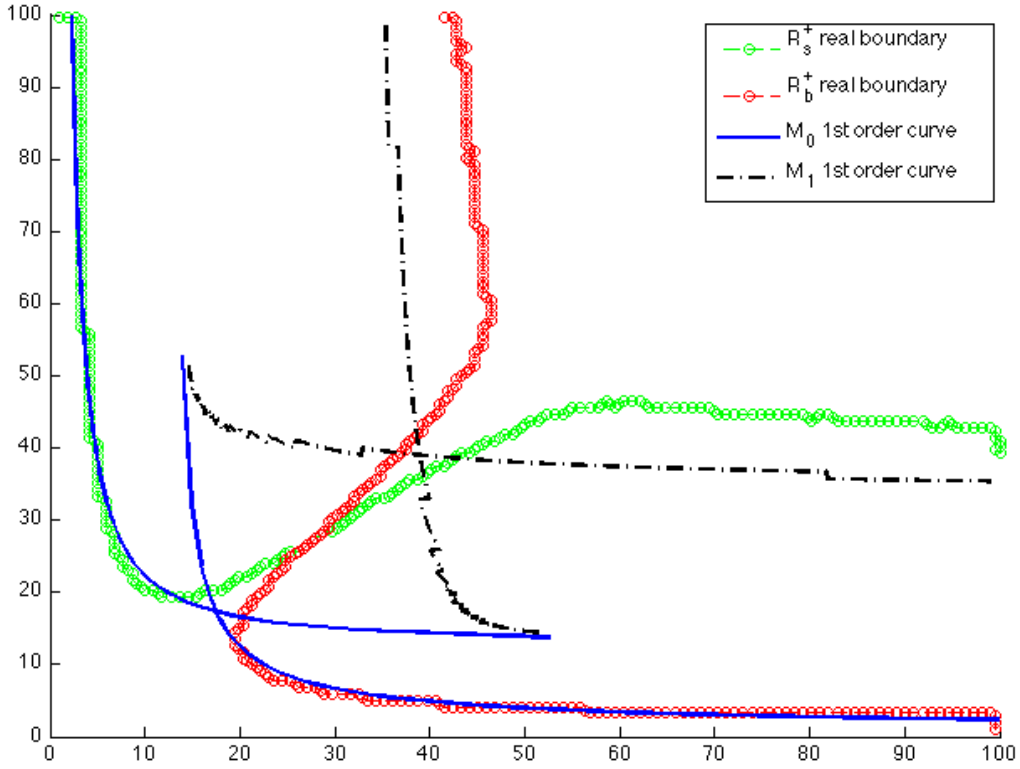

FIGURE 27 – Test 1 : the numerically computed switching curves (red and green dots) tries to conciliate the curves analytically computed at order 1 (dotted and solid lines).

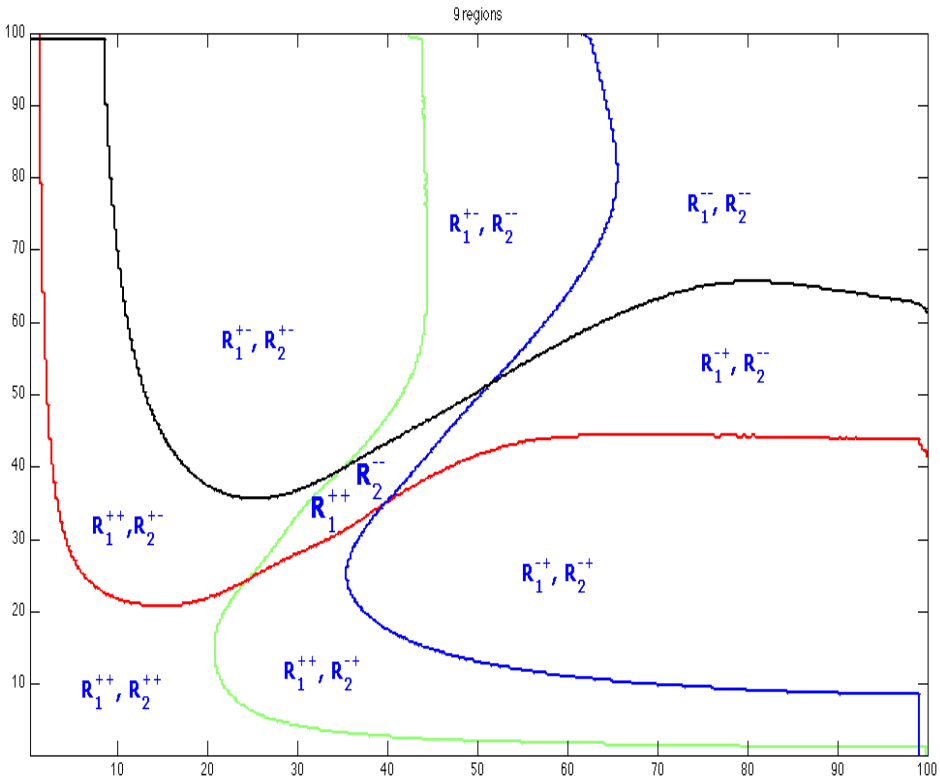

FIGURE 28 – There are nine regions in terms of trader type (HFT*vs* Institutional Investor) and trader action (LP *vs* LC)

2.3 Market Impact Modelling : some Microscopic Statistics and a Macroscopic Understanding

This Section has been published as a part of *Market Impacts and the Life Cycle of Investors Orders* (in review process at MML : Market Microstructure and Liquidity) [Bacry *et al.* , 2014]

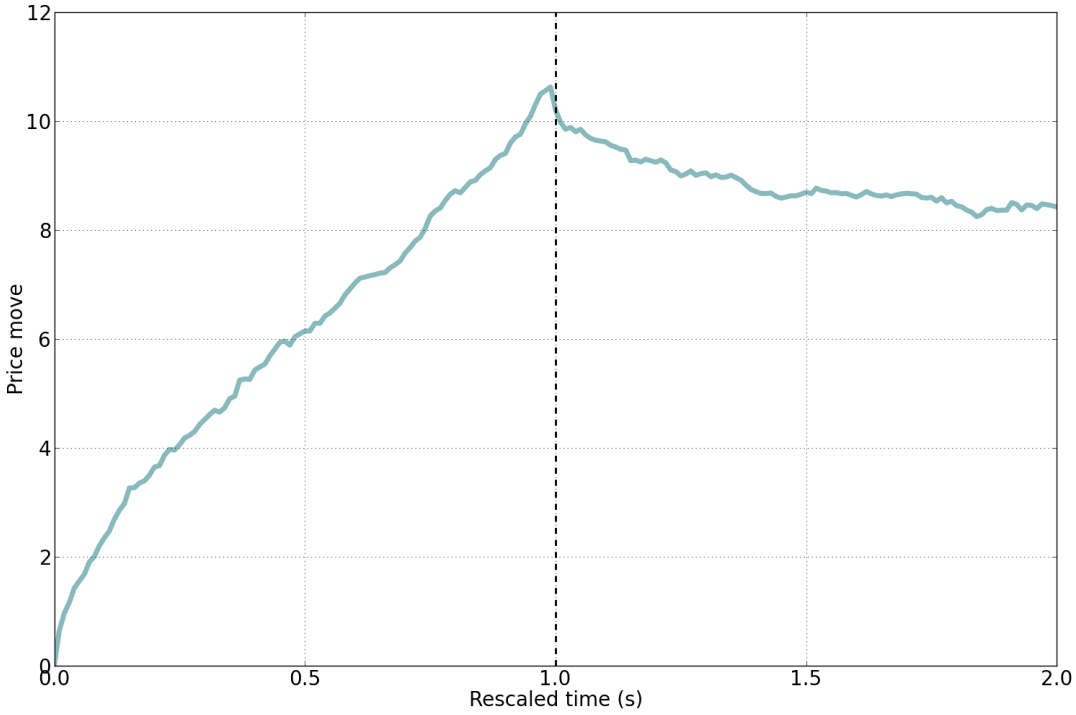

FIGURE 29 – The profile obtained averaging the prices on the 61,000 metaorders of database $\Omega^{(de)}$.

2.3.1 The main database

The database Ω is made of nearly 400,000 metaorders. The selected metaorders have been traded electronically by a single large broker during year 2010 on European markets. We built different databases from the original 400,000 database (see Table 2 for complete list of all database and associated filters) in order to adapt to the part of the market impact curve under study and to have as much orders as possible for each time scale :

- For intraday studies, we only kept orders traded by trading algorithms whose trading rate is as much as possible independent from the market conditions. Note in [Zarinelli *et al.* , 2014] , authors underline the potential influence of the trading rate on the market impact components. It allows to avoid sudden accelerated trading rates having an hidden influence on the price moves. Hence we kept VWAP, TWAP, PoV and only few Implementation Shortfall instances.
- We kept orders large enough to protect our estimation from noise.
- For the decay market impact curve, we needed the metaorder execution to stop halfway to the market close in order to observe prices relaxation long enough between the end of the metaorder and the closing time.

We ended up with five main databases : $\Omega^{(te)}$ to study temporary impact, $\Omega^{(tr)}$ to study transient impact and impact decay, and $\Omega^{(de)}$ to study decay impact and $\Omega^{(day)}$ to study daily effects. The filter *consistency with market data* means that we reject metaorders such that we cannot identify their child orders in the market data within the same second (i.e. when a child order is sent or is partially fill or fill, we try to find the correposning modification in the market data : same price, same quantity, same second ; if we cannot, we reject the metaorder).

2.3.2 Temporary Market Impact

The temporary market impact has been mainly studied from three viewpoints :

- As the main source of trading costs. The obtained model can then be used in an optimal trading scheme (see [Almgren *et al.* , 2005], [Gatheral, 2010b] and [Lehalle & Dang, 2010] and for link with optimal trading see [Almgren & Chriss, 2000], [Gatheral & Schied, 2012] and [Bouchard *et al.* , 2011]), or used by an investment firm to understand its trading costs (like in [Engle *et al.* , 2012], [Bershova & Rakhlin, 2013], [Brokmann *et al.* , 2014] or [Mastromatteo *et al.* , 2013] written by author involved in investment firms).
- It can be viewed as an important explanatory variable of price discovery and studied as such, often by economists, like the seminal work of Kyle [Kyle, 1985] or later in [Hautsch & Huang, 2012] or [Farmer *et al.* , 2004].

Databases Ω	Filters	Number of meta-orders
Original database Ω	$\emptyset$	**398.812**
Daily database $\Omega^{(day)}$	smooth execution condition	**299.824**
Database for temporary impact $\Omega^{(te)}$	10 atomic orders minimum	191.324
	consistency with HF market data	**157.061**
Database for transient impact $\Omega^{(tr)}$	number of daily trades on the stock ≥ 500	150.100
	$T > 3$ minutes	134.529
	$r \in [3\%, 40\%]$	94.818
	$R \in [0.1\%, 20\%]$	**92.100**
Database for decay impact $\Omega^{(de)}$	$t_0(\omega) + 2T(\omega) <$ closing time	**61.671**

TABLE 2 – Different databases and filters used to obtain them.

– Last but not least, statistical tools have been built to be able to estimate the temporary market impact at the scale of one trade (see [Bacry & Muzy, 2013] or [Eisler *et al.* , 2011]). The implicit conditioning of such "atomic" orders by metaorders is sometimes discussed in such papers, but it is not their main goal.

To study temporary market impact and its dependence to explanatory variables, we use the $\Omega^{(te)}$ database (see Table 2) and followed a direct regression approach. As an example, we tested the dependence in the daily participation $X = R$, since it has been identified as significant by other papers.

$$\epsilon(\omega)\Delta P_T(\omega) = a \cdot R(\omega)^\gamma + \epsilon(\omega)\Delta W_T(\omega),$$

and found an exponent $\gamma \simeq 0.449$ using the L^2 distance and a lower exponent (around 0.40) using the L^1 distance (see regression **(R.1)** in Table 3). The fact that we obtain different estimation when using the two distances L^1 and L^2 can be explained by the fact that the joint distribution of $\epsilon\Delta P$ and R is skewed to large values of ΔP. The L^2 distance has no other choice than to render this skewness by setting an high value to γ, while the L^1 distance focuses more on the center of the distribution. The source of this skewness could stem from an informational effect as a dependence between ϵ and W_T. However, we do not have enough elements to conclude on this point.

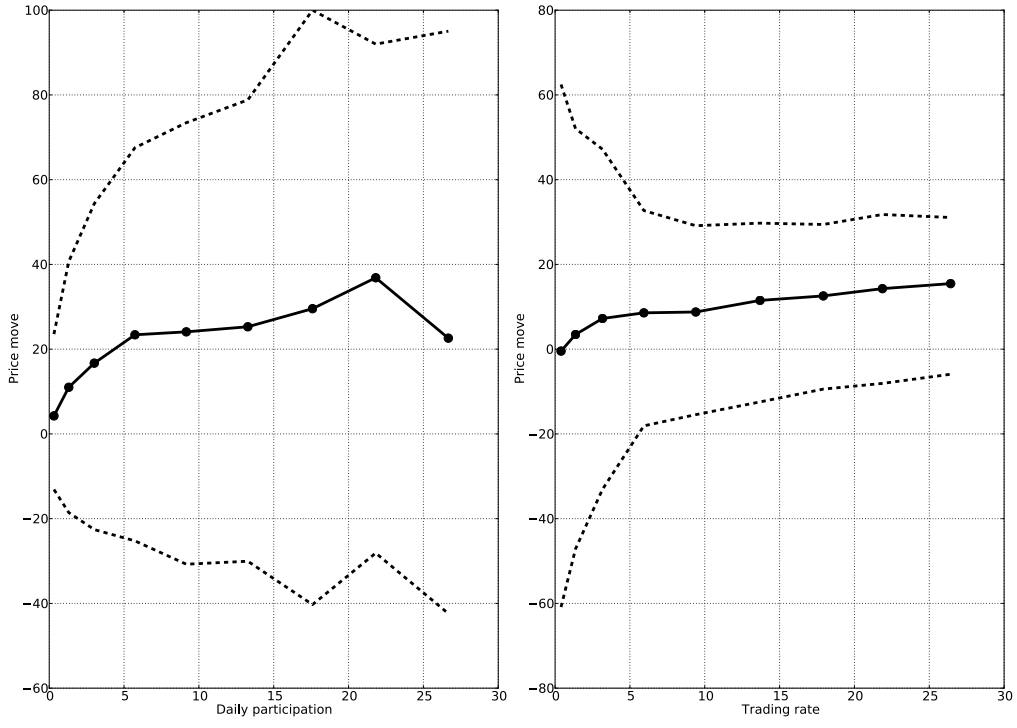

FIGURE 30 – The estimated market impact $\hat{\eta}_X$ (defined as a function of the daily participation $X = R$ (left) or of the trading rate $X = r$ (right). Each point is the average of one decile of the X variable, dotted lines are 25% and 75% quantiles, showing the amplitude of market moves.

2.3.3 Transient Market Impact

Previous section confirms, as many other empirical studies before us, that the temporary market impact of a metaorder of size v is proportional to v^γ (with γ close to 0.5). Thus, we expect the temporary market impact of the first half of the execution (the first $v/2$ contracts) to be more important than the one of the second half (the last $v/2$ contracts). Generalizing this argument to any portion of the metaorder, we expect the transient market impact curve to be a concave

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	–	–

TABLE 3 – Direct regression approach algorithm of the temporary market impact for various sets of explanatory variables. For each set, the power exponent estimations are given using L^1 distance, L^2 distance and regular log-log regressions. There is an horizontal line − when there is not significant difference between the three regressions.

function of the time. The first empirical study confirming this intuition is due to Moro et *al.* ([Moro *et al.* , 2009]), lately, it has also been confirmed by the work of Bershova & Rakhlin ([Bershova & Rakhlin, 2013]). In both cases, behavior close to power-laws were found. Let us point out that the *latent order book model* of [Mastromatteo *et al.* , 2013] can be seen as a possible qualitative explanation of this well established stylized fact. In this model the agents place limit orders only when the price is close enough to their vision of the price. Thus, more and more liquidity is revealed as the price is trending, it results in "slowing down" this trend.

In this section we use our $\Omega^{(tr)}$ database (see Table 2.3.1) and first confirm the concavity of the transient market impact curve. Apart from this well known stylized fact, we study the link between the curvature of the transient market impact curve and the execution duration.

Let us recall that the transient market impact curve corresponds to restricting the market impact curve to time $t \leq T$. In practice, we compute

$$\hat{\eta}_s = \langle \epsilon(\omega) \Delta P_{sT}(\omega) \rangle_{T=1}. \tag{56}$$

Let us recall that $\langle \ldots \rangle_T$ means that averaging is performed after rescaling in time for each metaorder ω so that all durations correspond to $T = 1$. So the function $\hat{\eta}_s$ is a function of the rescaled time s (let us recall that in the paper we use the letter s to refer to the rescaled time, whereas t stands for the physical time). We sampled this estimation in $s \in [0, 1]$ on 100 points using a uniform sampling grid. One can show on the dataset that a power law behavior

$$\hat{\eta}_s \propto s^{\gamma^{(tr)}} \tag{57}$$

with $\gamma^{(tr)} = 0.64$ fits the curve according to a log-log regression when $s \leq 1$ (the (tr) subscript in $\gamma^{(tr)} = 0.68$ stands for *transient*). Our empirical findings are compatible whit Moro *et al.* ([Moro *et al.* , 2009]) which found an exponent equal to 0.62 for metaorders executed on London Stock Exchange (LSE) and 0.71 for metaorders executed on Spanish Stock Exchange (BME, Bolsas y Mercados Espãnoles).

2.3.4 Convexity of market impact decay

During the execution of the metaorder the price is pushed in the adverse direction making it less attractive as time goes by reaching higher level (temporary impact) at the end of the execution. After the execution a reversal effect is expected. as seen in Fig. 29. This is the decay part or relaxation of the market impact. Some qualitative explanations can be sketches. For instance, consider the whole market impact cycle (transient, temporary and decay phases) in the spirit of a continuous version of Kyle's model [Kyle, 1985] : a *stylized market maker* makes a price from 0 to T and unwinds its accumulated inventory after T at a price compatible with the risk he usually take to accept the position. Another potential explanation is mechanical : if the orderbooks refill with a constant rate independent of the metaorder, the metaorder will first consume the limit orderbook on one side, pushing the price in its direction,

Subset name	Execution time	Average $\gamma^{(tr)}$	Q5%	Q25%	Q25%	Q75%	Q95%
$\mathcal{T}_1$	$T = [3, 15]$	0.80	0.76	0.78	0.80	0.82	0.85
$\mathcal{T}_2$	$T = [15, 30]$	0.66	0.62	0.65	0.66	0.68	0.70
$\mathcal{T}_3$	$T = [30, 60]$	0.62	0.58	0.60	0.62	0.64	0.66
$\mathcal{T}_4$	$T = [60, 90]$	0.55	0.49	0.52	0.56	0.58	0.62
$\mathcal{T}_5$	$T = [90, 300]$	0.54	0.48	0.52	0.55	0.57	0.62

TABLE 4 – Statistics (mean and quantiles) on the distribution of the power-law exponent $\gamma^{(tr)}$ of the transient market impact estimation of metaorders with a participation rate $R \in [1\%, 3\%]$. The exponent is estimated using log-log regression conditioned on different duration intervals. The larger T the larger the curvature of the transient market impact and the smaller the temporary market impact.

and then let an orderbook imbalanced enough at T such that symmetric and random consumption of liquidity on the two sides of the books will implement a mechanical decay. It is typically supported by a *latent orderbook model* like the one developed in [Mastromatteo *et al.* , 2013]. Last but not least, a model of *impact plus decay at an atomic size* (i.e. for each child order generated by the metaorder, like in [Farmer *et al.* , 2013]) will generate a transient phase as soon as the decay does not end when the next child order is generated (this will typically be the case for power law decay) ; after T, the "cumulated decay" will express itself, generating an observable reversal.

The existing empirical literature of decay metaorders market impact is limited ([Moro *et al.* , 2009], [Bershova & Rakhlin, 201) since the difficulty of obtaining data is very high. In the first study, Moro *et al.* are the first showing a decay of the impact to a level roughly equal to $0.5 \sim 0.7$ of its highest peak. In the second study, Bershova and Rakhlin shows the decay is a two-regime process : slow initial power decay followed by a faster relaxation.

In this section we confirm that the transient market impact curve is convex and that it seems to have a slow initial regime.

Numerical results. To have a chance to observe intraday decay, it is needed to restrict the numerical study to orders ending long enough before the close. We chose to follow [Moro *et al.* , 2009] and selected metaorders ending before the close, i.e. such that $t_0(\omega) + 2T(\omega)$ takes place before closing time. This exactly corresponds to the database $\Omega^{(de)}$ defined in Table 2.

Following the same lines as in the previous section, we use an averaging methodology to compute $\hat{\eta}_s$ for $s \leq 2$ for different intervals of duration (as previously, s corresponds here to the rescaled time, so $s = 2$, corresponds to the physical time $t = 2T$).

Fig. 31 shows such estimations (of both transient and market impact decay curves) for first four intervals.

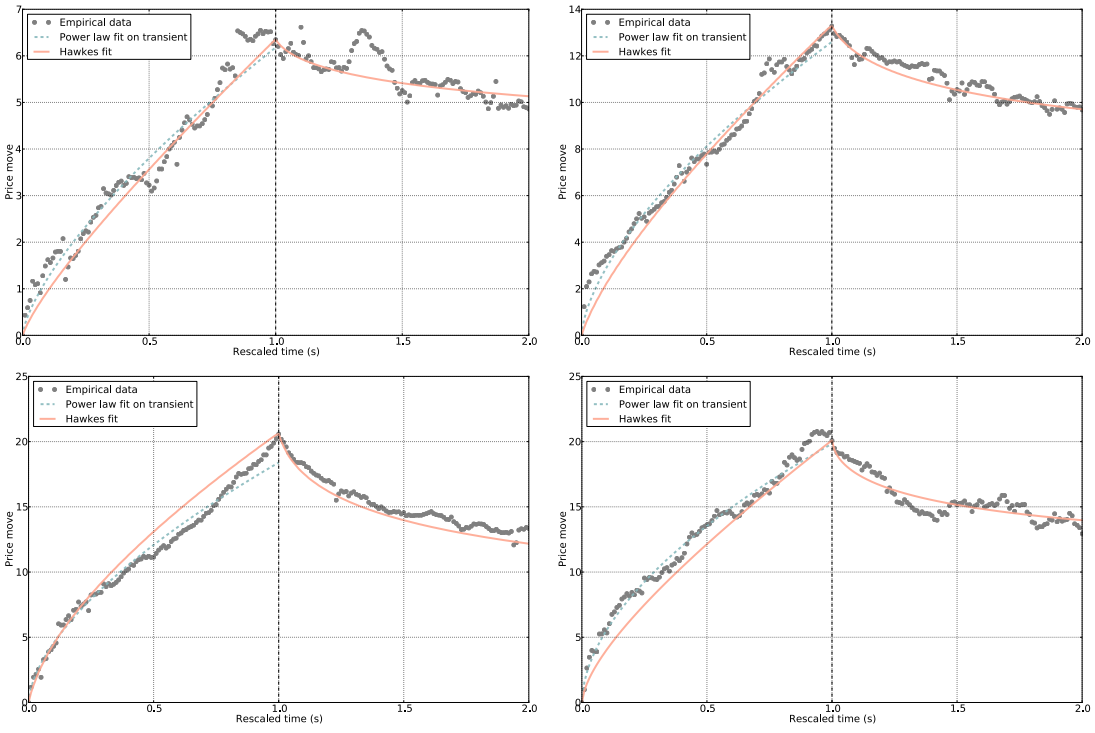

FIGURE 31 – Transient and market impact decay curve estimations $\hat{\eta}_s$ for $R \in [1\%, 3\%]$ and for different T ranges Power-law fit of the transient part are shown. Fits with the HIM model (see Section 2.3.5) are also displayed. (a) $T \in [3, 15[$, (b) $T \in [15, 30[$, (c) $T \in [30, 60[$ and (d) $T \in [60, 90[$.

In order to study the decay rate (towards the permanent impact value), we displayed log-log plots of $\hat{\eta}_s - \hat{\eta}_{s=2}$ as functions of $s-1$ for $s \in]1,2]$. They clearly show that the decay is much slower at the very beginning (i.e., right after the end of the execution of the metaorder). We have checked that changing the daily participation R does not affect qualitatively the result. This result confirms the ones obtained previously by [Bershova & Rakhlin, 2013] and [Waelbroeck & Gomes, 2013].

2.3.5 The Hawkes Impact Model (HIM) for market impact of a metaorder

Hawkes based models for microstructure. Hawkes processes have already been proved successful for modeling high frequency financial time-series (see [Bacry *et al.* , 2012, Bacry & Muzy, 2013, Hewlett, 2006]) and in optimal trading schemes (see [Alfonsi & Blanc, 2014]). Hawkes processes are point processes with a stochastic intensity which depends on the past of the process.

Following [Bacry *et al.* , 2012], we consider the following price model. Let P_t be a proxy for the high-frequency price of an asset (e.g., last-traded price, mid-price, ...). For the sake of simplicity, we shall not consider the size of the jumps in the price and consider that they are only of size 1. Let $(J_t^+,\ J_t^-)$ be the point processes representing respectively upward and downward jumps of P_t.

$$P_t = J_t^+ - J_t^-. \tag{58}$$

Let λ^+ and λ^- the respective intensities of $(J_t^+,\ J_t^-)$. It is well known that at microstructure level, the price is highly mean reverting (at least for large tick-size assets). It has been shown ([Bacry *et al.* , 2012]) that this mean-reversion property is well mimicked using a 2-dimensional Hawkes process using only a single "cross" kernel $\varphi(t)$:

$$\lambda_t^+ = \mu + \varphi \star dJ_s^- \quad \text{and} \quad \text{and} \quad \lambda_t^- = \mu + \varphi \star dJ_s^+ \tag{59}$$

where $\varphi(t)$ is a causal (i.e., supported by $\mathbb{R}^+$), positive function and where $\star$ stands for the convolution product $\varphi \star dJ_t = \int_{-\infty}^t \varphi(t-s)dJ(t)$. The mean reversion property reads clearly from these last two equations : the more P_t goes up (resp. down), the greater the intensity λ_t^- (resp. λ_t^+) will be. A criteria for the price increments and the intensities to be stationary is given by $||\varphi||_1 < 1$, where $||.||_1$ denotes the $\mathcal{L}^1(\mathbb{R})$ norm (for a complete mathematical study of Hawkes process, see [Daley & Vere-Jones, 2003]).

The Hawkes Impact Model. We model the impact of a metaorder starting at time t_0, ending at time t_0+T and corresponding to a continuous flow of buying orders with a trading rate r_t supported by $[t_0, t_0+T]$ ($r_t \neq 0$ only for $t \notin [t_0, t_0+T]$) by a perturbation of the intensities.

For the sake of simplicity, we will follow the microstructure model above and just consider mean-reversion reaction of the market (e.g., [Bacry & Muzy, 2013], [Da Fonseca & Zaatour, 2014]). Let us point out that this is clearly not a realistic hypothesis if one is interested in mimicking precisely the microstructure. However, this is not our goal. In this section, we want to build a structural model that allows to explain the main dynamics of the market impact curve. In the same line as Bouchaud [Bouchaud *et al.* , 2004], Gatheral [Gatheral, 2010b] and [Bacry & Muzy, 2013], we shall build a *linear* model, in the sense that the impact of the metaorder is nothing but the sum of the impact of its child order.

