Empirical study, modelling and applications of multiple limits trades in limit order books

Fabrizio Pomponio

Thèse de doctorat de l'Ecole centrale paris

Spécialité :
Finance Quantitative

Laboratoire :
Mathématiques Appliquées aux Systèmes

Présentée par :
Fabrizio POMPONIO
Chaire de finance quantitative,
Laboratoire MAS, Ecole Centrale Paris

En vue d'obtenir le grade de
Docteur de l'Ecole Centrale Paris

Etude empirique, modélisation et applications des trades à limites multiples dans les carnets d'ordres

Sous la direction de Frédéric ABERGEL

Rapporteurs : Janos KERTESZ
Fabrizio LILLO

Thèse soutenue publiquement le 14 décembre 2012 devant le jury composé de

Frédéric ABERGEL	Ecole Centrale Paris	Directeur
Jean-Philippe BOUCHAUD	Ecole Polytechnique	Examinateur
Janos KERTESZ	Budapest University of Technology and Economics	Rapporteur
Boris LEBLANC	BNP-Paribas	Examinateur invité
Fabrizio LILLO	Scuola Normale Superiore di Pisa	Rapporteur
Mathieu ROSENBAUM	Université Pierre et Marie Curie	Président
Stéphane TYC	BNP-Paribas	Examinateur

A mes parents, à mon frère.

Remerciements

A quelques jours de soutenir cette thèse, ma première pensée va à mon directeur de thèse Frédéric ABERGEL. Je le remercie chaleureusement pour son encadrement ; il a su me laisser suffisamment de liberté dans mon travail pour explorer des pistes nouvelles, tout en me guidant davantage lorsque cela s'est avéré nécessaire. Au final, même si cette thèse fut loin d'être une partie de plaisir, je le remercie d'avoir fait son possible pour m'aider à la faire aboutir.

Je voudrais également remercier Stéphane TYC et Boris LEBLANC pour m'avoir permis de réaliser cette thèse au sein de l'équipe de Recherche quantitative à BNP-Paribas Paris. J'ai beaucoup appris à leur contact.

I would also like to thank the referees Prof. Fabrizio LILLO and Prof. Janos KERTESZ for their careful reading of the manuscript and their interesting comments. I tried to include and possibly answer all of their comments, the best I could with the time and current conditions available. I am pretty sure that the remarks they made would lead to significative new and interesting developments for future research. Thanks to all the members of the jury, I am very honoured by their presence.

Au sein de l'équipe Recherche à BNP, je tiens à remercier toute l'équipe Histo pour m'avoir aidé dans l'utilisation des données tick à tick, et particulièrement Cédric JOULAIN, Mehdi ABOULKASSIM et Olivier SAINGRE. De même, je remercie ceux avec qui j'ai été amené à travailler et qui ont répondu à mes questions, en particulier Alexandre DAVROUX et Jacques-Olivier MOUSSAFIR. Merci de manière générale à toute l'équipe pour son accueil, et notamment aux "fous du volant" (Julien NGUYEN, Geoffrey SCOUTHEETEN, ...) pour les bons moments passés ensemble.

J'ai bien sûr une pensée amicale pour toute l'équipe de la Chaire de finance quantitative BNP-Paribas de Centrale Paris, et plus particulièrement pour Ioane MUNI-TOKE, avec qui j'ai eu le plaisir de travailler, mais aussi Anirban CHAKRABORTI, Laurent GUERBY, et les doctorants, post-doc et stagiaires de la première période de la chaire, et que j'ai pu davantage cotoyer durant ma thèse (Nicolas et Rémi, Mauro, Aymen et Riadh, Nicolas et Ban, Mehdi, ...).

Concernant la décision d'entreprendre une thèse, je dois reconnaître que le premier à m'avoir tout particulièrement encouragé dans cette voie est Erick HERBIN lors d'un projet de maths en 2ème année à Centrale. Merci Erick, c'était a posteriori une très bonne suggestion.

Je remercie également Lionel GABET pour son soutien amical, d'abord quand j'étais étudiant ingénieur, puis en tant que doctorant en maths appliquées à Centrale.

Toujours à Centrale, je voudrais remercier le personnel administratif pour leur bonne humeur et toute l'aide qu'il nous apporte au quotidien, en particulier Annie et Sylvie du labo MAS, et aussi Catherine et Géraldine de l'école doctorale. Merci à Dany qui sait résoudre tous les mystères liés aux connexions réseau, Internet et imprimantes.

Quelques très bons professeurs ont marqué ma scolarité et je tenais à les remercier, notamment Mme SEVERIN en maths au lycée, ainsi que Jean VOEDTS et Xavier COINTAULT, en maths et en physique en maths spé.

Un grand, grand merci à mes amis, notamment ceux de l'époque prépa (Erwann, Augustin, Benjamin, Christophe,...) et ceux de l'époque Centrale (Wafaa, Sébastien, David, ...).

Enfin, une pensée de remerciement et de gratitude envers ma famille, en particulier mes parents et mon frère.

Résumé

Cette thèse étudie certains évènements particuliers des carnets d'ordre - les "trades traversants". Dans le premier chapitre, on définit les trades traversants comme étant ceux qui consomment la liquidité présente dans le carnet d'ordres sur plusieurs limites, sans laisser le temps à la meilleure limite de se remplir par l'arrivée de nouveaux ordres limites. On étudie leurs propriétés empiriques en fournissant des statistiques de liquidité, de volume, de distribution de leurs temps d'arrivées, de clustering et de relaxation du spread. Leur impact de marché est supérieur à celui des trades classiques, et ce même à volume comparable : les trades traversants présentent donc un contenu informationnel plus grand. On propose deux applications au problème du lead-lag entre actifs/marchés, d'abord pour répondre à la question de savoir quel actif bouge en premier, et ensuite pour mesurer la force du signal des trades traversants dans le cadre d'une stratégie d'investissement basée sur le lead-lag entre actifs.

Le chapitre suivant approfondit l'étude empirique du clustering de l'arrivée des trades traversants. On y modélise leur arrivée par des processus stochastiques auto-excités (les processus de Hawkes). Une étude statistique de la calibration obtenue avec des modèles à noyaux exponentiels pour la décroissance temporelle de l'impact est menée et assure une modélisation satisfaisante avec deux processus indépendants, un pour le bid et un pour l'ask. La classe de modèles proposée à la calibration est bien adaptée puisqu'il n'existe pas d'effet inhibiteur après l'arrivée d'un trade traversant. On utilise ces résultats pour calculer un indicateur d'intensité basé sur l'arrivée des trades traversants, et améliorer ainsi une stratégie d'investissement de type "momentum". Enfin, une calibration non-paramétrique du noyau de décroissance temporel d'impact fournit une décroissance empirique encore plus forte qu'une loi exponentielle, et davantage proche d'une loi-puissance.

Le dernier chapitre rappelle une méthode générale de détection statistique de sauts dans des séries temporelles de prix/rendements qui soit robuste au bruit de microstructure. On généralise les résultats empiriques connus à de nouveaux indices financiers. On adapte cette méthode de détection statistique de sauts à des trajectoires intraday afin d'obtenir la distribution de la proportion de sauts détectés au cours de la journée. Les valeurs extrémales et les plus grandes variations de cette proportion se déroulent à des heures précises de la journée (14 :30, 15 :00 et 16 :30, heure de Paris), déjà rencontrées dans l'étude des trades traversants. Grâce à eux, on propose une explication des caractéristiques principales du profil intraday de la proportion de sauts détectés par le test, qui s'appuie sur une modification de la part relative de chacune des composantes de sauts dans la trajectoire des actifs considérés (la composante des mouvements continus et celle liée aux mouvements de sauts purs).

Mots-clés : Données financières tick-à-tick, statistiques empiriques, lead-lag, processus de Hawkes, calibration de modèle, détection de sauts, test statistique.

Abstract

This thesis aims at studying particular events occurring in the limit order books - the 'trades-through'. In the first chapter, we define trades-through as those who consume the liquidity available on several limits of the limit order book, without waiting for the best limit to be filled with new incoming limit orders. We study their empirical properties and present statistics about their liquidity, their volume, their arrival time distribution, their clustering and the spread relaxation that follows their arrival. Their market-impact is higher than the one of the other trades, even with a comparable trading volume : trades-through have a higher informational content. We present two applications linked to the lead-lag between assets/markets : to find which asset moves first, and also to measure the trades-through intensity signal in a simple trading strategy based on lead-lag.

The next chapter goes into more detail about the trades-through arrival time clustering. We model their arrival time with self-excited stochastic processes (Hawkes processes). A statistical study of the calibration obtained with models based on exponential-decay kernels for the temporal impact ensures a satisfactory modelling with two independent processes, one for the bid and one for the ask. The model class under scrutiny for the calibration is well-adapted as no inhibitory effects are measured after trades-through arrival. We use those results to compute an intensity indicator based on trades-through arrival, and thus we enhance a simple trading strategy that relies on them. Finally, a non-parametric calibration of the empirical decay kernel for the temporal impact of trades-through indicates a decrease faster than exponential, and closer to a power-law.

The last chapter recalls a general statistical method robust to market microstructure noise to find jumps in prices/returns time series. We generalize the empirical results already known in the literature to new financial indices and we adapt this statistical jump detection method to intraday trajectories in order to obtain the intraday proportion of detected jumps. Extreme values and biggest intraday variations of this jump proportion occurs at very specific hours of the day (14 :30, 15 :00 and 16 :30, Paris time reference), already linked with trades-through. Using trades-through, we explain the main characteristics of the intraday proportion of detected jumps with the test using a modification in the relative importance of each jump component in the assets trajectories (the continuous moves component and the pure-jumps component).

Keywords : Tick-by-tick financial data, empirical statistics, lead-lag, Hawkes processes, model calibration, jump detection, statistical test.

Table des matières

1 Introduction (French version) 15
- 1.1 Contexte général d'étude, présentation et unité de la thèse . . . 15
- 1.2 Propriétés empiriques des trades traversants et deux applications au lead-lag . . . 16
  - 1.2.1 Résultats obtenus . . . 18
- 1.3 Modélisation des trades traversants par des processus de Hawkes . . . 21
  - 1.3.1 Résultats obtenus . . . 22
- 1.4 Comparaison entre les trades traversants et la méthode statistique de Aït-Sahalia et al. [2012] pour la détection de sauts dans les trajectoires intraday . . . 23
  - 1.4.1 Résultats obtenus . . . 26

2 Introduction 29
- 2.1 A reminder on financial markets microstructure . . . 29
- 2.2 Outline of the thesis . . . 30

3 Trades-through : empirical facts and applications to lead-lag 33
- 3.1 Introduction . . . 34
- 3.2 Data presentation . . . 36
  - 3.2.1 General TRTH data presentation and processing . . . 36
  - 3.2.2 Data used in this study . . . 36
- 3.3 Elementary statistical properties of trades-through . . . 37
  - 3.3.1 On which limit of the order book is liquidity taken from ? . . . 38
  - 3.3.2 Frequency and volume . . . 38
  - 3.3.3 Links between trades-through and large trades . . . 41
  - 3.3.4 Clustering . . . 42
  - 3.3.5 Intraday distribution . . . 43
- 3.4 Market impact . . . 44
  - 3.4.1 Empirical response function . . . 46
  - 3.4.2 Response function conditioned on trade volume . . . 46
- 3.5 Spread relaxation . . . 47
- 3.6 Lead-lag estimation from trades-through time-series . . . 49
  - 3.6.1 The estimation technique . . . 50
  - 3.6.2 Empirical results . . . 51
- 3.7 Empirical lead-lag signal from trades-through . . . 52
  - 3.7.1 Data description and summary statistics . . . 53
  - 3.7.2 Empirical lead-lag signal based on trades-through . . . 54
- 3.8 Conclusion and further research . . . 54

4 Modelling trades-through in a limit order book using Hawkes processes 59
4.1 Introduction . . . 60
4.2 Trades-through . . . 60
4.2.1 Orders splitting and trades-through . . . 60
4.2.2 Definition of trades-through . . . 61
4.2.3 Empirical data . . . 61
4.2.4 Occurrences of trades-through . . . 61
4.2.5 Clustering . . . 61
4.2.6 Intraday timestamp distribution . . . 63
4.3 Hawkes processes . . . 64
4.3.1 Definition . . . 64
4.3.2 Stationarity condition . . . 65
4.3.3 Maximum-likelihood estimation . . . 66
4.3.4 Testing the calibration . . . 67
4.4 A simple Hawkes model for trades-through . . . 68
4.4.1 Model . . . 68
4.4.2 Calibration . . . 69
4.4.3 Goodness-of-fit . . . 72
4.5 Toy-trading strategy based on trades-through autocorrelated arrival times . . . 72
4.5.1 Hawkes signal computation . . . 72
4.5.2 Results . . . 73
4.6 Searching for inhibitory effects in trades-through . . . 75
4.6.1 Models . . . 75
4.6.2 Results . . . 75
4.7 Calibration of the empirical decay kernel of trades-through arrival times by a non-parametric method . . . 76
4.7.1 A reminder on nonparametric Hawkes process calibration . . . 76
4.7.2 Application to trades-through arrival times . . . 77
4.8 Conclusion . . . 77

5 Comparison of trades-through microstructural jumps with macroscopic jumps (detected by a statistical test robust to microstructure noise) 85
5.1 Introduction . . . 86
5.2 Testing for jumps in a discretely observed noisy process *a la* Aït-Sahalia and Jacod . 87
5.2.1 The case with no microstructure noise . . . 87
5.2.2 The case with microstructure noise . . . 88
5.3 Empirical comparison of high occurrence of jumps, detected by trades-through and by a statistical test method . . . 90
5.3.1 Empirical data and parameters used . . . 90
5.3.2 Testing for jumps in daily trajectories . . . 91
5.3.3 Testing for jumps in intraday trajectories . . . 91
5.3.4 Comparing the intraday distributions of jumps obtained with trades-through and with the statistical test method . . . 92
5.4 Conclusion . . . 93

Table des figures

1.1 Histogramme intraday moyen des temps d'arrivée des trades traversants pour les 30 composantes les plus liquides de l'indice Footsee en 2011 (avec des intervalles de temps de 1 minute). 19
1.2 Comparaison des fonctions de réponse globale pour le stock BNP Paribas (en échelle log-linéaire) entre tous les trades et les trades traversants, incluant la condition sur le volume initial du trade (ramené au volume de trade moyen, noté MTV). 20
1.3 Décroissance en loi-puissance de l'excès de spread après un trade traversant sur le stock BNP Paribas (en échelle log-log) 20
1.4 Comparaison des résultats, entre l'algorithme MISD et celui du maximum de vraisemblance, pour la calibration empirique du noyau des trades traversants du côté bid du carnet d'ordres, pour le stock BNP Paribas entre Juin et Octobre 2010 (10 :00 à 14 :00). 24
1.5 Comparaison de la distribution de la statistique de test normalisée t avec $N_{\text{daily}} = 10000$, entre $T = 1$ jour (trait plein) et $T = 15$ min (pointillé), avec une distribution normale standard (zone grisée). Une ligne verticale indique la médiane de la distribution. Etude menée sur l'indice Footsee (future et les 30 stocks les plus liquides) en 2011. 26
1.6 *(Haut)* Variations intraday de la médiane du test t avec $N_{\text{daily}} = 10000$ et $T = 15$ min ; *(Bas)* Histogramme moyen des temps d'arrivée des trades traversants avec des intervalles de temps de 1 minute. Etude menée sur l'indice Footsee (future et les 30 stocks les plus liquides) en 2011. 27

3.1 Example of a trade-through : *(up)* Limit order book configuration before the trade-through ; *(middle)* Trade-through ; *(down)* Limit order book configuration after the trade-through. 39
3.2 Fraction of total trading volume taken from each limit of the order book of BNP Paribas stock 40
3.3 Conditional probability for a trade to be a trade-through conditioned on his volume for BNP Paribas stock (in linear-logarithmic scale) 42
3.4 Distribution of waiting-time (in trade-time) until next trade-through for BNP Paribas stock (in double logarithmic scale) 43
3.5 Trades (above), 2nd-limit trades-through (middle) and 3rd-limit trades-through (below) intraday histograms for US equity index future E-mini S&P500 45
3.6 Global response functions for BNP Paribas stock (in linear-logarithmic scale). *(up)* All trades vs trades-through ; *(down)* All trades vs trades-through, with the condition on the initial trade volume with respect to the mean trade volume MTV. 48
3.7 Excess-spread decay after trades for BNP Paribas stock (in linear-logarithmic scale) 49
3.8 Power-law decay of excess-spread after trades-through for BNP Paribas stock (in double logarithmic scale) 50

3.9 Each trade-through of Asset 1 is put in relationship with the trade-through of Asset 2 that is the closest in time. 51
3.10 An example of lead-lag parameter distribution for two simulated Poisson processes with different intensities (one curve for each time grid reference) 52
3.11 Lead-lag parameter empirical distribution (above) and positive and negative cumulative distributions (below) of 2nd-limit trades-through for two groups of European and American futures 56
3.12 *(left)* Leader : BNPP | Lagger : SOGN ; *(right)* Leader : SOGN | Lagger : BNPP ; *(up)* Lead-lag signals across time ; *(down)* PnL distribution after 10s 57
3.13 *(left)* Leader : FDX | Lagger : STXE ; *(right)* Leader : STXE | Lagger : FDX ; *(up)* Lead-lag signals across time ; *(down)* PnL distribution after 0.5s 58

4.1 2nd-limit trade-through example 61
4.2 Global trades-through clustering for BNP Paribas : Empirical distribution of the time intervals between a trade or a trade-through and the next trade-through. 63
4.3 Intraday distribution of timestamps of trades-through for the stock BNP Paribas on June–October 2010, using one-minute bins. 64
4.4 "Mid-to-mid" strategy using trades-through on DAX Future FDX@ *(down)* and BNPP stock *(up)*. Benchmark (simple) : invest 1 after an observed trade-through. Hawkes (oracle) : invest $\phi(\delta)$ after an observed trade-through. 74
4.5 Comparison of trades-through kernel calibration (nonparametric with 20 bins vs exponential) for the stock BNP Paribas on June–October 2010 (10 a.m. to 2 p.m.) on the bid side *(up)* and the ask side *(down)* of the limit order book. 78
4.6 Power-law fit of the empirical nonparametric trades-through kernel for the stock BNP Paribas on June–October 2010 (10 a.m. to 2 p.m.) on the bid side *(up)* and the ask side *(down)* of the limit order book (on a log-log scale). 79

5.1 Triangular kernel g chosen for the robustification of the statistical test in presence of microstructure noise by the pre-averaging method 91
5.2 Daily normalized statistic distribution in 2011 (solid line) compared to standard normal distribution (shaded area). *(up)* DJIA index (30 stocks composition) with $N_{\text{daily}} = 5000$; *(middle)* DAX index (future and 30 stocks composition) with $N_{\text{daily}} = 10000$; *(down)* Footsee index (future and 30 most liquid stocks composition) with $N_{\text{daily}} = 10000$. 95
5.3 Normalized statistic distribution comparison in 2011 with $N_{\text{daily}} = 10000$, between $T = 1$ day (solid line) and $T = 15$ min (dotted line), with standard normal distribution (shaded area). A vertical line indicates the median of the distribution. *(up)* DJIA index (30 stocks composition) ; *(middle)* DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition). 96
5.4 Intraday distribution of the normalized statistic median in 2011 with $N_{\text{daily}} = 10000$ and $T = 15$ min. *(up)* DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition). 97
5.5 Intraday distribution of the normalized statistic boxplot in 2011 with $N_{\text{daily}} = 10000$ and $T = 15$ min. *(up)* DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition). 98
5.6 Mean intraday histograms of timestamps of trades-through for the 30 most liquid components of FDX and FFI indices on 2011, using one-minute bins (*2 graphs up*) and 15-minutes bins (*2 graphs down*) 99

Liste des tableaux

3.1 TRTH's most important trades flags . . . 36
3.2 Data summary table . . . 38
3.3 Basic statistics on 2nd-limit trades-through (estimations based on March 2010 data) 38
3.4 Basic statistics on 3rd-limit trades-through (estimations based on March 2010 data) 39
3.5 Stability of the definition of trades-through with respect to the Mean Trade Volume (estimations based on March 2010 data) . . . 41
3.6 Rarefaction of data for French stock in the estimation of conditional probability of trades-through . . . 43
3.7 Daily liquidity statistics on 2nd and 3rd-limit trades-through for Jan-Feb 2011 . . . 53
3.8 Lead-lag results summary . . . 55
3.9 Comparison of best mid-to-mid returns with transaction costs . . . 55

4.1 Occurrences of trades-through at bid and ask sides for BNP Paribas. . . . 62
4.2 Clustering of trades-through on bid and ask sides (on BNP Paribas data). . . . 62
4.3 Statistics summary for the maximum-likelihood estimates of the ask side of model (4.19). . . . 69
4.4 Statistics summary for the maximum-likelihood estimates of the bid side of model (4.19). . . . 69
4.5 Statistics summary for the maximum-likelihood estimates of the simplified model (4.20) with λ_0^A and λ_0^B constant. . . . 70
4.6 Statistics summary for the maximum-likelihood estimates of the ask side of model (4.20) with λ_0^A and λ_0^B piecewise-linear continuous functions. . . . 71
4.7 Statistics summary for the maximum-likelihood estimates of the bid side of model (4.20) with λ_0^A and λ_0^B piecewise-linear continuous functions. . . . 71
4.8 Performance of the calibration of the Hawkes models. For each model, this table gives the number of trading days (out of 109) where 4, 3, or 2 or less tests out of for where successfully passed. The four tests are two independence Ljung-Box tests and two Kolmogorov-Smirnov tests for the exponential distribution. Values in parentheses are percentages. . . . 72
4.9 Statistics summary for the expectation-maximization estimates of the bid side of model (4.20) with 20 bins. . . . 80
4.10 Statistics summary for the expectation-maximization estimates of the ask side of model (4.20) with 20 bins. . . . 80

Chapitre 1

Introduction (French version)

Cette thèse s'inscrit dans l'étude de la microstructure des marchés financiers au niveau du carnet d'ordres. Elle vise à décrire, à modéliser et à utiliser en vue d'applications certains évènements particuliers, dénotés ci-après "trades traversants" et caractérisés par une consommation agressive de la liquidité présente dans le carnet d'ordres.