The HIM model. This model consists in replacing (59) by the two equations :

$$\lambda_t^+ = \mu + \varphi \star dJ_t^- + \int_{t_0}^t f(r_s)g^+(s-t_0)ds \quad \text{and} \quad \lambda_t^- = \mu + \varphi \star dJ_t^+ + \int_{t_0}^t f(r_s)g^-(s-t_0)ds, \tag{60}$$

where $f(r_s)ds$ (with $f(0)=0$) codes the infinitesimal impact of a buy order of volume $r_s ds$. The f function corresponds to the *instantaneous* impact function and g^+ and g^- are the *impact kernel* functions. As empirical found in [Bouchaud & Potters, 2004] and used by others authors before us ([Gatheral, 2010b, Bouchaud *et al.* , 2004]), we suppose the market impact can be separated in a factorized form : one depending on volume (or volume per time) and the other depending only on time.

The impulsive-HIM model : a particular choice for the kernels. Following [Bacry & Muzy, 2013] , it is reasonable to consider that the only "upward" impact of a single buying order is instantaneous, i.e., either the corresponding order ate up the whole first limit (in which case there is an instantaneous jump in the price) or it did not (in which case a limit order fills up the missing volume). This case corresponds to consider that g^+ is "purely" impulsive, i.e.,

$$g^+(t) = g_i^+(t) = \delta(t), \tag{61}$$

where $\delta(t)$ stands for the Dirac distribution. As for the "downward"component, we shall consider that the market reacts to the newly arrived order as if it triggered an upward jump. Doing so leads to the choice

$$g^-(t) = g_i^-(t) = C\frac{\varphi(t)}{||\varphi||_1}, \tag{62}$$

where $C > 0$ is a very intuitive parameter that quantifies the ratio of contrarian reaction (i.e. impact decay) and of the "herding" reaction (i.e. impact amplification). Indeed the L^1 norm of the herding reaction to an impulsive buying order is $f(r_t)||g_i^+||_1 = f(r_t)||\delta||_1 = f(r_t)$ (with the notation r_t for the instatanous trading rate), whereas the contrarian reaction to the same order is $f(r_t)||g_i^-||_1 = f(r_t)||\varphi||C/||\varphi|| = Cf(r_t)$. Thus, one can distinguish 3 cases of interests (see (67) of Proposition 10 for analytical expressions) :

- $C = 0$: no contrarian reaction ; we expect a permanent effect on prices from metaorders,
- $C = 1$: the contrarian reaction is as "strong" (in terms of the norm $||.||_1$) as the herding one. So we expect the two to compensate asymptotically, i.e., we expect the permanent effect of the metaorder on prices to be 0 (see Eq. (67) of Proposition 10 for confirmation),
- $C \in]0, 1[$: the contrarian reaction is not zero but strictly smaller than the herding reaction.

Thus the impulsive-HIM model corresponds to the equations

$$\lambda_t^+ = \mu + \varphi \star dJ_t^- + f(r_t) \quad \text{and} \quad \lambda_t^- = \mu + \varphi \star dJ_t^+ + C\int_{t_0}^t f(r_s)\varphi(s - t_0)ds, \tag{63}$$

where C is a positive constant that controls the contrarian vs. herding reaction of the market.

Market impact curve within HIM. According to our definition of market impact : *the difference between the observed price moves and what it would have been without this specific order*, within HIM, the market impact of a metaorder (starting at time t_0) writes :

$$\eta_t = \mathbb{E}[P_t], \quad \forall t \geq t_0. \tag{64}$$

Then, one can prove

Proposition 10 *(***Transient, decay curves and permanent effect***)*
In the framework of the HIM model (60)*, for all $t \geq t_0$ (t_0 is the starting time of the metaorder), one has :*

$$\eta_t = \int_{t_0}^{\infty} f(r_s)\Big(G(t-s) - (\kappa \star G)(t-s)\Big)ds, \tag{65}$$

where

- $G(t) = \int_0^t (g^+(u) - g^-(u))du$
- $\kappa = \sum_{n=1}^{\infty}(-1)^{n-1}\varphi^{(\star n)}$*, where* $\varphi^{(\star 1)} = \varphi$ *and* $\varphi^{(\star n)} = \varphi^{(\star n-1)} \star \varphi$.

In the case of the impulsive-HIM model (63) *, this formula gives*

$$\eta_t = \int_{t_0}^t f(r_s)H_\varphi^C(t-s)ds, \;\; t \geq t_0, \tag{66}$$

where $H_\varphi^C(t) = 1-(1+C/||\varphi||_1)\int_0^t \kappa(s)ds$*. Moreover in the case of a constant rate strategy (i.e., $r_t = r$, $\forall t \in [t_0, t_0+T]$ and $r_t = 0$ otherwise), the permanent effect of the metaorder on prices is*

$$\eta_\infty = \lim_{t\to+\infty} \eta_t = f(r)T\frac{1-C}{1+||\varphi||_1} \tag{67}$$

Let us point out that several recent empirical results ([Bacry *et al.* , 2011] and [Hardiman *et al.* , 2013]) seem to show that the Hawkes kernel φ decays as a power-law. Both studies found the exponent in the interval $[-1.5, -1]$. The following corollary shows that, in the framework of the impulsive-HIM model, and in the case of a constant rate strategy, then if φ is power-law then the market impact curve asymptotically decays (to the limit permanent effect) as a power-law, with an exponent which is related to the exponent of φ. More precisely :

Corollary 1 *In the framework of the impulsive-HIM model, let us consider a constant rate strategy, i.e., $r_t = r$, $\forall t \in [t_0, t_0 + T]$ and $r_t = 0$ otherwise. Assume that φ is such that*

- $\varphi \geq \varphi^{(\star 2)}$ *and,*
- $\exists K > 0, \quad \lim_{t\to\infty} \varphi(t)t^{-b} = K$*, with* $b \in]-2, -1[$.

Then, the market impact curve decays to the permanent market impact η_∞ asymptotically as a power-law with exponent $b + 1$, in the sense that

$$\inf\left\{\gamma, \; : \int_1^\infty (\eta_t - \eta_\infty)t^{-\gamma-1}dt < \infty\right\} = b + 1. \tag{68}$$

Qualitative understanding of the impulsive-HIM model. One can give a qualitative understanding of the impulsive-HIM model, especially of the meaning of C. Assuming an idealized market where the metaorder is only traded against one market maker and potentially noise traders (i.e. no other metaorders), this impulsive-HIM model can be seen as modelling the market maker inventory $\mathcal{I} = \lambda^+ - \lambda^-$. This market maker accepts to provide liquidity to metaorders under his own risk limits. He tunes the level of interaction he accepts with metaorders (in short : the distance of his quotes to the mid-price and the inventory thresholds he uses to unwind his risk, stopping any interaction

with the market for a while) using backtests or experience given his risk budget. As a result : with a small risk budget he will provide very attractive quotes until his inventory crosses a given threshold.

It is typically the case for high frequency market makers (HFMM) ; they are confident their technology investment allow them to capture most of the flows on both sides of the book, and be able to detect fast that they are adversely selected. In this model they "provide" to the market a small value for the parameter C of the impulsive-HIM model. On the other hand, market makers with a large risk budget will provide lazier quotes and accept a larger inventory before unwinding it. It is typically the traditional behaviour of market makers in a quote driven market. In this model such market makers will provide a large value for the parameter C.

Thus, a market maker with a small C does not provide a lot of resistance to the metaorder pressure and does not generate a large herding effect when he unwinds its inventory. A market maker with a large C will generate large contrarian pressure to metaorders and, once his risk thresholds are crossed, will unwinding a large position, generating a large herding move (partly compensated by slower market makers with larger inventory).

Moreover, since market makers calibrate their C to earn money under their risk constraint (and potentially other operational and regulatory constraints), each of them will "specialize" its activity around a given duration of metaorders $T(C)$. They will earn money providing liquidity to metaorders with a duration smaller than $T(C)$, and loose money interacting with longer metaorders.

Consequently, in the impulsive-HIM model, we have used a single market maker only with a given value of C but in practice we should consider an extension of this model with a continuum of market makers, implementing a distribution of C reflecting the distribution of metaorders in a given market.

This interpretation of the impulsive-HIM model can explain the concavity of transient impact stems from the market makers "triggered" by a given metaorder. Since $T(C)$ is increasing in C : at the start of the trading, the metaorder interacts with market makers with small C (i.e. typically HFMM). After a while, such participants consider they are adversely selected, stop trading and unwind their inventory in front of market makers having a largest C. The longer the metaorder, the more market makers with large C, the more contrarian pressure and concavity to the transient impact. Once the metaorder stops, it is currently interacting with market makers with a perfectly adapted C (from their viewpoint) ; the later now have to unwind slowly their inventory to realize their gain.

3 L'Apport des Algorithmes Stochastiques et de l'Apprentissage Statistique

Sans même trop réfléchir, les algorithmes stochastiques et l'apprentissage statistique semblent pouvoir apporter des réponses dans le cadre de l'étude et du contrôle du processus de formation des prix :

- on sait l'apprentissage statistique adapté à la construction à la volée d'estimateurs sur de larges bases de données ;
- les algorithmes stochastiques ont quant à eux apporté des réponses aux problèmes d'exploration - exploitation comme celui des bandits manchots (cf. [Lamberton & Pagès, 2008]).

Trading Optimal "Forward". Bien entendu le contrôle stochastique répond précisément à question lorsqu'il s'agit de planifier le trading sur plusieurs heures (pourvu que le modèle de dynamique renferme suffisamment d'information à ces horizons de temps), mais l'étude de la dynamique des carnets d'ordres montre qu'une échelle de temps entre "le prochain mouvement de carnet" et les "quelques heures" peut être utile.

La Section 1.2 tente d'offrir un cadre pour conjuguer ces deux échelles :

- à moyen terme, le contrôle stochastique offre une planification des interactions avec les autres participants ;
- à court terme, il se repose sur des "*robots de trading*" qui vont faire leur possible pour obtenir au meilleur prix la quantité souhaité dans le delais spécifié.

Nous allons nous reposer sur les algorithmes stochastiques pour construire des tels robots :

- dans la Section 3.1 il s'agira de répartir la quantité à acheter ou vendre sur plusieurs plateformes de négociation. Ce genre de répartition répond à la *fragmentation des marchés* apparue depuis une petite dizaine d'année (cf. [Lehalle *et al.* , 2013] pour des détails).
- Quant à la Section 3.2, elle montrera comment adapter le prix auquel on poste un ordre d'achat ou de vente, non pas à une probabilité d'exécution espérée d'après une loi en $A \exp -k\delta$ vérifiée en moyenne (comme cela est fait dans la Section 1.3), mais en fonction des transactions obtenues tout récemment.

Il s'agit donc dans les deux cas de mettre en œuvre un processus d'exploration-exploitation, visant à optimiser non pas l'espérance future suivant la loi jointe bien spécifiée des contrôles et de l'état, mais à l'espérance *sur le long terme*, en se reposant sur une hypothèse de stationnarité de la loi sous jascente jamais spécifiée.

Ces deux approches très complémentaires sont difficiles à articuler, ici nous choisirons l'échelle de temps : l'optimisation rétrograde (i.e. *backward* en anglais) du contrôle stochastique restera dédiée aux grandes échelles, alors que dans les section qui suivent nous spécialiserons l'optimisation progressive (i.e. *forward*) des algorithmes stochastiques pour les petites échelles (sur quelques dixaines de minutes tout au plus).

Pour rester, comme lors de l'introduction de la Section 1, dans un formalisme approximatif, on peut résumer ainsi le passage du trading optimal stochastique au trading optimal par apprentissage :

$$\underbrace{\left|\begin{array}{rcl} \mathsf{EV}(t,x) & = & \max_{\rho(t,x)} \mathbb{E}\big(V(t, X_t = x, \rho(t,x)) \\ & & \qquad + \mathsf{EV}(t+1, X_{t+1}|x, \rho(t,x))\big) \\ \mathsf{EV}(T,x) & = & G(x) \end{array}\right.}_{\text{Stochastic Optimal Trading}} \longrightarrow \underbrace{\left|\begin{array}{rcl} \rho(t+1) & = & \rho(t) - \gamma_t \partial_\rho h(X_t, \rho(t)) \\ & & \\ H(\rho) & = & \mathbb{E}_X h(X, \rho) \end{array}\right.}_{\text{Optimal Trading by Learning}} .$$

Ici on cherche à maximiser un critère en espérance $H(\rho)$ qui s'exprime en fonction d'une valeur instantanée d'une action $h(X, \rho)$. Cette dernière dépend à la fois de l'état du système et du contrôle ρ. Plutôt que de maximiser explicitement le critère, on va construire une suite de contrôle, qui prend en compte les actions passées et leurs valeurs, de telle sorte de la limite de cette suite maximisera le critère : $\rho(\infty) = \arg\max_\rho H(\rho)$. Classiquement, cela recours à une *descente de gradient stochastique* d'un pas γ_t qui vérifie les hypothèses habituelles de Robbins-Monro[Robbins & Monro, 1951] (tend vers zéro mais de somme divergente). Il s'agit donc ici de maximiser *le long des trajectoires*, ce qui, via le théorème ergodique, permettra d'obtenir le manimum de l'espérance selon la bonne mesure sur X.

Monitoring d'un grand nombre d'ordres automatisés. L'apprentissage statistique n'est pas seulement utile pour la construction d'algorithmes de trading, il permet aussi de comprendre en temps réel la cause d'une baisse de performance d'un sous ensemble d'algorithmes au sein d'un très grand nombre. C'est ce à quoi s'attache la Section 3.3 : elle montre comment on peut construire un système de focalisation d'attention pour la surveillance d'un très grand nombre d'algorithmes de trading en cours de fonctionnement.

En effet, en pratique (cf. [Lehalle *et al.* , 2013]), les intermédiaires (courtiers et banques d'investissement) qui fournissent aux investisseurs finaux des algorithmes de trading paramétrables constituent des équipes de 2 à 12 traders chargés de surveiller chacun entre 200 et 1,000 algorithmes. Leur rôle est d'ajuster manuellement un ou deux paramètres de règlage (comme l'*aggressivité* –qui correspond à un paramètre d'avertion au risque–), voire d'arrêter le trading.

Ces *execution traders* ont un temps et une attention limités, et ils se doivent de les consacrer aux algorithmes qui ont "*le plus de difficulté*". La Section 3.3 construit à la volée des prédicteurs de bonne ou mauvaise qualité du trading. Ces prédicteurs ayant la capacité de sélectionner à chaque instant t les variables explicatives qu'ils utilisent comme régresseurs, on va considérer que les meilleurs prédicteurs pointent en fait du doigt *les meilleures variables explicatives de mauvaise performance*. Les traders peuvent ainsi accéder à une première forme d'interprétation et un regrouppement des baisses de performance.

English introduction to the section

The two preceding sections deal respectively with optimal trading (in the sense of the minimization of an expected cost) and market microstructure (modelling and understanding the price formation process in general). They both model averaged or usual behaviours. Nevertheless the reality is not that predictable, hence the *noise* (or the unexpected fluctuations around dynamics that are, on average, understood) around the expected behaviour by previous section, is high.

An answer to this would be to use recent data to build on the fly models capturing the current state of liquidity, or the current relationship between the efficiency of trading algorithms and some observables. By chance electronic markets are transparents, since regulators or any market participant have access to an order-by-ordrer update of the state of matching engines. Regulators have even more information since they have the identity of each owner of orders. The typical number of trades on a liquid stock per day being around 40,000 and the number of orderbook update being around ten times this number, we have access to a lot of data. Enough to put in place statistical learning approaches in attempts to capture oscillations around average behaviours.

This section presents two ways to use statistical learning on high frequency data :
- improve the efficiency of optimal trading scheme of Section 1,
- monitor a large number of trading algorithms and provide online decision support to human operators.

The work presented here is not just an application of known techniques, but it goes back to the root of stochastic algorithms (see [Benveniste *et al.* , 1991] and [Duflo, 1997]) and of statistical learning and non parametric modelling (see [Vapnik, 2006]) to provide an accurate description (in terms of bound or convergence rate) of the obtained tools.

The use of an exploration-exploitation learning approach inside robots synchronized thanks to stochastic control has been underlined in the introduction of Section 1. The *learning by trading* approaches presented here are the first attempts to give a mathematical framework to *tricks* and *tips* used by traders to improve the efficiency of their trading algorithms.

The alternatives to this stochastic algorithm approach would be :
- to put in place estimators as robust as possible of needed variables and plug them as if they were deterministic into a deterministically optimal strategy,
- to use a minimum regret approach, being more afraid by doing mistakes than being confident in the efficiency of a strategy designed to be optimal "on expectation".

By chance these two other approaches have been proposed in the scope of "Dark Pool trading" as alternatives of the work exposed in Section 3.1 and published in [Pagès *et al.* , 2011]. [Ganchev *et al.* , 2010] for a sensored statistics proposal and [Agarwal *et al.* , 2010] for a minimum regret one. As usual in learning, there is no best algorithm with respect to any criterion, just that :
- our stochastic algorithm approach is suitable for a participant having enough trading flow to attain the "expected optimal" (typically a broker) ;
- the sensus statistics is good for an opportunistic hedge fund primarily interested to know what is available in the dark rather than trading massively in it,
- the minimum regret is nice for an investor who trades not that often but would like to systematically send a fraction of its orders in Dark Pools.

The monitoring framework of Section 3.3 is the only paper providing a mathematical grasp to offer decision support to human operating or supervising trading algorithms. It can be adapted to the use of investment banks, market operators or regulators.

3.1 Trading in the Dark : an Exploration-Exploitation Solution

This Section has been published as a part of *Optimal split of orders across liquidity pools : a stochastic algorithm approach* (2011 SIAM Journal on Financial Mathematics 2, 1042–1076) [Pagès *et al.* , 2011]

3.1.1 A simple model for the execution of orders by dark pools

The execution model. As mentioned in the introduction, we will focus in this paper on the splitting order problem in the case of (competing) *dark pools*. The execution policy of a dark pool differs from a primary market : thus a dark pool proposes bid/ask prices with no guarantee of executed quantity at the occasion of an over the counter transaction. Its bid and ask prices are slightly different from those offered on the primary market (practically, it is often a small modification of the mid-price). Let us temporarily focus on a buy order to be split across several dark pools. We do not aim at specifying a given dynamics for the bid and ask prices, but just to make some reasonable generic assumptions. We mean assumptions not too far from reality which makes possible to connect the problem with a stochastic optimization that can be solved *on line* (hence in a forward way).

One can model the impact of the existence of N dark pools ($N \geq 2$) on a given transaction started at time t as follows : let $V(t) > 0$ be the random volume to be executed between time t and $t + \Delta t$. If S_t denotes the (bid) price of the asset at time t on the primary market and S_t^i the one of dark pool i, we assume that

$(H_1) \equiv \forall\, t \geq 0, \quad \mathbb{E}(S_t^i \,|\, S_t) = \theta_i S_t,\ i = 1, \ldots, N.$

where $\theta_i \in (0, 1)$ can be seen as the mean *discount factor* proposed by the dark pool $i \in \{1, \ldots, N\}$ with respect to the best opposite. Let $r_i(t)$ denote the percentage of $V(t)$ sent to the dark pool i for execution and let $D_i(t) \geq 0$ be the quantity of securities that can be delivered (or made available) by the dark pool i at price S_t^i. In case of of a remaining quantity on the order, what is left is sent at time $t + \Delta t$ for agressive execution on the primary market, at price $S_{t+\Delta t}$. Then the total cost $C(t)$ of the executed order is given by

$$C(t) \;=\; \sum_{i=1}^{N} S_t^i \min(r_i(t)V, D_i(t)) + S_{t+\Delta t}\left(V(t) - \sum_{i=1}^{N} \min(r_i(t)V(t), D_i(t))\right).$$

The second assumption is that, at the scale we consider (say around 10 minutes with an execution every few seconds) :

$(H_2) \equiv$ The price process $(S_t)_t$ and $(V(t), D_1(t), \ldots, D_{_N}(t))_t$ are independent.

$(H_3) \equiv$ The process $(V(t), D_1(t), \ldots, D_{_N}(t))_t$ is stationary with marginal distribution ν.

(From a mathematical viewpoint, we only need this process to be mean-reverting and to converge toward a steady regime.) We will also need in the next section that this shares a slightly more stringent ν-ergodic/averaging property under this steady regime.

Finally we assume that $(S_t)_t$ is either a martingale or more generally

$(H_4) \equiv \forall\, t \geq 0, \quad \mathbb{E}\big(S_t \,|\, \mathcal{F}_s^S\big) = a_{t-s} S_s,\ a_u \in (0, \infty)$. where a_u satisfies $\lim_{u \to 0} a_u = 1$ (Typically, one may think at a_u as

a discount factor like $e^{-\mu u}$, $\mu > 0$).

Then elementary computations show that

$$\mathbb{E}(C_t) = \mathbb{E}(S_{t+\Delta t})\mathbb{E}(V(t)) - a_{\Delta t}\mathbb{E}(S_t) \sum_{i=1}^{N} \Big(1 - \frac{\theta_i}{a_{\Delta t}}\Big)\mathbb{E}\Big(\min(r_i(t)V(t), D_i(t))\Big).$$

At this stage, one may wish to minimize the mean execution cost $\mathbb{E}(C_t)$. Taking into account the stationarity assumption (H_3), this amounts to solving the following (conditional) maximization problem

$$\max\left\{\sum_{i=1}^{N} \rho_i \mathbb{E}\left[\min(r_i V, D_i)\right],\ r \in \mathcal{P}_{_N}\right\}. \tag{69}$$

where the expectation is taken with respect to the marginal distribution ν and $\rho_i = 1 - \frac{\theta_i}{a_{\Delta t}}$, $i = 1, \ldots, N$. If one of the ρ_i is negative, dark pool i is out of the game and the problem is ill-posed. So we make the assumption that Δt is chosen small enough so that

$$\max_{1 \leq i \leq N} \theta_i < a_{\Delta t}$$

which implies in turn that $\rho_i > 0$, $i = 1, \ldots, N$.

If one considers symmetrically a sell order to be executed, the dark pool is supposed to propose a higher price $\theta_i S$, $\theta_i > 1$, than the best opposite of the visible order book. The seller aims at maximizing the execution global (mean)

price of the transaction. This leads to the same formal optimization problem, this time with $\rho_i = \theta_i - a_{\Delta t}$, $i = 1, \ldots, N$, all assumed to be positive which requires the assumption $\min_i \theta_i > a_{\Delta_t}$.

In both cases, the restriction on Δt means that, if *e.g.* $(a_u)_{u\geq 0}$ is also nonincreasing, $[0,1]$-valued with $\lim_{u\to\infty} a_u \leq 0$, one must complete the execution of the order fast enough to limit the exposure to the market risk in order to take advantage of dark pool executions.

All these considerations lead us to focus on the abstract optimal allocation problem (69) in which the price variable S no longer appears in what follow which means in some sense that we can forget about the primary market. In fact this is a *sine qua non* condition to define an *on line* approach to the optimal splitting based on a learning procedure.

3.1.2 The Learning Scheme

In practice, there is no *a priori* assumption – or information available – on the joint distribution of $(V, D_1, \ldots, D_{_N})$ under $\mathbb{P}$. So the only reasonable way to provide a procedure to solve this allocation problem is to devise an *on-line learning algorithm based on historical data*, namely the results of former transactions with the dark pools on this security executed in the past. This underlines that our agent dealing with the dark pools is a financial institution like an investment bank, a broker or possibly a large investor which often – that means at least daily – faces some large scale execution problems on the same securities.

This means that we will have to make some assumptions on the dynamics of these transactions *i.e.* on the data input sequence $(V^n, D_1^n, \ldots, D_{_N}^n)_{n\geq 1}$ supposed to be defined on the same probability space $(\Omega, \mathcal{A}, \mathbb{P})$ (where superscript n stands for an execution launched at time t_n (to be consistent we assume that the execution are made successively that is $t_n > t_{n-1} + \Delta t$).

Our basic assumption on the sequence $(D_i^n, V^n, i = 1, \ldots, N)_{n\geq 1}$ is of statistical – or ergodic – nature : we ask this sequence to be ν-averaging (*a.s.* and in $L^p(\mathbb{P})$), at least on bounded continuous functions, where ν is a distribution on $(\mathbb{R}_+^{N+1}, \mathcal{B}or(\mathbb{R}_+^{N+1}))$. This leads to the following formal assumption :

$$(ERG)_\nu \equiv \begin{cases} (i) & \text{the sequence } (V^n, D_i^n, i = 1, \ldots, N)_{n\geq 1} \text{ is averaging } i.e. \\ & \mathbb{P}\text{-}a.s. \quad \dfrac{1}{n}\displaystyle\sum_{k=1}^{n} \delta_{(V^k, D_1^k, \ldots, D_N^k)} \overset{(\mathbb{R}_+^{N+1})}{\Longrightarrow} \nu, \\ (ii) & \sup_n \mathbb{E}(V^n)^2 < +\infty. \end{cases}$$

where $\overset{(\mathbb{R}_+^{N+1})}{\Longrightarrow}$ denotes the weak convergence of probability measures on $\mathbb{R}_+^{N+1}$. For convenience, we will denote $(V, D_1, \ldots, D_{_N})$ the canonical random vector on $\mathbb{R}_+^{N+1}$ so that we can write $\nu = \mathcal{L}(V, D_1, \ldots, D_{_N})$.

Assumption (ii) on the marginal distribution of the sequence $(V^n)_{n\geq 1}$ is mainly technical. In fact standard arguments from weak convergence theory show that combining (i) and (ii) implies

$$\frac{1}{n}\sum_{k=1}^{n} V^k \longrightarrow \mathbb{E}\, V \quad \text{as} \quad n \to \infty$$

($\sup_n \mathbb{E}(V^n)^{1+\varepsilon} < +\infty$ would be enough). An important sub case is the the (IID) setting

$$(IID) \equiv \begin{cases} (i) & \text{the sequence } (V^n, D_1^n, \ldots, D_{_N}^n)_{n\geq 1} \text{ is i.i.d. with distribution } \nu = \mathcal{L}(V, D_1, \ldots, D_{_N}), \\ (ii) & V \in L^2(\mathbb{P}). \end{cases}$$

This more restrictive assumption is undoubtedly less realistic from a modeling point of view but it remains acceptable as a first approximation. It is the most common framework to apply the standard Stochastic Approximation machinery (*a.s.* convergence, asymptotically normal fluctuations, etc). So, its interest may be considered at least as pedagogical. The (ERG) setting is slightly more demanding in terms of assumptions and needs more specific methods of proof. It will be investigated as a second step, using some recent results established in [Laruelle & Pagès, 2012] which are well suited to the specificities of our problem (in particular we will not need to assume the existence of a solution to the Poisson equation related to the procedure like in the reference book [Benveniste *et al.* , 1991])

The mean execution function of a dark pool In view of the modeling section, we need to briefly describe the precise behaviour of the mean execution function $\varphi : [0,1] \to \mathbb{R}_+$ of a single dark pool.

Let (V, D) be an $\mathbb{R}_+^2$-valued random vector defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ representing the *global volume* to be executed and the *deliverable quantity* (by the dark pool) respectively. Throughout this paper we will assume the following consistency assumption

$$V > 0 \quad \mathbb{P}\text{-}a.s. \quad \text{and} \quad \mathbb{P}(D > 0) > 0. \tag{70}$$

The *a.s.* positivity of V means that we only consider true orders. The fact that D is not identically 0 means that the dark pool does exist in practice. The "rebate" coefficient ρ is specific to the dark pool.

To define in a consistent way the mean execution function of a dark pool we only need to assume that $V \in L^1(\mathbb{P})$ (although more stringent integrability assumptions are made throughout the paper).

Here the *mean execution function* $\varphi : [0,1] \to \mathbb{R}_+$ of the dark pool is defined by

$$\forall\, r \in [0,1], \qquad \varphi(r) = \rho\, \mathbb{E}(\min(rV, D)) \tag{71}$$

where $\rho > 0$. The function φ is finite, non-identically 0. It is clearly a concave non-decreasing bounded function. Furthermore, one easily checks that its right and left derivatives are given at every $r \in [0,1]$ by

$$\varphi'_l(r) = \rho\, \mathbb{E}\left(\mathbf{1}_{\{rV \le D\}} V\right) \qquad \text{and} \qquad \varphi'_r(r) = \rho\, \mathbb{E}\left(\mathbf{1}_{\{rV < D\}} V\right). \tag{72}$$

In particular,

$$\varphi'(0) = \rho\, \mathbb{E}(V \mathbf{1}_{\{D>0\}}) > 0$$

and if

$$\textit{the (right continuous) distribution function of } \tfrac{D}{V} \textit{ is continuous on } \mathbb{R}_+, \tag{73}$$

then

$$\varphi \textit{ is everywhere differentiable on the unit interval } [0,1] \textit{ with } \varphi' = \varphi'_l \textit{ on } (0,1].$$

Assumption (73) means that the distribution of $\frac{D}{V}$ has no atom except possibly at 0. It can be interpreted as the fact that a dark pool has no "quantized" answer to an order.

More general models of execution functions in which the rebate ρ and the deliverable quantity D may depend upon the quantity to be executed rV are briefly discussed further on.

Design of the stochastic Lagrangian algorithm. Let V be the quantity to be executed by N dark pools. For every dark pool $i \in \mathcal{I}_N$ the available quantity D_i is defined on the same probability space $(\Omega, \mathcal{A}, \mathbb{P})$ as V. We assume that all couples (V, D_i) satisfy the consistency assumption (70).

To each dark pool $i \in \mathcal{I}_N$ is attached a (bounded concave) mean execution function φ_i of type (71), introduced in Section 3.1.2.

Then for every $r = (r_1, \ldots, r_N) \in \mathcal{P}_N$,

$$\Phi(r_1, \ldots, r_N) := \sum_{i=1}^{N} \varphi_i(r_i). \tag{74}$$

In order to design the algorithm we will need to extend the mean execution function φ (whatever its form is) as a concave function on the whole real line by setting

$$\varphi(r) = \left(r - \frac{r^2}{2}\right)\varphi'(0) \quad \text{if } r < 0 \quad \text{and} \quad \varphi(r) = \varphi(1) + \varphi'(1)\log r \quad \text{if } r > 1. \tag{75}$$

Based on the extension of the functions φ_i defined by (75), we can formally extend Φ on the whole affine hyperplane spanned by $\mathcal{P}_N$ *i.e.*

$$\mathcal{H}_N := \{r \in \mathbb{R}^N \mid \sum_i r_i = 1\}.$$

As announced, we aim at solving the following maximization problem

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

but we will deal for algorithmic purpose with the same maximization problem when r runs over $\mathcal{H}_N$.

Before stating a rigorous result, let us a have a look at a Lagrangian approach that only takes into account the affine constraint that is $\max_r \Phi(r) - \lambda \sum_i r_i$. Straightforward formal computations suggest that

$$r^* \in \operatorname{argmax}_{\mathcal{P}_N} \Phi \text{ iff } \varphi'_i(r_i^*) \text{ is constant when } i \text{ runs over } \mathcal{I}_N$$

or equivalently if

$$\forall\, i \in \mathcal{I}_N, \qquad \varphi'_i(r_i^*) = \frac{1}{N}\sum_{j=1}^{N} \varphi'_j(r_j^*). \tag{76}$$

In fact this statement is not correct in full generality because the Lagrangian method does not provide a necessary and sufficient condition for a point to be a maximum of a (concave) function ; thus, it does not take into account the case where Φ reaches its maximum on the boundary $\partial\mathcal{P}_N$ where the above condition on the derivatives may fail. So, an additional assumption is necessary to make it true as established in the proposition below.

Proposition 11 *Assume that (V, D_i) satisfies the consistency assumptions (70) and (73) for every $i \in \mathcal{I}_N$.*

(a) Assume that the functions φ_i defined by (71) satisfy the following assumption

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

Then $\mathrm{argmax}_{\mathcal{P}_N}\Phi$ *is a compact convex set and*

$$\mathrm{argmax}_{\mathcal{P}_N}\Phi = \{r \in \mathcal{P}_N,\, |\, \varphi_i'(r_i) = \varphi_1'(r_1),\ i = 1, \ldots, N\}.$$

Furthermore $\mathrm{argmax}_{\mathcal{H}_N}\Phi = \mathrm{argmax}_{\mathcal{P}_N}\Phi$.

(b) If the functions φ_i satisfy the slightly more stringent assumption,

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

then

$$\mathrm{argmax}_{\mathcal{H}_N}\Phi = \mathrm{argmax}_{\mathcal{P}_N}\Phi \subset int(\mathcal{P}_N).$$

Remarks. • If $N = 2$, one checks that Assumption $(\mathcal{C})$ is also necessary to derive the conclusion of item (a).
• As a by-product of the proof below we have the following more precise result on the optimal allocation r^* : if $r^* \in \mathrm{argmax}_{\mathcal{P}_N}$ and $\mathcal{I}_0(r^*) := \{i \in \mathcal{I}_N \,|\, r_i^* = 0\}$, then

$$\max_{i\in\mathcal{I}_0(r^*)} \varphi_i'(0) \leq \min_{i\in\mathcal{I}_0(r^*)^c} \varphi_i'(0).$$

Interpretation and comments : • In the case of a "regular" mean execution function, Assumption $(\mathcal{C})$ is a kind of *homogeneity assumption on the rebates* made by the involved dark pools. If we assume that $\mathbb{P}(D_i = 0) = 0$ for every $i \in \mathcal{I}_N$ (all dark pools buy or sell at least one security with the announced rebate), then $(\mathcal{C})$ reads

$$\min_{i\in\mathcal{I}_N} \rho_i \geq \max_{i\in\mathcal{I}_N} \left(\rho_i \frac{\mathbb{E}\, V \mathbf{1}_{\{\frac{V}{N-1}\leq D_i\}}}{\mathbb{E}\, V}\right)$$

since $\varphi_i'(0) = \rho_i\, \mathbb{E}\, V$. In particular,

Assumption $(\mathcal{C})$ is always satisfied when all the ρ_i's are equal

i.e. all dark pools propose the same rebates which usually is usually the mid-price (defined as the middle between bid and ask prices).

• Assumption $(\mathcal{C})$ is in fact our main assumption in terms of modeling. It may look somewhat difficult to satisfy when the rebates are not equal. But the crucial fact in order to preserve the generality of what follows is that it contains *no assumption about the dependence between the volume V and the "answers" D_i from the dark pools.*

Design of the stochastic algorithm. Now we are in position to devise the stochastic algorithm for the optimal allocation among the dark pools, taking advantage of the characterization of $\mathrm{argmax}_{\mathcal{P}_N}\Phi$. In fact we will simply use the obvious remark that N numbers a_1,..., a_N are equal if and only if they are all equal to their arithmetic mean $\frac{a_1+\cdots+a_N}{N}$.

We consider the mean execution function as defined by (71). We assume from now on that the continuity assumption (73) holds so that the representation (72) of its derivative can be taken as its right or its left derivative on $(0, 1]$ (and its right derivative only at 0).

Using this representation (72) for all the derivatives φ_i' yields that, if Assumption $(\mathcal{C})$ is satisfied, then $\mathrm{argmax}_{\mathcal{H}_N}\Phi = \mathrm{argmax}_{\mathcal{P}_N}\Phi$ and

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

However, the set $\mathcal{P}_N$ is not stable for the "naive" zero search algorithm naturally derived from the above characterization, we are led to devise the procedure on the hyperplane $\mathcal{H}_N$.

Consequently, this leads to devise the following zero search procedure

$$r^n = r^{n-1} + \gamma_n H(r^{n-1}, V^n, D_1^n, \ldots, D_N^n),\ n \geq 1, \quad r^0 \in \mathcal{P}_N, \tag{77}$$

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

$$H_i(r, V, D_1, \ldots, D_N) = V\Big(\rho_i \mathbf{1}_{\{r_i V \leq D_i\} \cap \{r_i \in [0,1]\}} - \frac{1}{N}\sum_{j=1}^{N} \rho_j \mathbf{1}_{\{r_j V \leq D_j\} \cap \{r_j \in [0,1]\}} + R_i(r, V, D_1, \ldots, D_N)\Big) \tag{78}$$

and the "innovation" $(V^n, D_1^n, \ldots, D_N^n)_{n \geq 1}$ is a sequence of random vectors with non negative components such that, for every $n \geq 1$, $(V^n, D_i^n, i = 1, \leq, N) \stackrel{d}{=} (V, D_i, i = 1, \leq, N)$ and the remainder terms R_i have a mean-reverting effect to pull back the algorithm into $\mathcal{P}_N$. They are designed from the extension (75) of the derivative functions φ_i' outside the unit interval $[0, 1]$; to be precise, for every $i \in \mathcal{I}_N$,

$$R_i(r, V, D_1, \ldots, D_N) = \rho_i \left((1 - r_i)\mathbf{1}_{\{D_i > 0\} \cap \{r_i < 0\}} + \frac{1}{r_i}\mathbf{1}_{\{V \leq D_i\} \cap \{r_i > 1\}}\right) - \frac{1}{N}\sum_{j=1}^{N} \rho_j \left((1 - r_j)\mathbf{1}_{\{D_j > 0\} \cap \{r_j < 0\}} + \frac{1}{r_j}\mathbf{1}_{\{V \leq D_j\} \cap \{r_j > 1\}}\right).$$

3.1.3 Interpretation and implementability of the procedure

▷ Implementability. The vector $(r_i^n)_{1 \leq i \leq N}$ in (77) represents the dispatching of the orders among the N dark pools to be sent at time $n + 1$ by the investor. It is computed at time n. On the other hand V^n represents the volume to be executed (or its monetary value if one keeps in mind that we "plugged" the price into the volume) and the D_i^n the "answer" of dark pool i, still at time n.

The point is that the investor does have no access to the quantities D_i^n. However, he/she knows what he/she receives from dark pool i, *i.e.* $\min(D_i^n, r_i^{n-1} V^n)$. As a consequence, the investor has access to the event

$$\{\min(D_i^n, r_i^{n-1} V^n) = r_i^{n-1} V^n\} = \{r_i^{n-1} V^n \leq D_i^n\}$$

which in turn makes possible the updating of the procedure although he/she has no access to the true value of D_i^n.

So, except for edge effects outside the simplex $\mathcal{P}_N$, the procedure as set can be implemented on real data.

▷ Interpretation. As long as r is a true allocation vector, *i.e.* lies in the simplex $\mathcal{P}_N$, the interpretation of the procedure is the following : assume first that all the factors ρ_i are equal (to 1). Then the dark pools which fully executed the sent orders ($r_i V \leq D_i$) are rewarded proportionally to the numbers of dark pools which *did not fully executed* the request they received. Symmetrically, the dark pools which could not execute the whole request are penalized proportionally to the number of dark pools which satisfied the request.

Thus, if only one dark pool, say dark pool 1, fully executes the request at time n, its pourcentage will be increased for the request at time $n + 1$ by $\gamma_n(1 - \frac{1}{N})V^n$ *i.e.* it will asked to execute $r_1^n = r_1^{n-1} + \gamma_n(1 - \frac{1}{N})V^n$ % of the total order V^{n+1}. The other $N - 1$ dark pools will be penalized symmetrically : the pourcentage r_i^n of the total request V^{n+1} each of them will receive at time $n + 1$ will be reduced by $\gamma_n \frac{1}{N} V^n$.

If k dark pools totally execute their request at time n and the $N - k$ other fail, the pourcentages of V^{n+1} that the "successful" dark pools will receive for execution at time $n + 1$ will be increased by $\gamma_n(1 - \frac{k}{N})V^n$, each of the $N - k$ "failing dark pools" being reduced by $\gamma_n \frac{k}{N} V^n$.

If no dark pool was able to satisfy their received request at time n, none will be penalized and if all dark pools fully execute the received orders, none will be rewarded.

In short, the dark pools are rewarded or penalized by comparing their mutual performances. When the "attractivity" coefficients ρ_i are not equal, the reasoning is the same but weighted by these attractivities.

▷ Practical implementation. One may force the above procedure to stay in the simplex $\mathcal{P}_N$ by projecting, once updated, the procedure on $\mathcal{P}_N$ each time it exits the simplex. This amounts to replace the possibly negative r_i by 0, the $r_i > 1$ by 1 and to renormalize the vector r by dividing it by the sum of its terms.

Furthermore, to avoid that the algorithm leaves too often the simplex, one may simply normalize the step γ_n by considering the predictable step

$$\tilde{\gamma}_n = \gamma_n \times \frac{n - 1}{V^1 + \cdots + V^{n-1}} \approx \frac{\gamma_n}{\mathbb{E}\, V}.$$

3.1.4 Simulations

The pseudo-real data setting. Firstly we explain how the data have been created. We have considered for V the traded volumes of a very liquid security – namely the asset BNP – during an 11 day period. Then we selected the N most correlated assets (in terms of traded volumes) with the original asset. These assets are denoted S_i, $i = 1, \ldots, N$

and we considered their traded volumes during the same 11 day period. Finally, the available volumes of each dark pool i have been modeled as follows using the mixing function

$$\forall 1 \leq i \leq N, \quad D_i := \beta_i \left((1-\alpha_i)V + \alpha_i S_i \frac{\mathbb{E}\, V}{\mathbb{E} S_i} \right)$$

where $\alpha_i \in (0,1),\ i = 1, \ldots, N$ are the recombining coefficients, $\beta_i,\ i = 1, \ldots, N$ are some scaling factors and $\mathbb{E}\, V$ and $\mathbb{E} S_i$ stand for the empirical mean of the data sets of V and S_i.

The shortage situation corresponds to $\sum_{i=1}^N \beta_i < 1$ since it implies $\mathbb{E}\left[\sum_{i=1}^N D_i\right] < \mathbb{E}\, V$.

The simulations presented here have been made with four dark pools ($N = 4$). Since the data used here covers 11 days and it is clear that unlike the simulated data, these pseudo-real data are not stationary : in particular they are subject to daily changes of trend and volatility (at least). To highlight these resulting changes in the response of the algorithms, we have specified the days by drawing vertical doted lines. The dark pool pseudo-data parameters are set to

$$\beta = \begin{pmatrix} 0.1 \\ 0.2 \\ 0.3 \\ 0.2 \end{pmatrix} \quad \text{and} \quad \alpha = \begin{pmatrix} 0.4 \\ 0.6 \\ 0.8 \\ 0.2 \end{pmatrix}$$

and the dark pool trading (rebate) parameters are set to $\rho = \begin{pmatrix} 0.01 \\ 0.02 \\ 0.04 \\ 0.06 \end{pmatrix}$.

The mean and variance characteristics of the data sets of $(V^n)_{n\geq 1}$ and $(D_i^n)_{n\geq 1}$, $i = 1, \ldots, 4$ are the following :

	V	D_1	D_2	D_3	D_4
Mean	955.42	95.54	191.08	286.63	191.08
Variance	2.01×10^6	9.05×10^3	4.29×10^4	4.73×10^5	5.95×10^4

Firstly, we benchmarked both algorithms on the whole data set (11 days) as though it were stationary without any resetting (step, starting allocation, etc.). In particular, the running means of the satisfactions ratios are computed from the very beginning for the first 1500 data, and by a moving average on a window of 1500 data. As a second step, we proceed on a daily basis by resetting the parameters of both algorithms (initial allocation for both and the step parameter γ_n of the optimization procedure) at the beginning of every day.

▷ **Long-term optimization** We observe that, except for the first and the fourth days where they behave similarly, the optimization algorithm is more performing than the reinforcement one. Its performance is approximately 30 % higher on average (see Figure 3).

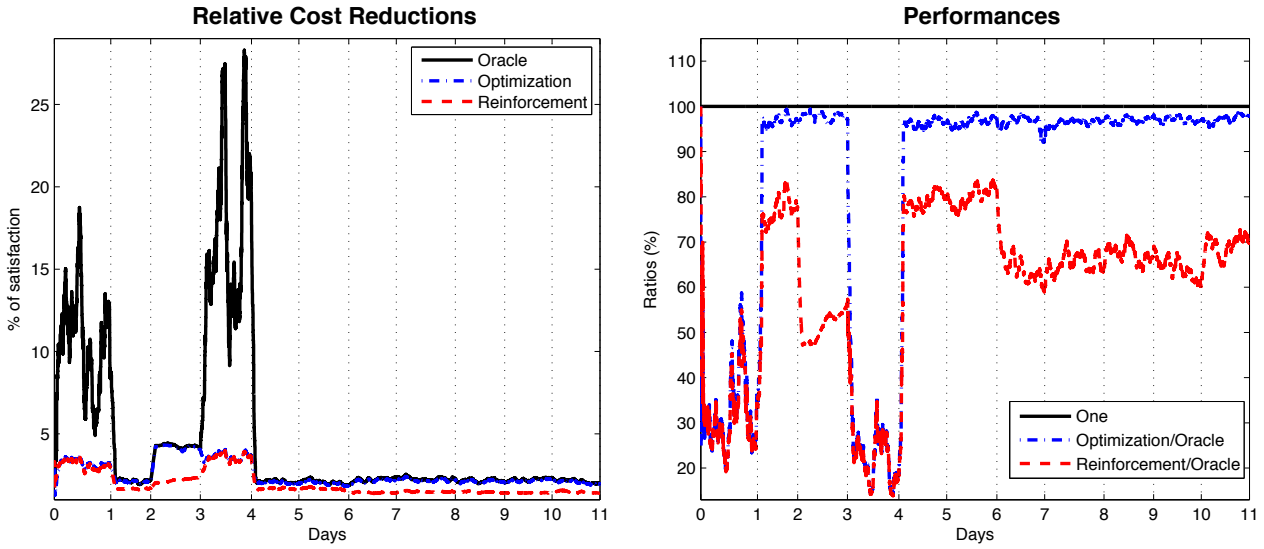

FIGURE 32 – *Long term optimization* : Case $N = 4$, $\sum_{i=1}^N \beta_i < 1$, $0.2 \leq \alpha_i \leq 0.8$ and $r_i^0 = 1/N$, $1 \leq i \leq N$.

This test confirms that the statistical features of the data are strongly varying from one day to another (see Figure 32), so there is no hope that our procedures converge in standard sense on a long term period. Consequently, it is necessary to switch to a short term monitoring by resetting the parameters of the algorithms on a daily basis as detailed below.

▷ **Daily resetting of the procedure** We consider now that we reset on a daily basis all the parameters of the algorithm, namely we reset the step γ_n at the beginning of each day and the satisfaction parameters and we keep the allocation coefficients of the preceding day. We obtain the following results

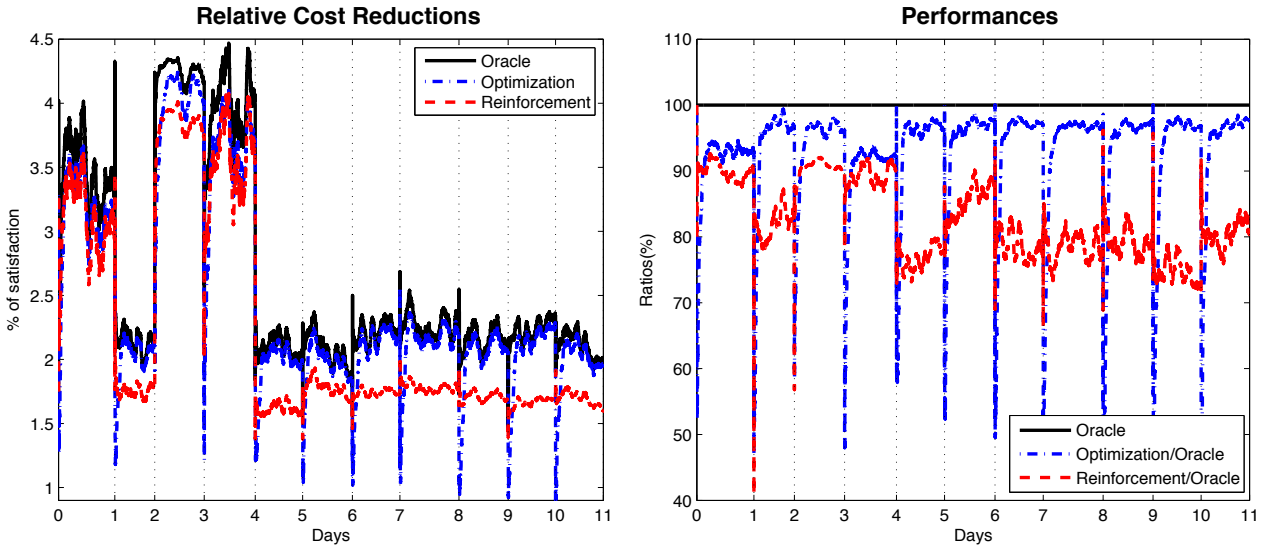

FIGURE 33 – *Daily resetting of the algorithms parameters* : Case $N = 4$, $\sum_{i=1}^{N} \beta_i < 1$, $0.2 \leq \alpha_i \leq 0.8$ and $r_i^0 = 1/N$ $1 \leq i \leq N$.

We observe (see Figure 33) that the optimization algorithm still significantly outperforms the reinforcement one, reaching more 95 % of the performance of the oracle. Furthermore, although not represented here, the allocation coefficients look more stable.

3.2 Choosing a Price to Post Without a Priori : Forward Optimality

This Section has been published as a part of *Optimal posting price of limit orders : learning by trading* (2013 Mathematics and Financial Economics 7, 359–403) [Laruelle *et al.* , 2013]

3.2.1 Modeling and design of the algorithm

We focus our work on the problem of optimal trading with limit orders on one security without needing to model the limit order book dynamics. To be more precise, we will focus on buy orders rather than sell orders in all that follows. We only model the execution flow which reaches the price where the limit order is posted with a general price dynamics $(S_t)_{t\in[0,T]}$ since we intend to use real data. However there will be two frameworks for the price dynamics : either $(S_t)_{t\in[0,T]}$ is a process bounded by a constant L (which is obviously an unusual assumption but not unrealistic on a short time scale see Section 3.2.2) or $(S_t)_{t\in[0,T]}$ is ruled by a Brownian diffusion model (see Section 3.2.2).

We consider on a short period T, say a dozen seconds, a Poisson process modeling the execution of posted passive *buy* orders on the market, namely

$$\left(N_t^{(\delta)}\right)_{0\leq t\leq T} \quad \text{with intensity} \quad \lambda(S_t-(S_0-\delta)) \text{ at time } t\in[0,T], \tag{79}$$

where $0\leq\delta\leq\delta_{\max}$ ($\delta_{\max}\in(0,S_0)$ is the depth of the order book), $\lambda:[-S_0,+\infty)\to\mathbb{R}_+$ is non-negative non-increasing function and $(S_t)_{t\geq 0}$ is a stochastic process modeling the dynamics of the "fair price" of a security stock (from an economic point of view). In practice one may consider that S_t represents the best opposite price at time t. Il will be convenient to introduce the cumulated intensity defined by

$$\Lambda_t(\delta,S):=\int_0^t\lambda(S_s-(S_0-\delta))ds \tag{80}$$

One way to build $N^{(\delta)}$ is to set

$$N_t^{(\delta)}=\widetilde{N}_{\int_0^t\lambda(S_s-(S_0-\delta))ds}$$

where $\widetilde{N}$ is a Poisson process with intensity 1 independent of the price $(S_t)_{t\in[0,T]}$.

This representation underlines the fact that for one given trajectory of the price S_t, the intensity of the Point process N is decreasing with δ : in fact the above representation for $N^{(\delta)}$ is even pathwise consistent in the sense that if $0<\delta<\delta'$ then

$$\mathbb{P}\text{-}a.s. \quad \left(\forall t\in[0,T], \quad N_t^{(\delta)}\leq N_t^{(\delta')}\right).$$

One natural question is how to account for simultaneously placing limit orders at multiple prices (which happens in reality) ; this case is much harder to handle and is not studied in this paper. In particular, due to interacting impact features, we would need a more sophisticated approach then simply considering $(N^{\delta(k)})_{1\leq k\leq K}$, processes as above with $\delta(1)<\delta(2)<\ldots<\delta(K)$ (with the same N).

We assume that the function λ is defined on $[-S_0,+\infty)$ as a finite non-increasing convex function. Its specification will rely on parametric or non parametric statistical estimation based on previously obtained transactions (see Figure 34 below. At time $t=0$, buy orders are posted in the limit order book at price $S_0-\delta$. Between t and $t+\Delta t$, the probability for such an order to be executed is $\lambda(S_t-(S_0-\delta))\Delta t$ where $S_t-(S_0-\delta)$ is the distance to the current fair price of our posted order at time t. The further the order is at time t, the lower is the probability for this order to be executed since λ is decreasing on $[-S_0,+\infty)$. Empirical tests strongly confirm this kind of relationship with a convex function λ (even close to an exponential shape, see Figure34 and [Avellaneda & Stoikov, 2008, Guéant *et al.* , 2013, Bayraktar & Ludkovski, 2012]).

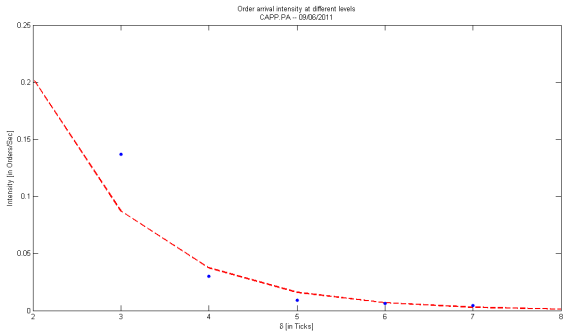

FIGURE 34 – Empirical probabilities of execution (blue stars) and its fit with an exponential law (red dotted line) with respect to the distance to the "fair price".

Over the period $[0,T]$, we aim to execute a portfolio of size $Q_T \in \mathbb{N}$ invested in the asset S. The execution cost for a distance δ is $\mathbb{E}\left[(S_0-\delta)\left(Q_T \wedge N_T^{(\delta)}\right)\right]$. We add to this execution cost a penalty function depending on the remaining quantity to be executed, namely we want to have Q_T assets in the portfolio at the end of the period T, so we buy the remaining quantity $\left(Q_T - N_T^{(\delta)}\right)_+$ at price S_T.