1.1 Contexte général d'étude, présentation et unité de la thèse

La mise à disposition d'immenses quantités de données financières regroupant les transactions et les carnets d'ordre d'un grand nombre d'actifs a permis un développement important de la compréhension du fonctionnement des marchés financiers au niveau du carnet d'ordres des actifs considérés. Au lieu de modéliser directement les séries de prix/rendements par des processus stochastiques, on s'est plutôt attaché à une mise en relation des différents phénomènes microscopiques à l'origine de la formation et de la modification des prix. Cette étude systématique des données de carnet d'ordres a mis en lumière des propriétés empiriques fondamentales concernant l'interdépendance de l'offre et de la consommation de la liquidité dans le phénomène de formation des prix. Il reste évidemment encore des questions en suspens afin d'accéder à une compréhension complète des mécanismes en jeu, incluant la fourniture et la consommation de liquidité, la dynamique des carnets d'ordre, le market-impact, la volatilité et le comportement des différents acteurs en présence. Toutefois, ce domaine a atteint un certain niveau de maturité, comme l'atteste le nombre croissant de revues disponibles du domaine (voir par exemple celles de Bouchaud et al. [2009] et de Chakraborti et al. [2011]). Cette maturité se traduit à présent par la recherche de nouveaux champs d'approfondissements dans la compréhension du domaine à l'étude. Parmi ceux-ci, deux semblent émerger tout particulièrement : les études empiriques basées sur des données permettant l'identification des acteurs à l'œuvre, et celles visant à comprendre des évènements particuliers les plus significatifs et les plus instructifs possibles. Au final, si les statistiques globales concernant la microstructure des marchés financiers ont d'ores et déjà permis une bonne compréhension du domaine, il semble que les statistiques conditionnelles (à l'identité des acteurs ou à certains évènements particuliers) ont ouvert un champ d'investigation intéressant et prometteur.

Cette thèse s'inscrit dans la tentative de meilleure compréhension de la microstructure des marchés financiers basée sur l'étude d'évènements particuliers qui ont trait, dans notre cas, à une consommation agressive de la liquidité présente dans le carnet d'ordres. Dans le chapitre 3, on rappelle d'abord certaines propriétés empiriques essentielles à la compréhension de la microstructure des marchés financiers et à la définition du sujet d'étude. La liquidité insuffisante présente dans le carnet d'ordres par rapport à celle désirée par les investisseurs, couplée à la volonté de ne pas révéler aux autres participants ses intentions d'investissements afin de ne pas décaler le prix dans le mauvais

sens, amènent à un découpage usuel des ordres marché en une succession de petits ordres adaptés à la liquidité disponible dans le carnet d'ordres. A l'inverse, il arrive que les investisseurs ne puissent pas étendre leur période d'investissements au-delà d'une certaine durée, mais doivent au contraire agir rapidement et sans considération de minimisation de leur impact de marché. Au cœur de ce compromis, les trades traversants sont justement ceux qui vident complètement plusieurs limites, en ne laissant pas le temps à la meilleure limite du carnet de se remplir à nouveau par l'arrivée de nouveaux ordres limites. Dans le premier chapitre, on réalise une étude approfondie de leurs propriétés empiriques et on les utilise dans le cadre de deux applications liées au lead-lag entre actifs/marchés.

Le chapitre 4 poursuit l'étude d'une des propriétés empiriques majeures des trades traversants, à savoir leur clustering. Une modélisation des temps d'arrivées des trades traversants par des processus stochastiques auto-excités (les processus de Hawkes) est proposée et sa calibration est validée statistiquement. Ce modèle permet le calcul d'un indicateur d'intensité espérée de l'arrivée future des trades traversants et son utilisation dans une stratégie d'investissement. On s'attache également à obtenir la forme empirique du noyau modélisant la décroissance temporelle de l'impact après un trade traversant, grâce à une méthode de calibration non-paramétrique.

Le chapitre 5 rappelle une méthode générale de détection statistique de sauts dans des séries temporelles de prix/rendements qui soit robuste au bruit de microstructure. On généralise les résultats empiriques connus à de nouveaux indices financiers. On adapte cette méthode à la détection statistique de sauts dans les trajectoires intraday afin d'obtenir la distribution de la proportion de sauts détectés au cours de la journée. On détermine les valeurs extrémales et les plus grandes variations de cette proportion de sauts détectés, et on les relie à la distribution intraday des trades traversants. Enfin, on propose une explication possible des résultats précédents basée sur une modification de la part relative de chacune des composantes de sauts dans la trajectoire des actifs considérés, à savoir la composante liée aux mouvements continus et celle liée aux mouvements de sauts purs.

Cette thèse vise donc à une meilleure compréhension de la microstructure des marchés financiers, via l'étude de ces évènements particuliers au sein des carnets d'ordres que sont les trades traversants. Des statistiques empiriques et leur modélisation par des processus de Hawkes mettent en lumière leurs principales propriétés et permettent des applications liées au lead-lag, par exemple. Enfin, ils aident à l'interprétation des résultats obtenus quand on généralise une méthode statistique connue de détection de sauts au cas de trajectoires intraday.

On donne à présent pour chaque chapitre un contexte, une problématique détaillée et les principaux résultats obtenus.

1.2 Propriétés empiriques des trades traversants et deux applications au lead-lag

On sait depuis les travaux de Bouchaud et al. [2004] et de Lillo and Farmer [2004] que le flux des trades présente une autocorrélation des signes particulièrement forte. Plus précisément, les auteurs ont proposé des modèles de type loi-puissance pour représenter la fonction empirique d'autocorrélation des signes des trades : $C(\tau) \sim \tau^{-\gamma}$, après un temps τ. Les exposants empiriques γ obtenus varient selon les marchés étudiés par ces auteurs, mais restent compris entre 0.2 et 0.7, ce qui traduit des décroissances temporelles très faibles. En d'autres termes, le signe d'un trade conserve un impact significatif sur le signe des trades suivants au bout de durées pouvant s'étendre jusqu'à deux semaines.

Plusieurs explications à ce phénomène ont été proposées (voir la revue de Bouchaud et al. [2009], et particulièrement les travaux de LeBaron and Yamamoto [2007], de Lillo et al. [2005] et de Gerig [2007]). Au final, la liquidité insuffisante qui est présente dans le carnet d'ordres par rapport à celle désirée par les investisseurs, couplée à la volonté de ne pas révéler aux autres participants ses intentions d'investissements afin de ne pas décaler le prix dans le mauvais sens, permettent de rendre compte de ce découpage usuel des ordres en une succession de petits ordres adaptés à la liquidité disponible dans le carnet d'ordres.

Les résultats de Gerig [2007] ont été obtenus grâce à l'utilisation de données incluant des informations sur l'identité des acteurs en présence, non pas l'identité individuelle des investisseurs comme ce serait idéalement le cas, mais le "membership code" identifiant celui qui exécute le trade. Cette information est partielle car à un "membership code" donné correspondent potentiellement plusieurs investisseurs individuels finaux, mais elle s'est avérée suffisante. Les auteurs ont conditionné le calcul de la fonction empirique d'autocorrélation des signes de trades à des trades provenant du même "membership code", de différents "membership codes" ou inconditionnellement. Le conditionnement par un même "membership code" aboutit à une fonction empirique d'autocorrélation des signes de trades un ordre de grandeur plus grand que celle inconditionnelle, tout en conservant une décroissance en loi-puissance avec un exposant plus petit. A l'inverse, quand on conditionne par des trades provenant de "membership codes" différents, la fonction empirique d'autocorrélation des signes de trades décroît très rapidement vers zéro et ne suit plus une loi-puissance. Sous l'hypothèse que les investisseurs n'utilisent qu'un faible nombre de brokers, ces statistiques confirment l'explication précédente due à Lillo et al. [2005].

Pourtant, il arrive parfois que les investisseurs ne puissent pas étendre leur période d'investissements au-delà d'une certaine durée, mais doivent au contraire agir rapidement et sans considération de minimisation de leur impact de marché. Ils ne peuvent se permettre de scanner le carnet d'ordres pour limiter la taille de leurs petits ordres à la liquidité présente en meilleure limite, mais sont alors contraints d'utiliser des trades caractérisés par une consommation particulièrement agressive de la liquidité. Ces trades vident complètement plusieurs limites, sans laisser le temps à la meilleure limite du carnet de se remplir à nouveau par l'arrivée de nouveaux ordres limites. Ce sont ces trades particuliers - les trades traversants - dont on s'applique à déterminer les principales propriétés empiriques dans ce premier chapitre.

Avant cela, rappelons que les propriétés statistiques des carnets d'ordres ont déjà été étudiées dans le détail, par exemple dans Bouchaud and Potters [2003], Bouchaud et al. [2009] et les références qui y sont mentionnées. Ainsi revisite-t-on dans ce chapitre certaines des statistiques déjà bien connues des carnets d'ordres, mais dans le cas particulier du conditionnement aux trades traversants. Par exemple, une question récurrente étudiée dans la littérature concerne l'impact de marché (ie la réponse moyenne du prix aux trades), cf les travaux de Potters and Bouchaud [2003], de Weber and Rosenow [2005], de Almgren et al. [2005], de Hautsch and Huang [2012], de Eisler et al. [2012] et de Cont et al. [2011]. Le conditionnement de la statistique d'impact de marché aux trades traversants sera donc détaillé dans la suite. Notons dès à présent que certains travaux ultérieurs de modélisation microscopique de l'impact de marché effectués par Eisler et al. [2012] ont choisi de doubler la classification usuelle des évènements en trois catégories (trades, ordres limites et annulations d'ordres) par un conditionnement similaire à celui que nous présentons, à savoir si le prix de l'évènement considéré a été modifié. Dans le cas particulier des trades, cela revient à former deux catégories de trades en fonction de leur consommation complète, ou pas, de la liquidité présente en meilleure limite. On comparera plus en détail dans le corps de la thèse les résultats issus de cette définition proche des trades traversants, même si différente.

Les trades traversants sont évidemment reliés aux gros trades et aux grandes variations de prix. Il est à noter que des recherches précédentes s'étaient intéressées à déterminer la cause des grands mouvements des carnets d'ordres (cf Farmer et al. [2004] et Weber and Rosenow [2006]), entre manque de liquidité et grand volume de trades. Farmer et al. [2004] ont relié les grands changements de prix sur le London Stock Exchange à des trades qui consomment complètement le volume disponible en meilleure limite (comme le font les trades traversants), créant ainsi un décalage de prix égal à l'écart entre la meilleure limite et la suivante. Cela indique une contribution importante des fluctuations de liquidité des carnets d'ordres dans les grands changements de prix. Des résultats similaires ont été obtenus par Weber and Rosenow [2006] sur des données issues du NASDAQ et ont confirmé que les fluctuations extrêmes de prix étaient principalement reliées à une faible liquidité dans le carnet d'ordres.

La détermination des relations de lead-lag entre actifs et/ou marchés est d'un grand intérêt pratique. On parle de relation de lead-lag entre deux actifs quand la connaissance des rendements d'un actif donné aide à la prédiction des rendements futurs d'un autre actif. Cette question a reçu une attention particulière dans la littérature récente, comme en témoignent les travaux de Abergel and Huth [2011], de Hoffmann et al. [2011] et les références qui y sont mentionnées. Une approche classique de résolution de ce problème consiste à étudier une fonction de corrélation croisée des rendements laggués, et à mesurer une asymétrie entre les lags positifs et ceux négatifs, ou bien à exhiber un lag non nul qui maximise la corrélation croisée. Notre problématique est quelque peu différente puisqu'on cherche à déterminer quel actif/marché bouge en premier, ie si un mouvement important ayant lieu sur l'actif/marché 1 est, en moyenne, plutôt suivi ou plutôt précédé par un mouvement important sur l'actif/marché 2. On termine ce premier chapitre par une seconde application au problème du lead-lag, plus classique en pratique, qui évalue la force prédictive du signal des traversants dans une stratégie de lead-lag entre actifs.

1.2.1 Résultats obtenus

On réalise une étude approfondie des propriétés empiriques des trades traversants. On s'intéresse d'abord à savoir si les trades traversants représentent une population statistique suffisamment importante pour être étudiée. Concernant leur liquidité, les trades traversants sont des évènements relativement rares, mais ils ne sont pas du tout négligeables puisqu'ils peuvent contenir quasiment 20% du volume traité pour le future sur indice DAX, par exemple. Ils représentent environ une proportion de 5% des trades des stocks du CAC40, ce qui conduit pour BNP-Paribas ou Renault à une occurrence journalière de 250-300 évènements.

Il est intuitivement clair que plus le volume d'un trade donné est important et plus sa probabilité d'être un trade traversant est grande. On vérifie statistiquement cette situation en traçant la courbe de probabilité d'être un trade traversant en fonction du volume du trade, courbe qui est bien empiriquement croissante. Il est également intéressant de noter qu'il existe certains trades de très grands volumes qui ne sont pas des trades traversants, ce qui reflète la présence exceptionnelle en meilleure limite d'une quantité bien supérieure à la normale.

Une propriété empirique importante des trades traversants, sur laquelle nous reviendrons plus en détail dans le prochain chapitre, concerne leur clustering. Les trades traversants ont tendance à arriver par grappes ; autrement dit, on attend en moyenne moins le prochain trade traversant quand on vient d'en observer un, plutôt que quand on vient d'observer un trade quelconque. Et ce résultat est vrai que le délai entre les trades soit mesuré en nombre de trades (temps événementiel) ou en secondes (temps physique).

Les temps d'arrivée des trades traversants ont également tendance à se regrouper à des heures

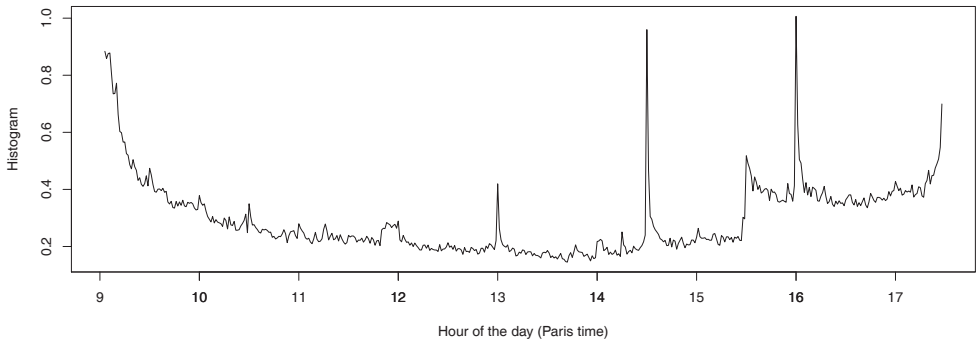

FIGURE 1.1 – Histogramme intraday moyen des temps d'arrivée des trades traversants pour les 30 composantes les plus liquides de l'indice Footsee en 2011 (avec des intervalles de temps de 1 minute).

précises de la journée, comme le révèle leur distribution intraday, particulièrement marquée par la présence de deux pics à des heures précises de la journée (14 :30 et 16 :00 [1]), liés à la tombée de nouvelles macro-économiques US particulièrement suivies, et par une hausse de régime à 15 :30 consécutive à l'ouverture des marchés américains NYSE et NASDAQ. On reproduit à la figure 1.1 l'histogramme moyen des temps d'arrivée des trades traversants des 30 stocks les plus liquides de l'indice Footsee pour illustrer ces résultats. [2]

L'impact de marché des trades traversants est supérieur à celui des trades classiques, et ce même à volume comparable, comme indiqué dans la figure 1.2. En résumé, les trades traversants présentent un contenu informationnel plus grand.

La figure 1.3 montre la relaxation du spread en temps physique après un trade traversant, avec une décroissance en loi puissance d'exposant 0.25 pour le stock BNP-Paribas.

Après ces différentes propriétés empiriques des trades traversants, on présente deux applications qui portent sur le même sujet d'étude (le lead-lag), mais qui diffèrent sensiblement dans leur esprit. La première application vise à utiliser les trades traversants pour répondre à la question de savoir quel actif/marché bouge en premier. Dans cette optique, on commence par définir deux marchés et à en extraire les actifs les plus représentatifs possibles (les plus liquides, par exemple). On met en relation les trades traversants les plus rapprochés entre les deux groupes d'actifs. Puis, on mesure le décalage temporel pour chaque paire de trades traversants ainsi mis en relation, avant d'en tracer la distribution. En comparant les cumulées des parties positives et négatives de la distribution de ce décalage temporel entre les trades traversants des deux groupes d'actifs, on peut déterminer quel marché est en avance sur l'autre. En utilisant cette technique, on montre par exemple que le marché actions US est en avance sur le marché actions européen, et que le stock français Société Générale est en avance sur BNP-Paribas.

Quant à la seconde application, elle mesure la force du signal des trades traversants dans le cadre d'une stratégie d'investissement de type lead-lag. On considère une paire d'actifs donnée qui soit susceptible de former un couple leader-lagger intéressant. [3] On utilise alors les trades traversants

1. heure de Paris

2. La même statistique a été faite sur des stocks européens, dont français et allemands, ainsi que sur des futures européens majeurs avec des conclusions similaires.

3. On peut par exemple déterminer cette paire d'actifs en choisissant celle dont la corrélation des returns laggués est empiriquement élevée.

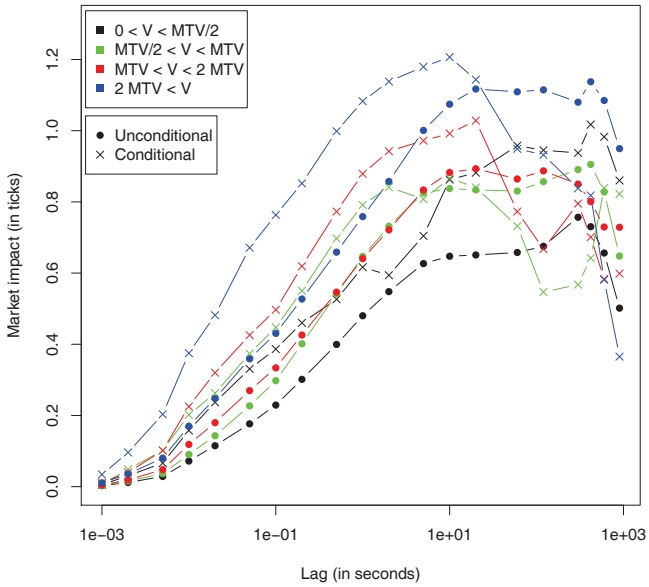

FIGURE 1.2 – Comparaison des fonctions de réponse globale pour le stock BNP Paribas (en échelle log-linéaire) entre tous les trades et les trades traversants, incluant la condition sur le volume initial du trade (ramené au volume de trade moyen, noté MTV).

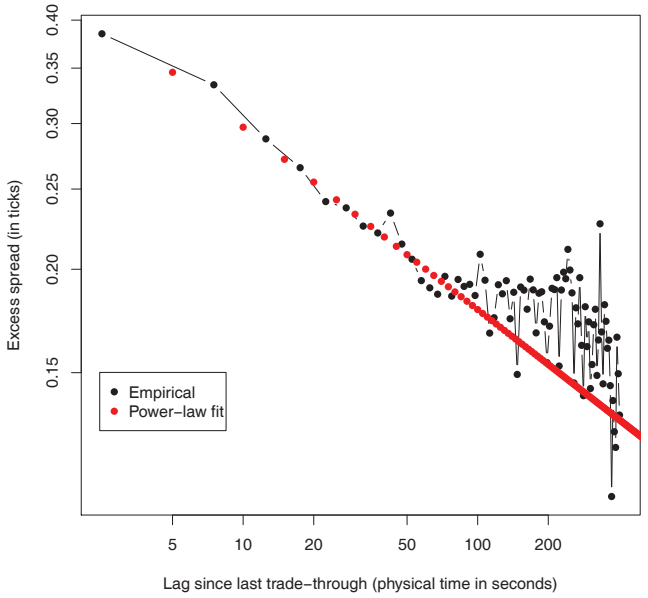

FIGURE 1.3 – Décroissance en loi-puissance de l'excès de spread après un trade traversant sur le stock BNP Paribas (en échelle log-log)

comme un indicateur local d'une relation de lead-lag. La stratégie d'investissement consiste à prendre position sur l'actif lagger après chaque trade traversant sur l'actif leader, et ce dans le même sens que ce dernier, ce qui en fait un exemple de stratégie 'momentum'. Les résultats obtenus montrent des valeurs moyennes de réponse en rendement midprice-to-midprice allant jusqu'à 1.4 bp après 0.5 secondes pour les futures et jusqu'à 1.2 bp après 10 secondes pour les stocks ; aucun signal ne parvient à dépasser les coûts de transaction, ce qui oblige à ne pas limiter la stratégie à l'utilisation d'ordres marché.

1.3 Modélisation des trades traversants par des processus de Hawkes

Les processus de Hawkes sont des processus auto-excités initialement introduits par Hawkes [1971]. On parle de processus auto-excités quand l'arrivée d'un évènement augmente la probabilité d'arrivée d'un nouvel évènement. L'étude de Hawkes [1971] se concentre sur le cas fondamental d'une augmentation instantanée et d'une décroissance exponentielle de la probabilité d'arrivée d'un évènement futur, après l'arrivée d'un évènement donné.

Dans un cas simple (unidimensionnel à paramètres constants avec un unique noyau), l'intensité d'un processus de Hawkes est décrite comme suit :

$$\lambda(t) = \lambda_0 + \int_0^t \alpha e^{-\beta(t-s)} dN_s = \lambda_0 + \sum_{t_i<t} \alpha e^{-\beta(t-t_i)}$$

Le processus $(N_t)_{t\geq 0}$ du nombre cumulé de sauts avant l'instant t est défini par son intensité stochastique $(\lambda(t))_{t\geq 0}$. A l'instant $t > 0$, cette intensité instantanée $\lambda(t)$ est la somme d'une intensité d'arrière-plan $\lambda_0 > 0$ et de la cumulée résiduelle (contre le noyau exponentiel) des hausses d'intensité consécutives aux évènements passés ayant eu lieu aux instants $t_i < t$. L'auto-excitation vient du fait qu'un évènement arrivant à l'instant t^* va provoquer une hausse instantanée de l'intensité d'une valeur $\alpha > 0$. Cette hausse de l'intensité va décroître au cours du temps vers zéro selon le noyau exponentiel $e^{-\beta(t-t^*)}$. Le paramètre $\beta > 0$ gère donc la vitesse de décroissance temporelle de l'impact d'un évènement passé sur l'intensité actuelle.

Si les processus de Hawkes ont trouvé un domaine d'application évident en sismologie (dans la prévision des tremblements de terre [4]), la finance quantitative a également tiré grand profit de l'introduction de ces nouveaux outils de modélisation. Nombre de contributions récentes ont souligné les propriétés intéressantes de ces processus auto- et mutuellement excités. Ils sont, par exemple, à la base des travaux de Large [2007] et de Muni Toke [2011] pour modéliser les temps d'arrivées des ordres dans des modèles de carnet d'ordres. On les retrouve dans l'étude de Bowsher [2007], accompagnés d'un cadre économétrique adapté, pour calibrer des arrivées de trades et de changements de midquotes. Plus récemment, Al Dayri et al. [2011] les utilise pour modéliser des arrivées de trades. Enfin, les travaux de Bacry et al. [2012] montrent qu'ils permettent une modélisation microscopique simple de l'effet Epps [5].