At this stage, we introduce a *market impact penalty function* $\Phi : \mathbb{R} \mapsto \mathbb{R}_+$, non-decreasing and convex, with $\Phi(0)=0$ to model the additional cost of the execution of the remaining quantity (including the market impact). Then the resulting cost of execution on a period $[0,T]$ reads

$$C(\delta) := \mathbb{E}\left[(S_0-\delta)\Big(Q_T \wedge N_T^{(\delta)}\Big) + \kappa\, S_T\, \Phi\Big(\big(Q_T - N_T^{(\delta)}\big)_+\Big)\right] \tag{81}$$

where $\kappa > 0$ is a free tuning parameter (the true cost is with $\kappa = 1$, but we could overcost the market order due to the bad estimation of the market impact of this order or conversely). When $\Phi(Q) = Q$, we assume we buy the remaining quantity at the end price S_T. Introducing a market impact penalty function $\Phi(x) = (1+\eta(x))x$, where $\eta \geq 0$, $\eta \not\equiv 0$, models the market impact induced by the execution of $\big(Q_T - N_T^{(\delta)}\big)_+$ at time T while neglecting the market impact of the execution process via limit orders over $[0,T)$. Our aim is then to minimize this cost by choosing the distance to post at, namely to solve the following optimization problem

$$\min_{0\leq\delta\leq\delta_{\max}} C(\delta). \tag{82}$$

Our strategy to solve (82) numerically using a large enough dataset is to take advantage of the representation of C and its first two derivatives as expectations to devise a recursive stochastic algorithm, specifically a stochastic gradient procedure, to find the minimum of the (penalized) cost function (see below). Furthermore we will show that under natural assumptions on the quantity Q_T to be executed and on the parameter κ, the function C is twice differentiable and *strictly convex on* $[0,\delta_{\max}]$ *with* $C'(0) < 0$. Consequently,

$$\operatorname{argmin}_{\delta\in[0,\delta_{\max}]} C(\delta) = \{\delta^*\}, \quad \delta^* \in (0,\delta_{\max}]$$

and

$$\delta^* = \delta_{\max} \quad \text{iff} \quad C \text{ is non-increasing on } [0,\delta_{\max}].$$

Criteria involving κ and based on both the risky asset S and the trading process especially the execution intensity λ can be established. We specify representations as expectations of the function C and its derivatives C' and C''. In particular we will exhibit a Borel functional

$$H : [0,\delta_{\max}] \times \mathbb{D}\left([0,T],\mathbb{R}\right) \longrightarrow \mathbb{R}$$

such that

$$\forall \delta \in [0,\delta_{\max}], \quad C'(\delta) = \mathbb{E}\Big[H\big(\delta,(S_t)_{t\in[0,T]}\big)\Big].$$

The functional H has an explicit form involving integrals over $[0,T]$ of the intensity $\lambda(S_t - S_0 + \delta)$ of the Poisson process $(N_t^{(\delta)})_{t\in[0,T]}$. In particular, any quantity $H\left(\delta,(S_t)_{t\in[0,T]}\right)$ can be simulated, up to a natural time discretization, either from a true dataset (of past executed orders) or from the stepwise constant discretization scheme of a formerly calibrated diffusion process modeling $(S_t)_{t\in[0,T]}$ (see below). This will lead us to replace, for practical implementations, the continuous time process $(S_t)_{t\in[0,T]}$ over $[0,T]$ with either a discrete time sample, *i.e.* a finite dimensional $\mathbb{R}^{m+1}$-valued random vector $(S_{t_i})_{0\leq i\leq m}$ (where $t_0 = 0$ and $t_m = T$) or with a time discretization scheme with step $\frac{T}{m}$ (typically the Euler scheme when $(S_t)_{t\in[0,T]}$ is a diffusion).

A THEORETICAL STOCHASTIC LEARNING PROCEDURE : Based on this representation of C', we can formally devise a recursive stochastic gradient descent *a.s.* converging toward δ^*. However to make it consistent, we need to introduce constraints so that it lives in $[0,\delta_{\max}]$. In the classical literature on Stochastic Approximation Theory (see [Kushner & Yin, 2003] and [Kushner & Clark, 1980]) this amounts to using a variant *with projection on the "order book depth interval"* $[0,\delta_{\max}]$, namely

$$\delta_{n+1} = \operatorname{Proj}_{[0,\delta_{\max}]}\Big(\delta_n - \gamma_{n+1} H\left(\delta_n, \left(S_t^{(n+1)}\right)_{t\in[0,T]}\right)\Big),\ n \geq 0,\ \delta_0 \in (0,\delta_{\max}), \tag{83}$$

where

- $\operatorname{Proj}_{[0,\delta_{\max}]} : (x \mapsto 0) \vee x \wedge \delta_{\max}$ denotes the projection on the (nonempty closed convex) $[0,\delta_{\max}]$,
- $(\gamma_n)_{n\geq 1}$ is a positive step sequence satisfying (at least) the minimal *decreasing step assumption* $\sum_{n\geq 1}\gamma_n = +\infty$ and $\gamma_n \to 0$.
- the sequence $\Big\{(S_t^{(n)})_{t\in[0,T]}, n \geq 0\Big\}$, is the "innovation" sequence of the procedure : ideally it is either a sequence of simulable independent copies of $(S_t)_{t\in[0,T]}$ or a sequence sharing some ergodic (or averaging) properties with respect to the distribution of $(S_t)_{t\in[0,T]}$.

The case of independent copies can be understood as a framework where the dynamics of S is typically a Brownian diffusion solution to a stochastic differential equation, which has been calibrated beforehand on a dataset in order to be simulated on a computer. The case of ergodic copies corresponds to a dataset which is directly plugged into the procedure *i.e.* $S_t^{(n)} = S_{t-n\Delta t}$, $t \in [0,T]$, $n \geq 0$, where $\Delta t > 0$ is a fixed shift parameter. To make this second approach consistent, we need to make the assumption that at least within a laps of a few minutes, the dynamics of the asset S (starting in the past) is *stationary* and shares *e.g.* mixing properties.

The resulting implementable procedure : In practice, the above procedure cannot be implemented since the full path $(S_t(\omega))_{t\in[0,T]}$ of a continuous process cannot be simulated nor a functional $H(\delta, (S_t(\omega))_{t\in[0,T]})$ of such a path can be computed. So we are led in practice to replace the "copies" $S^{(n)}$ by copies $\bar{S}^{(n)}$ of a time discretization $\bar{S}$ of step, say $\Delta t = \frac{T}{m}$, $(m \in \mathbb{N}^*)$. The time discretizations are formally defined in continuous time as follows

$$\bar{S}_t = \bar{S}_{t_i},\ t \in [t_i, t_{i+1}),\ i = 0, \ldots, m-1 \ \text{ with } t_i = \frac{iT}{m},\ i = 0, \ldots, m,$$

where

– $(\bar{S}_{t_i})_{0\leq i\leq m} = (S_{t_i})_{0\leq i\leq m}$ when $(S_{t_i})_{0\leq i\leq m}$ can be simulated exactly at a reasonable cost (see *e.g.* [Beskos & Roberts, 2005] for 1D-Brownian diffusions processes).

– $(\bar{S}_{t_i})_{0\leq i\leq m}$ is a time discretization scheme (at times t_i) of $(S_t)_{t\in[0,T]}$, typically an Euler scheme with step $\frac{T}{m}$.

Then, with an obvious abuse of notation for the function H, we can write the *implementable procedure* as follows :

$$\delta_{n+1} = \mathrm{Proj}_{[0,\delta_{\max}]}\Big(\delta_n - \gamma_{n+1} H\left(\delta_n, \left(\bar{S}_{t_i}^{(n+1)}\right)_{0\leq i\leq m}\right)\Big),\ n \geq 0,\ \delta_0 \in [0, \delta_{\max}] \tag{84}$$

where $\left(\bar{S}_{t_i}^{(n)}\right)_{0\leq i\leq m}$ are copies of $\left(\bar{S}_{t_i}\right)_{0\leq i\leq m}$ either independent or sharing "ergodic" properties, namely some averaging properties in the sense of [Laruelle & Pagès, 2012]. In the first case, one will think about simulated data after a calibration process and in the second case to a direct implementation using a historical high frequency database of best opposite prices of the asset S (with *e.g.* $\bar{S}_{t_i}^{(n)} = S_{t_i - n\frac{T}{m}}$).

3.2.2 Main convergence results

The following theorems give *a.s.* convergence results for the stochastic procedure (83) : the first one for i.i.d. sequences and the second one for "averaging" sequences (see [Laruelle & Pagès, 2012]).

I.i.d. simulated data from a formerly calibrated model. In this section, we consider innovation process $\left\{\left(\bar{S}_{t_i}^{(n)}\right)_{0\leq i\leq m}, n \geq 0\right\}$ coming from a diffusion model beforehand *calibrated on real data*, which can be simulated at time t_i, $0 \leq i \leq m$, either exactly or via a stepwise constant time discretization scheme.

Theorem 3 (*a*) Theoretical procedure. *Assume that C is strictly convex $[0, \delta_{\max}]$ with $C'(0) < 0$. Let $\left(S_t^{(n)}\right)_{t\in[0,T]}$, $n \geq 1$, be a sequence of i.i.d. copies of $(S_t)_{t\in[0,T]}$. Furthermore, assume that the decreasing step sequence satisfies the standard "decreasing step assumption"*

$$\sum_{n\geq 1} \gamma_n = +\infty \quad \textit{and} \quad \sum_{n\geq 1} \gamma_n^2 < +\infty. \tag{85}$$

Then the recursive procedure defined by (83) converges a.s. towards its target $\delta^* = \mathrm{argmin}_{\delta\in[0,\delta_{\max}]} C(\delta)$:

$$\delta_n \xrightarrow{a.s.} \delta^*.$$

(*b*) Implementable procedure. *Assume the cost function $\bar{C}$ related to the discretization scheme $(\bar{S}_t)_{t\in[0,T]}$ is strictly convex $[0, \delta_{\max}]$ with $\bar{C}'(0) < 0$ and the step sequence satisfies the "decreasing step" assumption. Let $\left(\bar{S}_{t_i}^{(n)}\right)_{0\leq i\leq m}$, $n \geq 1$, be a sequence of i.i.d. copies of $\left(\bar{S}_{t_i}\right)_{0\leq i\leq m}$, then the recursive procedure defined by (84) converges a.s. towards its target $\bar{\delta}^* = \mathrm{argmin}_{\delta\in[0,\delta_{\max}]} \bar{C}(\delta)$.*

This theorem is a straightforward application of the classical *a.s.* convergence for constrained stochastic algorithms. In particular, the fact that in the original theorem the innovation process takes values in a finite dimensional space $\mathbb{R}^q$ plays no role in the proof.

Direct implementation on a historical high frequency dataset sharing averaging properties. In this framework we will focus on the time discretized procedure *i.e.* on $(\bar{S}_t)_{t\in[0,T]}$ rather than on $(S_t)_{t\in[0,T]}$ itself. Keep in mind that, when directly implementing a high frequency dataset, then

$$\bar{S}_t = S_{t_i},\ t \in [t_i, t_{i+1}), i = 0, \dots, m \ \text{ and } \bar{S}_T = S_T.$$

and that the sequence $(\bar{S}^{(n)}_{t_i})_{0\le i\le m}$, $n \ge 1$, is usually obtained by shifting the data as follows : if $\Delta t > 0$ denotes a fixed time shift parameter such that $t_i - t_{i-1} = \Delta t = \frac{T}{m}$, we set

$$\forall\, t \in [0,T], \quad \bar{S}^{(n)}_{t_i} = \bar{S}_{t_i - n\Delta t} = \bar{S}_{t_{i-n}}.$$

We will assume that the sequence $(\bar{S}^{(n)}_{t_i})_{0\le i\le m}$ shares an averaging property with respect to a distribution ν as developed in [Laruelle & Pagès, 2012]. The definition is recalled below.

Definition 1 *Let $m \in \mathbb{N}$ and ν be a probability measure on $([0,L]^{m+1}, \mathcal{B}or([0,L]^{m+1}))$. A $[0,L]^{m+1}$-valued sequence $(\xi_n)_{n\ge1}$ is ν-averaging if*

$$\frac{1}{n}\sum_{k=1}^{n} \delta_{\xi_k} \overset{(\mathbb{R}^{m+1})}{\Longrightarrow} \nu \quad \text{as } n \to \infty.$$

Then $(\xi_n)_{n\ge1}$ satisfies

$$D^*_n(\xi) := \sup_{x\in[0,L]^{m+1}} \Big|\frac{1}{n}\sum_{k=1}^{n} \mathbf{1}_{\{0,x\}}(\xi_k) - \nu(\{0,x\})\Big| \longrightarrow 0 \quad \text{as } n \to \infty,$$

where $D^*_n(\xi)$ is called the discrepancy at the origin or star discrepancy.

The resulting execution cost function $\bar{C}$ is defined by (81) where S is replaced by $(\bar{S}_t)_{t\in[0,T]}$ whose distribution is entirely characterized by the distribution ν. In some sense this function $\bar{C}$ is *the best possible approximation of the true execution function C* that we can get from the high frequency database.

In this setting, we apply the previous results to the price sequence $\big\{\big(\bar{S}^{(n)}_{t_i}\big)_{0\le i\le m}, n \ge 0\big\}$, *i.e.* we set for every $n \ge 1$, $\xi_n = \big(\bar{S}^{(n)}_{t_i}\big)_{0\le i\le m}$. In particular we will make the assumption that the dataset is bounded by a real number $L \in (0,+\infty)$ so that $\xi_n \in [0,L]^{m+1}$ for every $n \ge 1$. Moreover, we will need to prove the existence of a *pathwise* Lyapunov function, which means in this one dimensional setting that $H(\cdot, ((s_{t_i})_{0\le i\le m})$ is non-decreasing for every $(s_{t_i})_{0\le i\le m} \in \mathbb{R}^{m+1}_+$, $n \ge 1$.

Theorem 4 IMPLEMENTABLE PROCEDURE. *Let $\lambda(x) = Ae^{-ax}$, $A > 0$, $a > 0$. Assume $\big(\bar{S}^{(n)}\big)_{n\ge1}$ is an $[0,L]^{m+1}$-valued ν-averaging sequence where ν is a probability measure on $(\mathbb{R}^{m+1}, \mathcal{B}or(\mathbb{R}^{m+1}))$. Assume that the execution cost function C is strictly convex over $[0,\delta_{\max}]$ with $\bar{C}'(0) < 0$ and $\bar{C}'(\delta_{\max}) > 0$. Finally assume that the step sequence $(\gamma_n)_{n\ge1}$ is a positive non-increasing sequence satisfying*

$$\sum_{n\ge1}\gamma_n = +\infty, \quad nD^*_n(\bar{S})\gamma_n \underset{n\to\infty}{\longrightarrow} 0, \quad \text{and} \quad \sum_{n\ge1} nD^*_n(\bar{S}) \max\big(\gamma_n^2, |\Delta\gamma_{n+1}|\big) < +\infty. \tag{86}$$

Furthermore (having in mind that $\bar{S}_0 = S_0$), assume that

$$Q_T \ge 2T\lambda(-\bar{S}_0) \quad \text{and} \quad \kappa \le \frac{1 + a(\bar{S}_0 - \delta_{\max})}{a\,\big\|\bar{S}\big\|_\infty (\Phi(Q_T) - \Phi(Q_T - 1))} \tag{87}$$

Then the recursive procedure defined by (84) converges a.s. towards its target $\bar{\delta}^* = \operatorname{argmin}_{\delta\in[0,\delta_{\max}]}\bar{C}(\delta)$:

$$\delta_n \overset{a.s.}{\longrightarrow} \bar{\delta}^*.$$

Practical comments on the needed bounds.

- The constraint $Q_T \ge 2T\lambda(-S_0)$ is structural : it only involves parameters of the model and the asked quantity Q_T. It means that Q_T is likely not to be fully executed before the end of a slice of duration T (*i.e.* the intensity of trades obtained very far away from the current price is smaller than $Q_T/2$).
- The criterion involving the free parameter κ may be interpreted in two different ways, depending on the modeling of the "market impact".
  - *The market impact does not depend on the remaining quantity to be traded* (*i.e.* when $\Phi = \mathrm{id}$ or $\eta \equiv 0$ which implies that $\Phi(Q_T) - \Phi(Q_T - 1) = 1$). This setting is appropriate for executing very small quantities or trading very liquid assets (like equity futures). Then the criterion on the free parameter κ reads

$$\kappa \le \frac{\bar{S}_0 - \delta_{\max}}{\big\|\bar{S}\big\|_\infty} + \frac{1}{k\,\big\|\bar{S}\big\|_\infty}.$$

It states that in this case, the constant premium to pay for the remaining quantity (*i.e.* κ, in basis points) has to be lower than the price range inside which we wish to trade (*i.e.* $(S_0 - \delta_{\max})/\|S\|_\infty$) plus a *margin*, namely $1/(a\|S\|_\infty)$. It can be seen as a *symmetry argument* : a model where one cannot imagine buying at a lower price than a given threshold, one cannot accept paying (on the other side) more market impact than this very threshold either.

– *The market impact of the remaining quantity is a function of the quantity.* The interpretation is very similar to the previous one. In this case a quantity homogeneous to the market impact (*i.e.* $\kappa(\Phi(Q_T) - \Phi(Q_T - 1))$ in basis points) should not exceed $(\bar{S}_0 - \delta_{\max})/\left\|\bar{S}\right\|_\infty + 1/(a\left\|\bar{S}\right\|_\infty)$. Here again it is a symmetry argument : the trader cannot accept paying more market impact for almost one share (*i.e.* $\Phi(Q_T) - \Phi(Q_T - 1)$ plays the role of $\Phi(1)$ but "taken around Q_T") than his reasonable trading range ($(\bar{S}_0 - \delta_{\max})/\left\|\bar{S}\right\|_\infty$ in basis points again), plus a margin. Looking more carefully at this margin : $1/(a\left\|\bar{S}\right\|_\infty)$ (the same as in the constant market impact case), we see that it is in basis points, and means that if one considers large intensities of fill rates (*i.e.* a is small) then the trader can be lazy on the market impact constraint (because his margin, proportional to $1/a$, is large in such a case), mainly because he will have to pay market impact not that often. If on the contrary, a is large (*i.e.* he will have remaining quantities) then he needs to really fulfill his constraint on market impact.

The needed bounds to obtain the convergence of δ_n toward δ^* are thus not only fulfilled naturally. They also emphasis the consistency of the model and the mechanisms of the proofs.

3.2.3 Simulations

Setting1 : *With market impact* ($\Phi(x) = (1 + A'e^{a'x})x$)

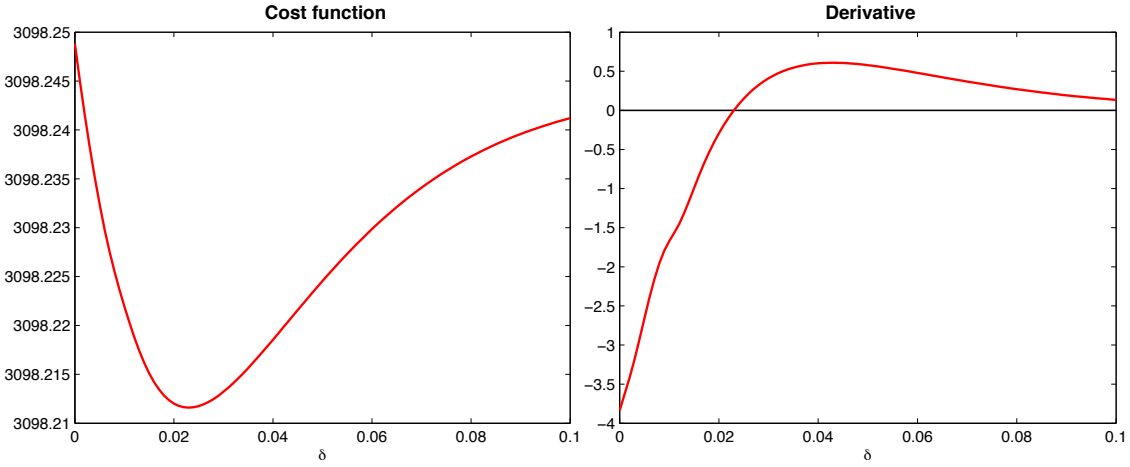

FIGURE 35 – $\Phi \not\equiv \mathrm{id}$: $A = 1/50$, $a = 50$, $Q = 100$, $\kappa = 1$, $A' = 0.001$, $a' = 0.0005$ and $N_{\text{cycles}} = 220$.

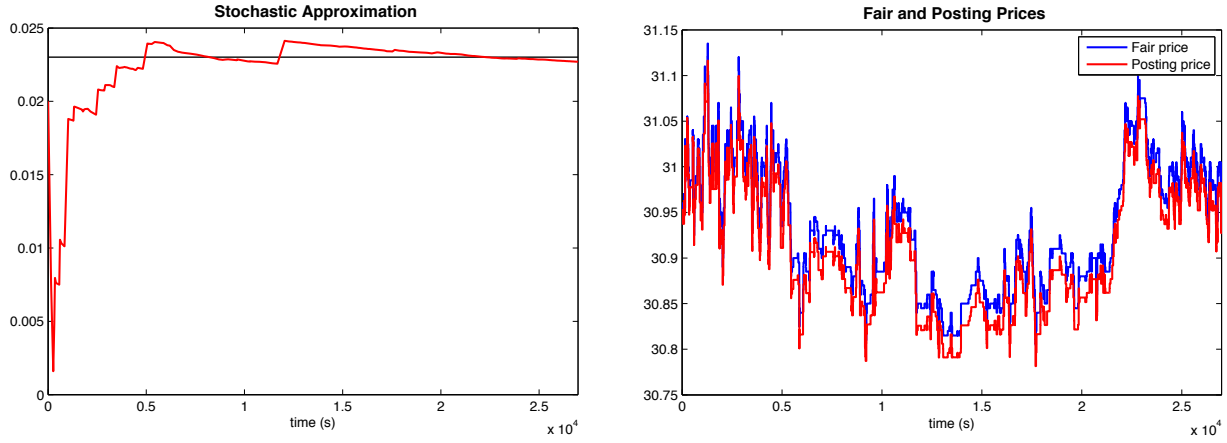

FIGURE 36 – *Crude algorithm* : $A = 1/50$, $a = 50$, $Q = 100$, $\kappa = 1$, $A' = 0.001$, $a' = 0.0005$, $N_{\text{cycles}} = 220$ and $\gamma_n = \frac{1}{550n}$, $1 \le n \le N_{\text{cycles}}$.

Once we checked that our algorithm converges using real data for the optimization on a whole day, we can implement the on-line procedure for practical implementation. We use the same data (Accor SA on 11/11/2010) and the same parameters for the cost and penalty functions, but also for the algorithm. Thus we consider a period T=15 trades and we use $N_{\text{cycles}} = 220$ for the recursive stochastic procedure. The difference is that we do not consider distinct periods of 15 trades, but we use a sliding window of 15 trades along the path of the fair price $(S_t)_{t\in[0,T]}$ moving back a trade at each step of the algorithm. Consequently, this procedure takes into account all the possible periods of 15 trades over the day and requires less data than the previous one to converge.

We consider now that we reset the step γ_n at the beginning of each period and we keep the posting price of the end of the previous period. To highlight these resulting changes in the response of the algorithm, we have specified the periods by drawing vertical doted lines. We present the results on both market impact frameworks (namely $\Phi \neq \mathrm{id}$ and $\Phi = \mathrm{id}$) and for crude and averaged algorithms as in the long-term optimization section.

Setting 1 : *With market impact* $(\Phi(x) = (1 + A'e^{a'x})x)$

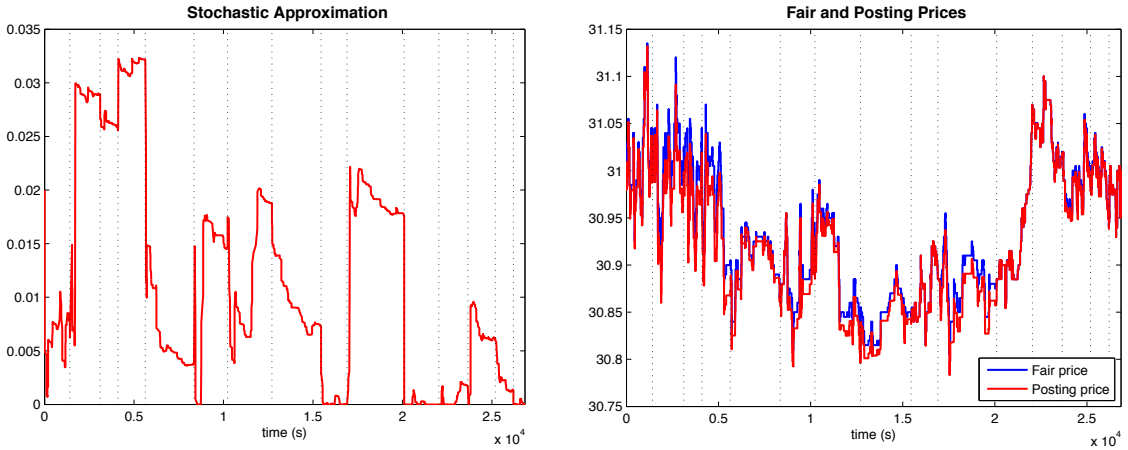

FIGURE 37 – *Crude algorithm* : $A = 1/50$, $a = 50$, $Q = 100$, $\kappa = 1$, $A' = 0.001$, $a' = 0.0005$ and $\gamma_n = \frac{1}{550n}$.

3.3 On the Fly Monitoring of Large Automated Portfolios

This Section has been published as a part of *Realtime market microstructure analysis : online Transaction Cost Analysis* (2014 Quantitative Finance , 0–19) [Azencott *et al.* , 2014]

Motivated by the practical challenge in monitoring the performance of a large number of algorithmic trading orders, this paper provides a methodology that leads to automatic discovery of causes that lie behind poor trading performance. It also gives theoretical foundations to a generic framework for real-time trading analysis. The common acronym for investigating the causes of bad and good performance of trading is TCA (Transaction Cost Analysis [Rosenthal, 2009]). Automated algorithms take care of most of the traded flows on electronic markets (more than 70% in the US, 45% in Europe and 35% in Japan in 2012). Academic literature provides different ways to formalize these algorithms and show how optimal they can be from a mean-variance [Almgren & Chriss, 2000], a stochastic control [Guéant *et al.* , 2013], an impulse control [Bouchard *et al.* , 2011] or a statistical learning [Laruelle *et al.* , 2013] viewpoint. This paper is agnostic about the way the algorithm has been built and provides a theoretical formalism to identify in real-time the market conditions that *influenced* its efficiency or inefficiency. For a given set of characteristics describing the market context, selected by a practitioner, we first show how a set of additional derived explanatory factors, called *anomaly detectors*, can be created for each market order (following for instance [Basseville & Nikiforov, 1993]). We then will present an online methodology to quantify how this extended set of factors, at any given time, predicts (i.e. have *influence*, in the sense of predictive power or information defined in [Cristianini & Shawe-Taylor, 2000], [Shannon, 1948] and [Alkoot & Kittler, 1999]) which of the orders are underperforming while calculating the predictive power of this explanatory factor set. Armed with this information, which we call *influence analysis*, we intend to empower the order monitoring user to take appropriate action on any affected orders by re-calibrating the trading algorithms working the order through new parameters, pausing their execution or taking over more direct trading control. Also we intend that use of this method in the post trade analysis of algorithms can be taken advantage of to automatically adjust their trading action.

We consider a portfolio of at most K trading orders $\mathbf{T}(k), k = 1, \dots, K$ driven by automatic trading algorithms, and supervised by one or more traders. Each trading order is defined by a few "static" variables such as buy/sell label, order size, trading place, section, country, capitalisation, free-float, benchmark type (VWAP, arrival price, etc), etc. In our intra-day benchmark studies a portfolio typically involves $200 \leq K_t \leq 700$ active orders at any arbitrary 5 minute time slice.

3.3.1 The Monitoring Framework

Market descriptors. Each trading order $\mathbf{T}(k)$ focuses on a specific asset whose dynamics is recorded at each time point t through a fixed number of basic "market descriptors" $M_t^1(k), M_t^2(k), M_t^3(k), \dots$; in our benchmark study below, we have focused on a subset of the following market descriptors :

- M^1 = Volatility
- M^2 = Spread
- M^3 = Momentum in bid-ask spread
- M^4 = Momentum in bp

This list can be augmented by the *rarity scores* $\text{Score}(M^i)$ of the market variables M^i. These scores are defined by $\text{Score}(M^i) = F^i(M^i)$ where F^i is the cunulative distribution function of M_i. Each such score necessarily has a uniform distribution [Borovkov, 1998]. If poor performance on a given set of stocks is due to a strong increase in the volatility level, the concrete cause may either be due to volatility reaching an "absolute" psychological threshold, or to volatility being high *relatively to its usual levels*. In this last case the volatility score will be a better explanatory factor for poor performance. We will use here the following scores, increasing our number of market variables :

- M^5 = Volume Rarity Score
- M^6 = Volatility Rarity Score
- M^7 = Spread Rarity Score

Figure 40 displays typical intraday (January 14, 2011) plots of the time series corresponding to the 7 basic market variables listed above for an anonymous stock.