Dans le cadre d'une modélisation des trades traversants, cette auto-excitation des processus de Hawkes incite à les utiliser pour rendre compte du clustering empirique observé. En y ajoutant la possibilité de généralisation au cas multi-dimensionnel, l'interprétabilité physique des différents paramètres utilisés, la disponibilité d'une méthode performante de calibration par maximum de vraisemblance et enfin, le calcul possible par transformée de Laplace d'un indicateur d'intensité

4. où l'on sait bien que la survenue d'un tremblement de terre initial est généralement accompagnée de tremblements de terre ultérieurs, appelés répliques.

5. décorrélation empirique des rendements de deux actifs aux échelles microscopiques.

future espérée, les processus de Hawkes sont donc des candidats privilégiés à la modélisation des trades traversants.

1.3.1 Résultats obtenus

Dans le chapitre 4, on approfondit d'abord l'étude empirique du clustering de l'arrivée des trades traversants. On sait déjà que les trades traversants ont tendance à arriver par grappes. On se demande à présent si ce résultat reste vrai quand on tient également compte dans l'étude du côté du carnet qui est touché. On observe alors qu'un clustering est à l'œuvre si on regarde le même côté du carnet ; il y a bien auto-excitation d'arrivée des trades traversants sur un côté donné du carnet, après l'arrivée d'un trade traversant sur ce même côté. En revanche, l'excitation empirique mutuelle des deux côtés du carnet après un trade traversant est plus faible.

On rappelle ensuite la définition et les propriétés principales des processus de Hawkes utiles à la suite de l'étude (condition de stationnarité, estimation paramétrique récursive par maximum de vraisemblance et tests de la calibration obtenue). Le modèle bidimensionnel d'arrivée des trades traversants au bid et à l'ask que l'on propose est le suivant :

$$\begin{aligned}
\lambda^A(t) &= \lambda_0^A + \int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A + \int_0^t \alpha_{AB} e^{-\beta_{AB}(t-s)} dN_s^B, \\
\lambda^B(t) &= \lambda_0^B + \int_0^t \alpha_{BA} e^{-\beta_{BA}(t-s)} dN_s^A + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B.
\end{aligned}$$

Par exemple, l'intensité du processus d'arrivée des trades traversants à l'ask $\lambda^A(t)$ est somme d'une intensité d'arrière-plan λ_0^A, de la composante d'auto-excitation due aux trades traversants précédents à l'ask ($\int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A$) et de la composante d'excitation mutuelle due aux trades traversants précédents au bid ($\int_0^t \alpha_{AB} e^{-\beta_{AB}(t-s)} dN_s^B$).

La calibration des différents paramètres par maximum de vraisemblance conduit à des valeurs très différentes entre les paramètres d'auto-excitation et ceux d'excitation mutuelle. Si on compare leur influence cumulée, représentée par le ratio $\frac{\alpha}{\beta}$ égal à l'intensité totale pour le noyau exponentiel $\int_0^{+\infty} \alpha e^{-\beta u} du$, on obtient des valeurs moyennes bien plus faibles dans le cas de l'excitation mutuelle ($\frac{\alpha_{AB}}{\beta_{AB}} = 0.017$, $\frac{\alpha_{BA}}{\beta_{BA}} = 0.038$) que dans le cas de l'auto-excitation ($\frac{\alpha_{AA}}{\beta_{AA}} = 0.111$, $\frac{\alpha_{BB}}{\beta_{BB}} = 0.120$).

Ce résultat confirme les statistiques empiriques de clustering précédemment mentionnées et nous amène à comparer le modèle complet à un modèle simplifié, sans excitation mutuelle :

$$\begin{aligned}
\lambda^A(t) &= \lambda_0^A + \int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A, \\
\lambda^B(t) &= \lambda_0^B + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B.
\end{aligned}$$

En suivant la méthode de Bowsher [2007], on détaille deux tests statistiques de validation de la calibration des modèles de Hawkes aux données empiriques. Grâce à une propriété de changement de temps pour les processus de Hawkes, on peut calculer une intensité intégrée empirique, censée décrire les intervalles de temps d'un processus de Poisson homogène. Il reste alors à vérifier si ces intensités intégrées sont bien exponentiellement distribuées (grâce à un test standard de Kolmogorov-Smirnov, par exemple) et si elles sont bien indépendantes (avec un test de Ljung-Box, par exemple).

Il y a donc 4 tests à effectuer au total : deux pour les variables empiriques calculées à partir de la modélisation au bid, de même pour celles issues de la modélisation à l'ask. On montre alors que

ces 4 tests sont tous validés dans 76% des cas, si on inclut l'excitation mutuelle dans le modèle. On reste sinon à un niveau de 71% de cas validés, ce qui confirme la faiblesse de l'excitation mutuelle des trades traversants comparée à leur auto-excitation. Si on offre, en outre, la possibilité aux paramètres λ_0^A et λ_0^B de varier au cours du temps via un profil intraday linéaire par morceaux, on parvient même à valider statistiquement ces 4 tests dans plus de 87% des cas. Dès lors, il n'y a plus aucun apport statistique de la composante d'excitation mutuelle dans la modélisation, par rapport à celle d'auto-excitation.

Il est naturel de vouloir utiliser la modélisation des données financières par des processus de Hawkes dans le cadre de stratégies d'investissements (voir, par exemple, la note de Carlsson et al. [2007], dans laquelle un signal d'achat-vente est basé sur le ratio d'intensités des trades signés). Ici, on profite de la structure exponentielle du noyau pour calculer un indicateur d'intensité future moyenne d'arrivée des trades traversants, dans le cadre du modèle sans excitation mutuelle validé précédemment. Grâce à une transformation de Laplace, on obtient l'expression de l'indicateur recherché, sous la forme d'une somme de l'intensité moyenne de sauts ($\lambda_0 \times \frac{\beta}{\beta-\alpha}$) légèrement perturbée, et d'un terme d'impact dû aux sauts précédents qui décroît exponentiellement vite :

$$E(\lambda(t+\delta)|\mathcal{F}_t^N) = \lambda_0 \times \frac{\beta - \alpha e^{-(\beta-\alpha)\delta}}{\beta-\alpha} + e^{-(\beta-\alpha)\delta} \int_0^t \alpha e^{-\beta(t-u)} dN_u$$

On montre que l'utilisation de cet indicateur permet une amélioration sensible du rendement moyen obtenu dans une stratégie simple de type 'momentum' basée sur les trades traversants, ce qui confirme l'intérêt d'une modélisation permettant de rendre compte du clustering de ces évènements.

On s'assure également que la classe de modèles proposée précédemment à la calibration, et qualifiée de "Generalized linear models" dans Bowsher [2007], est adaptée et qu'il n'existe pas d'effet inhibiteur après l'arrivée d'un trade traversant. On pourrait en effet imaginer une situation où l'arrivée d'un type d'évènements donné diminue la probabilité d'occurrence d'un autre type d'évènements, phénomène qui ne pourrait être intégré au cadre de modélisation précédent. La calibration de deux modèles d'intensité inspirés des processus de Hawkes et permettant la mise en lumière d'un effet inhibiteur (un modèle log-linéaire et un modèle à seuil) montre que cet effet n'est pas significativement présent dans les données de trades traversants.

Enfin, on prolonge l'étude de la forme empirique du noyau de décroissance temporel d'impact. A cette fin, on utilise une calibration non-paramétrique de type EM ("Expectation Maximization") pour les processus de Hawkes, issue des travaux de Marsan and Lengline [2008] et appelée MISD ("Model Independent Stochastic Declustering"). Celle-ci met en évidence, sur la figure 1.4, une décroissance plus forte que la loi exponentielle précédemment utilisée, et davantage proche d'une loi-puissance.

1.4 Comparaison entre les trades traversants et la méthode statistique de Aït-Sahalia et al. [2012] pour la détection de sauts dans les trajectoires intraday

De nombreux problèmes provenant de l'industrie financière, allant de la modélisation d'actifs financiers à leur couverture, en passant par le gestion des risques, sont fortement dépendants de la présence de sauts dans les actifs considérés. Cela explique probablement qu'une question majeure dans la littérature de la finance quantitative soit de déterminer, à partir d'observations empiriques de trajectoires, si ces données sont compatibles avec un processus sous-jacent continu, ou bien si une composante de sauts doit être incorporée à la modélisation. Dans la réalité, on ne dispose que

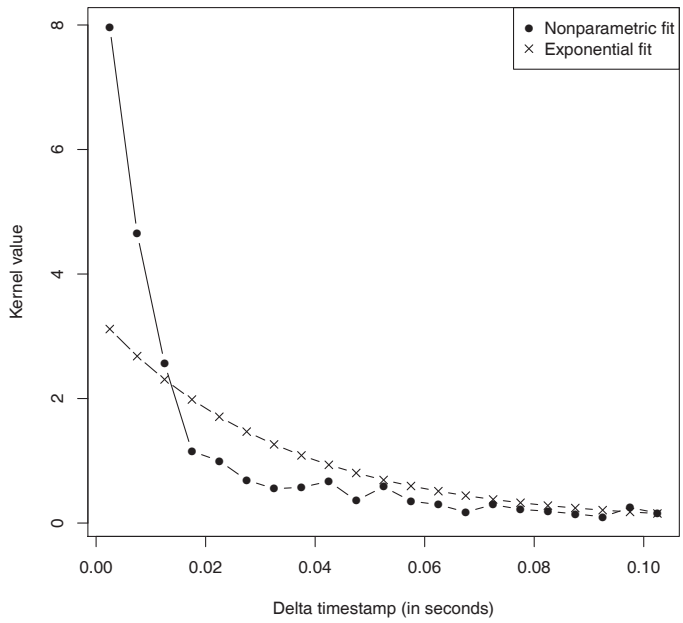

FIGURE 1.4 – Comparaison des résultats, entre l'algorithme MISD et celui du maximum de vraisemblance, pour la calibration empirique du noyau des trades traversants du côté bid du carnet d'ordres, pour le stock BNP Paribas entre Juin et Octobre 2010 (10 :00 à 14 :00).

d'observations discrètes qui sont composées entièrement de sauts, par construction. Il faut alors dire si ces sauts empiriques sont dûs à l'échantillonnage d'un processus continu, ou s'ils reflètent la présence de vrais sauts de l'actif.

Plusieurs approches ont été proposées pour détecter la présence de sauts dans les trajectoires discrètes observées sur les marchés financiers. Aït-Sahalia [2002] teste une inégalité de continuité basée sur la fonction de transition d'un processus supposé suivre une diffusion. Le test de Carr and Wu [2003] est basé sur une propriété des options maturant sous peu ; quand la durée avant échéance de l'option décroît vers zéro, la valeur-temps de l'option converge vers zéro à une vitesse dépendant à la fois de la structure du processus du sous-jacent (en clair, de la présence potentielle de sauts) et de la position par rapport à la monnaie (dans la monnaie ou hors de la monnaie). Barndorff-Nielsen and Shephard [2004] ont initié une nouvelle série de tests reposant sur des généralisations de la variance réalisée, appelées variations bi-puissances et, plus généralement, variations multi-puissances. Jiang and Oomen [2008] ont développé un test qui s'appuie sur l'erreur de réplication cumulée d'un swap de variance couvert de manière discrète. Lee and Mykland [2008] et Lee and Hannig [2010] calculent, quant à eux, un quotient de test comparant les rendements et les volatilités instantanées (estimées avec des variations bi-puissances réalisées pour l'un, et des variations quadratiques tronquées pour l'autre). Enfin, Aït-Sahalia and Jacod [2009] développent un test à partir du quotient de variations puissances échantillonnées à des fréquences différentes.

Le cadre d'estimation usuel en statistique assure une meilleure convergence des estimateurs quand la fréquence d'échantillonnage augmente. Mais les données financières ont ceci de particulier qu'une augmentation de la fréquence d'échantillonnage vers le domaine de la haute-fréquence amène

la présence d'un bruit statistique, appelé bruit de microstructure et dû, entre autres, à la discrétude des changements de prix et aux rebonds de prix entre le bid et l'ask. Utilisés dans ce cas précis, les tests précédents perdent donc de leur puissance statistique quand on augmente trop la fréquence d'échantillonnage, à cause de ces nouvelles erreurs de mesures.

Les problèmes liés au bruit de microstructure ont été étudiés dans la littérature à l'aide de deux modèles principaux, celui des erreurs additives et celui des erreurs d'arrondi (cf Aït-Sahalia et al. [2012] et les références qui y sont mentionnées, pour plus de détails). Les tests précédents ont été robustifiés afin de circonvenir ce problème, principalement par le biais de méthodes de pré-moyennes et de post-moyennes des données et/ou des mesures statistiques. Dans la suite de l'étude, on se focalise sur une méthode de test aux hypothèses particulièrement générales, qui a été introduite dans Aït-Sahalia and Jacod [2009] et robustifiée dans Aït-Sahalia et al. [2012] pour fonctionner en présence de bruit de microstructure.

Le chapitre 5 présente d'abord les travaux de Aït-Sahalia and Jacod [2009] et Aït-Sahalia et al. [2012] qui décrivent une méthode générale de détection statistique de sauts dans des séries temporelles de prix/rendements qui soit robuste au bruit de microstructure.

Dans le cas simple de la détection de la présence de sauts en absence de bruit de microstructure, on considère un processus de log-prix X très général, dont on observe à intervalle régulier $\Delta_n = \frac{T}{n}$ les accroissements discrets (ie les log-rendements) : $\Delta_i^n X = X_{i\Delta_n} - X_{(i-1)\Delta_n}$. Pour $p > 0$, on définit l'estimateur $\hat{B}$ à partir des p-variations des accroissements à l'échelle Δ_n : $\hat{B}(p, \Delta_n)_t = \sum_{i=1}^{[\frac{t}{\Delta_n}]} |\Delta_i^n X|^p$. Le test statistique utilisé pour discriminer entre la présence et l'absence de sauts est obtenu en calculant le quotient de ces variations puissances échantillonnées à deux fréquences différentes : $\hat{S}(p, k, \Delta_n)_t = \dfrac{\hat{B}(p, k\Delta_n)_t}{\hat{B}(p, \Delta_n)_t}$, pour $p > 2$ et $k \geq 2$. On dispose de deux convergences en probabilité quand $\Delta_n \to 0$: $\hat{S}(p, k, \Delta_n)_t \to 1$ si X a des sauts, et $\hat{S}(p, k, \Delta_n)_t \to k^{p/2-1}$ si X est continu. On a également un estimateur $\hat{V}^c_{n,t}$ de la variance asymptotique du test, sous l'hypothèse de continuité. Au final, le théorème de limite centrale suivant permet de conclure quant à la question posée : si X est continu et $p \geq 2$, les variables $\dfrac{\hat{S}(p, k, \Delta_n)_t - k^{p/2-1}}{\sqrt{\hat{V}^c_{n,t}}}$ convergent en loi vers une variable normale centrée réduite, conditionnellement à la filtration historique.

La présence du bruit de microstructure détériore la puissance discriminante du test précédent en faisant converger le quotient test $\hat{S}(p, k, \Delta_n)_t$ vers la même valeur dans les deux cas (présence ou absence de sauts). Les auteurs réalisent donc des pré-moyennes des données grâce à un noyau, afin de gérer le bruit de microstructure (dans un modèle valide jusqu'à une fréquence d'environ une donnée toutes les 5 secondes). Les accroissements du processus X sont remplacés par une version localement moyennée des accroissements du processus bruité Z. Combiné à un estimateur local de la variance du bruit de microstructure, ils en déduisent une version robustifiée de l'estimateur des variations p-puissances $\bar{V}(Z, g, p)^n_t$, ainsi qu'une version robustifiée du quotient test : $S_{RJ}(g, h, p)_n = \dfrac{\bar{V}(Z, g, p)^n_T}{\gamma' \bar{V}(Z, h, p)^n_T}$. Ce nouveau quotient test possède des propriétés de convergence similaires au précédent. On peut également estimer sa variance asymptotique $\Sigma^c_{R,J,n}$ sous l'hypothèse de continuité, pour en tirer un théorème central limite : le test statistique standard t défini par $t = \dfrac{S_{RJ}(g, h, p)_n - \gamma''}{(\Delta_n)^{1/4}\sqrt{\Sigma^c_{R,J,n}}}$ converge, sous l'hypothèse de continuité du processus X, en loi vers une variable gaussienne centrée réduite, conditionnellement à la filtration historique.

Au final, la validité de l'hypothèse de continuité est caractérisée par une distribution du test t proche de celle d'une loi normale centrée réduite, alors qu'une distribution décalée vers des valeurs négatives indique la présence de sauts. Empiriquement, plus la distribution de la statistique standard de test t est décalée vers la gauche, plus les sauts sont présents dans les données.

On veut à présent utiliser cette méthode de détection de sauts et la relier aux trades traversants, qui sont des sauts microscopiques détectés à partir des données tick-à-tick. Quelles sont les similarités et les différences de ces deux méthodes de détection de sauts, quand on les teste sur des données empiriques ? Plus précisément, peut-on relier et comprendre les changements dans la proportion de sauts détectés au cours de la journée par la méthode de Aït-Sahalia et al. [2012], à la lumière des variations dans la présence des trades traversants ?

1.4.1 Résultats obtenus

On confirme d'abord avec nos données les résultats empiriques obtenus par ces auteurs sur la détection de sauts pour l'indice DJIA, puis on les généralise aux indices Footsee et DAX. On décrit empiriquement comment ce test se comporte à plus grande fréquence d'échantillonnage dans le cadre de la détection statistique de sauts intraday. La figure 1.5 confirme que le passage de la détection de sauts dans des trajectoires journalières à des trajectoires de 15 minutes décale significativement la moyenne du test vers des valeurs plus positives (signe d'une plus faible proportion de sauts détectés).

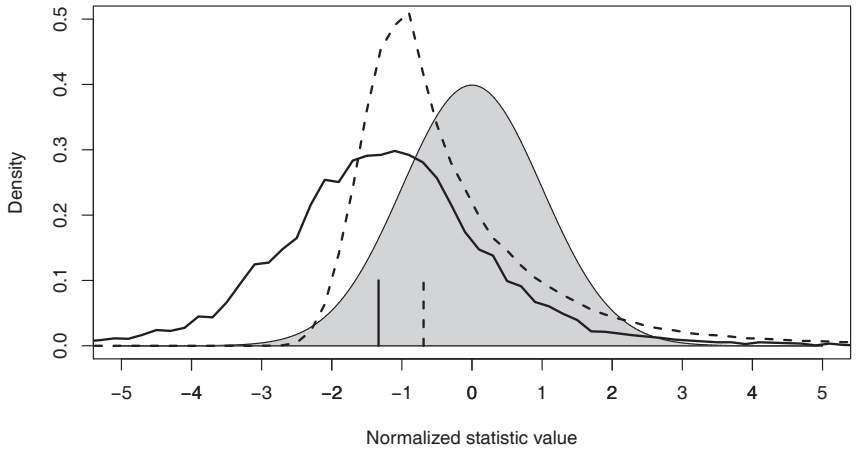

FIGURE 1.5 – Comparaison de la distribution de la statistique de test normalisée t avec $N_{\text{daily}} = 10000$, entre $T = 1$ jour (trait plein) et $T = 15$ min (pointillé), avec une distribution normale standard (zone grisée). Une ligne verticale indique la médiane de la distribution.
Etude menée sur l'indice Footsee (future et les 30 stocks les plus liquides) en 2011.

La figure 1.6 met en relation la distribution intraday de la médiane du test t avec celle des temps d'arrivée des trades traversants. Les boxplots de la distribution du test statistique standard t pour le Footsee montrent un intervalle de variation de la médiane au cours de la journée assez faible (de l'ordre d'un tiers de l'écart interquartile), ce qui nous invite à concentrer l'étude sur les

valeurs extrêmes (minimum et maximum) et sur les grandes variations intraday de cette médiane. On observe en particulier que le minimum de cette médiane est obtenue à 14 :30 (signe de la plus grande proportion de sauts) et son maximum à 15 :30 (signe de la plus faible proportion de sauts). Les plus grands mouvements intraday de la médiane ont lieu à 15 :30 (atteinte d'un minimum global de la proportion de sauts détectés) et à 16 :00 (atteinte d'un maximum local de la proportion de sauts détectés). On sait également que les horaires de 14 :30 et 16 :00 correspondent aux grands pics de la distribution intraday des trades traversants, et qu'à 15 :30 a lieu un changement de régime à la hausse du nombre de trades traversants.

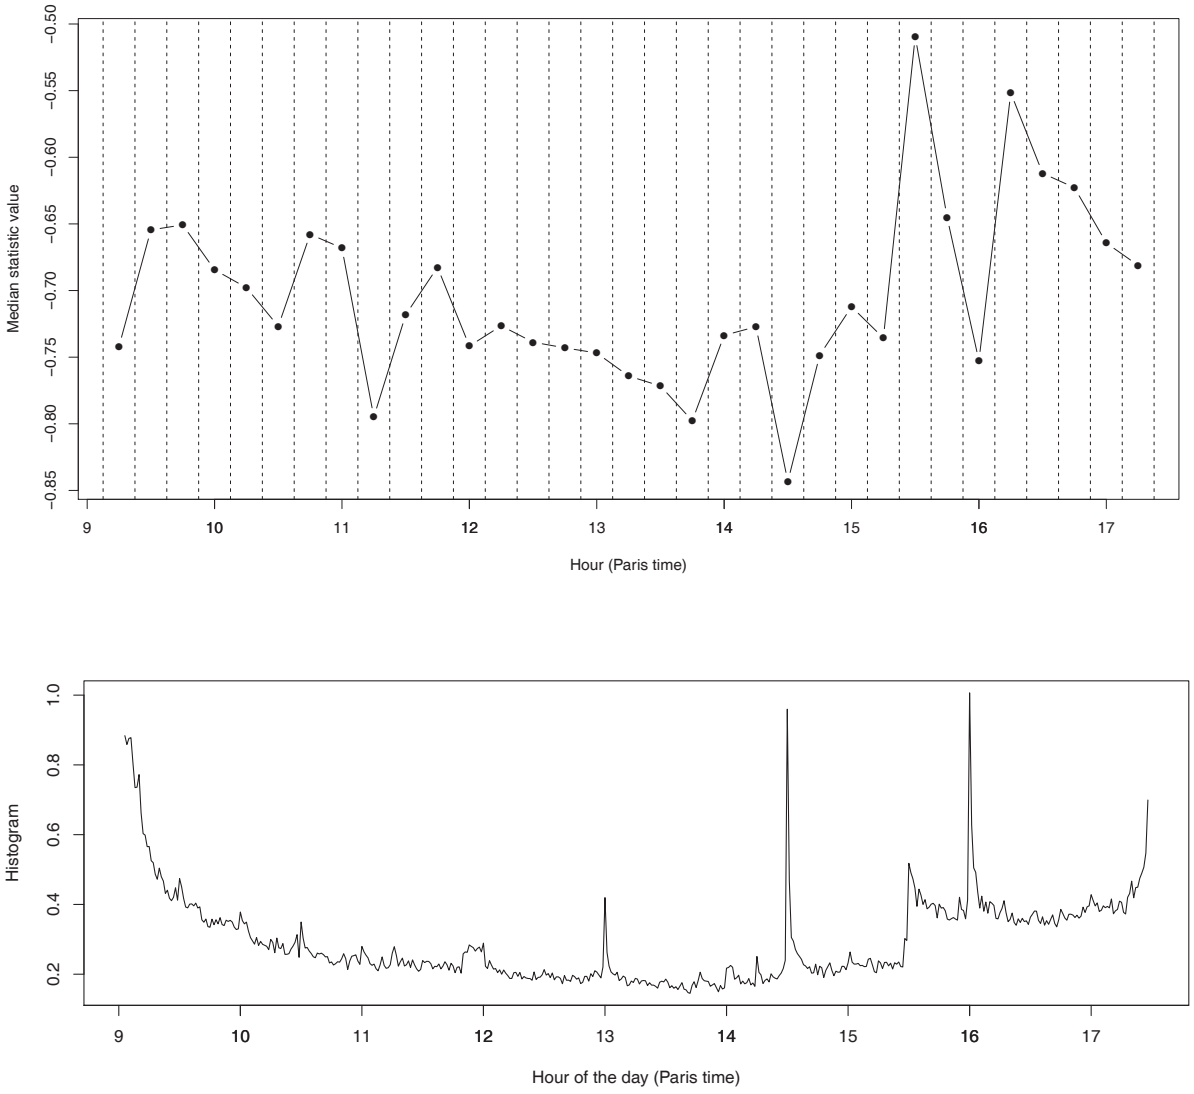

FIGURE 1.6 – *(Haut)* Variations intraday de la médiane du test t avec $N_{\text{daily}} = 10000$ et $T = 15$ min ; *(Bas)* Histogramme moyen des temps d'arrivée des trades traversants avec des intervalles de temps de 1 minute. Etude menée sur l'indice Footsee (future et les 30 stocks les plus liquides) en 2011.

On rappelle que les mouvements globaux des actifs étudiés sont supposés être la somme de

mouvements continus (dûs aux variations browniennes pondérées d'un facteur volatilité σ_t) et de mouvements de sauts purs. On propose à présent une explication possible des résultats précédents pour les variations intraday dans la détection de sauts par le test statistique à l'aide des trades traversants. A partir de 15 :30, le changement de régime à la hausse du nombre moyen de trades traversants se reflète dans une hausse du facteur volatilité σ_t, ce qui permet à davantage de mouvements de l'actif d'être considérés comme des mouvements acceptables de la composante continue, d'où une plus faible proportion de sauts détectés après 15 :30. A 14 :30, le pic particulièrement fort et bref du nombre de trades traversants est source d'une forte hausse de la composante de sauts purs, d'où le maximum intraday de la proportion de sauts détectés. Enfin, à 16 :30, ce pic similaire du nombre de trades traversants aurait également pu conduire à un maximum intraday de la proportion de sauts détectés, mais le facteur volatilité σ_t plus élevé après 15 :30 amène certains de ces sauts à être acceptés comme mouvements valides de la composante continue, d'où un maximum de la détection de sauts à 16 :00 moins marqué que celui de 14 :30.

Chapitre 2

Introduction

The current availability of huge financial databases including all transactions and limit order books of a large number of financial assets lead to an important footstep in the comprehension of financial markets at the microstructure level. Instead of directly modelling price/returns series with stochastic processes, the different microscopic phenomena at the origin of the price formation and modification were investigated. This systematic study of order books data highlighted some fundamental empirical properties about the interconnection between liquidity supply and consumption in the price formation process. Of course, many open questions remain to access to a complete understanding of the mechanisms at play, including liquidity supply, liquidity consumption, order books dynamics, market-impact, volatility and the different behaviours of the actors involved. But, this domain reached a maturity level, as indicated by the increasing number of reviews available (see for example those of Bouchaud et al. [2009] and Chakraborti et al. [2011]). Consequently, there is now a research for new development fields in this domain, and two of them seem to emerge : empirical studies based on data that include the identification of the different actors involved, and empirical studies of particular events meant to be the most significant and interesting possible. To sum up, global statistics about financial market microstructure already permit a good comprehension of this domain, but it seems that conditional statistics (on the identity of the actors, or on particular events) are opening a particularly interesting and promising path.

This thesis is part of this global effort for a better understanding of the financial markets microstructure based on the study of particular events, which are, in our case, related to an aggressive consumption of the liquidity available in the limit order books.

2.1 A reminder on financial markets microstructure

We quickly review the main characteristics of financial markets microstructure that will be useful in the sequel. We only consider here the case of continuous trading double auction markets. A market is continuously trading when any market participant (also called a trader) can enter and exit the market whenever he wants to. In double auction markets, the investors have access to three main order types : limit orders, market orders and cancelations. The limit order book is the structure in which the different limit orders are stored. Inside the limit order book, the interactions between those three types of orders give birth to the different microstructure prices.

Limit orders are used by patient traders who want to contribute to the liquidity present in the limit order book, by offering the possibility to the other participants to buy or sell a given quantity of the asset at a given price. Limit orders are sorted by ascendant (respectively descendant) prices for the sell (respectively buy) limit orders. Inside the queue of orders with a given price, they are

sorted by time precedence. Buy limit orders are called bids and sell limit orders are called asks. The minimum (respectively maximum) price at which one can sell (respectively buy) using a limit order is called the best ask (respectively bid). The difference between the best ask and the best bid is strictly positive and is called the spread. The average between the best prices is called the midprice. The limit orders prices can only be chosen on a grid of prices, all multiples of the minimal price increment called the tick.

On the contrary, impatient traders use market orders, which are orders to buy or sell a given quantity of the asset at the best price immediately available in the limit order book. A trade occurs when a market order sent by a trader consumes the liquidity available in the limit order book, by being matched with the limit orders of the opposite side. The matching algorithm of limit orders is simple : the limit orders consumed first by a sell (respectively buy) market order are the bids (respectively asks) with the biggest (respectively smallest) prices and, for a given price, the one which arrived first. The total volume available at the best limit may be completely consumed by an incoming trade and this will change the best price. But if the volume of the trade is smaller than the liquidity available at the best limit, the best price will not be changed.

The liquidity contributors are also given the right to remove an existing limit order from the book with a cancelation. Consequently, the limit order book and the price are completely determined by the arrival of limit orders, market orders and cancelations. Limit orders increase the liquidity of the order book, whereas market orders and cancelations consume it.

When a trader wants to quickly buy a large given quantity of an asset, he may post a buy market order of the corresponding size. The asks of the limit order book will be matched with this market order until the given quantity of asset is reached to fill the market order. This may consume several best limits of the order book and move the best ask upwards. After the trade, other traders may want to post new limit orders at the previous best ask. They may also consider that this trade contained important information about the stock, and they will also buy the stock considering that the price will increase in the future. So, a market order may move the price of an asset, due its own impact on the limit order book, and also to the impact of the reactions of other traders to this initial market order. The total impact following a market order is called the market impact.

2.2 Outline of the thesis

In the chapter 3, we first recall some empirical properties essential to the comprehension of financial markets microstructure and to the definition of our study. The insufficient liquidity available in the order book with respect to what is needed by the investors, added to the will to hide one's investments intentions to the others investors in order not to move the price in the wrong direction, directly lead to the usual splitting of market orders in a series of small orders adapted to the liquidity available in the limit order book. However, speed is sometimes more important than minimizing market-impact. At the heart of this trade-off, trades-through are precisely the trades that stand outside the usual trading pattern. In this chapter, we give empirical statistics about them and we use them in two different lead-lag applications.

The first results of this chapter (from 3.1 to 3.6) have been presented at 'Econophys-Kolkata V conference', published in its proceedings 'Econophysics of order-driven markets' at Springer Verlag, co-authored with Frédéric Abergel, and submitted as an article for publication with Frédéric Abergel as co-author. The last part of the chapter (3.7 on lead-lag trading signal) has been presented at 'QMF 2011 Conference' in Sydney.

The chapter 4 goes into more detail about the trades-through arrival times clustering, which is one of their main empirical properties. We model their arrival times with self-excited stochastic processes (Hawkes processes) and this calibration is statistically validated. This model allows to compute an intensity indicator based on trades-through that will be used in a simple trading strategy. A non-parametric calibration of the empirical decay kernel for the temporal impact of trades-through indicates a decrease faster than exponential, and closer to a power-law.

The first results of this chapter (from 4.1 to 4.4) have been published as an article, co-authored with Ioane Muni-Toke, in 'Economics : The Open-Access, Open-Assessment E-Journal', Vol. 6, 2012-22, and presented at the conference 'Fourth European summer school in financial mathematics' in Zürich, 2011. The last results (4.4, 4.5 and 4.7) have been presented by Ioane Muni-Toke at 'QMF 2011 Conference' in Sydney.

The chapter 5 recalls a general statistical method robust to market microstructure noise to find jumps in prices/returns time series. We generalize the empirical results already known in the literature to new financial indices and we adapt this statistical jump detection method to intraday trajectories in order to obtain the intraday proportion of detected jumps. Extreme values and biggest intraday variations of this jump proportion occur at very specific hours of the day (14 :30, 15 :00 and 16 :30, Paris time reference), already linked with trades-through arrival times. Using trades-through, we offer a possible explanation of the main characteristics of the intraday proportion of detected jumps using a modification in the relative importance of each jump component of the assets trajectories (the continuous moves component and the pure-jumps component).

This chapter will hopefully be submitted as an article for publication soon, with Frédéric Abergel as co-author.

To sum up, this thesis is an attempt to contribute to a better comprehension of the microstructure of financial markets, with the study of particular events happening in the limit order books : the trades-through. Empirical statistics and their modelling by Hawkes processes reveal their main properties and open the way to lead-lag applications, for example. Finally, they help to provide some understanding of the results obtained when generalizing a statistical method for detecting jumps inside intraday trajectories.

Chapitre 3

Trades-through : empirical facts and applications to lead-lag

Résumé

Order splitting is a standard practice in trading : traders constantly scan the limit order book and choose to limit the size of their market orders to the quantity available at the best limit, thereby controlling the market impact of their orders. In this article, we focus on the other trades, multiple-limits trades that go through the best available price in the order book, or "trades-through". We provide various statistics on trades-through : frequency, volume, intraday distribution, market impact and clustering. We present a new method for the measurement of lead-lag parameters between assets, sectors or markets, and we measure the trades-through intensity signal in a simple trading strategy based on lead-lag.

Contents

3.1 Introduction 34
3.2 Data presentation 36
3.2.1 General TRTH data presentation and processing 36
3.2.2 Data used in this study 36
3.3 Elementary statistical properties of trades-through 37
3.3.1 On which limit of the order book is liquidity taken from ? 38
3.3.2 Frequency and volume 38
3.3.3 Links between trades-through and large trades 41
3.3.4 Clustering 42
3.3.5 Intraday distribution 43
3.4 Market impact 44
3.4.1 Empirical response function 46
3.4.2 Response function conditioned on trade volume 46
3.5 Spread relaxation 47
3.6 Lead-lag estimation from trades-through time-series 49
3.6.1 The estimation technique 50
3.6.2 Empirical results 51
3.7 Empirical lead-lag signal from trades-through 52
3.7.1 Data description and summary statistics 53
3.7.2 Empirical lead-lag signal based on trades-through 54
3.8 Conclusion and further research 54

3.1 Introduction

It is a well-documented fact that the order flow is a highly autocorrelated long-memory process, see e.g. Bouchaud et al. [2004], Lillo and Farmer [2004] [1]. In Lillo and Farmer [2004], the empirical autocorrelation function of trades signs for 20 stocks traded at the London Stock Exchange is fitted to a power-law $C(\tau) \sim \tau^{-\gamma}$ with exponent $\gamma = 0.6$. On the Paris Stock Exchange, Bouchaud et al. [2004] measure power-law exponents γ ranging from 0.2 to 0.7. In other words, the sign of a trade has an impact on futures trades signs, with levels that stay statistically significant over a period of time as long as two weeks.

Different models have been suggested in the literature to explain this phenomenon, see the review in Bouchaud et al. [2009]. LeBaron and Yamamoto [2007] consider an evolutive market of heterogeneous investors allowed to learn and adapt their trading, and manage to replicate the long memory of order signs thanks to the imitative behavior of investors. Lillo et al. [2005] argue that long memory is mainly due to a delay in market clearing. A trader facing a large order will split it in several orders for two main reasons : the available liquidity in the order book may not be sufficient and, even if it were, revealing her/his intention to the market will cause the price to move too much in the wrong direction.

Empirical statistics that discriminate between those two explanations are not easy to make, because of the lack of available data on individual investors. However, some exchanges (like the London Stock Exchange, the Spanish Stock Exchange and the Australian Stock Exchange) provide a partial information (the membership code) that identify the member of the exchange who executes the trade. Gerig [2007] computes the empirical autocorrelation function of orders signs considering market orders coming from the same membership code, different membership codes, and all market orders. When conditioning on the same membership code, the autocorrelation function is one order of magnitude bigger than the unconditional one, and is also decaying as a power-law, but with a smaller exponent. On the contrary, when conditioning on different membership codes, the empirical autocorrelation function of orders signs is fast decaying to zero and is not a power-law anymore. If we assume that most investors trade only through a small number of brokers, these statistics support the explanation of Lillo et al. [2005] rather than that of LeBaron and Yamamoto [2007] : the long-memory of orders signs is due to delay in market clearing.

It is really worth noticing that some recent work presented in Toth et al. [2011] has shown that the autocorrelation function of the order flow can be decomposed into the sum of one splitting and one herding components, allowing to directly compare the relative strength of those two effects. The empirical results presented confirmed that the splitting of orders by individuals investors was by far the main cause for the long-memory of the order flow.

So, in practice, the splitting of orders is an important pattern of market microstructure : traders wanting to trade a large order constantly scan the limit order book and split their orders to restrict their size to the quantity available at the best limit. However, speed is sometimes more important than minimizing market-impact. In this article, we study trades-through, which are precisely the trades that stand outside the usual trading pattern.

The statistical properties of limit order books have been extensively studied, see for example Bouchaud and Potters [2003], Bouchaud et al. [2009] and the references therein. We will revisit some well-known statistics of limit order books, albeit restricted to trades-through. For example, we will study whether market impact (the average response of prices to trades), which has received considerable attention over the last years, see e.g. Potters and Bouchaud [2003], Weber and Rosenow [2005], Almgren et al. [2005], Hautsch and Huang [2012], Eisler et al. [2012], Cont et al. [2011], changes when one considers the response to trades-through.

For obvious reasons that become even clearer in the bulk of the paper, trades-through are

1. A centered process X is said to exhibit long-memory behavior when its series of auto-covariances is not summable, i.e. $\sum_{h\in\mathbb{N}} |\gamma(h)| = +\infty$ where $\gamma(h) = E[X_t X_{t-h}]$

naturally related to large trades and large price moves. Previous research works such as Farmer et al. [2004] or Weber and Rosenow [2006] investigate large moves in limit order books, mainly focusing on their cause : are they due to large trade volume, or lack of liquidity ? Farmer et al. [2004] found that large price changes on the London Stock Exchange occur when a market order removes all the volume at the best limit (as trades-through do), thereby creating a change in the corresponding best price equal to the size of the first gap. This indicates that a major contribution of large price changes is due to fluctuations in the liquidity inside the limit order books. Similar results have been obtained by Weber and Rosenow [2006] on NASDAQ data and confirmed that extreme price fluctuations were mainly caused by a small liquidity in the limit order book. In this paper, we somehow make this interpretation more complete, by studying the rather peaked distribution of the occurrences of trades-through, connecting them with exogenous events such as market opening and closing or economic news announcements.

Also of great practical interest is the study of lead-lag relationships between two assets or market places. A lead-lag relationship occurs when knowledge of the return of one asset helps predict the future return of a second asset. Lead-lag relationships have received much attention in the recent literature, see Abergel and Huth [2011] and Hoffmann et al. [2011] and the references therein. A standard approach is to measure the lagged cross-correlation of the returns, and to study a possible asymmetry between positive and negative lags, or even, to find a non-zero lag that maximizes the cross-correlation. Our approach is different, and we believe it is new. We try and answer the following question : which of the two assets moves first ? In other words, when an important move occurs on Asset 1, is it followed or is it caused by an important move on Asset 2 ? When focusing only on trades-through as we do, and due to the relative rarity of such events, we will show that one can answer that question in a statistically significant way.

The paper is organized as follows : in Section 2, we give a general presentation of the TRTH database used in this article. We also present the set of data we focus on (major US and EU equity futures and major French stocks). In Section 3, we give a precise definition of a trade-through and study some elementary statistical properties [2] : arrival time distribution, volume, seasonality, clustering. Section 4 is devoted to the study of the market impact of trades-through. In Section 5, we examine the behaviour of the spread after a trade-through, showing a typical power-law relaxation of the excess spread. Section 6 deals with the characterization of lead-lag relationships between US and EU equity markets and between pairs of French stocks of the same sector. Finally, in Section 7, we measure the trades-through intensity signal in a simple trading strategy based on lead-lag.

2. All computations are made using the free statistical software R, available at http ://cran.r-project.org.

Flag	Information and signification
normal	Trades occurring during the continuous trading session.
auction	Trades occurring at the end of the auction phase.
OTC	Over-the-counter trades occurring directly between two parties. They are opposed to trades occurring on centralized exchanges.
off book	Trades occurring outside the usual trading system of the considered exchange. They may be trades reported from a broker, for example. In Euronext, trades outside of the NSC (Nouveau Système de Quotation) fall inside this category.
block trade	Reporting of block trades (very large trades).
rck	TRTH/exchange's threshold break alerting that price and/or volume of this trade (as reported by the exchange) seem too different from an usual behavior and should not be considered as relevant.
market closed	Trades occurring before or after the regular trading session.
cancelled	Cancelled trades.
late0day & lateNdays	Trades that were reported the same day or N days after.
late report	Trades reported later in market data feeds. They include late0day and lateNdays trades.
unknown	All the other trades.

TABLE 3.1 – TRTH's most important trades flags

3.2 Data presentation

3.2.1 General TRTH data presentation and processing

The data used in this study come from the TRTH (Thomson Reuters Tick History) databases. There are two different databases, one for the quotes (grouped in the 'Quotes' file) and one for the trades (grouped in the 'Time And Sales' file). Both quotes and trades are timestamped in milliseconds by Thomson-Reuters timestamping machines located in London. Quotes entries are composed of Bid/Ask/BidSize/AskSize. Trades entries contain Price/Volume of each transaction.

An important point to be mentioned in the data presentation is that TRTH data are flagged. Each entry of both quotes and trades files has a flag indicating information to be taken into account in the data analysis. Those flags are market- and exchange-dependent in the sense that specific knowledge of each market and exchange is necessary to correctly interpret each TRTH/exchange flag. After processing the flags [3], we end up with trades tagged within a limited number of flags' categories. The most important trades flags (normal, auction, OTC, offbook, block trade, rck, market closed, cancelled, late0day & lateNdays, late report) are detailed in the table 3.1.

Note that some data sent by exchanges are corrections of previous entries, such as the cancellation of a previous trade and its replacement by another one. Corrections are the only case when TRTH data are modified at BNP Paribas before being made accessible to users.

3.2.2 Data used in this study

Tick-by-tick market data used in this study are the data after all corrections have been taken into account. We consider only trades that are flagged as 'normal' trades. In particular, we do not consider any block-trade or off-book trade. Moreover, we restrict ourselves to data coming from a

3. The flag processing is done by BNP Paribas Equities & Derivatives Quantitative R&D Histo team. It is a reliable process including, for example, matching with data published by the exchanges (such as extreme prices -open, high, low, close- and total volumes).

single venue, the main exchange where the considered assets are traded. [4]

EU-US equity futures

Here is how we choose the data for the comparison of the US and EU equity market : first, we select the most representative instruments in the US and EU equity markets. To this end, we rank all equity financial instruments available in TRTH according to their ADV (Average Daily Volume), and then pick the most liquid ones (3 from US equity markets and 3 from EU equity markets). By doing so, we end up with a small number of financial instruments representing the most liquid instruments of the markets under scrutiny. This choice is based on the rationale that market moves are first seen in the most liquid instruments, as those most liquid assets tend to incorporate information faster.

This selection process provides a set composed of E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures (for US equity markets) and Eurostoxx, DAX and Footsie futures (for European equity markets). These assets trade on the CME Exchange (for E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures), the Eurex Exchange (for Eurostoxx and DAX futures) and the NYSE Liffe Exchange (for the Footsie future). Moreover, in order to consider the most liquid instruments, we focus on the futures with the nearest maturity.

For the basic statistics on frequency and volumes of trades-through, we use the data of March 2010 (from 16 :00 to 21 :30, Paris Time reference). For the intraday distribution, we focus on the first half of March 2010 (to avoid difference in daylight saving time between Paris and the USA). For the lead-lag study, we use data from the beginning of December 2009 to mid-March 2010 and restrict our data time-frame to the period of the day when both EU and US equity markets are open and widely trading (from 15 :30 to 17 :30, Paris Time reference), as the lead-lag phenomenon is particularly relevant at that moment.

French stocks

For French stocks, the assets we select are BNP Paribas, Société Générale, Renault and Peugeot during March 2010, with a daily time frame from 9 :30 to 17 :00, Paris Time reference, so as not to be impacted by auction phases.

Data summary table

We present in table 3.2 a summary of the different data used in this part.

3.3 Elementary statistical properties of trades-through

Traders usually scan the limit order book and restrict the size of their orders to the available liquidity. If necessary, they split a large order into several smaller orders to control the trade size. As explained in the introduction, we are interested in the trades that deviate from this usual behavior. We focus on trades that consume the liquidity available in the order book in an aggressive way, namely the *trades-through*.

We define as an x-th limit trade-through, any trade that consumes at least one share at the x-th limit in the order book. For example, a 2nd-limit trade-through completely consumes the first limit

4. It would also be interesting to extend this study by including data coming from alternative exchanges, and also from the 'block' trades and OTC transactions.

Asset	Market	Exchange	Period (liquidity stat.)	Period (intraday distrib.)	Period (lead-lag)
E-mini S&P500 - ES@	US Equity future	CME	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
Nasdaq E-mini - NQ@	US Equity future	CME	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
Dow Jones E-mini - YM@	US Equity future	CME	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
Eurostoxx - STXE@	EU Equity future	Eurex	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
Footsee - FFI@	EU Equity future	Liffe	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
DAX - FDX@	EU Equity future	Eurex	March 2010	1/3/10-12/3/10	1/12/09-12/3/10
Peugeot - PEUP.PA	French stock	Euronext Paris	March 2010	March 2010	March 2010
BNP Paribas -BNPP.PA	French stock	Euronext Paris	March 2010	March 2010	March 2010
Renault - RENA.PA	French stock	Euronext Paris	March 2010	March 2010	March 2010
Société Générale - SOGN.PA	French stock	Euronext Paris	March 2010	March 2010	March 2010

TABLE 3.2 – Data summary table

and begins to consume the second limit of the order book. In the figure 3.1, we show an example of such a trade.