We also display heat map representations of rarity scores for multiple stocks in Figures 38 and 41, where each row displays the time series of rarity scores for one single stock (associated to one trading order in our benchmark data), and each column represents one time slice. Clearly "extreme" rarity scores tend to appear in clusters, and to co-occur across multiple stocks.

As will be seen below, high co-occurrence frequency of "dynamic anomalies" such as peaks, jumps, etc within a group of stocks tend to "explain" simultaneous lack of performance for the corresponding trading orders.

Trading performance evaluation. We also select a "trading performance evaluator" PE providing at each time point t, and for each active trading order $\mathbf{T}(k)$, a quantitative evaluation $\text{PE}_t(k)$ for the current performance of $\mathbf{T}(k)$.

In our study we have selected by default PE = "slippage in bid-ask spread" (*slippage* being the average price of the order minus the benchmark –VWAP, arrival price, close price, etc– for a sell order, and the opposite for a buy order), but there are no restrictions on the user choice for this PE variable. In particular, other examples of PE include "slippage in bp" (basis points), "Slippage in Dollars", "Absolute value of slippage in bid-ask spread", etc. We are

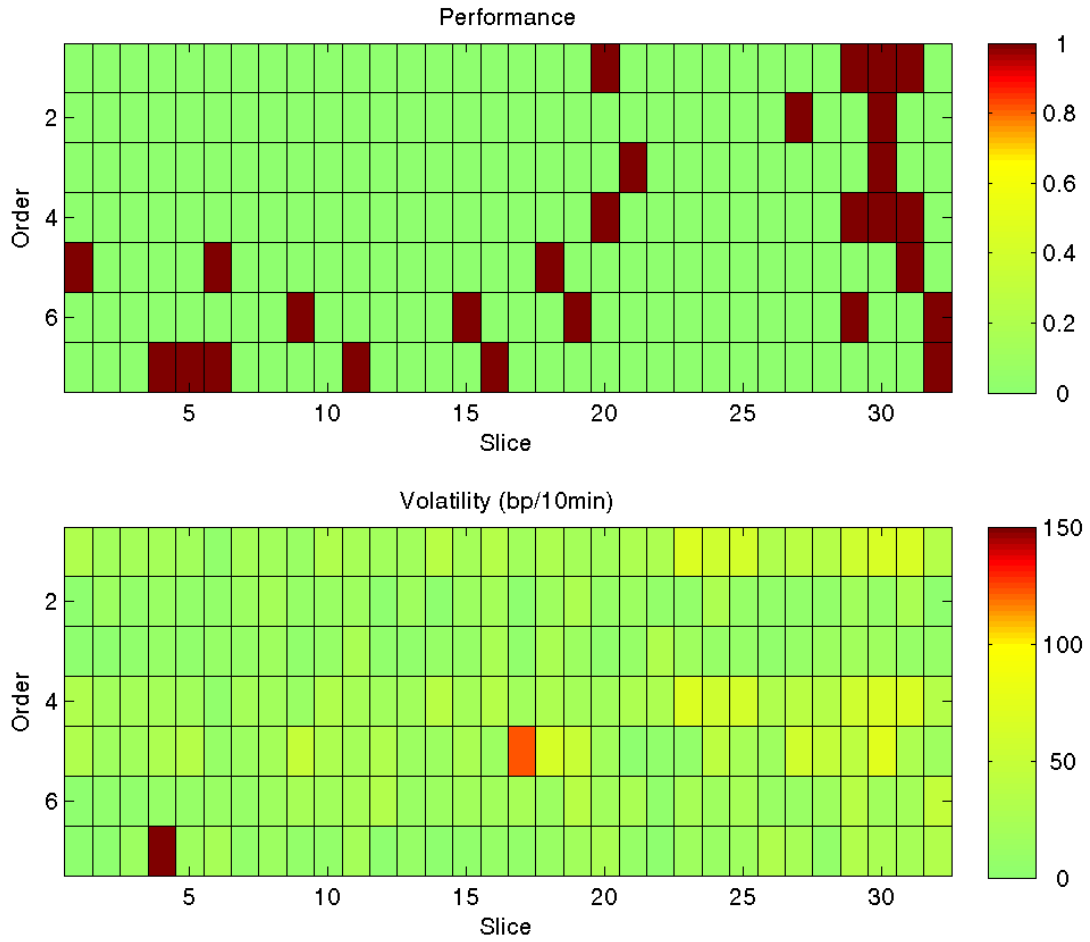

FIGURE 38 – Heatmap of online performances of a traded basket of 7 orders (top) in conjunction with values of one explanatory variable (the volatility ; bottom), the correspondance between the two heatmaps is not obvious.The abcissa represents the evolution over five-minute time slices.

assuming that degraded performances are associated to low values of PE. For each trading order $\mathbf{T}(k)$ the performance evaluator and the 7 market descriptors are volume averaged over successive time slices of arbitrary duration (set at 5 minutes for our benchmark study). Thus we generate 8 time series $M_t^1(k), \ldots, M_t^7(k)$ and $\mathrm{PE}_t(k)$ indexed by time slices t. These time series generically have missing values since orders do not necessarily begin or end at the same time. Fix a low percentile threshold q such as $q = 5\%$ or $q = 3\%$ to binarise the performance evaluator.

At time slice t, call $K_t \leq K$ the number of currently active trading orders $\mathbf{T}(k)$. The $q\%$- quantile of the corresponding K_t performance evaluations $\mathrm{PE}_t(k)$ is denoted by ℓ_t. We consider ℓ_t as a PE-threshold, separating "bad trading performances" (tagged "1") from "normal trading performances" (tagged "0"). We then *binarize the performance evaluations* $\mathrm{PE}_t(k)$ by setting

$$\begin{aligned} Y_t(k) &= 1 \quad \text{if} \quad \mathrm{PE}_t(k) < \ell_t, \\ Y_t(k) &= 0 \quad \text{if} \quad \mathrm{PE}_t(k) \geq \ell_t. \end{aligned} \tag{88}$$

In Figure 40, we plot an example of the intraday behaviour of a trading algorithm. Its trading performance evaluation PE_t can observed within a trading day, in real-time, like some of the market context variable we used. A trading order may or may not be active at a given time slice as observed in Figure 38.

Figure 41 displays synchronous intraday plots of trading performance evaluations in conjunction with the values of a few selected market variables. An essential goal of our methodology is, for each fixed time slice, to quantify on line the current influence of a market variable on trading performance degradation. Our automated online influence quantification replaces expert visual inspection of current trading orders performances, to identify critical market variables explaining trading performance degradations. For instance, visual inspection of Figures 38 and 41 will naturally "explain" the low performances observed at time slices $t = 39, 40, 41$ by the obvious trend changes simultaneously observed on rarity scores as well as by the volatility peak.

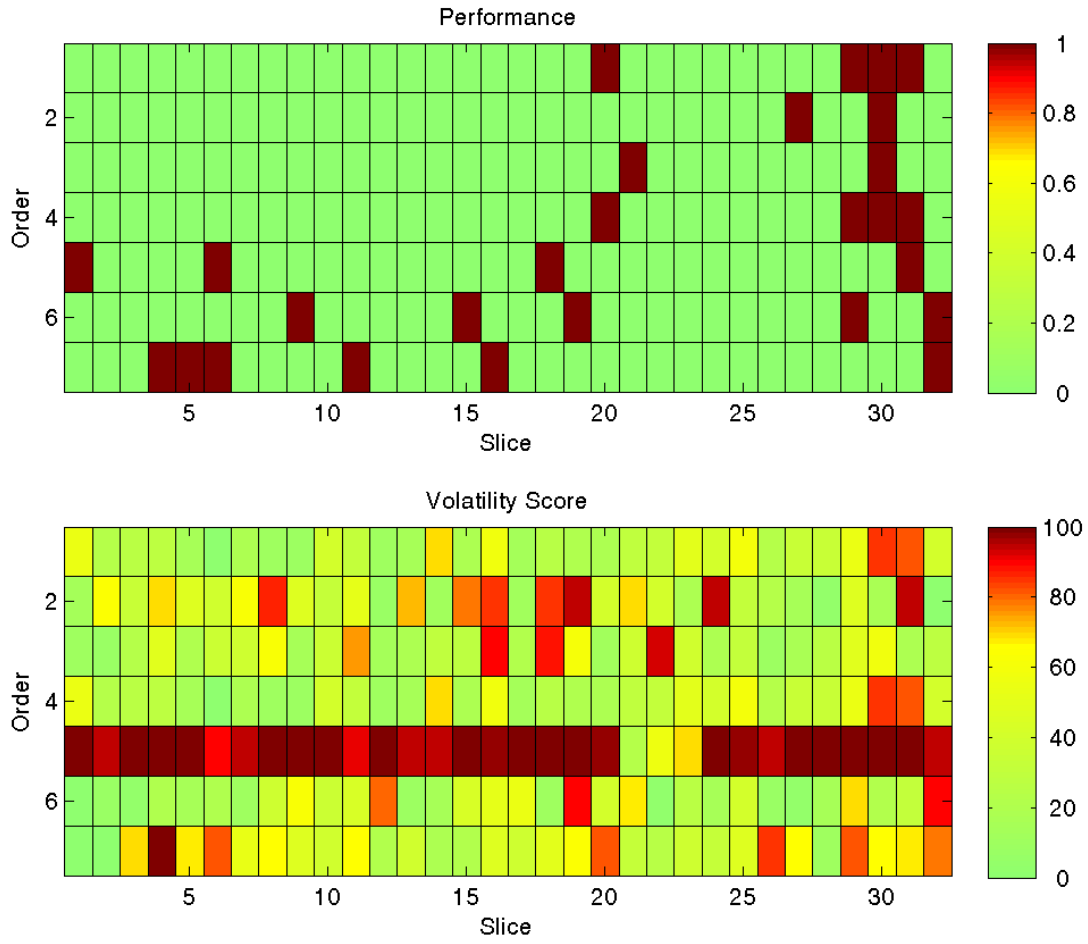

FIGURE 39 – Heatmap of online performances of the some basket of 7 orders (top) in conjunction with the scores of the same explanatory variable (the volatility ; bottom), compared to Figure 38, one can see that the scores change the high and low values of the explanatory variables, giving birth to more potential conjunctions with bad trading performances.

3.3.2 Online anomaly detectors

Anomaly detection. Online anomaly detection is a critical step in many applications, such as safety of complex systems, safety monitoring in automotive or aeronautics industries, remote health monitoring in biomedicine, real-time quality control for industrial production lines, etc (See [Aviv, 1991, Basseville & Nikiforov, 1993, Basseville, 1988, Gustafsson, 2000] [Lehalle & Azencott, 2004]).
In the context of trading performance online monitoring, it is also quite natural to systematically enrich raw market descriptors by automated detection of anomalies affecting their dynamics. We have thus developed algorithmics dedicated to the online implementation of this processing step.

The occurrence of visually evident anomalies can be detected by algorithmic tracking of local regime changes in market descriptors dynamics, and may have potentially strong influence on performance degradation for the corresponding trading orders. We have hence developed and implemented a set of 3 parametrized *anomaly detectors*, dedicated to the online identification of "significant" *Peaks or Crenels, Jumps, and Trend Changes* on generic time series. These 3 detectors automatically locate emerging anomalies, quantify their intensities, and filter them through adjustable gravity thresholds.

Building online detectors. Consider a generic discrete time series U_t. A smoothed "baseline" BU_t is generated as a moving local median of U_t. One then computes the local standard deviation σ_t of the "noise" $U_t - BU_t$, and in turn, this defines "outlier" values of U_t. Our three online anomaly detectors are based on local trend extractions at each time slice t by fitting linear or quadratic regression models on short moving time windows to the left and the right of t. The detector parameters have simple geometric interpretations for the users and are kept fixed during online influence analysis. Each anomaly detector is dedicated to a fixed type of anomaly, and generates a binary time series encoding the presence or absence of this anomaly type at successive time slices of U_t.

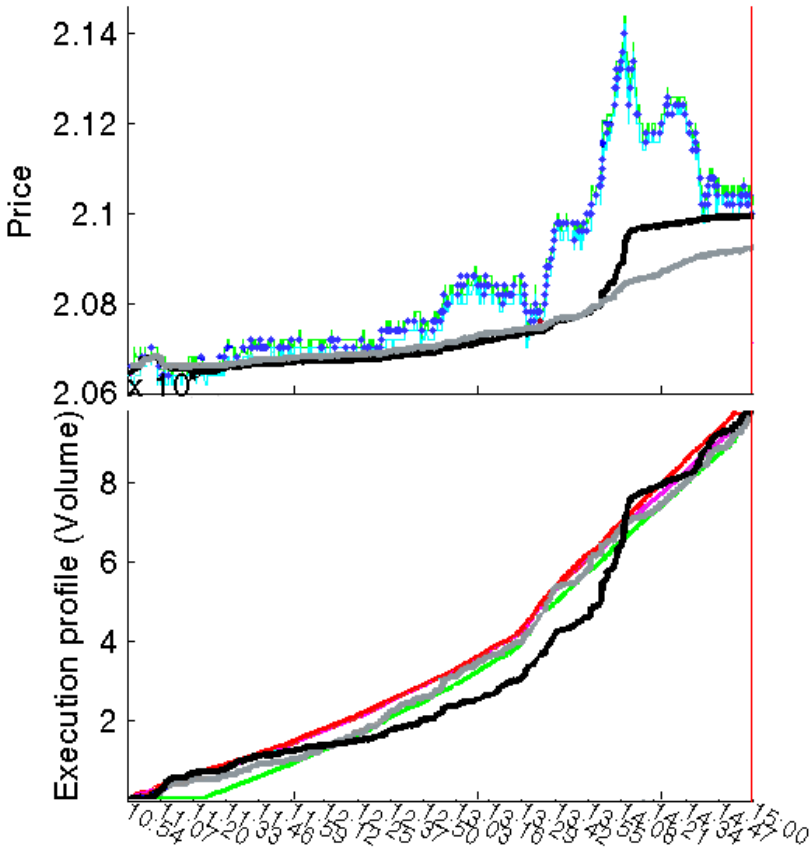

FIGURE 40 – Example of intraday behaviour of performance variables for a given order. Top : the prices (the variations of the average obtained price is in gray, the market VWAP in dark) ; bottom : the cumulated traded volume (grey) and market volume (dark).

Peaks/ crenels detector. A "peak" is the sudden occurrence of a high "outlier" value of U_t. More generally, a "crenel" is a cluster of successive high "outliers" with approximately equal values. Each crenel is described by 3 geometric "crenel features", namely, its *time duration*, its *thickness* (i.e., absolute difference between highest and lowest crenel points), and its *height* above the baseline BU_t. Minimal threshold values are imposed on these 3 features, as well as a minimal *time gap* between successive crenels.

To detect peaks and/or crenels on the series U_t, one first extracts outliers with respect to the baseline BU_t ; then one applies simple filters to detect local geometric configurations of outliers which satisfy the threshold constraints imposed on the three "crenel features" described above. If a peak or crenel is detected at time t, then "peak/crenel intensity" $Peak_t$ is set equal to "height" of the peak/crenel above the baseline. If no peak or crenel is detected at time t, one sets $Peak_t = 0$.

Jumps detector. A "jump" at time t is a sudden level change between the U_s values on finite time windows to the left and to the right of t. Bona fide jumps are described by 2 features, namely, a *duration* $2L$ and a minimal *jump size* Δ. For each t, one fits two distinct quadratic regressions to the baseline BU_s, namely Reg^- for $(t-1-L) \leq s \leq (t-1)$ and Reg^+ for $t \leq s \leq (t+L)$, where L is a fixed parameter.

A jump is detected at t if the "jump size" $JS(t) = |Reg^+(t) - Reg^-(t)|$ is larger than Δ, provided the two regressions have small enough residuals. If a jump is detected at time t on the series U_t, then "jump intensity" $Jump_t$ is set equal to the "jump size" $JS(t)$. If no trend change is detected at time t, one sets $Jump_t = 0$.

Trend changes detector. Bona fide "trend changes" are described by 3 features, a *duration* $2L$, a minimal *slope change* λ, and a *continuity modulus* ε. For each t, one fits as above two quadratic regressions to the baseline BU_s, namely Reg^- to the left of t and Reg^+ to the right of t. Call α^+, α^- the slopes of Reg^+, Reg^-, and define the "trend change size"

$$TCS(t) = |\alpha^+ - \alpha^-|.$$

A local "trend change" is detected at time t if Reg^+, Reg^- have sufficiently small residuals and verify

$$TCS(t) > \lambda \, ; \quad \text{and} \quad |Reg^+(t) - Reg^-(t)| < \varepsilon.$$

If a Trend Change is detected at time t on the series U_t, then "trend change intensity" $Trend_t$ is set equal to the "trend change size" $TCS(t)$. If no jump is detected at time t, one sets $Trend_t = 0$.

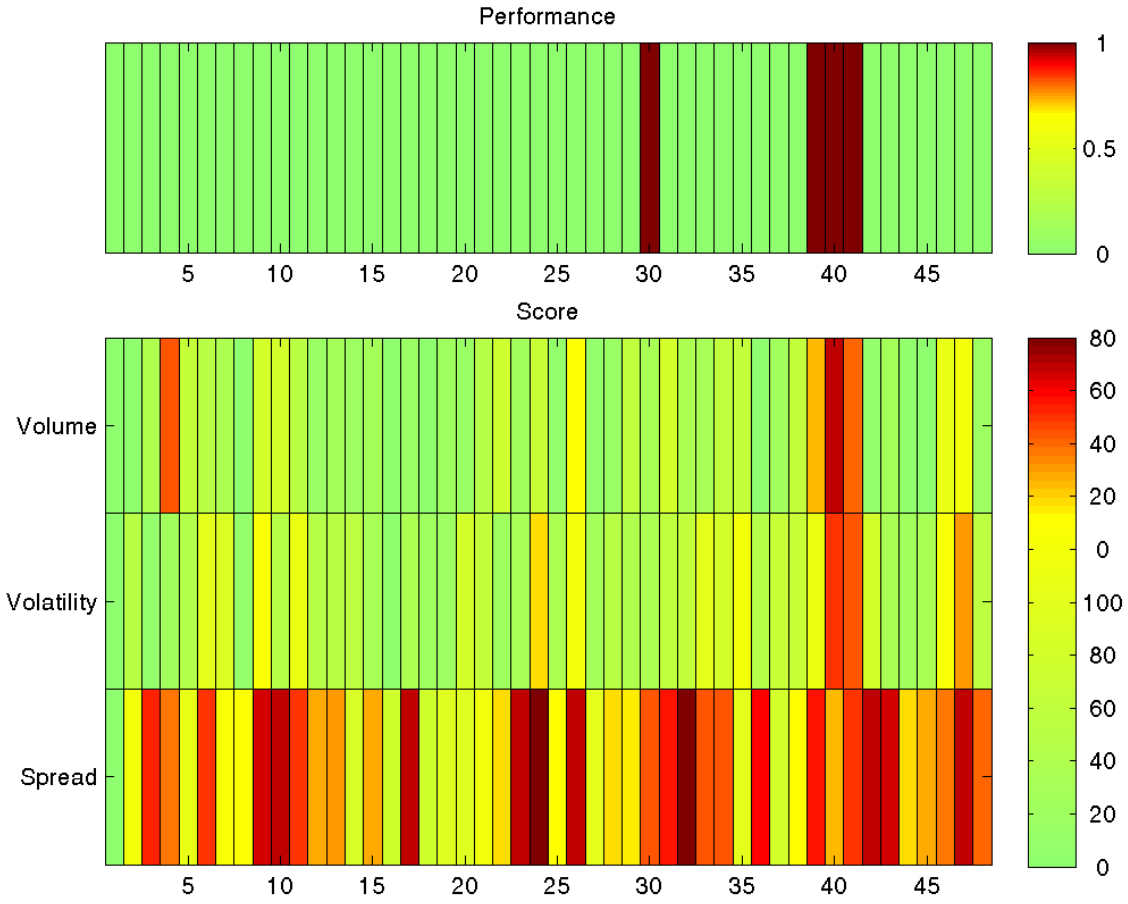

FIGURE 41 – Conjunction of the performance (top curve) of one traded order (first line of Figure 41) with the market context (bottom heatmap).

3.3.3 Online Influence Computation : Result

Influence computation for a single explanatory factor. Consider any real valued single explanatory factor Z having a continuous conditional density function $w(z)$ given $Y = 1$. To predict Y given Z the best binarization of Z should select a subset B of S maximizing the predictive power of the predictor $\mathbf{1}_B(Z)$. It can be shown that any optimal B should be the set of all $z \in \mathbb{R}$ such that $w(z) > c$ for some $c > 0$. Thus an optimal B must be a closed level set L of the unknown conditional density function $w(z)$. The family of all closed sets in $\mathbb{R}$ is well known to have infinite Vapnik-Cervonenkis dimension (see [Vapnik & Chervonenkis, 1971],[Vapnik & Chervonenkis, 1971],[Vapnik, 2010]). So, in view of Vapnik's theorems on automated learning (see [Vapnik, 2006],[Cristianini & Shawe-Taylor, 2000]), empirical optimal selection of B among all closed sets will have weak generalization capacity, increasing extremely slowly with the number K_t of data. This has naturally led us to select sub-optimal but much more robust classes of predictors, whith radically reduced Vapnik-Cervonenkis dimension.
In the cases where $w(z)$ can be considered as roughly unimodal or monotonous, the level sets of w are unions of at most two disjoint intervals. We thus deliberately restrict our class of binary predictors of Y to two-sided ones :

Definition 2 (Two-sided binary predictor) Two-sided binary predictors *are of the form* $h_\theta = \mathbf{1}_B(Z)$ *where* B *is the union of the two disjoint intervals* $(-\infty, \theta^-)$ *and* $(\theta^+, +\infty)$ *, indexed by the vector* $\theta = (\theta^-, \theta^+) \in \mathbb{R}^2$*, with* $\theta^- < \theta^+$.

Note that h_θ predicts bad trading performances if and only if the explanatory factor Z takes sufficiently large or sufficiently small values. Hence these estimators of trading performance degradation have *an immediate interpretability* for natural users of online trading performance monitoring.

At time t, given the current K_t joint observations of the explanatory factor Z and of the binarized trading performance Y, an immediate counting provides for each θ the empirical estimates $\hat{P}^1$ and $\hat{P}^0$ of the probabilities of correct prediction P^1, P^0 for the estimator h_θ, given by

$$P^1 = \mathbb{P}\{(Z < \theta^-) \cap (Y = 1)\} + \mathbb{P}\{(Z > \theta^+) \cap (Y = 1)\}$$

$$P^0 = \mathbb{P}\{(\theta^- < Z < \theta^+) \cap (Y = 0)\}$$

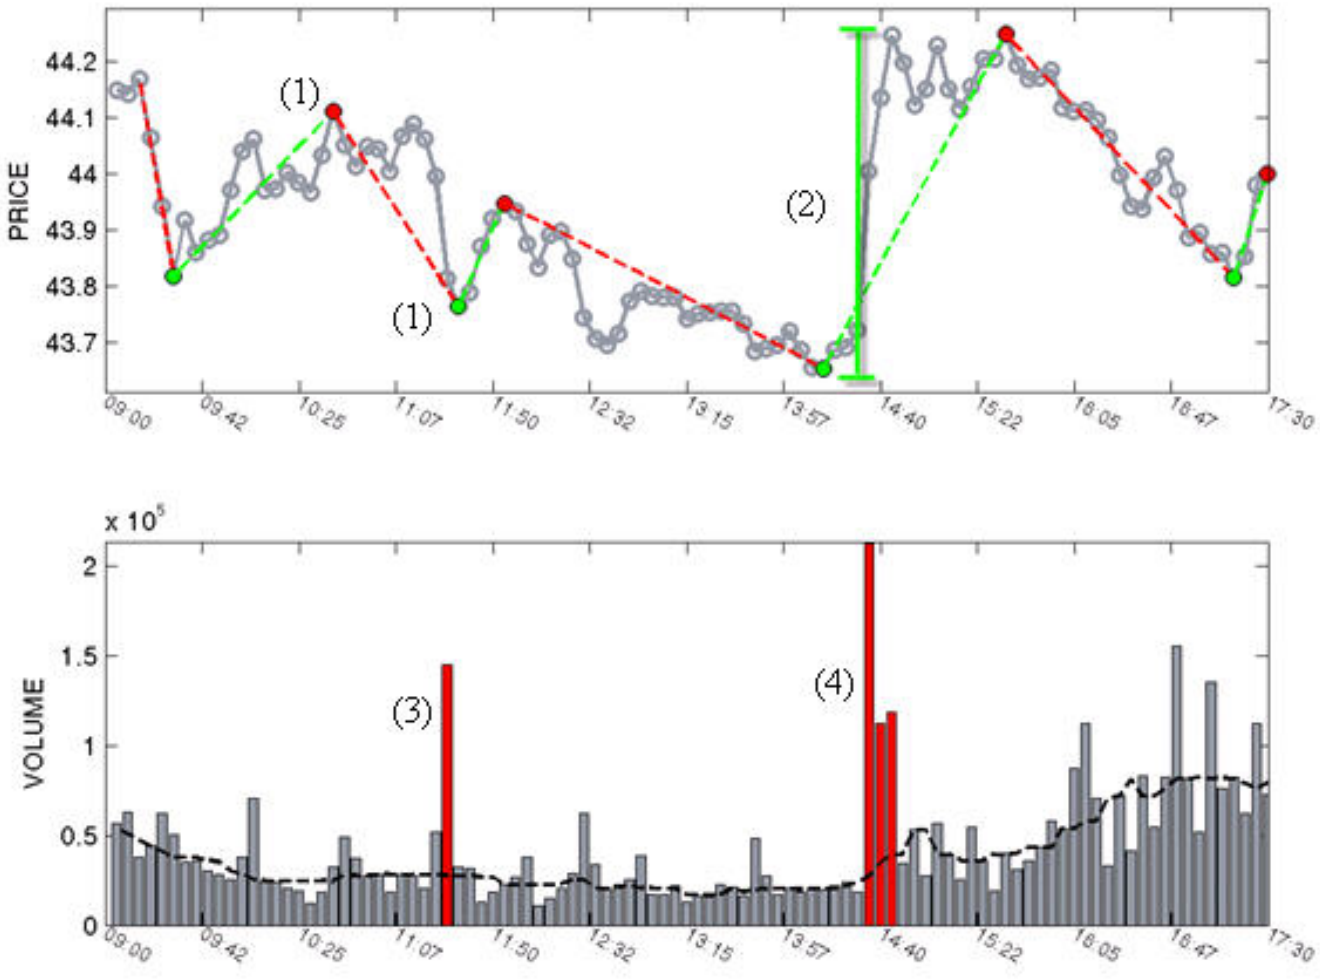

FIGURE 42 – The three abnormal patterns targeted by our three anomaly detectors : (1) price trends, (2) price jump, (3) volume peak, (4) volume crenel.

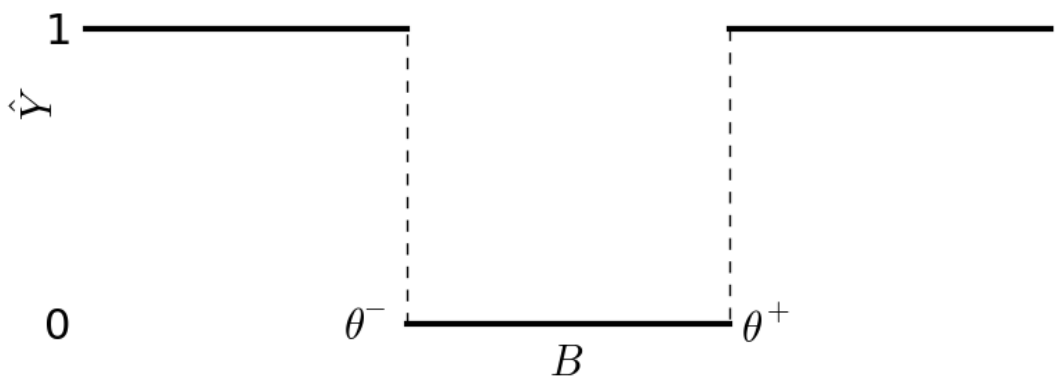

FIGURE 43 – Typical predictor defined by two intervals $(-\infty, \theta^-)$ and $(\theta^+, +\infty)$.