3.3.1 On which limit of the order book is liquidity taken from ?

We study the location where liquidity is taken from the order book. More precisely, we want to measure the fraction of the total trade volume that is taken from each limit of the order book. In the figure 3.2, we plot this fraction against the algebraic limit number (strictly positive for the ask side, strictly negative for the bid side and zero for the trades we could not find on which limit they occurred) for the French stock BNP Paribas. For example, approximately 7% of the trading volume is taken outside the first limit of the order book for this stock (84% of which is taken from the second limit). This is clearly non-negligible and confirms that traders may sometimes consume the liquidity in a more aggressive way, rather than wait for new liquidity to be provided.

3.3.2 Frequency and volume

Here, we present basic statistics on occurrences and volumes of trades-through, in order to better measure the significance of this phenomenon.

Financial asset considered	2nd-limit TT Occurrence (in %| daily number)	2nd-limit TT Volume (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	1,43|638	2,66	2,2
Nasdaq E-mini - NQ@	1,91|404	5,73	1,3
Dow Jones E-mini - YM@	3,17|479	9,65	0,9
Eurostoxx - STXE@	2,2|237	5,4	3,5
Footsee - FFI@	4,3|419	12,7	0,9
DAX - FDX@	8,57|724	22,6	0,8
Peugeot - PEUP.PA	4,3 | 112	12,2	2,3
BNP Paribas -BNPP.PA	4,3 | 298	11,3	1,8
Renault - RENA.PA	5,0 | 224	12,8	1,5
Société Générale - SOGN.PA	6,0 | 451	14,8	1,1

TABLE 3.3 – Basic statistics on 2nd-limit trades-through (estimations based on March 2010 data)

Note that even if trades-through are rare events, they form a sizeable part of the daily volume (up to 20% for the DAX index future).

An important remark should be made at this stage : the smaller the relative tick value (which is the absolute tick value divided by the value of the asset), the more important trades-through are,

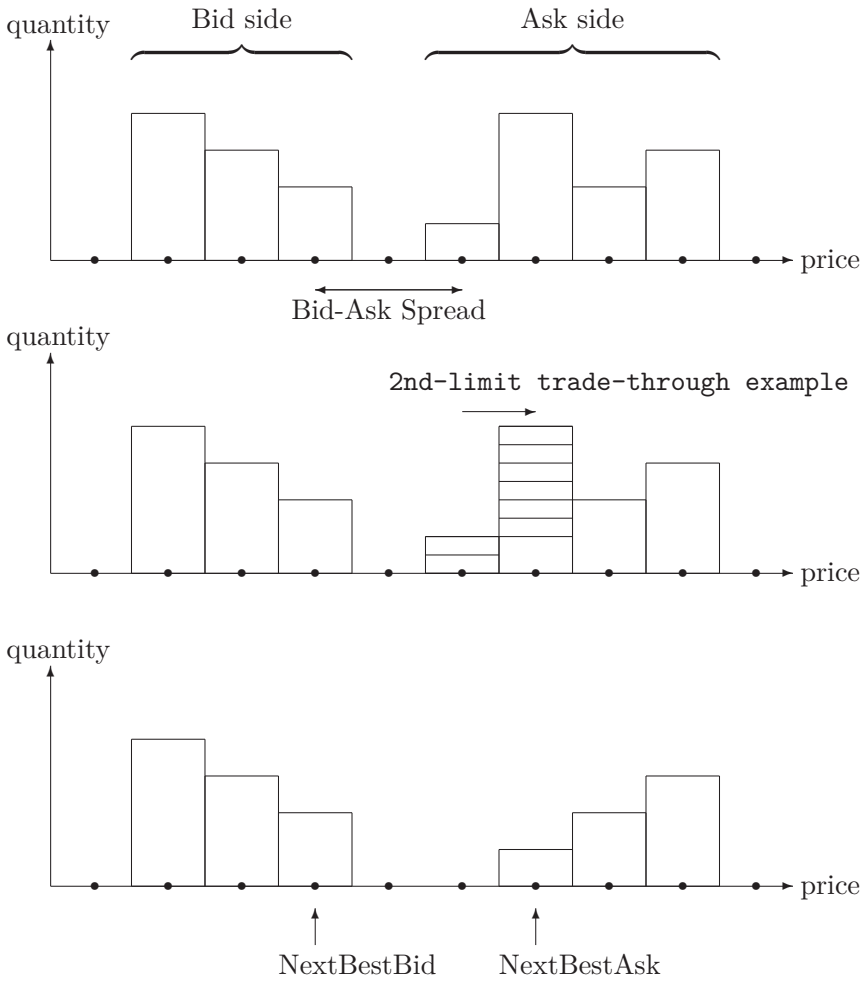

FIGURE 3.1 – Example of a trade-through : *(up)* Limit order book configuration before the trade-through ; *(middle)* Trade-through ; *(down)* Limit order book configuration after the trade-through.

Financial asset considered	3rd-limit TT Occurrence (in %| daily number)	3rd-limit TT Volume (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	0,12 | 0,6	0,0067	2,2
Nasdaq E-mini - NQ@	1,4|2,9	0,26	1,3
Dow Jones E-mini - YM@	6,2|8,5	0,79	0,9
Eurostoxx - STXE@	0,41 | 0,47	0,096	3,5
Footsie - FFI@	27|26	2,3	0,9
DAX - FDX@	49|42	3,58	0,8
Peugeot - PEUP.PA	59 | 16	3,47	2,3
BNP Paribas -BNPP.PA	48 | 33	2,49	1,8
Renault - RENA.PA	81 | 37	3,8	1,5
Société Générale - SOGN.PA	127|94	5,2	1,1

TABLE 3.4 – Basic statistics on 3rd-limit trades-through (estimations based on March 2010 data)

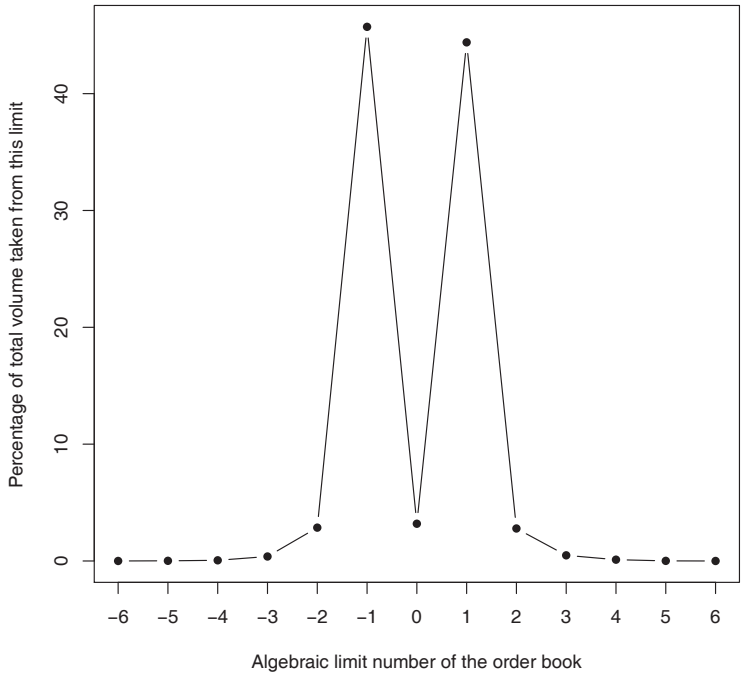

FIGURE 3.2 – Fraction of total trading volume taken from each limit of the order book of BNP Paribas stock

both in occurrence and volume. This result seems natural in the sense that, the smaller the relative tick value on an asset is, the more aggressively this asset is traded.

From a liquidity point of view, we also know that limit order books of assets with a big relative tick value tend to be more filled with high liquidity at the best, whereas order books of assets with a small relative tick value are less filled with liquidity at the best and may also present gaps. This explains why assets with small relative tick values are more often hit by trades-through. So, this liquidity pattern for assets with small relative tick values is consistent with the larger number of occurrences of trades-through in tables 3.3 and 3.4 on those assets.

The typical daily number of such 2nd-limit trades-through is about a few hundreds (from about 100 for a French stock like Peugeot to about 600 for the E-mini S&P500 future). When looking at the statistics on 3rd-limit trades-through, we can see that this daily number falls to less than 10 events for US futures (which shows that almost all trades-through on those assets only reach the 2nd-limit of the order book), to less than 50 events for EU futures and 100 events for French stocks.

On the volume part, we can see that except for US futures and the Eurostoxx future, the volume fraction due to 3rd-limit trades-through is in the range of 2 % to 5%. Clearly, such an analysis shows that trades-through, and especially 2nd-limit trades-through, are significant as a fraction of the number of trades and the trade volume.

Let us now check that our definition of trades-through is stable with respect to the volume of the trade. To do so, we first compute the mean trade volume of the usual trades, i.e., trades that consume less than the quantity available at the first limit. Then, we check if the statistical set of data we defined as trades-through is stable when adding a volume condition to their definition,

Financial asset considered	2nd-limit TT Occurrence fraction (in %)	3rd-limit TT Occurrence fraction (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	31	84	2,2
Nasdaq E-mini - NQ@	62	91	1,3
Dow Jones E-mini - YM@	69	90	0,9
Eurostoxx - STXE@	40	94	3,5
Footsee - FFI@	70	97	0,9
DAX - FDX@	80	99	0,8
Peugeot - PEUP.PA	75	96	2,3
BNP Paribas -BNPP.PA	74	96	1,8
Renault - RENA.PA	78	96	1,5
Société Générale - SOGN.PA	76	95	1,1

TABLE 3.5 – Stability of the definition of trades-through with respect to the Mean Trade Volume (estimations based on March 2010 data)

namely, that they consume more than the mean trade volume.

In the case of French stocks in table 3.5, a fraction of 75% of the data considered as 2nd-limit trades-through is stable with the new definition, and this fraction goes to 95% for 3rd-limit trades-through.

For the US and EU futures, the situation is more complex. If we look only at the 3rd-limit trades-through, the fraction of trades-through stable with the new definition is over 84%. But for 2nd-limit trades-through, there exists a bigger difference in this fraction, especially for the E-mini S&P500 and the Eurostoxx futures (which are both under 40%).

For these two particular assets, a fraction of 60 to 70% of 2nd-limit trades-through are caused by relatively small trades (with a volume smaller than the mean trade volume of the usual trades). We notice that they have the biggest relative tick values among the considered futures. As the best limit is crossed by those trades, this means that trades-through on futures with big relative ticks tend to happen in regimes when liquidity is lacking on the best limits.

In conclusion, except for these two particular futures in the case of 2nd-limit trades-through, the definition is stable when we add a volume condition.

3.3.3 Links between trades-through and large trades

In this section, we examine the link between trades-through and large trades. Intuitively, the larger a trade, the more probable the fact that this trade is a trade-through. On figure 3.3, we show the conditional probability for a trade to be a trade-through if its volume is higher than a given threshold, where this threshold is expressed in multiples of the mean trade volume. For French stocks, this conditional probability roughly increases from 0 to 0.4 for trade volumes ranging from 0 to 20 mean trade volumes.

This shows that trades-through are related to big trades as expected. However, empirically, this relation is not as clear as one might expect, since conditional probabilities are not reaching levels close to 1 which would reflect that every big trade is a trade-through. A statistical flaw may explain this difference : the bigger the trade volume limit used to measure the conditional probability, the less data we have (as shown in the table 3.6). Hence, the most reliable part of the graph is that close to the origin, with a rapidly increasing conditional probability for a trade to be a trade-through as its volume increases. But even taking this issue into account, we believe based on our data that some differences between trades-through and large trades will remain, as that some very large trades simply reflects the presence of large quantities at the best limit.

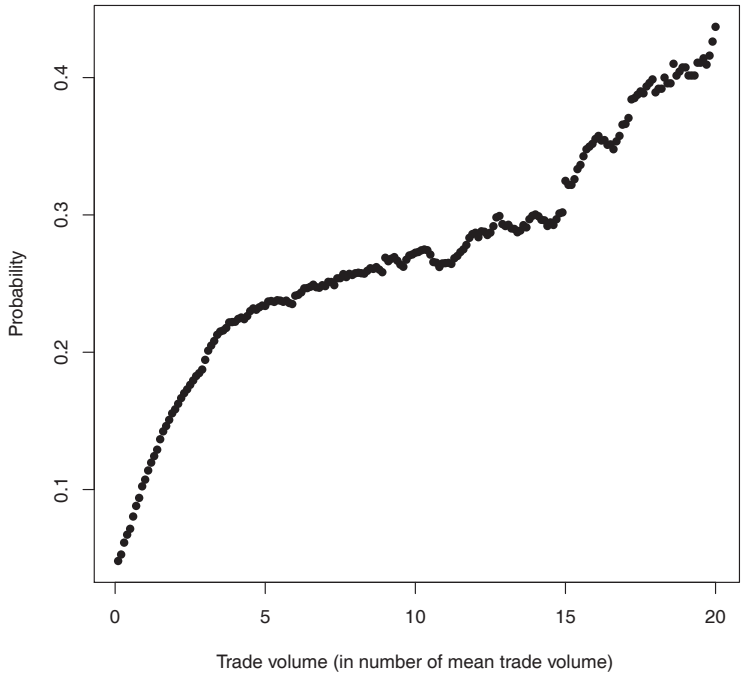

FIGURE 3.3 – Conditional probability for a trade to be a trade-through conditioned on his volume for BNP Paribas stock (in linear-logarithmic scale)

3.3.4 Clustering

Clustering of trades-through is studied by looking at the arrival time of the next trade-through : if there is a clustering of trades-through, the next trade-through should arrive faster after a trade-through than after any trade. To verify this, we compute the empirical arrival time distribution of the next trade-through, conditioned or not by the fact that the current trade is a trade-through.

We can see on the figure 3.4 that the distribution is more peaked for short waiting-times (measured in trades number) when the current trade is a trade-through. A similar graph is obtained when the lag is measured in physical time. We obtain a mean of 30 trades (respectively 113 seconds) to wait for the next trade-through in the unconditional case and a mean of 23 trades (respectively 91 seconds) when already on a trade-through.

To verify that this difference is statistically significant, we use a Welch two samples t-test to compare the means of the distributions in the conditional and unconditional cases and obtain an equality test p-value of less than 2.2×10^{-16} in both cases. So, trades-through are both more likely followed by trades-through (in trades number), and more closely in time followed by trades-through. [5]

5. It would also be interesting to compare the autocorrelation functions obtained in the unconditional and conditional cases to test for the presence of clustering in trades-through data.

Volume of the trade (in MTV)	Proportion of trades with a volume higher than (in %)			
	BNPP.PA	SOGN.PA	RENA.PA	PEUP.PA
1	29	30	32	31
2	12	11	12	12
5	2	2	2	2
10	0.4	0.3	0.4	0.4
20	0.07	0.06	0.06	0.05

TABLE 3.6 – Rarefaction of data for French stock in the estimation of conditional probability of trades-through

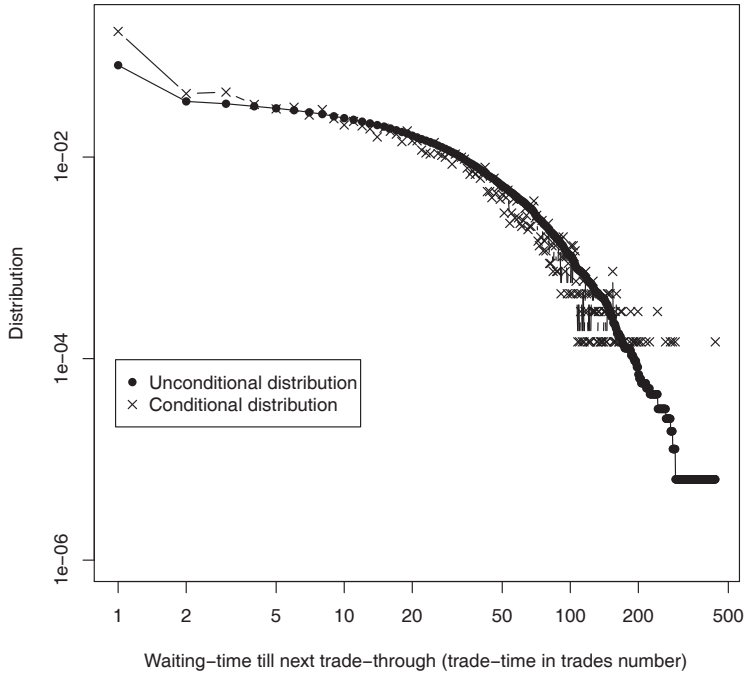

FIGURE 3.4 – Distribution of waiting-time (in trade-time) until next trade-through for BNP Paribas stock (in double logarithmic scale)

3.3.5 Intraday distribution

Figure 3.5 shows the intraday distributions of timestamps [6] of trades in the unconditional case (all trades) and the conditional case (trades-through) for the US equity future E-mini S&P500. Comparing the distribution for 2nd-limit trades-through with the unconditional one shows that the U-shaped central part is less important, and that the relative size of the peaks drastically increases. Those peaks are present at very specific hours, observed for both US and European equity futures, as will be detailed later. Note that if we restrict the study to 3rd-limit trades-through, the U-shape part is almost completely to completely removed from the distribution, in which only the peaks

6. All timestamps presented in this article are referenced in the time-reference of Europe/Paris = CET = UTC/GMT + 1h.

remain.

Using a two-sample Kolmogorov-Smirnov test to compare the unconditional and conditional distributions of the timestamps of trades-through yields a p-value less than 2.2×10^{-16}, which indicates that both distributions are statistically different.

The peaks of the intraday distribution of trades-through timestamps are more pronounced at specific hours of the day :

- 07 :50 : Eurex trading phase beginning (FESX, FDAX).
- 09 :00 : Euronext trading phase beginning (FTSE).
- 14 :30 : US major macro news releases (Jobless claims, Employment situation, International trade or GDP, for example). CME open-outcry trading phase beginning (major equity index futures). [7]
- 15 :30 : NYSE regular trading phase beginning.
- 16 :00 : US major macro news releases (ISM Manufacturing Index, Philadelphia Fed survey or New Homes Sales, for example).
- 17 :30 : End of the calculation of the DAX index using Xetra electronic trading system.
- 22 :00 : Eurex trading phase end (FESX, FDAX). Euronext trading phase end (FTSE).

Similar results are obtained for European Equity futures and French stocks, but with smaller peaks for French stocks. This suggests a weaker dependence of French stocks on major macro-economic news, compared with US and EU equity futures.

3.4 Market impact

In double auction limit order books, the price is completely determined by the liquidity available in the order book and the way liquidity is consumed by market orders. An interesting and important question is to assess the market impact, measured as the correlation between the sign of an incoming market order and the subsequent price changes. An empirical measure of market impact is given in Bouchaud et al. [2004] by the impact function R as a function of the lag l :

$$R(l) = \langle (m_{n+l} - m_n).\epsilon_n \rangle$$

where ϵ_n is the sign of the n-th trade (1 if the trade is on the ask side, -1 if it is on the bid side) and m_n is the midprice before the n-th trade.

Eisler et al. [2012] studied the market impact of 14 of the most liquid stocks traded on the NASDAQ in 2008. They obtained a response function that increases for lags up to 100 events and then, depending on the tick size, remains roughly constant for large tick stocks, or slowly decreases to a non-zero value for small tick stocks. In this section, we adapt the methodology in Eisler et al. [2012] to measure the market impact of trades-through. One important point has to be taken into account : by definition, there is a mechanical, instantaneous price change after a trade-through that has to be removed to fairly compare the price impact of a trade-through with that of any trade. To this end, we start counting the lags one instant [8] after the considered trade-through. That way, we make sure that the reference level with respect to which the returns are calculated is the price after the trade-through, and not before it.

7. The possible links between news and major events in stocks have already been studied, for example by Joulin et al. [2008]

8. In the data, we use a time lag of one millisecond.

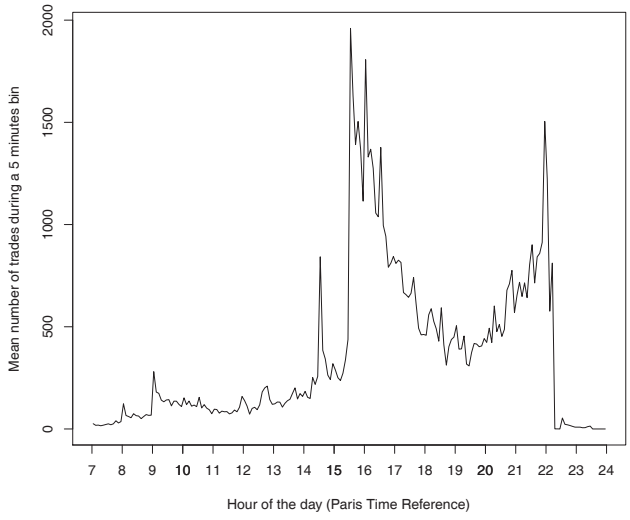

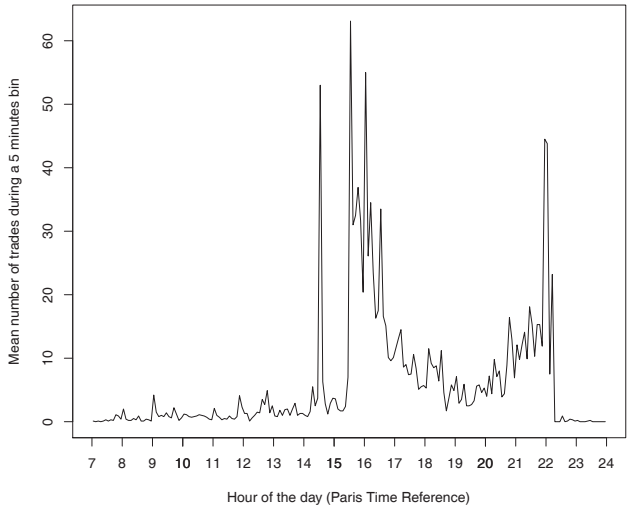

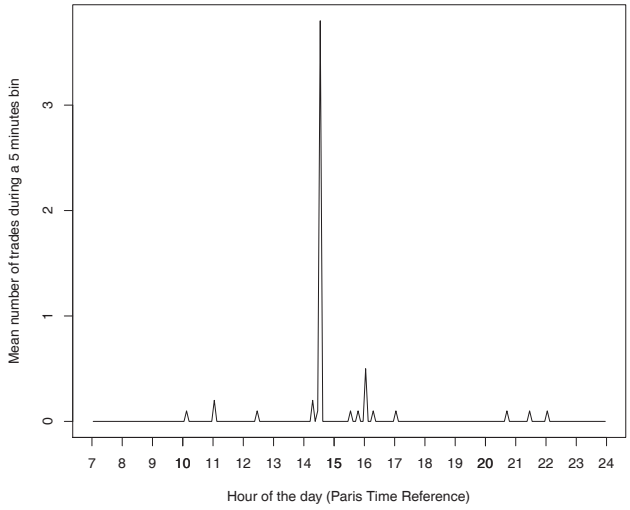

FIGURE 3.5 – Trades (above), 2nd-limit trades-through (middle) and 3rd-limit trades-through (below) intraday histograms for US equity index future E-mini S&P500

3.4.1 Empirical response function

In the upper figure 3.6, are plotted the response functions in the unconditional case (all trades are considered) and in the conditional case (only trades-through), for lags smaller than 15 minutes. One can see that the two response functions are similar : they first sharply increase from 0 to a maximum value (equal to 1 tick for trades-through and to 0.8 tick for all trades) reached after approximately 10 seconds, and then slightly decrease to half of its maximum value for trades-through, and 75% of its maximum value in the unconditional case as the time lag goes from 10 seconds to 15 minutes. Also note that the response function of trades-through is larger than the unconditional one for lags less than one minute, whereas both response functions become approximately equal for lags ranging from 1 minute to 15 minutes.

The global response function clearly measures the correlation between the sign of a trade and that of the next return. One can see that for lags less than 1 minute, the sign of a trade-through is a better indicator of the sign of the next midprice return than the sign of any trade. In an economic perspective where information is displayed in the market through the combination of events affecting the limit order book, one can say that trades-through have a higher informational content than usual trades.

Comparing our results with those in Eisler et al. [2012], we first notice a difference in the order of magnitude of the price impact function. After removing the instantaneous price impact in their graphs, their response functions take values ranging from 0 to 0.4 tick for large tick stocks and from 0 to 0.8 tick for small tick stocks, whereas in our example, the response function is more pronounced and takes values ranging from 0 to 1 tick. Apart from the instantaneous price impact, they obtain very similar response functions for market orders that do not change the best price and for those that change the best price. The only difference between both cases is, in the case of small tick stocks and for lags larger than 2000 events, they find a constant impact for market orders that change the price, and a response function that continues to decay for market that do not change the price. In the study we perform on trades-through, the response functions for small and large tick stocks both decay for lag values of 10 seconds to 15 minutes. Given the number of events per day, of the order of 70 000 events per day on the French stock BNP Paribas, this is most likely caused by the fact that we do not reach the regime where Eisler et al. [2012] began to measure a constant impact function, as we do not concatenate successive days when performing our statistics.

3.4.2 Response function conditioned on trade volume

We saw in sections 3.3.2 and 3.3.3 that trades-through tend to be trades with volume bigger than the mean trade volume. It is also clear that trade volume influences market impact : the larger the trade size, the greater the impact. The results shown on the upper figure 3.6, where we compared global response functions for trades-through and all trades, may arguably be explained by the fact that trades-through generally have bigger volume. To rule out this simplistic explanation, we include the volume of trades-through in the response function and study the difference in market impact.

Following Bouchaud et al. [2004], we use a generalized version of market impact to take the trade volume into consideration :

$$R(l, V) = \langle (m_{n+l} - m_n).\epsilon_n \rangle \,|_{V_n = V}$$

where ϵ_n is the sign of the n-th trade, m_n is the midprice before the n-th trade and V_n is the volume of the n-th trade.

Trades are divided into four volume categories related to the mean trade volume MTV : $[0, \frac{MTV}{2}], [\frac{MTV}{2}, MTV], [MTV, 2 \times MTV], [2 \times MTV, +\infty]$. Figure 3.6 shows that each response

function of trades-through is still bigger than the unconditional one of the same trade volume category. This property holds true for all lags for the first trade volume category $[0, \frac{MTV}{2}]$ and for lags ranging from 0 to 10-20 seconds for the other trade volume categories.

It is worth noticing that this bigger market impact is especially valid for lags smaller than 10 seconds, which are those with the smaller standard deviation of results for each measure indicated in the graph. For lags bigger than 10 seconds, as the standard error for each measure is higher, both the decline of response and the relative magnitude between the trades-through and the unconditional responses are less significant.

One must therefore conclude that, independently of the trade volume, the sign of trades-through is a better indicator of the sign of the next return than the sign of an unconditional trade, at least over timescales ranging from 0 to 10 seconds. [9]This confirms what may be termed as the higher informational content of trades-through as opposed to that of usual trades [10].

3.5 Spread relaxation

As trades-through are rather rare and informationally rich events, it is natural to consider a trade-through as an excitation of the limit order book. The results of the previous sections essentially addressed the relaxation of the mid-price, and we now focus on the behaviour of the spread. Hence, we study the spread of the limit order book in order to measure :

– the spread level in the excited state (on a trade-through)

– the evolution of the spread after the trade-through.

Figure 3.7 shows the behaviour of the excess spread for the French stock BNP Paribas after all trades and after trades-through. The excess spread is defined as the difference between the value of the spread after a trade and its value just before. The excited value (in ticks) of the excess-spread is approximately 0.5 after a trade-through, and 0.25 after a standard trade. Since trades-through instantaneously increase the spread by one tick, this result indicates that approximately half of the relaxation takes place during the first lag interval, namely, 5 seconds.

The unconditional excess spread decay seems to decay almost linearly for lags smaller than 100 seconds. For longer lags, one must bear in mind that the scarcity of available data where two successive trades-through are separated by a large lag makes a large portion of this unconditional curve not statistically reliable. It is plotted it here mostly for indicative purpose, and is mostly reliable for the very first lags.

For the sake of completeness, one can fit a power-law decay of the excess spread in physical time with an exponent close to 0.25, as shown on figure 3.8.

At this stage, a comment is in order : the section 3.3.5 shows that there is approximately one trade-through per minute for BNP Paribas stock. However, there are periods during the day when more events occur, and the maximum number of trades-through reaches a frequency of 1 event every 25 seconds. Therefore, the excess-spread relaxation is in fact statistically reliable for a period of the order of 25 seconds. For lags larger than 25 seconds, there are significantly less data used in the conditional statistics and one should be very careful drawing any conclusion from that part of the graph.

9. It is worth keeping in mind that a part of trades-through might be due to liquidity fluctuations, as a usual trade incoming in an order with few liquidity at the best limit will mechanically become a trade-through and, as a consequence, will not necessarily have a higher informational content. But, as the definition of trade-through is stable when including a condition of minimal volume (see table 3.5), we know that in general trades-through are not due to liquidity fluctuations. As a consequence, most trades-through are really linked with a fast and big consumption of the liquidity available at the best limit. This helps to better understand the higher informational content of them.

10. Note that similar empirical results are obtained using other time definitions, e.g. tick time, or trade volume time where the lag is measured by the fraction of trading volume with respect to the daily trading volume.

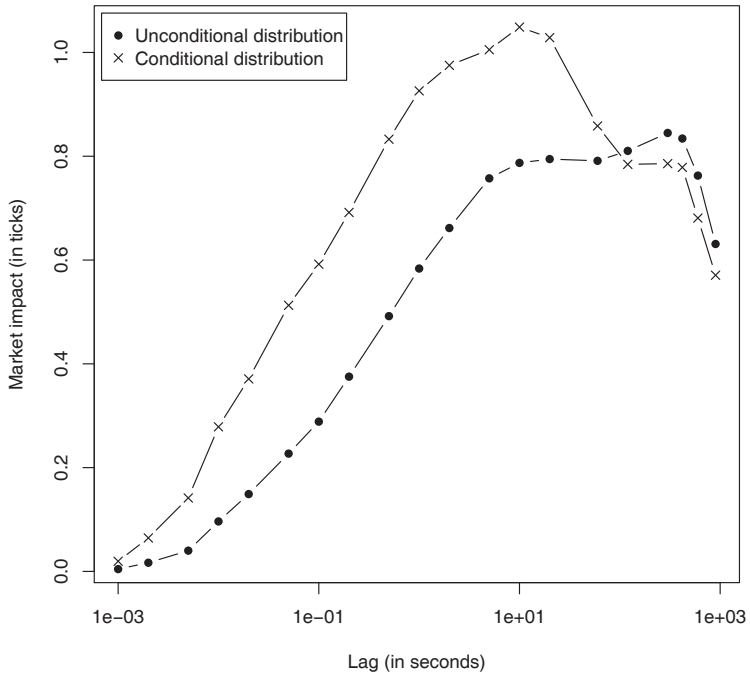

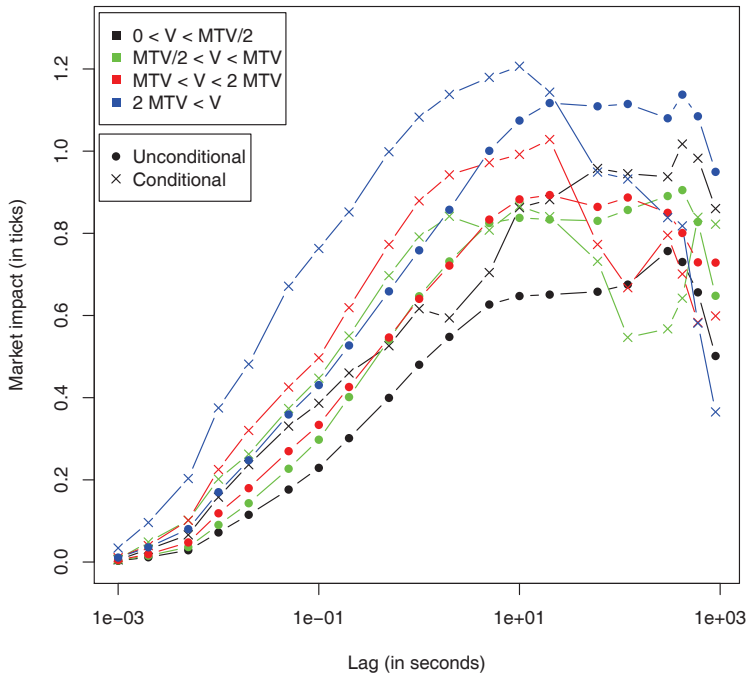

FIGURE 3.6 – Global response functions for BNP Paribas stock (in linear-logarithmic scale). *(up)* All trades vs trades-through ; *(down)* All trades vs trades-through, with the condition on the initial trade volume with respect to the mean trade volume MTV.

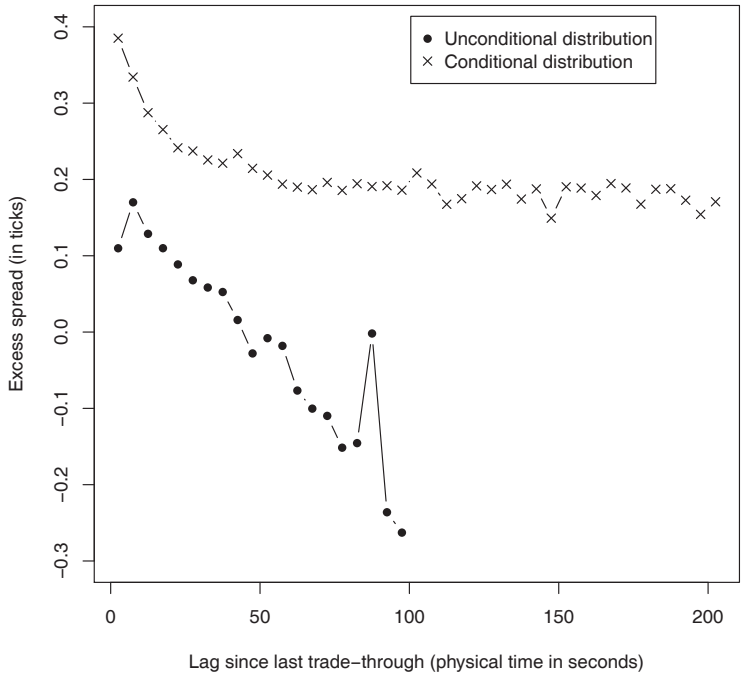

FIGURE 3.7 – Excess-spread decay after trades for BNP Paribas stock (in linear-logarithmic scale)

Also note that we consider physical time in this section. Tick time could be relevant, as it incorporates all the changes in the order book. However, it seems clear that trade time, especially trade-through time, is not the correct clock, since most of the spread dynamics takes place before a next trade occur, as the order book replenishes itself.

Zawadowski et al. [2004] and Zawadowski et al. [2006] first study and measure the changes occurring after large intraday 15-minutes price changes and they found that the spread was decaying as a power-law on the NYSE exchange.

Ponzi et al. [2009] study a similar problem in the relaxation of the spread. Conditioned on a move of the spread, they measure a relaxation and obtain a power-law behavior of the excess spread in trade time with exponent between 0.4 and 0.5). They provide no explanation for this empirical observation.

Another example is given by the work of Toth et al. [2009] where relaxation after large price moves is studied. They show a power-law relaxation of the excess bid-ask spread in physical time, with exponent 0.38.

In both cases, the methodologies and results are similar to ours, in that they provide evidence of a slow relaxation of the excess-spread, with power-law fits showing exponents of the same order of magnitude.

3.6 Lead-lag estimation from trades-through time-series

It is a well-known fact amongst financial practitioners that some pairs of assets, not only move in a correlated way, but also that sometimes one asset (called the leader) tends to move in advance

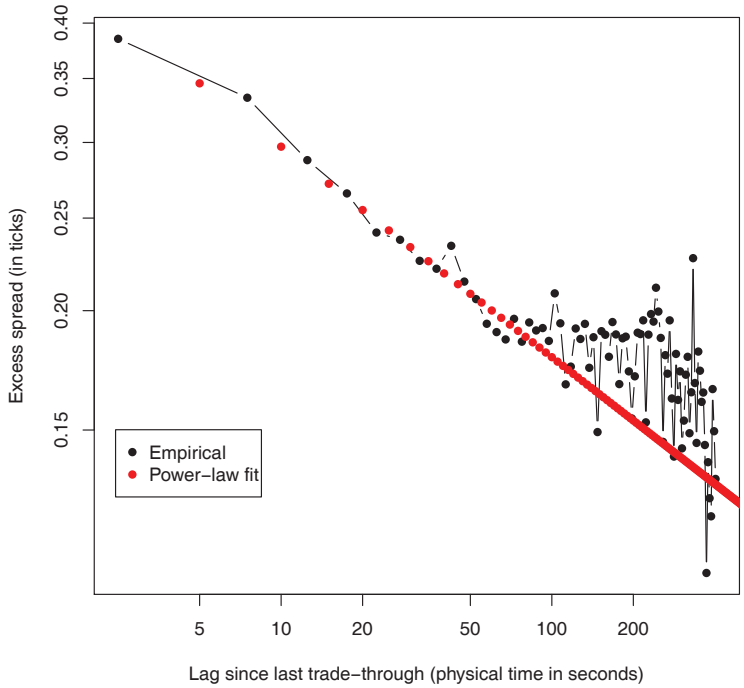

FIGURE 3.8 – Power-law decay of excess-spread after trades-through for BNP Paribas stock (in double logarithmic scale)

with respect to the other asset (the lagger). Over the years, this empirical phenomenon has received considerable attention and traders tried to build profitable leader-lagger pairs of assets, in order to trade the lagger with the information coming from the moves of the leader. The recent availability of high-frequency financial data confirmed this interest in lead-lag relationships, as reflected in the financial literature (see Lo and MacKinlay [1990], De Jong and Nijman [1997], De Jong and Donders [1998], Abergel and Huth [2011] and the references therein, for example).

Our approach is different, as we directly address the following question : which asset moves first ? In other words, when a trade-through occurs on Asset 1, is it followed or is it preceded by a trade-through on Asset 2 ? Thanks to the relative rarity of trades-through, we are able to answer this question in a statistically significant way, as will now be described in details.

3.6.1 The estimation technique

Assume that we have two grids representing the timestamps of trades-through for two different assets, and that we want to get an empirical distribution of the lead-lag parameter between the two assets. Our method is quite simple : we connect every timestamp on one grid with its nearest neighbour on the other grid. At the end of this process, every trade-through on Asset 1 is linked with the closest trade-through on Asset 2. Figure 3.9 shows how different timestamps are connected pairwise. Each trade-through of asset 1 is put in relationship with the trade-through of asset 2 that is closest in time.

The empirical distribution of the lead-lag parameter between the two assets is then obtained by calculating the difference between two connected timestamps and plotting this distribution.

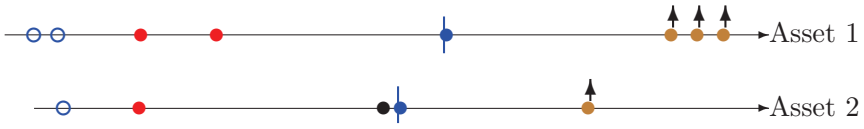

FIGURE 3.9 – Each trade-through of Asset 1 is put in relationship with the trade-through of Asset 2 that is the closest in time.

This method can be generalized to two groups of more than one asset by first merging the timestamps of trades-through in each group. For example, in the study of the lead-lag between EU and US equity markets, we first build two grids, one for the US equity market and one for the EU equity market, by merging the timestamps of the trades-through of E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures (for the US grid) and those of Eurostoxx, DAX and Footsie futures (for the EU grid).

In order to better understand this estimation technique, and also provide numerical evidence that it is sound, we test it with a very basic model for the arrival times of trades-through as two independent Poisson processes with different intensities. In such a model, there is no lead-lag relationship between the two assets. On the figure 3.10 are plotted the two empirical lead-lag distributions using either one of the assets as a reference grid. Clearly, the two distributions are different, but they are both perfectly symmetric with zero mean.

An important feature of this method is that it is purely empirical, totally model-free, and returns a full distribution of lead-lag parameter, not only one value that maximizes some contrast criterion. Again, let us insist on the fact that only the relative rarity of trades-through enables us to use this simple method : should the average lead-lag parameter be of the order of the average frequency of trades-through, no conclusion could be drawn. [11]

3.6.2 Empirical results

In this section, we apply the methodology just introduced to empirical lead-lag estimates to the EU and US equity futures. Results are displayed on figure 3.11, where two useful statistics are plotted : the empirical distribution of the lead-lag parameter, and the positive and negative cumulative distributions of the lead-lag parameter. On the left side of each graph, the EU futures lead the US futures and on the right side, the US lead Europe.

Define the lag l between the timestamps TS of the two stocks as $l = TS_{EU} - TS_{US}$. We have to compare $P_- = \mathbf{P}(l < 0) \equiv \mathbf{P}(\text{EU leads US})$ and $P_+ = \mathbf{P}(l > 0) \equiv \mathbf{P}(\text{US lead EU})$. Empirically, there holds that $P_+ > P_-$. Let us now check that this difference is statistically significant, by comparing it to its standard deviation σ. Since

$$\sigma^2 = \mathbf{Var}(P_+ - P_-) = \mathbf{Var}(2P_+ - 1) = 4\mathbf{Var}(P_+) = \frac{4P_+P_-}{N_{data}}, \tag{3.1}$$

the empirical values $P+ \approx 0.5126$, $P- \approx 0.4874$, $N_{data} = 118207$ yield $P_+ - P_- \approx 0.025$ and $\sigma \approx 0.0029$, and therefore $P_+ - P_- \approx 8\sigma$. One can then conclude that the empirical lead of US futures on EU futures is statistically significant at an 8 standard deviations level on the time period

11. It would be interesting to test our method on the analytically tractable case presented in Kullmann et al. [2002].

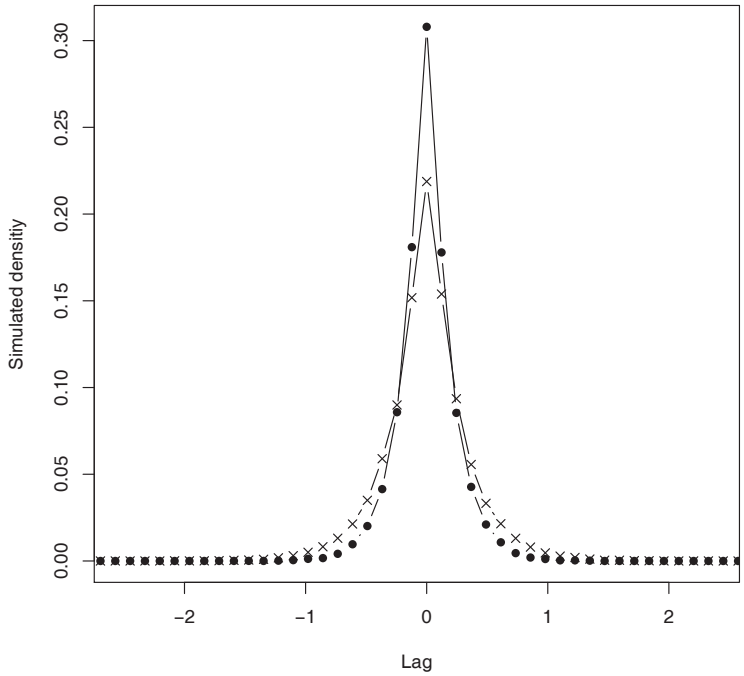

FIGURE 3.10 – An example of lead-lag parameter distribution for two simulated Poisson processes with different intensities (one curve for each time grid reference)

considered.

We also perform a similar analysis on French stocks, choosing the pair Société Générale and BNP Paribas. Defining similarly the lag l between the timestamps TS of the two stocks as $l = TS_{SOGN} - TS_{BNP}$, we obtain $P_- = \mathbf{P}(l < 0) \equiv \mathbf{P}(\text{SOGN leads BNP})$ and $P_+ = \mathbf{P}(l > 0) \equiv \mathbf{P}(\text{BNP leads SOGN})$. Clearly, $P_- > P_+$. The empirical values $P+ \approx 0.489$, $P- \approx 0.511$, $N_{data} = 6816$ yield $P_- - P_+ \approx 0.022$ and $\sigma \approx 0.012$, which means that the lead of Société Générale on BNP Paribas is statistically significant at approximately a 2 standard deviations level.

It is interesting to notice that our method gives access to the full distribution of the lead-lag parameter between each pair of assets studied. In particular, one can have information about the usual size of the lag. For the French stocks under scrutiny for example, we obtain average lag values of approximately 1 second : Société Générale leads on average by 0.76 seconds BNP Paribas (financial sector) and Peugeot leads on average by 1.23 seconds Renault (automotive sector).

3.7 Empirical lead-lag signal from trades-through

It has been shown previously (see Section 3.4) that for a single asset, one can look at trades-through as a local indicator of future market moves. So, for a given pair of assets, it seems interesting to generalize this approach and look at trades-through as as indicator of local lead-lag relationship. As a consequence, we will investigate what happens when one takes position on a given lagger asset,

based on input signals from the trades-through of the leader asset.