The predictive power $\pi(\theta)$ of h_θ is then readily estimated by the explicit formula

$$\hat{\pi}(\theta) = Q_r(\hat{P}^1, \hat{P}^0).$$

The influence coefficient $\mathcal{I}(Z, Y)$ of Z on Y at time t is then estimated by maximizing $\hat{\pi}(\theta)$ over all θ in $\mathbb{R}^2$. At time t, the set S of currently observed values of Z has cardinal inferior or equal to K_t. The previous formulas show that to maximize $\hat{\pi}(\theta)$, we may in fact restrict both θ^- and θ^+ to belong to S, so one needs only explore at most $K_t^2/2$ values

of θ.

Clearly this computation tends to underestimate the influence $\mathcal{I}(Z, Y)$. Nevertheless in our benchmark studies we have systematically applied this approach for the following reasons.

- the set of binary predictors h_θ has the merit of having finite Vapnik-Cervonenkis dimension equal to 2 , so that our empirical estimate of the maximum of $\pi(\theta)$ will be statistically robust even for moderate realistic values of $K_t \equiv 700$;
- the immediate interpretability of the predictors h_θ enables user friendly online graphic displays of the explanatory factors currently having high influence on performance degradation.
- at each time t, at most K_t^3 basic operations suffice to implement the brute force maximization of $\pi(\theta)$ which generates our current evaluation of $\mathcal{I}(Z, Y)$.

Note that the two optimal thresholds (θ^-, θ^+) will of course strongly depend on the time slice t.

When the explanatory factor Z is one of the 21 *smoothed anomaly detectors* $[[A_t]] \geq 0$ introduced above in section 3.3.2, the preceding implementation can be simplified. Recall that $[[A_t]]$ records the maximal gravity of very recent anomalies of fixed type affecting the dynamics of a fixed market variable, and that for "most" time slices, $[[A_t]]$ takes the values 0. Thus it is natural to expect that only higher values of $[[A_t]]$ to be potential explanations for currently degraded trading performances. So for practical applications to $Z = [[A_t]]$ of the preceding approach, we may actually impose the constraint $\theta^- = 0$, with essentially no loss of predictive power.

A few examples for single explanatory factors. The empirical strategy just presented to estimate the influence $J(G)$ when cardinal(G) equals 1 has been numerically validated on our benchmark set of intraday data. We now outline a few examples. Recall that our benchmark study used the predictive power functional $\pi = Q_r(P^1, P^0)$, which is specifically sensitive to predictors capacity to detect *degradations* of trading performances. On our intraday data sets, we have methodically tested the values $r = 70\%, 75\%, 80\%, 85\%, 90\%, 95\%$ for the "floor predictive power" r ; the value $r = 85\%$ turned out to be the best choice for these data sets, and was adopted for all results presented below.

Figure 44 illustrates for the fixed time slice $t = 45$, the predictive power of the predictors h_θ based on the single market variable $Z =$ "Momentum in Bid-Ask Spread" .

The trading performance evaluator PE is the "slippage in bid-ask spread". The PE-thresholds $\ell_t = \ell_{45}$ determining low trading performance is fixed at the 3% -quantile of all performance evaluations observed at time $t = 45$.
The x and y axes in the graph (figure 44) indicate the threshold values (θ^-, θ^+) for the market variable $Z =$ "Momentum in Bid-Ask Spread". The z axis displays the predictive power $\pi(\theta)$ of h_θ. The red marker indicates at time $t = 45$, the estimated influence coefficient $\mathcal{I}(Z, Y)$ on Y for this specific market variable Z, which turns out to be equal to 100%. The threshold vector $\theta = \theta_{45}$ which achieves maximum predictive power at time 45 is equal to (66.76, 3.87). At each time t, our 7 basic market variables can then be ranked on the basis of their approximate influence values computed as above, which provides a ranking of their respective capacity to explain current bad trading performances.

Influence computation for pairs of explanatory factors. Again at fixed time t, we now sketch our online "optimized fusion" of predictors to estimate the influence coefficient $J(G)$ when G is a group of 2 explanatory factors $Z = [\ Z^1, Z^2\ ]$. Our statistical robustness analysis above indicates the necessity to consider only classes of trading performance predictors having radically low Vapnik-Cervonenkis dimension. So our predictive power maximization among predictors based on Z is deliberately restricted to the following class of predictors.

Let $\mathfrak{P}_2$ be the set of all 16 functions mapping $\{0; 1\}^2$ into $\{0; 1\}$. The class $\mathcal{H}$ will be the set of all predictors of the form

$$h(Z) = \mathfrak{m}(f(Z_1), g(Z_2))$$

where $\mathfrak{m} \in \mathfrak{P}_2$, and the indicator functions $f(z)$ and $g(z)$ are both two-sided binary predictors in the sense of definition 2. The class of binary predictors $\mathcal{H}$ has Vapnik-Cervonenkis dimension equal to 4. Hence the estimation of maximal predictive power within $\mathcal{H}$ by empirical estimation of probabilities P^1 and P^0 on the basis of the current K_t joint observations of Y, Z_1, Z_2 will be statistically robust. This provides at time t a stable estimator of $J(G)$, which as above tends to undervalue the true $J(G)$.

In concrete implementation of this approach at fixed time t, we first select only pairs of predictors $f(Z_1), g(Z_2)$ which already have reasonably high probabilities of correctly predicting Y. To maximize the predictive power of $h(Z) = \mathfrak{m}(f(Z_1), g(Z_2))$ we need to select the best binary polynomial $\mathfrak{m}$ among the 16 elements of $\mathfrak{P}_2$. We then impose $\mathfrak{m}(0, 0) = 0$ and $\mathfrak{m}(1, 1) = 1$, so that whenever the predictions of $f(Z_1)$, and $g(Z_2)$ agree, we also have $h(Z) = f(Z_1) = g(Z_2)$. This "accelerated fusion" is fairly classical in multi-experts fusion (see [Alkoot & Kittler, 1999]) and obviously provides an acceleration multiplier of 4 in the online computation of $J(G)$.

For groups G of $k = 3$ or $k = 4$ explanatory factors, one could estimate $J(G)$ by similar sub-optimal but implementable strategies. However the corresponding predictor classes have Vapnik-Cervonenkis dimensions 6 and 8, and their statistical robustness is hence much weaker in the concrete context of our intraday datas set, since at each time t the key Vapnik ratios $K_t/6$ and $K_t/8$ were resp. inferior to 120 and 70, values which are much too small and strongly suggested to avoid the estimation of $J(G)$ for cardinal(G) ≥ 3.

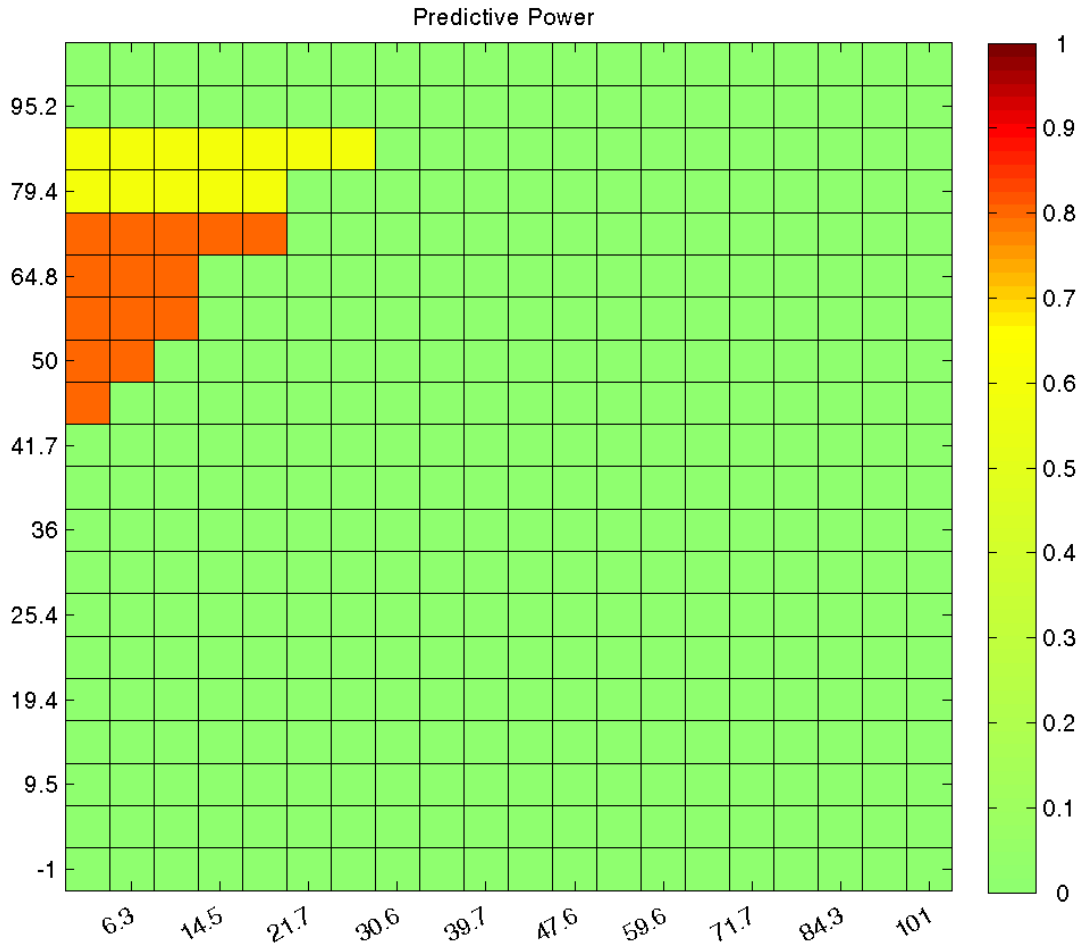

FIGURE 44 – Predictive power as a functional of the two thresholds θ^+ (y-axis) and θ^- (x-axis) of the market variable *Volume Score* at time slice $t = 72$. It can be seen that θ^- lower than 10 and θ^+ around 65 generate an efficient predictor of bad trading performance during this time slice.

3.3.4 Numerical results

We now present the numerical results obtained by applying the above methodology to our benchmark dataset of intra-day trading records.

Dataset (portfolio) description. Recall that our intra-day benchmark data involve a total of 79 time slices of 5 minutes each (i.e. this portfolio has been traded from 8 :55 to 15 :30 London time), and that we are monitoring a portfolio of 1037 trading orders, with a maximum of 700 trading orders active simultaneously at each time slice.

At each fixed time slice t, we compute the current influence coefficient for each one of our 28 explanatory factors, namely the 7 market descriptors M^j themselves and the 3×7 smoothed anomaly detectors monitoring the dynamic of these market descriptors.

These 28 explanatory factors generate $378 = 28 \times 27/2$ pairs of factors. By accelerated fusion as above, we compute, at each time slice t, the influence of each one of these 378 pairs of explanatory factors.

Among these $406 = (378 + 28)$ groups of explanatory factors, at each time t, we retain only those having both conditional probabilities of correct predictions (p^1, p^0) larger than $r\%$. Here $r\% > 70\%$ is the user selected "floor predictive power". Note that each single factor or pair of factors retained at time t can predict current degraded performance degradations with a false alarm rate $FAR = 1 - p^0$ inferior to $(100 - r)\%$.

Among the retained groups of explanatory factors, we compute the maximal influence $\max J_t$ achievable at time t. We also determine the set $\mathcal{D}_t$ of dominating groups of explanatory factors, defined as the groups of 1 or 2 factors having an influence equal to $\mathfrak{J}_t$ and achieving the minimal false alarm rate $(100 - r)\%$.

Predictive power of market descriptors. Figure 45 gives an heatmap of the predictive power of a selected subset of market descriptors on the whole portfolio. The display shows that :

– no market descriptor is used before slice 60 (i.e. 14 :00), meaning that there are no significant predictive links between bad trading performances and specific values of the descriptors.

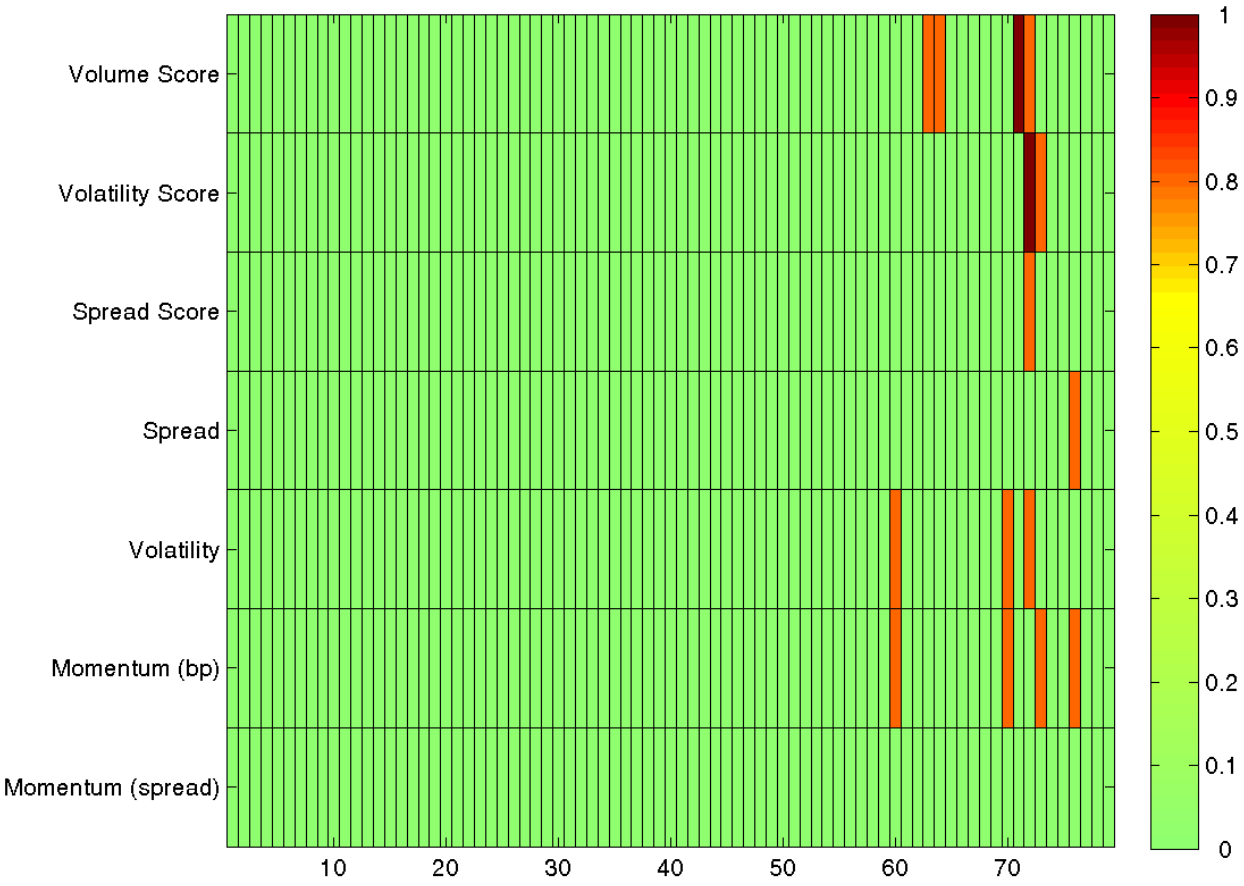

FIGURE 45 – Predictive powers of some explanatory variables (horizontal scale is time in slices of 5 minutes).

- Then the *Volume Score* has the capability to explain bad trading performance from 14 :10 to 14 :20 and from 14 :50 to 15 :00. It means that during these two time intervals, bad performances occurred simultaneously with quite unusual levels of traded volumes.
- The *Volatility Score* emerges as a complementary explanatory factor between 14 :55 and 15 :05 ; orders with bad trading performances focussed on stocks having unexpectedly high volatility levels during these 10 minutes.
- The *Bid-Ask Spread Score* conforts this automated diagnosis : a rare event did indeed degrade trading performances around 15 :00. Keeping in mind that scores are computed according to historical values during last weeks, it means that for this portfolio, the worst performances occurred on stocks for which volumes, volatility, and bid-ask spread had abnormal values.

It is interesting to note that, before scoring, market performances do not explain that well bad performances.

How alarm zones explain bad trading performance. Since the two-sided binary predictors are built to explain the degraded performances of the worst trading orders, it is easy to identify the most impacted orders one given market descriptor. Here we consider the trading order $\mathbf{T}(139)$ whose lifecycle is shown on Figure 41, in order to visualize the impact of its *Volume Score*, *Volatility Score*, and *Bid-Ask Score* on the order performance. First note that this order has been active from 10 :55 to 15 :30 London time, with a start time two hours after the launch of the portfolio. It means that 5 minutes slices on this order are numbered from 0 to 49 ; they have to be shifted by 30 to be synchronized with the time scale of the other portfolio orders.

Alarm zones on the Volume Score. Figure 46 shows alarm zones for the *Volume Score* of order $\mathbf{T}(139)$: the top subplot draws the value of the volume score through time and the associated alarm zones have been added on top of it, when triggered.

For the whole portfolio (Figure 45), alarm zones are triggered on the *Volume Score* from 14 :10 to 14 :20 and from 14 :50 to 15 :00 (i.e. slices 32 to 32 and 40 to 41 for this specific order). They are drawn like "gates" from the low threshold (θ_t^-) to the high one (θ_t^+) ; if the value of the Volume Score is outside bounds for the given order : we thus state that "*the Volume Score contributed to the degradation of trading performance for order* $\mathbf{T}(139)$".

It is important to note that if if our thresholds had not been adaptive as proposed by our "*influence analysis*"

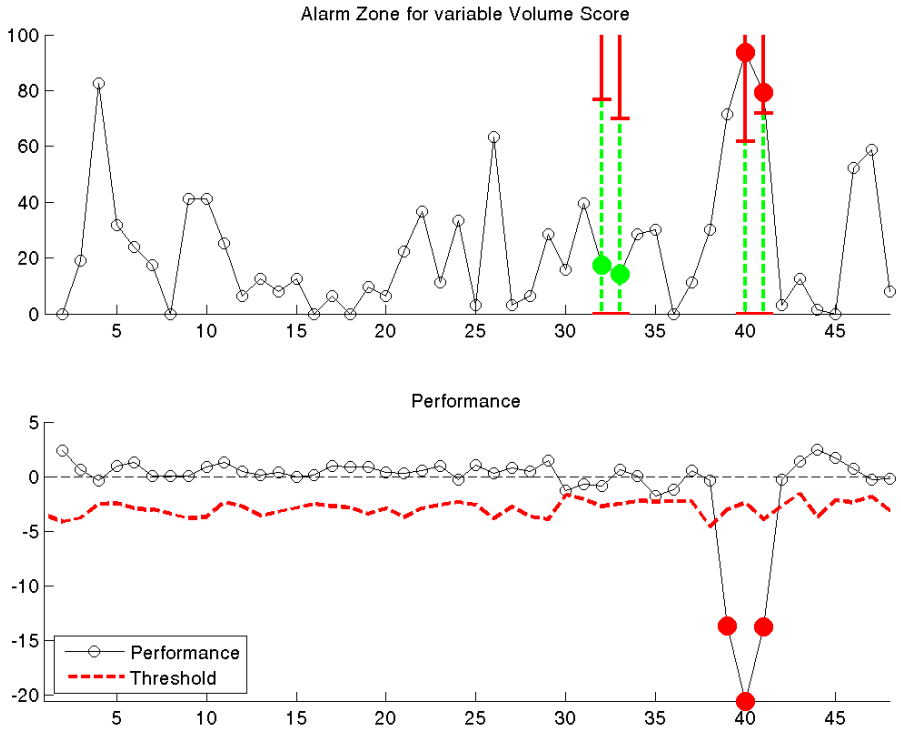

FIGURE 46 – Auto adaptive alarm zones on the *Volume Score* explanatory variable for the order $T(139)$ displayed in Figure 41 ; Top : four alarm zones are active, two realizations of the Volume Score exceed the auto adaptive thresholds and thus emerge as a highly likely explanation for the bad performance exhibited by this order (see Bottom graph.).

methodology, they would have generated false explanations at the start of the order (around slice 4, i.e. 11 :15 London time).

Specifically for this order, the Volume Score does not enter the first alarm zone (around 14.15) ; it is in line with the performance of order $\mathbf{T}(139)$ that is normal (bottom subplot of Figure 46). The boundaries of the next alarm zone (around 15 :00) are crossed by the order Volume Score ; indeed this order performs quite poorly at that time.

Alarm zones on the Bid-Ask Spread Score. Only one alarm zone has been activated on the portfolio (around 15 :00) and the advantage of the auto adaptive approach proposed in this paper is straightforward : a unique threshold for all the duration of the portfolio would clearly not have been able to separate the 41st slice of order $\mathbf{T}(139)$ from the others.

Alarm zones on the Volatility Score. Once again it is clear that the alarm zones succeeded in isolating slices to efficiently explain the bad trading performance around 15 :00.

To summarize. When applied applied to our real portfolio of 1037 orders traded during 6 hours and 35 minutes, the *automated influence analysis methodology* presented and studied from a theoretical viewpoint in this paper efficiently selects quite pertinent explanatory factors for degraded trading performance :

- our alarm zones use thresholds that are automatically adapted online to successive time slices, as computed via the *predicting power* of *two-sided binary predictors* (see Definition 2) based on *market descriptors*.
- Our approach generates generates auto adaptive thresholds taking into account currently observed synchronicity between the user selected performance criterion (chosen according to the trading goal)with market descriptors *market descriptors*.
- At each time slice, the computed adaptive thresholds on market descriptors apply to the whole portfolio, bad trading performance of orders for which market descriptors take values outside *alarm zones* are said to be explained or *influenced* by the given descriptors. We commented real examples to illustrate these automated selection of explanatory factors.
- The added value of augmenting the state space of market descriptors using *scores* has been illustrated on several examples within our benchmark data set o trading orders..

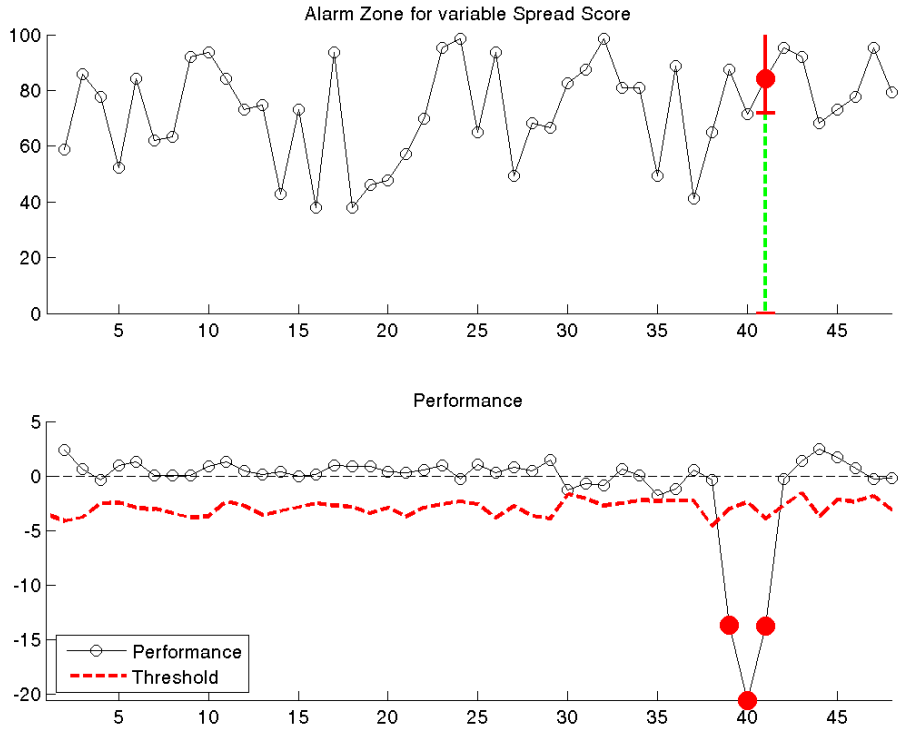

FIGURE 47 – Auto adaptive alarm zones on the *Bid-Ask spread Score* explanatory variable for the order of Figure 41 ; Top : one alarm zone is active, since the Bid-Ask Score exceeds the auto adaptive thresholds, and thus emerges as a quite likely explanation of bad trading performance (as validated on the bottom figure).

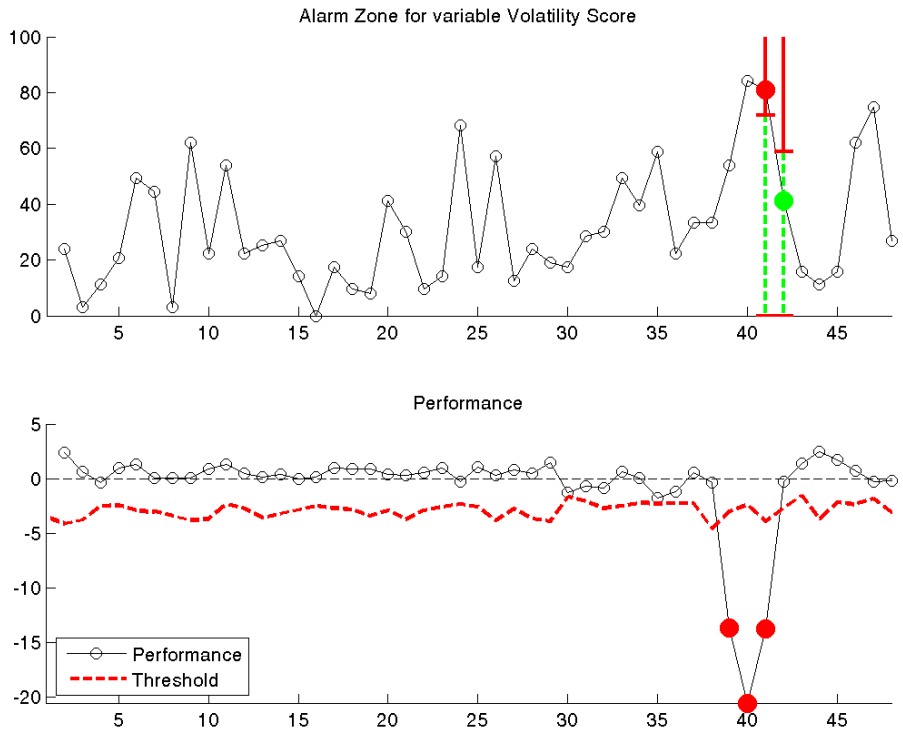

FIGURE 48 – Auto adaptive alarm zones for the *Volatility Score* explanatory variable for the order of Figure 41 ; Top : two alarm zones are active, one realization of the Volatility Score exceeds the auto adaptive thresholds giving and thus provides a highly likely explanation of the current bad trading performance (as displayed on the bottom graph).

Conclusion

Le trading optimal est désormais un champ à part entière des mathématiques financières :
- initié dans un cadre moyenne-variance, nous avons montré comment même dans ce cadre des explorations théoriques peuvent apporter des réponses aux ingénieurs financiers en charge des algorithmes de trading,
- ses prolongations à des critères plus réalistes fait apparaître des problématiques sophistiquées du contrôle stochastique, dont l'exploration et parfois la résolution éclaire la pratique des différents acteurs (teneurs de marché, banques d'investissement, investisseurs finaux).