3.7.1 Data description and summary statistics

Empirical data

We use TRTH data (see Section 3.2.1 for more detail). The perimeter we looked at is composed of four different assets : two major French stocks (BNP Paribas and Societe Generale) traded on Euronext-Paris exchange and two major Equity index futures (DAX and Eurostoxx) traded on Eurex exchange. We denoted them with their usual Reuters tags (BNPP(.PA), SOGN(.PA)GLE, FDX(@), STXE(@)). For a given day, we always focus on the most liquid traded future, which is the one maturing first. We only look at the data coming from the main exchange for each asset and only during simultaneous hours of trading.

We focused our study on the time period going from January 2011 to May 2011 included, which gives about 100 days of trading. Moreover, we excluded the first and last half-hours of the trading day in order not to be impacted by Open or Close phenomenons. We also excluded from the statistics all the days when Daylight Saving Time were present in the US and not in Europe (which corresponds to two weeks in March 2011).

Finally, we also know that trades-through are linked with macroeconomic news arrival (as detailed in 3.3.5). Indeed, some market players with given anticipations of the level of the news are fast readapting their position to the actual level of the news, and so may increase trades-through arrival. As we want to look at intrinsic lead-lag properties, and not at the relative speed reaction of market participants on each asset after a news release, we choose to remove all trades-through close to 2.5 pm and 4 pm (Paris time).

Liquidity difference and choice of limit number

In the table 3.7, we present liquidity statistics on trades-through (already looked at in the Section 3.3.2) in order to choose the best data source for trades-through detection, with sufficient liquidity and maximum predictive power. There is a trade-off in the data source selection : we want to find a data source which is liquid enough in order to be possibly traded in real life and, at the same time, which has enough predictive power and is not to fast to be traded, which means that we will focus on higher limit trades-through whenever possible. Moreover, for a given lead-lag couple of assets, we want pairs of signals with similar liquidities in order to be able to compare the relative strengths of the signals without including liquidity distortions.

In the end, we choose 2nd-limit trades-through for the (BNPP, SOGN) stocks pair, whereas we choose 2nd-limit trades-through for the STXE future, coupled with 3rd-limit trades-through for the DAX future signals. In those two cases, we have similar liquidities for both signals in the pairs of assets, and also maximum predictive power by choosing when possible (in our case, for the DAX) the highest limit number possible for trades-through.

Field	BNPP.PA 2	SOGN.PA 2	STXE@ 2	FDX@ 2
trades number	8 416	9 081	20 602	24 033
trades-through number	628	968	231	1 521
	BNPP.PA 3	SOGN.PA 3	STXE@ 3	FDX@ 3
trades number	8 416	9 081	20 602	24 033
trades-through number	82	195	3	131

Table 3.7 – Daily liquidity statistics on 2nd and 3rd-limit trades-through for Jan-Feb 2011

3.7.2 Empirical lead-lag signal based on trades-through

When looking at two different assets that presumably form a lead-lag couple (because of a high lagged correlation, for example), it seems interesting to look at trades-through as indicators of lead-lag trading opportunities. In practice, we extract all the timestamps and directions of the trades-through on the leader. We consider this event occurring on the leader asset as significant and we want to take advantage of it on the lagger asset. To this end, we take positions on the lagger by entering at the midprice, just one millisecond after the trade-through (to account for the trades-through detection process on the leader) in the same direction as the leader trade-through. We then keep this position for a fixed duration, before closing it at the current midprice of the lagger.

We plot in the following figures (3.12 and 3.13) the empirical returns (measured in bp) obtained with the previous lead-lag trading strategy. In abscissa, we reported the return horizon, measured as the duration from the moment we take position in the lagger asset to the moment when we liquidate the position (both being done at the midprice). [12] We also summarized the lead-lag results obtained in the table 3.8.

For the (BNP-SOGN) pair, we can see that the two mean response functions are very similar, whatever the leader chosen (BNP leader or SOGN leader). More precisely, we found that the predictive power of SOGN as a lagger on BNP is slightly higher for lags less than 5 seconds. [13] We can see that, after only a return horizon of 10 seconds, the response return reached about 90% of its maximum value and is about 1.2 bp. We choose this specific duration of 10 seconds as our objective duration and we compare the empirical distributions of returns on the lagger at this precise horizon time. In both choices of leader, we obtain similar distributions of returns, when comparing both the means and the standard deviation values. Finally, when looking at a given choice of leader, we see that the bigger the mean return, the bigger the standard deviation of returns, which indicates that a trade-off has to be made to obtain the best Sharpe ratio.

For the (FDX-STXE) pair, we obtain very different results when changing the choice of leader, with maximum response function of 1.5 bp when the DAX is leader and of 0.5 bp when the STXE is leader. Moreover, the maximum response function is also reached before (after about 0.5 seconds), when the DAX is chosen as the leader signal. We choose 0.5 seconds as our objective duration to compare the complete distributions of returns. We still have the mean-variance trade-off to find the best leader signal amongst the DAX and STXE futures, with higher mean return and standard deviation of return for the DAX future chosen as the leader.

Finally, we included transaction costs in this lead-lag signal analysis. In the table 3.9, we compare the best mid-to-mid PnL per trade obtained with the lead-lag strategy (traded on the lagger asset) with a minimal one-tick return cost on the lagger. We can see that there is no profitable lead-lag trading strategy based on trades-through signal, when using only market-orders. As a consequence, one has to include limit orders in its strategy.

3.8 Conclusion and further research

In this chapter, we perform an extensive study of multiple-limits trades, or trades-through. Various statistics are provided measuring their liquidity, volumes, arrival times, clustering and spread relaxation. Evidence is provided regarding important peaks in the arrival times distribution at 2.5 pm and 4pm (Paris time reference), the moments of the day when major US macro-economic news are released. Market-impact of trades-through is studied, and is demonstrated to be larger than that of other trades, a fact that we synthetize by saying that trades-through have a higher informational content. We also introduce a new methodology to assess lead-lag relationships, and apply

12. These curves are closely related to generalized cross market-impact function.
13. This is not contrary to the general idea amongst practitioners that SOGN leads BNP.

Leader	Lagger	Mean PnL (in bp/trade)	Standard deviation (in bp/trade)	Indicative Sharpe ratio (%/trade)
BNPP.PA	SOGN.PA	1.15	5.34	21.4
SOGN.PA	BNPP.PA	1.13	4.36	25.9
FDX@	STXE@	1.41	2.26	62.7
STXE@	FDX@	0.23	1.64	13.8

TABLE 3.8 – Lead-lag results summary

Traded asset	Best mid-to-mid PnL (in bp)	Tick	Asset value	1 Tick return (in bp)
BNPP.PA	≈ 1.1	0.005	30	1.6
SOGN.PA	≈ 1.2	0.005	20	2.5
FDX@	≈ 0.6	0.5	6000	0.8
STXE@	≈ 1.4	1	2000	5

TABLE 3.9 – Comparison of best mid-to-mid returns with transaction costs

it to pairs of financial markets (European vs US futures). Finally, we studied lead-lag entry signals based on trades-through, for French stocks and major European futures. We obtained good results with midprice-to-midprice strategies but no strategy was profitable when including transaction costs (with exit times after a fixed duration). It would be interesting in future research to also include limit orders in order to bypass some of the transaction costs.

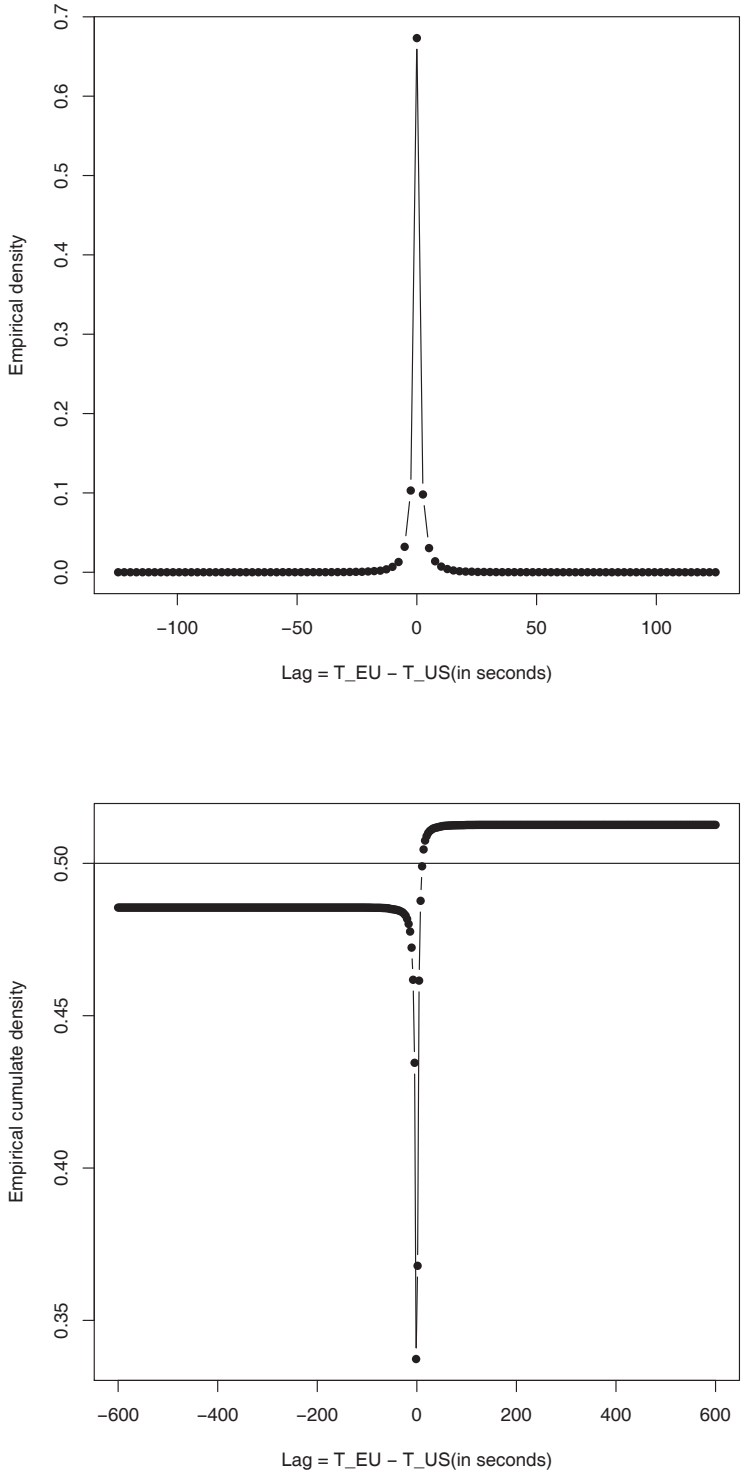

FIGURE 3.11 – Lead-lag parameter empirical distribution (above) and positive and negative cumulative distributions (below) of 2nd-limit trades-through for two groups of European and American futures

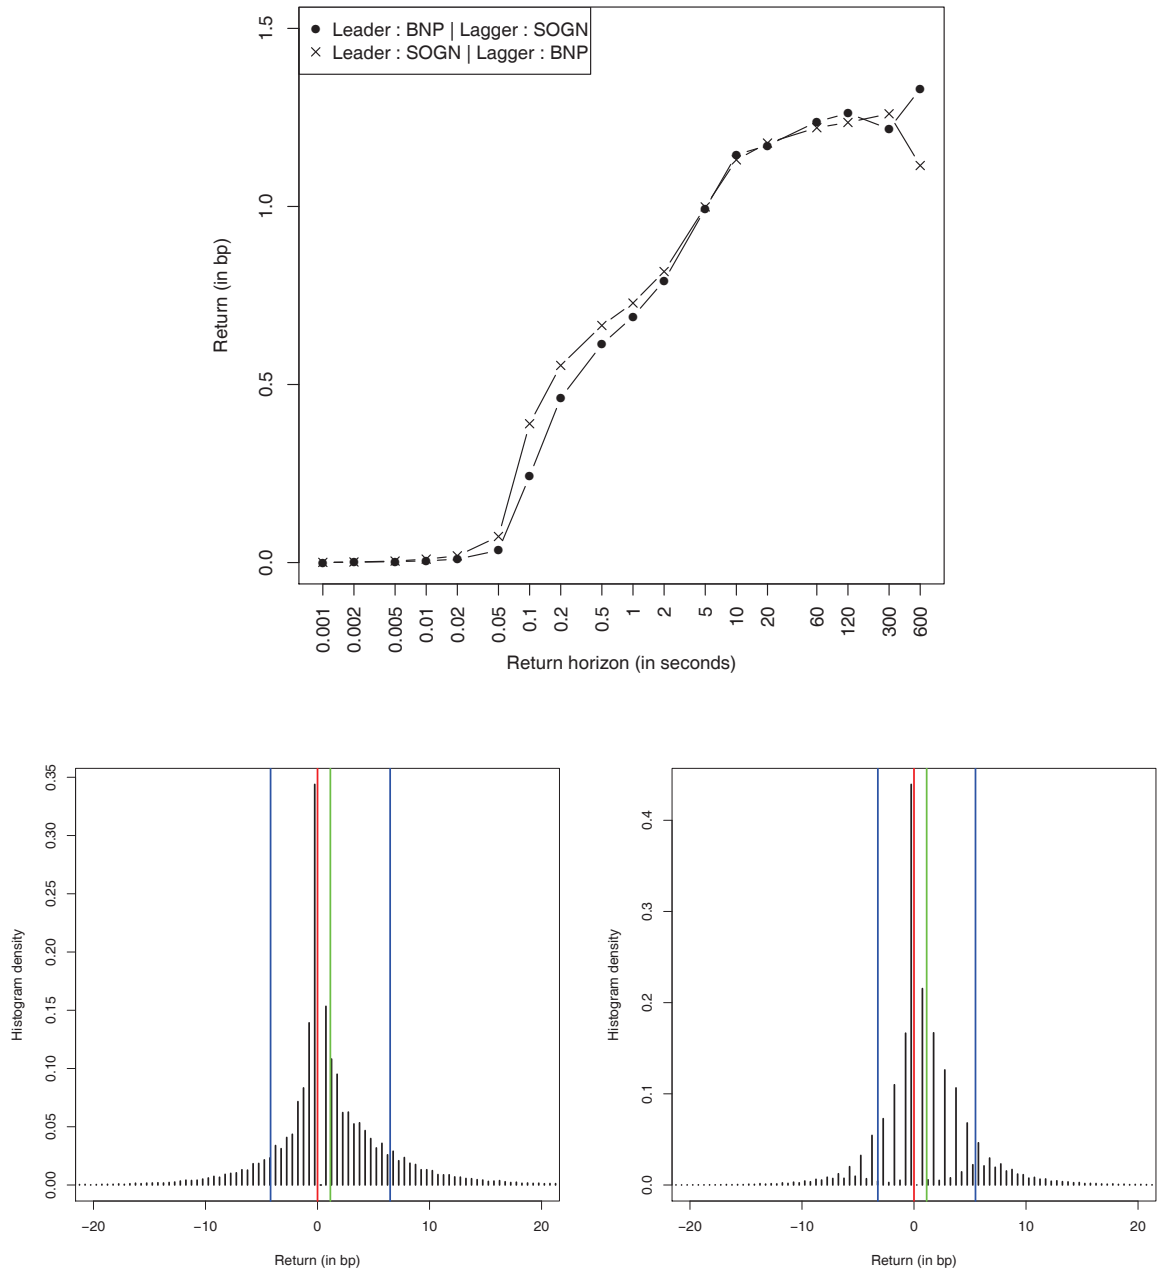

FIGURE 3.12 – *(left)* Leader : BNPP | Lagger : SOGN ; *(right)* Leader : SOGN | Lagger : BNPP ; *(up)* Lead-lag signals across time ; *(down)* PnL distribution after 10s

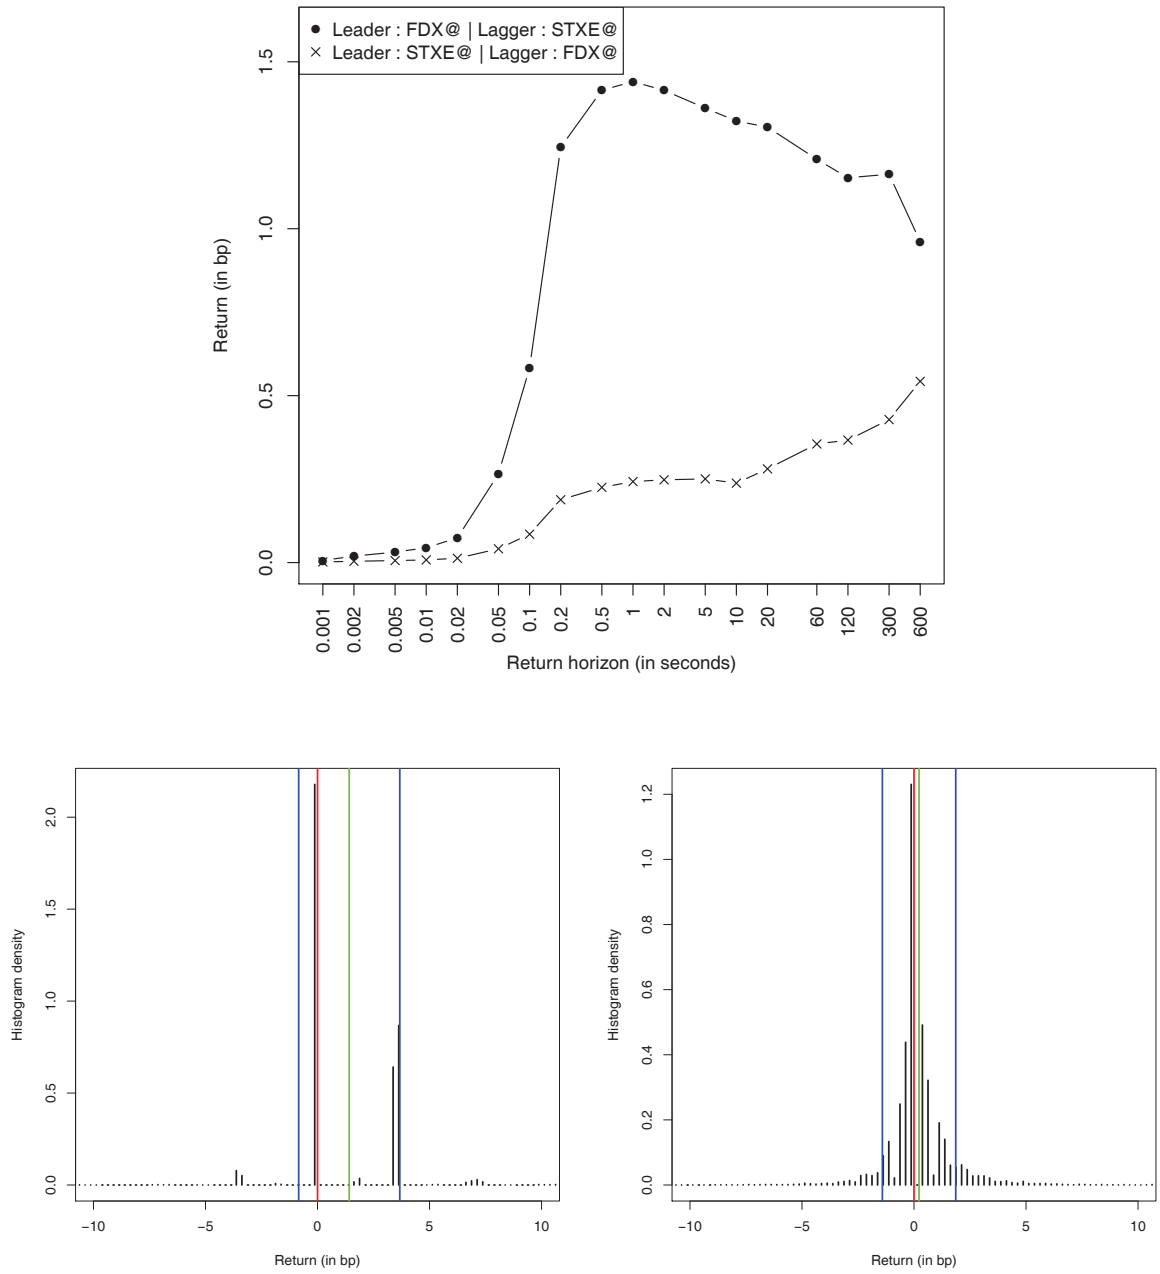

FIGURE 3.13 – *(left)* Leader : FDX | Lagger : STXE ; *(right)* Leader : STXE | Lagger : FDX ; *(up)* Lead-lag signals across time ; *(down)* PnL distribution after 0.5s

Chapitre 4

Modelling trades-through in a limit order book using Hawkes processes

Résumé

We model trades-through, i.e. transactions that reach at least the second level of limit orders in an order book. Using tick-by-tick data on Euronext-traded stocks, we show that a simple bivariate Hawkes process fits nicely our empirical observations of trades-through. We show that the cross-influence of bid and ask trades-through is weak. We also obtain that a toy-trading strategy based on the autocorrelation of trades-through arrival is improved by the use of our Hawkes model of trades-through, on a Euronext-traded stock and a Eurex-traded future. We validate our choice of a generalized Hawkes linear model for trades-through by confirming the absence of empirical inhibitory effects. Finally, we use a nonparametric statistical method to get the complete empirical profile of trades-through decay kernel. We found this decay to be faster than exponential, and closer to a power-law.

Contents

4.1 Introduction 60
4.2 Trades-through 60
4.2.1 Orders splitting and trades-through 60
4.2.2 Definition of trades-through 61
4.2.3 Empirical data 61
4.2.4 Occurrences of trades-through 61
4.2.5 Clustering 61
4.2.6 Intraday timestamp distribution 63
4.3 Hawkes processes 64
4.3.1 Definition 64
4.3.2 Stationarity condition 65
4.3.3 Maximum-likelihood estimation 66
4.3.4 Testing the calibration 67
4.4 A simple Hawkes model for trades-through 68
4.4.1 Model 68
4.4.2 Calibration 69
4.4.3 Goodness-of-fit 72
4.5 Toy-trading strategy based on trades-through autocorrelated arrival times 72
4.5.1 Hawkes signal computation 72
4.5.2 Results 73
4.6 Searching for inhibitory effects in trades-through 75
4.6.1 Models 75

4.6.2 Results 75
4.7 Calibration of the empirical decay kernel of trades-through arrival times by a non-parametric method 76
4.7.1 A reminder on nonparametric Hawkes process calibration 76
4.7.2 Application to trades-through arrival times 77
4.8 Conclusion 77

4.1 Introduction

Recent contributions have emphasized that Hawkes processes exhibit interesting features for financial modelling. For example, these self- and mutually exciting point processes can model arrival times of orders in an order book model (see Large [2007] and Muni Toke [2011]), or explain the Epps effect in a microstructure toy model by Bacry et al. [2012]. An econometric framework has been derived by Bowsher [2007].