Tout comme le reste des mathématiques financières, il apporte un point de vue "contrôle du risque" aux praticiens qui sans cela mettraient en production des algorithmes de trading guidés par d'éventuels paris sur les prix futurs, encadrés par des bornes approximatives et mal maîtrisées.

En permettant de poser clairement un critère mélangeant gains ou pertes attendus et risques courus, le trading optimal fait émerger des algorithmiques qui conjuguent au mieux les deux aspects. Il permet de bien comprendre comment des paramètres comme l'aversion au risque modifient le comportement d'un algorithme de trading.

L'étude de la microstructure des marchés est restée longtemps l'apanage des économistes qui utilisent des modèles grossiers pour obtenir des résultats d'équilibre généraux. Avec la disponibilité de larges bases de données précises, nous avons montré comment la dynamique des carnets d'ordres peut désormais être modélisée finement et de façon réaliste par des modèles mathématiques complexes, comme des files d'attente couplées, ou les processus de Hawkes.

Il est même possible d'envisager l'étude d'équilibres généraux, grâce au cadre des jeux à champ moyen, aboutissant à des systèmes d'EDP couplées.

La microstructure est finalement le lieu qui permet de confronter données empiriques et des mélanges de stratégies optimales tentant de rendre compte de comportements stylisés. Cette confrontation donne lieu à une meilleure compréhension de la relation entre les décisions d'investissement, le processus de trading, et les prix. Portée à l'attention des régulateurs ces travaux devraient permettre de mieux réguler la formation des prix.

Par ailleurs, les algorithmes stochastiques et l'apprentissage statistique permettent d'adapter le processus d'achat ou de vente aux variations fugaces de liquidité et de prix. Ils permettent aussi de mettre en place d'aide à la décision dans le cadre de la surveillance de la formation des prix.

Références

[Agarwal *et al.* , 2010] Agarwal, Alekh, Bartlett, Peter L., & Dama, Max. 2010 (May). Optimal Allocation Strategies for the Dark Pool Problem. *Pages 9–16 of :* Teh, Y. W., & Titterington, M. (eds), *Proceedings of The Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 9.

[Alfonsi & Blanc, 2014] Alfonsi, Aurélien, & Blanc, Pierre. 2014 (Apr.). *Dynamic optimal execution in a mixed-market-impact Hawkes price model.*

[Alfonsi & Schied, 2010] Alfonsi, Aurélien, & Schied, Alexander. 2010. Optimal Execution and Absence of Price Manipulations in Limit Order Book Models. *SIAM J. Finan. Math.*, **1**(Nov.), 490–522.

[Alkoot & Kittler, 1999] Alkoot, F. M., & Kittler, J. 1999. Experimental evaluation of expert fusion strategies. *Pattern Recognition Letters*, **20**(11), 1361–1369.

[Almgren & Chriss, 2000] Almgren, R. F., & Chriss, N. 2000. Optimal execution of portfolio transactions. *Journal of Risk*, **3**(2), 5–39.

[Almgren *et al.* , 2005] Almgren, Robert, Thum, Chee, Hauptmann, Emmanuel, & Li, Hong. 2005. Direct Estimation of Equity Market Impact. *Risk*, **18**(July), 57–62.

[Almgren, 2003] Almgren, Robert F. 2003. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, **10**(1), 1–18.

[Avellaneda & Stoikov, 2008] Avellaneda, Marco, & Stoikov, Sasha. 2008. High-frequency trading in a limit order book. *Quantitative Finance*, **8**(3), 217–224.

[Aviv, 1991] Aviv, David. 1991 (June). *Detection of Abrupt Changes in Statistical Models.* Ph.D. thesis.

[Azencott *et al.* , 2014] Azencott, Robert, Beri, Arjun, Gadhyan, Yutheeka, Joseph, Nicolas, Lehalle, Charles-Albert, & Rowley, Matthew. 2014. Realtime market microstructure analysis : online Transaction Cost Analysis. *Quantitative Finance*, Mar., 0–19.

[Bacry & Muzy, 2013] Bacry, E., & Muzy, J. F. 2013 (Jan.). *Hawkes model for price and trades high-frequency dynamics.*

[Bacry *et al.* , 2001] Bacry, E., Delour, J., & Muzy, J. F. 2001. Multifractal random walk. *Physical Review E*, **64**(2), 026103+.

[Bacry *et al.* , 2011] Bacry, E., Dayri, K., & Muzy, J. F. 2011 (Dec.). *Non-parametric kernel estimation for symmetric Hawkes processes. Application to high frequency financial data.*

[Bacry *et al.* , 2012] Bacry, E., Delattre, S., Hoffmann, M., & Muzy, J. F. 2012. Modeling microstructure noise with mutually exciting point processes. *Quantitative Finance (forthcoming)*, Jan.

[Bacry *et al.* , 2014] Bacry, Emmanuel, Iuga, Adrian, Lasnier, Matthieu, & Lehalle, Charles-Albert. 2014. Market Impacts and the Life Cycle of Investors Orders. *Social Science Research Network Working Paper Series*, Dec.

[Bacry *et al.* , 2015] Bacry, Emmanuel, Mastromatteo, Iacopo, & Muzy, Jean-François. 2015 (Feb.). *Hawkes processes in finance.*

[Basseville, 1988] Basseville, Michèle. 1988. Detecting changes in signals and systems - A survey. *Automatica*, **24**(3), 309–326.

[Basseville & Nikiforov, 1993] Basseville, Michele, & Nikiforov, Igor V. 1993. *Detection of Abrupt Changes : Theory and Application (Prentice Hall information and system sciences series).* Prentice Hall.

[Bayraktar & Ludkovski, 2012] Bayraktar, Erhan, & Ludkovski, Michael. 2012. Liquidation in Limit Order Books with Controlled Intensity. *Mathematical Finance*, June.

[Benveniste *et al.* , 1991] Benveniste, Albert, Metivier, Michel, & Priouret, Pierre. 1991. *Adaptive Algorithms and Stochastic Approximations.* Springer-Verlag.

[Bershova & Rakhlin, 2013] Bershova, Nataliya, & Rakhlin, Dmitry. 2013. The Non-Linear Market Impact of Large Trades : Evidence from Buy-Side Order Flow. *Social Science Research Network Working Paper Series*, Jan.

[Bertsimas & Lo, 1998] Bertsimas, Dimitris, & Lo, Andrew W. 1998. Optimal control of execution costs. *Journal of Financial Markets*, **1**(1), 1–50.

[Beskos & Roberts, 2005] Beskos, Alexandros, & Roberts, Gareth O. 2005. Exact simulation of diffusions. *The Annals of Applied Probability*, **15**(4), 2422–2444.

[Borovkov, 1998] Borovkov, A. A. 1998. *Probability Theory.* Gordon and Breach Science Publishers.

[Bouchard & Touzi, 2011] Bouchard, Bruno, & Touzi, Nizar. 2011. Weak dynamic programming principle for viscosity solutions. *SIAM Journal on Control and Optimization*, **49**(3), 948–962.

[Bouchard *et al.* , 2011] Bouchard, Bruno, Dang, Ngoc-Minh, & Lehalle, Charles-Albert. 2011. Optimal control of trading algorithms : a general impulse control approach. *SIAM J. Financial Mathematics*, **2**(1), 404–438.

[Bouchaud, 2010] Bouchaud, Jean-Philippe. 2010. Price Impact.

[Bouchaud & Potters, 2004] Bouchaud, Jean-Philippe, & Potters, Marc. 2004. *Theory of Financial Risk and Derivative Pricing : From Statistical Physics to Risk Management.* Cambridge University Press.

[Bouchaud *et al.* , 2004] Bouchaud, Jean-Philippe, Gefen, Yuval, Potters, Marc, & Wyart, Matthieu. 2004. Fluctuations and response in financial markets : the subtle nature of 'random' price changes. *Quantitative Finance*, **4**(2), 176–190.

[Brokmann *et al.* , 2014] Brokmann, X., Serie, E., Kockelkoren, J., & Bouchaud, J. P. 2014 (July). *Slow decay of impact in equity markets*.

[Cartea *et al.* , 2011] Cartea, Álvaro, Jaimungal, Sebastian, & Ricci, Jason. 2011. Buy Low Sell High : A High Frequency Trading Perspective. *Social Science Research Network Working Paper Series*, Nov.

[Chakraborti *et al.* , 2011] Chakraborti, Anirban, Toke, Ioane M., Patriarca, Marco, & Abergel, Frederic. 2011 (June). *Econophysics review : II. Agent-based models*.

[Cont & De Larrard, 2013] Cont, Rama, & De Larrard, Adrien. 2013. Price Dynamics in a Markovian Limit Order Book Market. *SIAM Journal for Financial Mathematics*, **4**(1), 1–25.

[Cont *et al.* , 2008] Cont, Rama, Stoikov, Sasha, & Talreja, Rishi. 2008. A Stochastic Model for Order Book Dynamics. *Social Science Research Network Working Paper Series*, Sept.

[Cristianini & Shawe-Taylor, 2000] Cristianini, Nello, & Shawe-Taylor, John. 2000. *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. 1 edn. Cambridge University Press.

[Da Fonseca & Zaatour, 2014] Da Fonseca, José, & Zaatour, Riadh. 2014. Clustering and Mean Reversion in a Hawkes Microstructure Model. *Journal of Futures Markets*, Apr., n/a.

[Daley & Vere-Jones, 2003] Daley, D. J., & Vere-Jones, D. 2003. *An introduction to the theory of point processes Vol. I, Elementary theory and methods.* 2nd edn. Springer.

[Donier & Bonart, 2014] Donier, Jonathan, & Bonart, Julius F. 2014. A Million Metaorder Analysis of Market Impact on the Bitcoin. *Social Science Research Network Working Paper Series*, Dec.

[Duflo, 1997] Duflo, Marie. 1997. *Algorithmes stochastiques*. Springer.

[Eisler *et al.* , 2011] Eisler, Zoltan, Bouchaud, Jean-Philippe, & Kockelkoren, Julien. 2011. Models for the impact of all order book events. July.

[El Aoud & Abergel, 2014] El Aoud, Sofiene, & Abergel, Frederic. 2014. A Stochastic Control Approach for Option Market Making. *Social Science Research Network Working Paper Series*, Sept.

[Embrechts, 2002] Embrechts, Paul. 2002. *Selfsimilar Processes (Princeton Series in Applied Mathematics)*. Princeton University Press.

[Engle *et al.* , 2012] Engle, Robert F., Ferstenberg, Robert, & Russell, Jeffrey R. 2012. Measuring and Modeling Execution Cost and Risk. *The Journal of Portfolio Management*, **38**(2), 14–28.

[Euronext, 2006] Euronext. 2006 (Apr.). *EuroNext rules - Book I.* Tech. rept. EuroNext.

[Fama & French, 2003] Fama, Eugene F., & French, Kenneth R. 2003. The Capital Asset Pricing Model : Theory and Evidence. *Social Science Research Network Working Paper Series*, Sept.

[Farmer *et al.* , 2004] Farmer, Doyne J., Gillemot, Laszlo, Lillo, Fabrizio, Mike, Szabolcs, & Sen, Anindya. 2004. What really causes large price changes ? *Quantitative Finance*, **4**(4), 383–397.

[Farmer *et al.* , 2013] Farmer, J. Doyne, Gerig, Austin, Lillo, Fabrizio, & Waelbroeck, Henri. 2013. How efficiency shapes market impact. *Quantitative Finance*, **13**(11), 1743–1758.

[Field & Large, 2008] Field, Jonathan, & Large, Jeremy. 2008. *Pro-rata matching and one-tick futures markets*. CFS working paper 2008,40. Frankfurt, Main.

[Fodra & Pham, 2013] Fodra, Pietro, & Pham, Huyên. 2013 (May). *Semi Markov model for market microstructure*.

[Foucault & Menkveld, 2008] Foucault, Thierry, & Menkveld, Albert J. 2008. Competition for Order Flow and Smart Order Routing Systems. *The Journal of Finance*, **63**(1), 119–158.

[Ganchev *et al.* , 2010] Ganchev, Kuzman, Nevmyvaka, Yuriy, Kearns, Michael, & Vaughan, Jennifer W. 2010. Censored exploration and the dark pool problem. *Commun. ACM*, **53**(5), 99–107.

[Gareche *et al.* , 2013] Gareche, A., Disdier, G., Kockelkoren, J., & Bouchaud, Jean-Philippe. 2013. A Fokker-Planck description for the queue dynamics of large tick stocks. *Phys. Rev. E*, **88**(3), 32809–32820.

[Gatheral, 2010a] Gatheral, Jim. 2010a. No-Dynamic-Arbitrage and Market Impact. *Quantitative Finance*, **10**(7).

[Gatheral, 2010b] Gatheral, Jim. 2010b. No-dynamic-arbitrage and market impact. *Quantitative Finance*, **10**(7), 749–759.

[Gatheral & Schied, 2012] Gatheral, Jim, & Schied, Alexander. 2012. Dynamical models of market impact and algorithms for order execution. *In :* Fouque, Jean-Pierre, & Langsam, Joseph (eds), *Handbook on Systemic Risk (Forthcoming)*. Cambridge University Press.

[Gatheral *et al.* , 2012] Gatheral, Jim, Schied, Alexander, & Slynko, Alla. 2012. Transient Linear Price Impact and Fredholm Integral Equations. *Mathematical Finance*, **22**, 445–474.

[Glosten & Milgrom, 1985] Glosten, Lawrence R., & Milgrom, Paul R. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, **14**(1), 71–100.

[Gökay *et al.* , 2011] Gökay, Selim, Roch, AlexandreF, & Soner, Mete H. 2011. Liquidity Models in Continuous and Discrete Time. *Pages 333–365 of :* Di Nunno, Giulia, & Øksendal, Bernt (eds), *Advanced Mathematical Methods for Finance*. Springer Berlin Heidelberg.

[Gomes & Waelbroeck, 2013] Gomes, C., & Waelbroeck, H. 2013. *Is Market Impact a Measure of the Information Value of Trades ? Market Response to Liquidity vs Informed Trades*. Tech. rept.

[Graewe *et al.* , 2013] Graewe, Paulwin, Horst, Ulrich, & Séré, Eric. 2013 (Dec.). *Smooth solutions to portfolio liquidation problems under price-sensitive market impact*.

[Guéant *et al.* , 2012] Guéant, O., Lehalle, Charles-Albert, & Fernandez-Tapia, J. 2012. Optimal Portfolio Liquidation with Limit Orders. *SIAM Journal on Financial Mathematics*, **13**(1), 740–764.

[Guéant & Pu, 2013] Guéant, Olivier, & Pu, Jiang. 2013 (Nov.). *Option pricing and hedging with execution costs and market impact*.

[Guéant *et al.* , 2013] Guéant, Olivier, Lehalle, Charles-Albert, & Fernandez-Tapia, Joaquin. 2013. Dealing with the inventory risk : a solution to the market making problem. *Mathematics and Financial Economics*, **4**(7), 477–507.

[Guilbaud & Pham, 2011] Guilbaud, Fabien, & Pham, Huyen. 2011 (June). *Optimal High Frequency Trading with limit and market orders*.

[Gustafsson, 2000] Gustafsson, Fredrik. 2000. *Adaptive Filtering and Change Detection*. 1 edn. Wiley.

[Hardiman *et al.* , 2013] Hardiman, Stephen J., Bercot, Nicolas, & Bouchaud, Jean-Philippe. 2013 (June). *Critical reflexivity in financial markets : a Hawkes process analysis*.

[Hautsch & Huang, 2012] Hautsch, Nikolaus, & Huang, Ruihong. 2012. The market impact of a limit order. *Journal of Economic Dynamics and Control*, **36**(4), 501–522.

[Hewlett, 2006] Hewlett, P. 2006. Clustering of order arrivals, price impact and trade path optimisation. *In : Workshop on Financial Modeling with Jump processes*. Ecole Polytechnique.

[Ho & Stoll, 1981] Ho, Thomas, & Stoll, Hans R. 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, **9**(1), 47–73.

[Ho & Stoll, 1983] Ho, Tsy, & Stoll, H. R. 1983. The dynamics of dealer markets under competition. *Journal of Finance*, **38**(4), 1053–1074.

[Hoffmann *et al.* , 2013] Hoffmann, Marc, Labadie, Mauricio, Lehalle, Charles-Albert, Pages, Gilles, Pham, Huyen, & Rosenbaum, Mathieu. 2013. Optimization and Statistical Methods for High Frequency Finance. *In : ESAIM Proceedings*. SMAI, Seignosse, France.

[Huang *et al.* , 2013] Huang, Weibing, Lehalle, Charles-Albert, & Rosenbaum, Mathieu. 2013 (Dec.). *Simulating and analyzing order book data : The queue-reactive model*.

[Jaisson, 2014] Jaisson, Thibault. 2014 (Feb.). *Market impact as anticipation of the order flow imbalance*.

[Karatzas & Shreve, 1998] Karatzas, I., & Shreve, S. E. 1998. *Methods of Mathematical Finance*. New York : Springer.

[Kharroubi & Pham, 2010] Kharroubi, Idris, & Pham, Huyen. 2010. Optimal portfolio liquidation with execution cost and risk. *SIAM Journal on Financial Mathematics*, **1**(1), 897–931.

[Kharroubi *et al.* , 2013] Kharroubi, Idris, Langrené, Nicolas, & Pham, Huyên. 2013 (Nov.). *A numerical algorithm for fully nonlinear HJB equations : an approach by control randomization*.

[Kushner & Clark, 1980] Kushner, H. J., & Clark, D. S. 1980. Stochastic Approximation Methods for Constrained and Unconstrained Systems. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift fur Angewandte Mathematik und Mechanik*, **60**(1), 63–64.

[Kushner & Yin, 2003] Kushner, Harold J., & Yin, George G. 2003. *Stochastic Approximation and Recursive Algorithms and Applications (Stochastic Modelling and Applied Probability) (v. 35)*. 2nd edn. Springer.

[Kyle, 1985] Kyle, Albert P. 1985. Continuous Auctions and Insider Trading. *Econometrica*, **53**(6), 1315–1335.

[Labadie & Lehalle, 2014] Labadie, Mauricio, & Lehalle, Charles-Albert. 2014. Optimal starting times, stopping times and risk measures for algorithmic trading. *The Journal of Investment Strategies*, **3**(2).

[Lachapelle *et al.* , 2013] Lachapelle, Aimé, Lasry, Jean-Michel, Lehalle, Charles-Albert, & Lions, Pierre-Louis. 2013 (May). *Efficiency of the Price Formation Process in Presence of High Frequency Participants : a Mean Field Game analysis*.

[Lamberton & Pagès, 2008] Lamberton, Damien, & Pagès, Gilles. 2008. A penalized bandit algorithm. *Electronic Journal of Probability*, **13**(0).

[Laruelle & Pagès, 2012] Laruelle, Sophie, & Pagès, Gilles. 2012. Stochastic approximation with averaging innovation applied to Finance. *Monte Carlo Methods and Applications*, **18**(1), 1–51.

[Laruelle *et al.* , 2013] Laruelle, Sophie, Lehalle, Charles-Albert, & Pagès, Gilles. 2013. Optimal posting price of limit orders : learning by trading. *Mathematics and Financial Economics*, **7**(3), 359–403.

[Lehalle, 2008] Lehalle, Charle-Albert. 2008. Rigorous optimisation of intra day trading. *Wilmott Magazine*, Nov.

[Lehalle, 2009] Lehalle, Charles-Albert. 2009. Rigorous Strategic Trading : Balanced Portfolio and Mean-Reversion. *The Journal of Trading*, **4**(3), 40–46.

[Lehalle, 2013] Lehalle, Charles-Albert. 2013. Market Microstructure knowledge needed to control an intra-day trading process. *In :* Fouque, Jean-Pierre, & Langsam, Joseph (eds), *Handbook on Systemic Risk*. Cambridge University Press.

[Lehalle & Azencott, 2004] Lehalle, Charles-Albert, & Azencott, Robert. 2004 (May). On the fly health monitoring of mechanical hazards from under sampled signals in formula one. FISITA World Automotive Congress.

[Lehalle & Dang, 2010] Lehalle, Charles-Albert, & Dang, Ngoc M. 2010. Rigorous post-trade market impact measurement and the price formation process. *Liquidity Guide*.

[Lehalle *et al.* , 2010] Lehalle, Charles-Albert, Guéant, Olivier, & Razafinimanana, Julien. 2010. High Frequency Simulations of an Order Book : a Two-Scales Approach. *In :* Abergel, F., Chakrabarti, B. K., Chakraborti, A., & Mitra, M. (eds), *Econophysics of Order-Driven Markets*. New Economic Windows. Springer.

[Lehalle *et al.* , 2012] Lehalle, Charles-Albert, Lasnier, Matthieu, Bessson, Paul, Harti, Hamza, Huang, Weibing, Joseph, Nicolas, & Massoulard, Lionel. 2012 (Aug.). *What does the saw-tooth pattern on US markets on 19 July 2012 tell us about the price formation process*. Tech. rept. Crédit Agricole Cheuvreux Quant Note.

[Lehalle *et al.* , 2013] Lehalle, Charles-Albert, Laruelle, Sophie, Burgot, Romain, Pelin, Stéphanie, & Lasnier, Matthieu. 2013. *Market Microstructure in Practice*. World Scientific publishing.

[Li & Almgren, 2014] Li, Tianhui M., & Almgren, Robert. 2014 (Oct.). *Option Hedging with Smooth Market Impact*. Tech. rept.

[Lillo *et al.* , 2003] Lillo, F., Farmer, Doyne J., & Mantegna, R. 2003. Econophysics - Master Curve for Price - Impact Function. *Nature*.

[Mandelbrot & Hudson, 2004] Mandelbrot, Benoit, & Hudson, Richard L. 2004. *The (Mis)behavior of Markets*. First printing edn. Basic Books.

[Mantegna & Stanley, 1994] Mantegna, Rosario N., & Stanley, H. Eugene. 1994. Stochastic Process with Ultraslow Convergence to a Gaussian : The Truncated Lévy Flight. *Physical Review Letters*, **73**(22), 2946–2949.

[Markowitz, 1952] Markowitz, Harry. 1952. Portfolio Selection. *The Journal of Finance*, **7**(1), 77–91.

[Mastromatteo *et al.* , 2013] Mastromatteo, Iacopo, Toth, Bence, & Bouchaud, Jean-Philippe. 2013 (Nov.). *Agent-based models for latent liquidity and concave price impact*.

[Mendelson & Amihud, 1991] Mendelson, Haim, & Amihud, Yakov. 1991. *How (Not) to Integrate the European Capital Markets*. Cambridge University Press.

[Moro *et al.* , 2009] Moro, Esteban, Vicente, Javier, Moyano, Luis G., Gerig, Austin, Farmer, Doyne J., Vaglica, Gabriella, Lillo, Fabrizio, & Mantegna, Rosario N. 2009. Market impact and trading profile of hidden orders in stock markets. *Physical Review E*, **80**(Dec.), 066102+.

[Muniesa, 2003] Muniesa, Fabian. 2003. *Des marchés comme algorithmes : sociologie de la cotation électronique à la Bourse de Paris*. Ph.D. thesis, Ecole Nationale Supérieure des Mines de Paris.

[Obizhaeva & Wang, 2005] Obizhaeva, Anna, & Wang, Jiang. 2005. Optimal Trading Strategy and Supply/Demand Dynamics. *Social Science Research Network Working Paper Series*, Feb.

[Pagès *et al.* , 2011] Pagès, Gilles, Laruelle, Sophie, & Lehalle, Charles-Albert. 2011. Optimal split of orders across liquidity pools : a stochastic algorithm approach. *SIAM Journal on Financial Mathematics*, **2**, 1042–1076.

[Robbins & Monro, 1951] Robbins, Herbert, & Monro, Sutton. 1951. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, **22**(3), 400–407.

[Robert & Rosenbaum, 2011] Robert, Christian Y., & Rosenbaum, Mathieu. 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, 2009] Roşu, Ioanid. 2009. A Dynamic Model of the Limit Order Book. *Review of Financial Studies*, **22**(11), 4601–4641.

[Rosenthal, 2009] Rosenthal, Dale W. 2009. Performance Metrics for Algorithmic Traders. *Social Science Research Network Working Paper Series*, July.

[Schied & Schöneborn, 2009] Schied, Alexander, & Schöneborn, Torsten. 2009. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, **13**(2), 181–204.

[Shannon, 1948] Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Technical Journal*, **27**(July), 379–423.

[Shiryaev, 1999] Shiryaev, Albert N. 1999. *Essentials of Stochastic Finance : Facts, Models, Theory*. 1st edn. World Scientific Publishing Company.

[Smith *et al.* , 2003] Smith, E., Farmer, Doyne J., Gillemot, L., & Krishnamurthy, S. 2003. Statistical Theory of the Continuous Double Auction. *Quantitative Finance*, **3**(6), 481–514.

[Stoikov & Saglam, 2009] Stoikov, Sacha, & Saglam, M. 2009. *Option market making under inventory risk.*

[Tauchen & Pitts, 1983] Tauchen, George E., & Pitts, Mark. 1983. The Price Variability-Volume Relationship on Speculative Markets. *Econometrica*, **51**(2), 485+.

[Vapnik, 2006] Vapnik, V. 2006. *Estimation of Dependences Based on Empirical Data (Information Science and Statistics).* Springer.

[Vapnik & Chervonenkis, 1971] Vapnik, V. N., & Chervonenkis, A. Y. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, **16**(2), 264–280.

[Vapnik, 2010] Vapnik, Vladimir. 2010. *The Nature of Statistical Learning Theory (Information Science and Statistics).* Softcover reprint of hardcover 2nd ed. 2000 edn. Springer.

[Vapnik, 1998] Vapnik, Vladimir N. 1998. *Statistical Learning Theory.* Wiley-Interscience.

[Waelbroeck & Gomes, 2013] Waelbroeck, Henri, & Gomes, Carla. 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*, July.

[Wang *et al.* , 2007] Wang, Hansheng, Li, Guodong, & Tsai, Chih-Ling. 2007. Regression Coefficient and Autoregressive Order Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **69**(1).

[Zarinelli *et al.* , 2014] Zarinelli, Elia, Treccani, Michele, Farmer, J. Doyne, & Lillo, Fabrizio. 2014 (Dec.). *Beyond the square root : Evidence for logarithmic dependence of market impact on size and participation rate.* lastlillo.

Modèles Mathématiques pour l'Etude et le Contrôle du Processus de Formation des Prix

—

Mathematical Models to Study and Control the Price Formation Process

Charles-Albert Lehalle

Université Pierre et Marie Curie

Habilitation à Diriger les Recherches, le 18 Décembre 2015

Outline

Motivations: The Context Changed

Following the financial crisis (2008):

- complex financial products have been replaced by more linear ones, the activity of investment banks moved from hedging a large inventory to trade massive flows;
- the cost of trading (i.e. the transaction costs) are now of paramount importance;

Simultaneously:

- regulators pushed towards more trading on electronic platforms;
- exchanges and clearing houses are now in competition, generating a more complex trading process, demanding real-time choices to be made;
- more data than ever are available to describe the price formation process, and participants have better capabilities to record, store, and process those data.

Motivations: A New Field for Financial Mathematics (?)

- Innovative frameworks allowing the application of optimization techniques to the trading process have been initiated by mathematicians [Almgren and Chriss, 2000, Bertsimas and Lo, 1998];
- The main fatures of the price formation process have been modelled a suitable way [Almgren et al., 2005, Smith et al., 2003]; especially the *market impact*;
- The transitory nature of liquidity and the high dimension of the trading state space needs statistical estimates to be built and used during the optimization process.

The conjunction of these three components: optimal trading, price formation modelling, and statistical learning, gave birth to a new area of financial mathematics.

For generic considerations, see *Market Microstructure knowledge needed to control an intra-day trading process* [L. 2013], *Market Microstructure in Practice* [L. and Laruelle, 2013], or *Market Microstructure Confronting Many Viewpoints* [Abergel, Bouchaud, Foucault, L., Rosenbaum, 2012]. For focussed papers, see *Market Microstructure and Liquidity* issues.

Outline

Optimal Trading Speed Against Temporary Impact

Optimal trading is about optimizing a trading activity taking into account "microscopic" effect.
The first of such effect is the **price impact**: when you send a liquidity consuming order (example: a sell order ⍗), you impact negatively the price of your trade, you can expect the liquidity to come back if you wait *enough* (empirics in [Bouchaud et al., 2004], example of models for trading in [Obizhaeva and Wang, 2005] and [Alfonsi et al., 2009]).

From a (large orders) trading perspective, this aggregates in **temporary impact**. At a mesoscopic scale [Almgren et al., 2005], you can afford to model this impact as a local price premium to pay as a function of your trading intensity.

Because of this penalty, it is not optimal to trade immediately after a decision. Even with no information on future prices, **you have to cut your large order into small pieces** and to wait for relaxation after each trade.

The Almgren-Chriss Framework And Its Analogy With Portfolio Construction

The first models proposed in [Bertsimas and Lo, 1998], [Grinold and Kahn, 1999, Chap. 16] and [Almgren and Chriss, 2000] can be seen as natural extensions of portfolio allocation. They are based on discretization of time into T time bins (of around 5 min for liquid instruments). You have to allocate a volume $\sum_n v_n = V^*$ over these T bins.
Replace portfolio components by time bins and expected returns by (negative) market impact. But it is a little more sophisticated since price at bin $t+h$ is conditioned by price at time t. Assume this Markov chain for the price $S_{t+1} = S_t + \sigma_{t+1}\,\xi_{t+1}$, and use $\eta_t(v_t)$ for the additive temporary market impact function. The price of buying over T reads

$$W(v_1, \ldots, v_T) = \sum_{t=1}^{T} v_t \cdot (S_t + \eta_t(v_t)).$$

Now you can use a mean-variance criterion on W, cancels its derivative with respect to $x_t = \sum_{s>t} v_t$. Provided the market impact is linear (i.e. $\eta_t(v_t) = \eta \cdot \sigma_t\, v_t / V_t$), you obtain this recurrence relation:

$$x_{t+1} = \left(1 + \frac{\sigma_{t-1}}{\sigma_t}\frac{V_t}{V_{t-1}} + \frac{\lambda}{\eta}\sigma_t^2\right) x_t - \frac{\sigma_{t-1}}{\sigma_t}\frac{V_t}{V_{t-1}} x_{t-1},\ x_0 = V^*,\ x_T = 0.$$

This is The Swiss Knife of Trading Curves

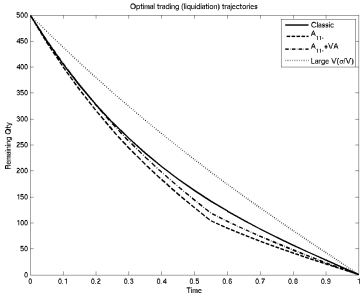

It introduced the idea of **optimal trading curves** ⍈→ crucial for risk control.

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

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

This is The Swiss Knife of Trading Curves

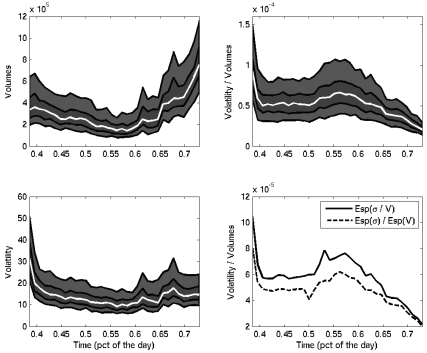

It introduced the idea of **optimal trading curves** $\underset{\downarrow}{\rightarrow}$ crucial for risk control.

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

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

Open questions (within this framework)

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

More Within the AC Framework

Why introduce a variance term in the cost function?

- you took the decision $\Rightarrow$ you want to realize it as soon as possible. With a small change of the criterion, Cartea-Jaimungal studied a lot of variations based on this cost function [Cartea et al., 2013]:

$$H_t(v) := \mathbb{E}_t \left(W_T^v + Q_T^v (S_T - \eta(Q_T^v)) - \lambda \int_t^T (Q_s^v)^2 ds \right).$$

- with $\lambda = 0$ and an explicit (exponential) form of orderbook relaxation, Alfonsi and Schied elaborated in this direction ([Alfonsi et al., 2009, Gatheral et al., 2012]). A U-shaped trading curve ⍗ generally stems from such choice (due to market impact relaxation and blindness after the last transaction).
- The choice of a variance term (instead of any p-variation) can be discussed, too [Labadie and L. 2014].Especially if you think intraday price dynamics are *more mean reverting* than daily ones.

Practical Use of Automated Trading: Adding Constraints

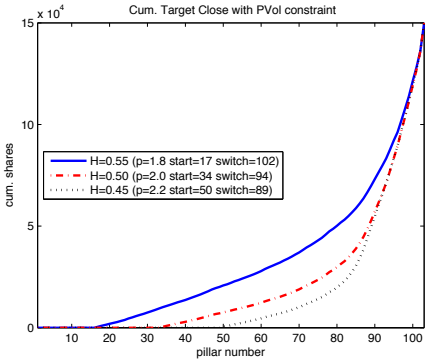

The Almgren-Chriss criterion is the **Implementation Shortfall** (i.e. $W(v)$), but other "trading styles" are possible, like **VWAP** (follow the usually traded volume [Cartea and Jaimungal, 2014]), **PoV** (follow the real-time market volume), and **Target Close**.

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

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

Outline

Optimal Trading Against Permanent Impact: stylized facts

On our database of 300,000 large orders

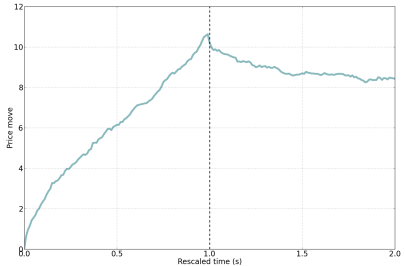

Market Impact takes place in different phases

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

In [Bacry, Iuga, Lasnier, L. 2014] we studied all the phases, using intraday and daily analysis (for the first time). We underlined the importance of some "normalization variables": the uncertainty on the price formation process , the capability of the orderbook to resist to volume pressure , and the duration of the metaorder.

Following [Waelbroeck and Gomes, 2013] and simultaneously with [Brokmann et al., 2014], we proposed an explanation of **permanent impact** .

Is it Really Possible to Minimize Permanent Impact?

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

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

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

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

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

Outline

Optimal Trading Against Orderbook Dynamics

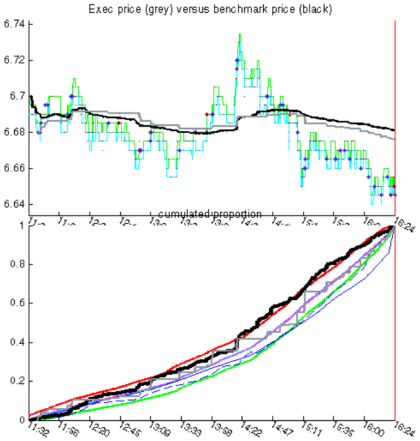

This idea of a decision process in two scales is close to reality:

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

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

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

Optimizing a Two-Scales Controller

The control

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

The dynamics

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

The gain

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

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

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

Optimizing a Two-Scales Controller

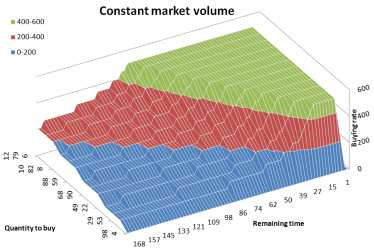

In [Bouchard, Dang, L. 2011], we developed a model describing this two scales process:

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

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

Outline

Using Stochastic Algorithms to Optimize Trading Robots

In 2008, R Berenstein's internship *Algorithmes d'exécution optimale pour les Dark Pools* has been the beginning of a research project with G Pagès and S Laruelle, initially inspired by R Almgren's whitepaper on Smart Order Routing embedding an hidden orders detector.
The goal was to use the **results of stochastic algorithms theory** to address questions raised during the **construction of the small trading robots** to be used by optimal schedulers focused on trading rates.

In such a framework

- You **do not focus on medium term** dynamics and on the usual terminal constraints (i.e. flat position at the end of the day).
- You strongly rely on **short term ergodicity** to minimize the expectation of the cost thanks to the realizations of this cost along the observed trajectory.
- Starting from the criterion, you obtain an **updating rule for your controls** thanks to a **stochastic descent** -like algorithm,
- you obtain CLT-like results to guaranteeing a given convergence rate.

Learning by Trading: an Online Approche to Optimal Trading

With S Laruelle and G Pagès, we mainly addressed two applications:

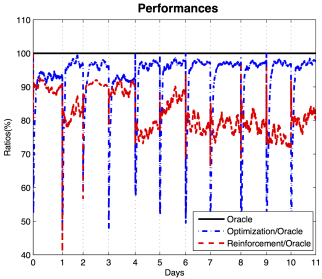

Dark Pool online learning

① Trading in Dark Pools [Pagès, Laruelle, L. 2011] is by essence an **exploration-exploitation** problem. The control is the split $(r_n)_n$ of a large order V across N Dark Pools, each of them accepting transactions up to invisible quantities $(D_n)_n$. The cost is hence

$$\mathcal{C}(r) = \sum_t \sum_n \theta_n(r_n V(t) \wedge D_n(t)) + [r_n V(t) - (r_n V(t) \wedge D_n(t))]_+.$$

It can be shown the associated updating rule for r is

$$r_n(t+1) = r_n(t) + \gamma_{t+1}\pi\big(\mathbf{1}_{r_n(t)>D_n(t)/V(t)} - \tfrac{1}{N}\textstyle\sum_k \mathbf{1}_{r_k(t)>D_k(t)/V(t)}\big).$$

where γ verifies the usual Robbins-Monro assumptions and π is a projection on the simplex.

Two other frameworks for optimal Dark Pool splitting exist: one using censored statictics [Ganchev et al., 2010] and the under minimum regret [Agarwal et al., 2010].

② Adjusting the **price of limit orders** into a visible orderbook [Laruelle, Pagès, L. 2013] can be a way to take into account liquidity dynamics around the mid price.

Outline

Towards a Better Understanding of Orderbook dynamics: The Queue Reactive Model (QRM)

The most commonly used framework for orderbook dynamics in stochastic control of trading algorithms [Avellaneda and Stoikov, 2008] is: "*when you insert an order at a distance δP of the mid price, it is consumed by a Poisson process of intensity $A \exp -k\delta P$*". It is in fact well suited as far as you cannot observe the orderbook dynamics with accuracy during the trading process. In [Huang, Rosenbaum, L. 2015] we went deeper in the details.

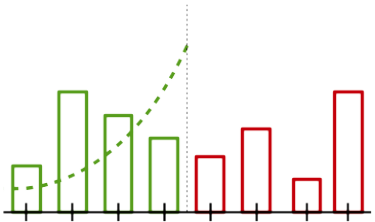

$y-$axis: quantities; $x-$axis: prices.

Towards a Better Understanding of Orderbook dynamics: The Queue Reactive Model (QRM)

The most commonly used framework for orderbook dynamics in stochastic control of trading algorithms [Avellaneda and Stoikov, 2008] is: "*when you insert an order at a distance δP of the mid price, it is consumed by a Poisson process of intensity $A \exp -k\delta P$*". It is in fact well suited as far as you cannot observe the orderbook dynamics with accuracy during the trading process. In [Huang, Rosenbaum, L. 2015] we went deeper in the details.

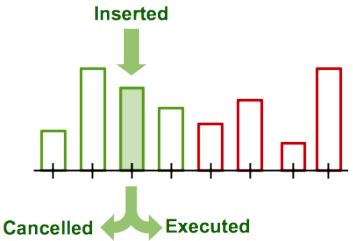

- Each orderbook queue can be filled by insertions and depleted via cancellations and transactions.

Towards a Better Understanding of Orderbook dynamics: The Queue Reactive Model (QRM)

The most commonly used framework for orderbook dynamics in stochastic control of trading algorithms [Avellaneda and Stoikov, 2008] is: "*when you insert an order at a distance δP of the mid price, it is consumed by a Poisson process of intensity $A \exp -k\delta P$*". It is in fact well suited as far as you cannot observe the orderbook dynamics with accuracy during the trading process. In [Huang, Rosenbaum, L. 2015] we went deeper in the details.

- Each orderbook queue can be filled by insertions and depleted via cancellations and transactions.
- The state flow affecting each queue is influenced by its neighbourhood.

Towards a Better Understanding of Orderbook dynamics: The Queue Reactive Model (QRM)

The most commonly used framework for orderbook dynamics in stochastic control of trading algorithms [Avellaneda and Stoikov, 2008] is: "*when you insert an order at a distance δP of the mid price, it is consumed by a Poisson process of intensity $A \exp -k\delta P$*". It is in fact well suited as far as you cannot observe the orderbook dynamics with accuracy during the trading process. In [Huang, Rosenbaum, L. 2015] we went deeper in the details.

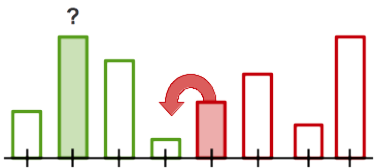

- Each orderbook queue can be filled by insertions and depleted via cancellations and transactions.
- The state flow affecting each queue is influenced by its neighbourhood.
- When a tick is fully depleted, the price changes. This can have a long range influence.

Typical Empirical Intensities

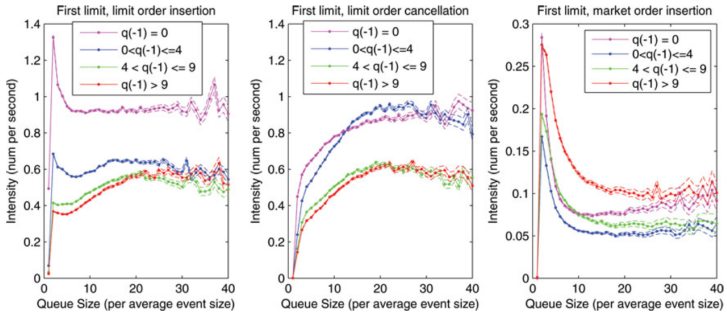

(a) (b) (c)

$$\mathbb{E}dq(1) = (\lambda^L(q(1), q(-1)) - \lambda^C(q(1), q(-1)) - \lambda^M(q(1), q(-1)))\, dt.$$

At left: The empirical intensities of the process of arrivals (a) and departures (b-c) at the first queue as a function of its own size (x-axis) and of the size of the opposite queue (colors). □

Typical Empirical Intensities

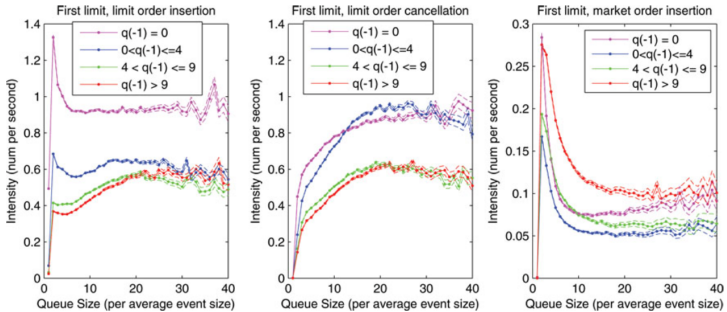

(a) (b) (c)

$$\mathbb{E}dq(1) = (\lambda^L(q(1), q(-1)) - \lambda^C(q(1), q(-1)) - \lambda^M(q(1), q(-1)))\, dt.$$

At left: The empirical intensities of the process of arrivals (a) and departures (b-c) at the first queue as a function of its own size (x-axis) and of the size of the opposite queue (colors).

From a **control perspective** (of the rank r of an infinitesimal order in the first queue):

$\mathbb{E}dr \simeq -(r/q(1)\,\lambda^C(q(1), q(-1)) + \lambda^M(q(1), q(-1)))dt$ and $\Pr(q(1) + dq(1) < 0)$ is known.

More importantly the **strength of adverse selection** $\mathbb{E}(q^{\text{new}}(1)|q(1) + dq(1) < 0)$ can be controlled.

Outline

Optimal Trading Against the Crowd

If one wants to take into account the fact other **participants are playing a similar game** in posting limit orders (waiting for a possible transaction at a better price) or sending market orders (paying the bid-ask spread and price impact to gain immediacy), you face to a very complex problem. In [Lachapelle, Lasry, L., Lions 2013] we have made an attempt to understand the outcome of such a (mean field) game.

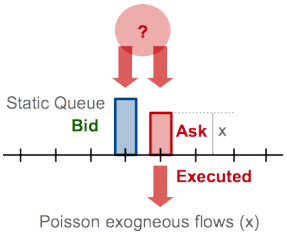

Optimal Trading Against the Crowd

If one wants to take into account the fact other **participants are playing a similar game** in posting limit orders (waiting for a possible transaction at a better price) or sending market orders (paying the bid-ask spread and price impact to gain immediacy), you face to a very complex problem. In [Lachapelle, Lasry, L., Lions 2013] we have made an attempt to understand the outcome of such a (mean field) game.

- We assumed (1) no exogenous price change, (2) pro-rata rule, (3) uniform shape of the book —i.e. focus on the first limit(s). We take profit of existing convergence results in Shared Risk MFG.
- The 2-dim **Mean Field** is made of the two first sizes (Q^a, Q^b).
- For a newcomer, the choice (i.e. control) is made by **waiting into its own queue** (and paying a waiting cost but potentially being rewarded by a better price), or **consuming the opposite queue** (and paying an immediate cost).
- Participants can anticipate others' behaviour.
- We are **interested by the stationary state** of the queues and the stationary value of waiting in the cost.

Main Components of the Liquidity Game

Note that thanks to the pro rata assumption all limit orders in the queue face similarly all the events.

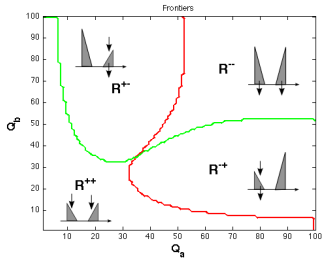

The events affecting the value functions u (sellers) and v (buyers) during an infinitesimal time dt:

- the queue is consumed: a fraction of its **value increases** by the price impact, the remaining quantity is valued with a smaller queue,
- the queue increases: its **value decreases**,
- in any case the value decreases by a waiting cost.

- the increase of the queue is conditioned by the value function being greater than the mid price,
- the decrease of the queue is conditioned by the opposite value function being greater than the mid price.

First order (forward) PDE satisfied by u and v (with the notation $x = Q^a$, $y = Q^b$, λs being intensities, and R the control, i.e. backward, components associated to the controls):

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

$\Rightarrow$ We found some **stable configurations of liquidity imbalance**, that can be reduced when the **diversity of players** increases. It explain some virtues of market making $\square$.

Our theoretical results are confirmed empirically by [Gareche et al., 2013].

Outline

Conclusion

The control of a trading process in an auction game lies at the crossing of a wide variety of fields:

- **Optimal Control**, from mean variance optimality to impulse or hybrid control.
- **Stochastic Algorithms**, to adapt in real-time to liquidity variations.
- **Statistics of Random Processes**, to capture limit orderbook dynamics.
- **Game Theory**, to go deeper in the details and assess robustness.

This research agenda has been driven by the necessity to answer to **practitioners' needs** and **regulators' questions**.

Financial Mathematics have a important role to play in providing results and tools in terms of **risk control**, thanks to a good understanding of the **cost functions** involved in the optimizations solved by practitioners.
It can prevent trading systems to become black boxes opaques for users, risk departments and regulators.

References I

Abergel, F., Bouchaud, J.-P., Foucault, T., Lehalle, C.-A., and Rosenbaum, M., editors (2012).
Market Microstructure Confronting Many Viewpoints.
Wiley.

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

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

Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005).
Direct Estimation of Equity Market Impact.
Risk, 18:57–62.

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

Avellaneda, M. and Stoikov, S. (2008).
High-frequency trading in a limit order book.
Quantitative Finance, 8(3):217–224.

References II


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. (2014).
Market Impacts and the Life Cycle of Investors Orders.
Social Science Research Network Working Paper Series.


Bacry, E. and Muzy, J. F. (2013).
Hawkes model for price and trades high-frequency dynamics.


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.


Bouchaud, J.-P., Gefen, Y., Potters, M., and Wyart, M. (2004).
Fluctuations and response in financial markets: the subtle nature of 'random' price changes.
Quantitative Finance, 4(2):176–190.


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

References III

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.

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.

Grinold, R. and Kahn, R. (1999).
Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk.
McGraw-Hill, 2 edition.

References IV

Guéant, O. and Lehalle, C.-A. (2013).
General intensity shapes in optimal liquidation.
Mathematical Finance, page n/a.

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.

Huang, W., Lehalle, C.-A., and Rosenbaum, M. (2015).
Simulating and analyzing order book data: The queue-reactive model.

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. (2013).
Efficiency of the Price Formation Process in Presence of High Frequency Participants: a Mean Field Game analysis.

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

References V


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.


Obizhaeva, A. and Wang, J. (2005).
Optimal Trading Strategy and Supply/Demand Dynamics.
Social Science Research Network Working Paper Series.


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.

References VI

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.

Monitoring a Large Collection of Automated Orders

With the increase of electronization of the trading process, we have **more transparency** of the price formation process (with the recording of any order, on any trading floor), but it became **more complex** to capture. It **does no more take place a localized way**. I.e. on few trading floors where you can stand and have a "sentiment" of what is happening. It takes place over a connected network of computers fed by human decisions incorporating real-time valuations of assets.

A great challenge is to have humans and machines **sharing information**:

- humans should have access to parameters to tune the behaviour of algorithms under controlled risk,
- machines should deliver synthetic description of their internal state.

In 2010, we started a research project to provide decision support to human traders **monitoring simultaneously hundreds of trading algorithms**. Our goal was to extract implicit clusters of algos.

Conjunction Detection as a way to Understand Market Behaviour

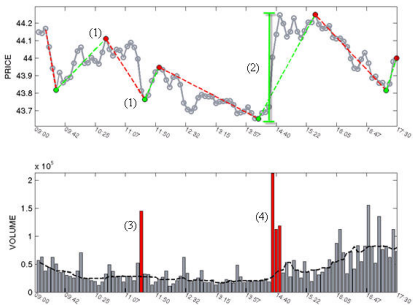

Say you succeed into modelling in real-time the relationship between the instantaneous performance of 1,000 trading algorithms and a collection of explanatory variables.

- First extract features of the price, volume, volatility, and bid-ask spread process (see at Left)
- Then normalize these features ("scoring" them) to obtain a collection of variables $Z_1, \ldots, Z_K$.
- Use **binary predictors**: $Z_k(t) \notin (\theta_k^-, \theta_k^+)$ predicts a bad performance at t and combinations of such predictors.

In [Azencott, Beri, Gadhyan, Joseph, L., Rowley 2014] we have shown how to select the best predictors very fast a reliable way (i.e. confidence intervals).

Conjunction Detection as a way to Understand Market Behaviour

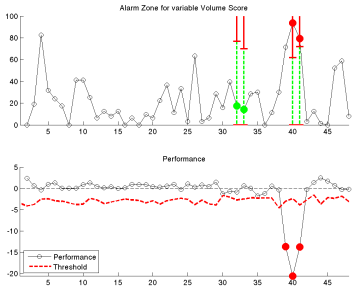

Say you succeed into modelling in real-time the relationship between the instantaneous performance of 1,000 trading algorithms and a collection of explanatory variables.

- First extract features of the price, volume, volatility, and bid-ask spread process (see at Left)
- Then normalize these features ("scoring" them) to obtain a collection of variables $Z_1, \ldots, Z_K$.
- Use **binary predictors**: $Z_k(t) \notin (\theta_k^-, \theta_k^+)$ predicts a bad performance at t and combinations of such predictors.

In [Azencott, Beri, Gadhyan, Joseph, L., Rowley 2014] we have shown how to select the best predictors very fast a reliable way (i.e. confidence intervals).

$\Rightarrow$ A good predictor at t provides **information on influencing variables**.
You can hence group trading algos by influencing variables, to help human traders to understand what is happening and take enlighten actions.

A First Attempt of Price Formation Modelling

In practice we need synthetic markets to test trading algorithms. In [Guéant, Razafinimanana, L. 2010]we develop a microfunded market simulator based on a Mean Field Game approach.
We introduced the concept of **orderbook of the views**, now best known as **latent orderbook**. We described the price formation process dynamics taking place in two stages:

- the mean field: a **latent orderbook** is shaped by the aggregated views of all investors (thanks to isolated actualization of anticipations),
- the exchanges **orderbooks**, where the price forms at high frequency.

Dynamics of the latent ordebook m, where p is the price, $p^*(t)$ the (temporary) equilibrium:

$$dm = \frac{\varepsilon^2}{2}\partial^2_{pp} m(t,p)dt - \frac{\varepsilon^2}{2}\partial_p m\left(t, p^*(t)\right)\left[source(t,p)\right]dt + \nu g(p, p^*(t))dW^p_t \quad (1)$$

With three types of agents: trend followers, mean reverters and noise traders. Each of the two first agent type has a gamma distribution for his time horizon.

- The distance between the observed orderbook and this latent one (the one of the *views*) conditions the occurrences of market events (insert or cancel) on the observed one.

Markov Chains on the Sizes in the Book is Not Enough

In [Huang et al., 2015] we have shown

- how to properly (theoretically) define a model in which the intensity of the point processes of arrival and departures (via cancellations and transactions) are functions of the state of all the queues of the book,
- using HF data to fit the conditional intensities of such a model enables to render with accuracy the **liquidity game** (i.e. dynamics of queues sizes),
- nevertheless **it does not fit well the price changes dynamics**.
- Adding an **exogenous component** allows to recover price dynamics. When a queue is completely deplete, reinitialize the orderbook state according to its usual distribution.

Note this modelling is orthogonal to the (usual) Hawkes process one [Bacry and Muzy, 2013], in which the past intensities of the flows are conditioning the future flows 무. Our **Queue Reactive Model** is a conditionning on the observed sizes of the queues.

Markov Chains on the Sizes in the Book is Not Enough

In [Huang et al., 2015] we have shown

- how to properly (theoretically) define a model in which the intensity of the point processes of arrival and departures (via cancellations and transactions) are functions of the state of all the queues of the book,
- using HF data to fit the conditional intensities of such a model enables to render with accuracy the **liquidity game** (i.e. dynamics of queues sizes),
- nevertheless **it does not fit well the price changes dynamics**.
- Adding an **exogenous component** allows to recover price dynamics. When a queue is completely deplete, reinitialize the orderbook state according to its usual distribution.

Note this modelling is orthogonal to the (usual) Hawkes process one [Bacry and Muzy, 2013], in which the past intensities of the flows are conditioning the future flows 🙂. Our **Queue Reactive Model** is a conditionning on the observed sizes of the queues.

What Drives Orderbook Dynamics?

Two different phases

- Usually a **liquidity game** takes place, driven by QRM-like dynamics.

→ this phase has a very short term memory and shapes orderbooks.

- Time to time, all agents seems to **agree on a price change**.

→ this phase has a long(er) term memory and shapes price trajectories.

⇒ In reality each participant adjusts his views according to a lot of external parameters.