In this paper, we are interested in modelling trades-through, i.e. transactions that reach at least the second level of limit orders in an order book. Trades-through are very important in price formation and microstructure. Since traders usually minimize their market impact by splitting their orders according to the liquidity available in the order book, trades-through may contain information. They may also reach gaps in orders books, which is crucial in price dynamics.

In a first part, we give basic statistical facts on trades-through, focusing on their arrival times and clustering properties. Our second part is a general introduction to Hawkes processes. In a third part, using tick-by-tick data on Euronext-traded stocks, we show that a simple bi-dimensional Hawkes process fits nicely our empirical data of trades-through. We show that the cross-influence of bid and ask trades-through is weak. Following Bowsher [2007], we improve the statistical performance of our maximum likelihood calibrations by enhancing the stationary model using deterministic time-dependent base intensity. A fourth part is devoted to the computation of intensity indicator based on trades-through arrival, and this signal is used to improve a simple momentum trading strategy based on trades-through. In a fifth part, we investigate if the models class under scrutiny for the calibration was well-adapted and if no inhibitory effects are measured after trades-through arrival. Finally, a sixth part uses a non-parametric calibration to measure the empirical decay kernel of the temporal impact.

4.2 Trades-through

4.2.1 Orders splitting and trades-through

It has been shown several times that the time series built from trading flows are long-memory processes (see e.g. Bouchaud et al. [2009]). Lillo and Farmer [2004] argue that this is mainly explained by the splitting of large orders. Indeed, let us assume that a trader wants to trade a large order. He does not want to reveal its intentions to the markets, so that the price will not "move against him". If he were to submit one large market order, he would eat the whole liquidity in the order book, trading at the first limit, then the second, then the third, and so on. When "climbing the ladder" this way, the last shares would be bought (resp. sold) at a price much higher (resp. lower) than the first ones. This trader will thus split its large order in several smaller orders that he will submit one at a time, waiting between each submitted order for some limit orders to bring back liquidity in the order book. We say that the trader tries to minimize its market impact.

In practice, this mechanism is widely used : traders constantly scan the limit order book and very often, if not always, restrict the size of their orders to the quantity available at the best limit.

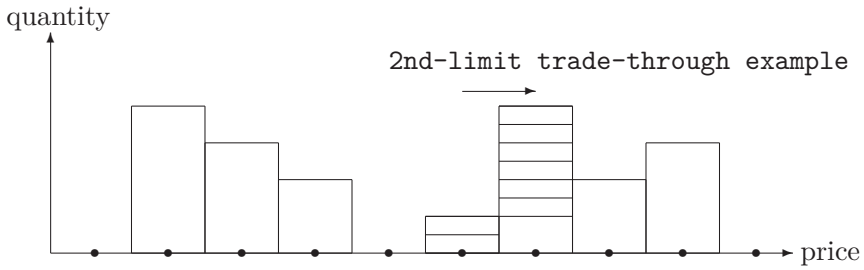

FIGURE 4.1 – 2nd-limit trade-through example

But sometimes speed of execution is more important than minimizing market impact. In this case, orders larger than the size of the first limit may be submitted : thus, trades-through are precisely the trades that stand outside the usual trading pattern, and as such are worth being thoroughly studied.

Trades-through have already been empirically studied by Pomponio and Abergel [2010] : their occurrences, links with big trades, clustering, intraday timestamps distribution, market impact, spread relaxation and use in lead-lag relation. In this paper, we model trades-through with Hawkes processes.

4.2.2 Definition of trades-through

In general, we call a n-th limit trade-through any trade that consumes at least one share at the n-th limit available in the order book. For example, a second limit trade-through completely consumes the first limit available and begins to consume the second limit of the order book. Our definition is inclusive in the sense that, if p is greater than q, any p-th limit trade-through is also a q-th limit trade-through. In this study, we will focus on second limit trades-through, and simply call them trades-through in what follows. Figure 4.1 shows an example of trade-through.

4.2.3 Empirical data

We now describe the empirical data that will be used in the remaining of the paper. We use Thomson-Reuters [1] tick-by-tick data of the Euronext-Paris limit order book for the stock BNP Paribas (BNPP.PA) from June 1st, 2010 to October 29th, 2010, i.e. 109 trading days. This data gives us trades (timestamp to the millisecond, volume and price) and quotes (volume, price, side of the order book) for the stock, from the opening to the close of the market. For each trading day, we extract the series of timestamps $\left(t_i^A\right)_{i\geq 1}$ and $\left(t_i^B\right)_{i\geq 1}$ of the trades-through.

4.2.4 Occurrences of trades-through

We look at the occurrences of trades-through on the different sides of the order book. Basic statistics are given in Table 4.1. We can see that for second limit trades-through, there are around 400 events per day on each side of the book.

4.2.5 Clustering

Trades-through are clustered both in physical time and in trade time (see Pomponio and Abergel [2010]). Here we study in detail several aspects of this problem that will be helpful for further modelling : is the global clustering of trades-through still true when looking only at one side of the

1. See 3.2.1 for more detail on TRTH data.

Limit considered	Number of trades-through per day (all)	Number of trades-through per day (bid side)	Number of trades-through per day (ask side)
2	829.0	401.8	427.2
3	124.1	59.0	65.1
4	30.5	14.6	15.9

TABLE 4.1 – Occurrences of trades-through at bid and ask sides for BNP Paribas.

Impact studied	Mean waiting time until next trade-through (in seconds)
$(\lambda^+ + \lambda^-) \to (\lambda^+ + \lambda^-)$	36.9
$(\Lambda^+ + \Lambda^-) \to (\lambda^+ + \lambda^-)$	51.8
$(\lambda^+) \to (\lambda^+ + \lambda^-)$	36.3
$(\Lambda^+) \to (\lambda^+ + \lambda^-)$	51.7
$(\lambda^-) \to (\lambda^+ + \lambda^-)$	37.5
$(\Lambda^-) \to (\lambda^+ + \lambda^-)$	51.7
$(\lambda^+) \to (\lambda^+)$	76.1
$(\Lambda^+) \to (\lambda^+)$	107.9
$(\lambda^-) \to (\lambda^-)$	71.6
$(\Lambda^-) \to (\lambda^-)$	98.1
$(\lambda^+) \to (\lambda^-)$	80.4
$(\Lambda^+) \to (\lambda^-)$	101.8
$(\lambda^-) \to (\lambda^+)$	91.1
$(\Lambda^-) \to (\lambda^+)$	111.6

TABLE 4.2 – Clustering of trades-through on bid and ask sides (on BNP Paribas data).

book ? If so, is there an asymmetry in trades-through clustering on the bid and on the ask sides ? Is there a cross-side effect for trades-through, in other words will a trade-through on one side of the book be followed more rapidly than usual by a trade-through on the other side of the book ? Which is the stronger from those different effects ?

In order to grasp the clustering of trades-though, we compute the mean of the distribution of waiting times between two consecutive trades-through, and we compare it with the mean waiting time between one trade (of any kind) and the next trade-through.

Table 4.2 summarizes our result on BNP Paribas stock in the considered period of study. We use the notation λ when looking at trades-through and Λ when looking at all the trades. When a specific side of the book is under scrutiny we mention it with a $+$ for ask side and a $-$ for bid side. For example, $(\Lambda^+) \to (\lambda^+ + \lambda^-)$ means that we look at the time interval between a trade at the ask side and the next trade-through, whatever its sign.

Analysing the first group of statistics ($(\lambda^+ + \lambda^-) \to (\lambda^+ + \lambda^-)$ and $(\Lambda^+ + \Lambda^-) \to (\lambda^+ + \lambda^-)$), we see that previous result on global clustering of trades-through is confirmed : you wait less the next trade-through when you already are on a trade-through, compared to when you are on a trade. Moreover, when looking at the second group of statistics, we see there is no asymmetry in this effect : both trades-through at the ask and at the bid are more closely followed in time by trades-through (whatever their sign), than trades at the bid and trades at the ask are.

The third group of statistics indicates that if you restrict the study to only one side of the book, the clustering is still valid. Finally, the fourth group of statistics shows that there seems to be a

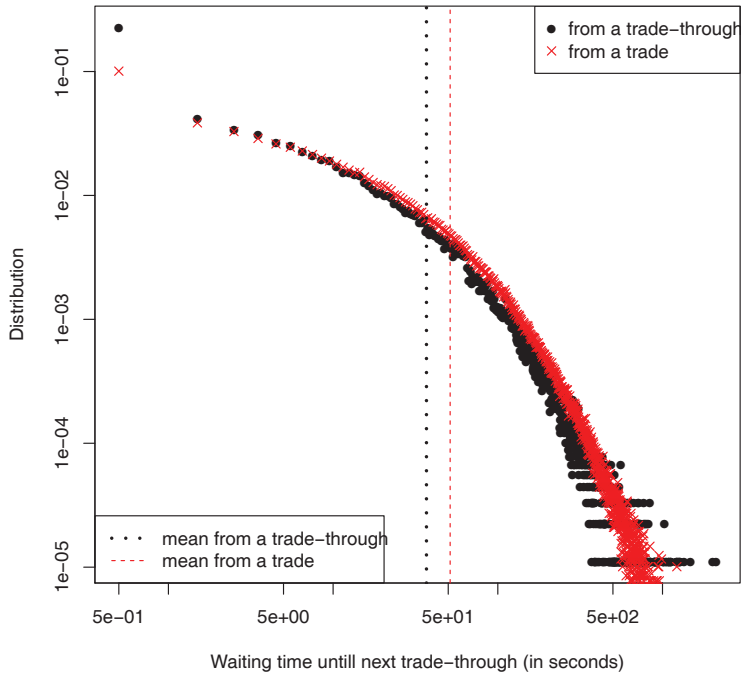

FIGURE 4.2 – Global trades-through clustering for BNP Paribas : Empirical distribution of the time intervals between a trade or a trade-through and the next trade-through.

cross-side effect of clustering of trades-through : a trade-through at one side of the book will be more closely followed in time by a trade-through on the other side of the book. But comparing the relative difference between mean waiting times of $(\lambda^+) \to (\lambda^+)$ and $(\Lambda^+) \to (\lambda^+)$, we have approximately a 30% decrease on the same side of the book. Whereas there is only a 20% decrease of mean waiting time between $(\lambda^+) \to (\lambda^-)$ and $(\Lambda^+) \to (\lambda^-)$, which reflects that cross-side clustering effect is weaker than same side clustering for trades-through.

Figure 4.2 plots the distributions of waiting times $(\lambda^+ + \lambda^-) \to (\lambda^+ + \lambda^-)$ and $(\Lambda^+ + \Lambda^-) \to (\Lambda^+ + \Lambda^-)$ studied in this paragraph.

In brief, looking at these distributions of durations gives us global tendencies on clustering and relative comparisons of the influences of trades-through with respect to limit order book sides. A more quantitative measurement of those effects will be done in the following part using the analysis of calibrated parameters of a fitted stochastic model, namely Hawkes processes.

4.2.6 Intraday timestamp distribution

We now look at the intraday distribution of timestamps for second-limit trades-through on the BNP Paribas stock, plotted on Figure 4.3. We can see that the distribution is globally the sum of two parts : a U-shape curve (linked to the global U-shape trading activity curve) and two peaks at very precise hours (2 :30 pm and 4 :00 pm – Paris time) reflecting the impact of major macro-economic news released at that moment of the day.

What is important for further modelling is to notice that it seems very difficult to find a pure stochastic model able to capture both the local behaviour and fluctuations of trades-through arrival

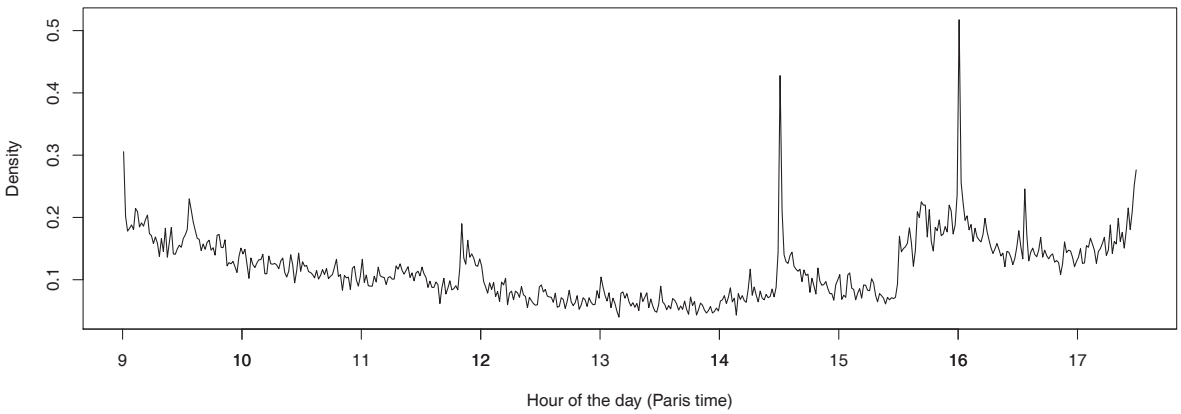

FIGURE 4.3 – Intraday distribution of timestamps of trades-through for the stock BNP Paribas on June–October 2010, using one-minute bins.

times and the two big peaks at very precise hours of the day. A first attempt may be to simply remove those peaks in the distribution. In the remaining of the paper, we will restrict ourselves to a two-hour interval, thus removing major seasonality effects.

4.3 Hawkes processes

Let us first recall standard definitions and properties of Hawkes processes. Hawkes processes are self-exciting processes which were introduced by Hawkes [1971]. Roughly speaking, in a self-exciting process, the occurrence of an event increases the probability of occurrence of another event. Hawkes [1971] specifically studies the case in which this increase is instantaneous and exponentially decaying, case which we mainly use in this paper. This section only gives results necessary to the remaining of the paper. Many references may provide more details : Bremaud [1981] and Daley and Vere-Jones [2003] are well-known general textbook treatments of point processes ; more precisely in our case, Bowsher [2007] provides a framework for "Generalized Hawkes models" which fully encompasses the processes used in this study.

In the following, we will note $\overline{\mathbb{R}}_+^* = \mathbb{R}_+^* \cup \{+\infty\}$ and $\overline{\mathbb{N}} = \mathbb{N} \cup \{+\infty\}$.

4.3.1 Definition

Let $M \in \mathbb{N}^*$. We start by defining a M-variate simple point process. Let $\{T_i, Z_i\}_{i\in\mathbb{N}^*}$ be a double sequence of random variables on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. $\{T_i\}_{i\in\mathbb{N}^*}$ is $\overline{\mathbb{R}}_+^*$-valued, (i) almost surely increasing, (ii) such that $T_i < T_{i+1}$ almost surely on $\{T_i < +\infty\}$, and (iii) such that $\lim_i T_i = +\infty$ almost surely. For all i, T_i represents the time of occurrence of the i-th event. $\{Z_i\}_{i\in\mathbb{N}^*}$ is $\{1, \dots, M\}$-valued. For all i, Z_i represents the type of the i-th event. For all $m = 1, \dots, M$, let us define $N^m(t) = \sum_{i\in\mathbb{N}^*} \mathbf{1}_{\{T_i \le t\}} \mathbf{1}_{\{Z_i = m\}}$ the counting process associated to events of type m. $N = (N^1, \dots, N^M)$ is the M-dimensional vector of counting processes. N (or equivalently $\{T_i, Z_i\}_{i\in\mathbb{N}^*}$) is called a M-variate simple point process.

Let us now precise the notion of (conditional) intensity. Let $\{\mathcal{F}_t\}_{t\in\mathbb{R}_+}$ be a filtration such that for all $t \in \mathbb{R}_+$, $\mathcal{F}_t \supset \mathcal{F}_t^N$, where $\mathcal{F}_t^N$ is the natural filtration of N. For all $m = 1, \dots, M$, a $\mathcal{F}_t$-intensity

of the monovariate simple point process N^m is a nonnegative $\mathcal{F}_t$-progressively measurable process $\{\lambda^m(t)\}_{t\in\mathbb{R}_+}$ such that

$$\forall 0 \leq s \leq t, \;\; \mathbf{E}\left[N^m(t) - N^m(s))|\mathcal{F}_s\right] = \mathbf{E}\left[\int_s^t \lambda^m(u)du\Big|\mathcal{F}_s\right] a.s. \tag{4.1}$$

With sufficient regularity conditions, one can derive that :

$$\lim_{t\to s, t>s} \frac{1}{t-s}\mathbf{E}\left[N^m(t) - N^m(s))|\mathcal{F}_s\right] = \lambda^m(s) \;\; a.s., \tag{4.2}$$

which is often taken as a definition for λ^m (see e.g. Hautsch [2004, Definition 2.1]). The intensity λ^m can thus be interpreted as the conditional probability to observe an event of type m per unit of time.

In this paper, we're interested in the special case of Hawkes processes with an exponentially-decaying kernel, i.e. processes for which the m-th coordinate N^m, $m \in \{1, \ldots, M\}$, admits an $\mathcal{F}_t^N$-intensity of the form

$$\lambda^m(t) = \lambda_0^m(t) + \sum_{n=1}^{M} \int_0^t \sum_{j=1}^{P} \alpha_j^{mn} e^{-\beta_j^{mn}(t-s)} dN_s^n, \tag{4.3}$$

where $\lambda_0^m : \mathbb{R}_+ \to \mathbb{R}_+$ is a deterministic function, the number P of exponential kernels is a fixed integer, and for all $m, n = 1, \ldots, M$, and $j = 1, \ldots, P$, α_j^{mn} and β_j^{mn} are positive constants.

In its simplest version with $P = 1$ and $\lambda_0^m(t) = \lambda_0^m$ a positive constant, the definition becomes for all $t \in \mathbb{R}_+$:

$$\lambda^m(t) = \lambda_0^m + \sum_{n=1}^{M} \int_0^t \alpha^{mn} e^{-\beta^{mn}(t-s)} dN_s^n, \tag{4.4}$$

Parameters α^{mn} and β^{mn} express the influence (scale and decay) of the past events of type n on the m-th coordinate of the process. It follows from this definition that two phenomena are present : self-excitation ($m = n$) and mutual excitation ($m \neq n$).

4.3.2 Stationarity condition

Taking here $P = 1$ and rewriting equation (4.4) using vector notations $\lambda = (\lambda^1, \ldots, \lambda^M)$, we have :

$$\lambda(t) = \lambda_0 + \int_0^t G(t-s) dN_s, \tag{4.5}$$

where G is a matrix-valued function of time defined by

$$G(t) = \left(\alpha^{mn} e^{-\beta^{mn}(t)}\right)_{m,n=1,\ldots,M}. \tag{4.6}$$

A multivariate point process is stationary if the joint distribution of any number of types of events on any number of given intervals is invariant under translation, i.e. if, forall integer r, the distribution of $\{N^{i_1}(t_1 + t) - N^{i_1}(s_1 + t), \ldots, N^{i_r}(t_r + t) - N^{i_r}(s_r + t)\}$ for $i_j \in \{1, \ldots, M\}, s_j \leq t_j, j = 1, \ldots, r$ does not depend on $t \in \mathbb{R}$. This of course cannot be the case if λ_0 is not constant. In the financial models of the following sections, such an assumption translates the fact that the level of market activity for all the events considered does not dramatically changes in the period considered, which may be far from obvious, as discussed in section 4.4.2.

In the monovariate case $M = 1$, Hawkes and Oakes [1974, Theorem 1] show that a sufficient condition for the existence of stationary point process with intensity as in equation (4.5) is that $\int_0^{+\infty} G(u)du < 1$. In our special case of exponential kernels, this gives the sufficient stationarity condition $\frac{\alpha^{11}}{\beta^{11}} < 1$.

Bremaud and Massoulie [1996, Theorem 7] generalizes the result to the multivariate case (and even much further to the non-linear case, which we don't need here [2]). The condition becomes that the spectral radius of the matrix

$$\Gamma = \int_0^{\infty} G(u)du = \left(\frac{\alpha^{mn}}{\beta^{mn}}\right)_{m,n=1,\dots,M} \tag{4.7}$$

is strictly smaller than 1. We recall that the spectral radius of the matrix G is defined as $\rho(G) = \max_{a\in\mathcal{S}(G)} |a|$, where $\mathcal{S}(G)$ denotes the set of all eigenvalues of G. In the two-dimensional case that will be used in the following sections, this can be written :

$$\frac{1}{2}\left(\frac{\alpha^{11}}{\beta^{11}} + \frac{\alpha^{22}}{\beta^{22}} + \sqrt{\left(\frac{\alpha^{11}}{\beta^{11}} - \frac{\alpha^{22}}{\beta^{22}}\right)^2 + 4\frac{\alpha^{12}}{\beta^{12}}\frac{\alpha^{21}}{\beta^{21}}}\right) < 1. \tag{4.8}$$

4.3.3 Maximum-likelihood estimation

In this section, we introduce the notation $\{T_i^m\}_{i\in\mathbb{N}^*}$ which is the re-enumerated sequence of events of type m. $\{T_i, Z_i\}_{i=1,\dots,N}$ is thus the ordered pool of all $\{\{T_i^m, m\}_{m=1,\dots,M}\}$. The log-likelihood of a multidimensional Hawkes process can be computed as the sum of the likelihood of each coordinate, and is thus written :

$$\ln \mathcal{L}(\{N(t)\}_{t\leq T}) = \sum_{m=1}^{M} \ln \mathcal{L}^m(\{N^m(t)\}_{t\leq T}), \tag{4.9}$$

where each term is defined by :

$$\ln \mathcal{L}^m(\{N^m(t)\}_{t\leq T}) = \int_0^T (1 - \lambda^m(s))\, ds + \int_0^T \ln \lambda^m(s) dN^m(s). \tag{4.10}$$

This partial log-likelihood can be computed as :

$$\begin{aligned} \ln \mathcal{L}^m(\{N^m(t)\}_{t\leq T}) &= T - \Lambda^m(0,T) \\ &+ \sum_{i:T_i\leq T} \mathbf{1}_{\{Z_i=m\}} \ln\left[\lambda_0^m(T_i) + \sum_{n=1}^{M}\sum_{j=1}^{P}\sum_{T_k^n<T_i} \alpha_j^{mn} e^{-\beta_j^{mn}(T_i - T_k^n)}\right] \end{aligned} \tag{4.11}$$

where $\Lambda^m(0,T) = \int_0^T \lambda^m(s)ds$ is the integrated intensity. Following Ozaki [1979], we compute this in a recursive way by observing that, thanks to the exponential form of the kernel :

$$\begin{aligned} R_j^{mn}(l) &= \sum_{T_k^n<T_l^m} e^{-\beta_j^{mn}(T_l^m - T_k^n)} \\ &= \begin{cases} e^{-\beta_j^{mn}(T_l^m - T_{l-1}^m)} R_j^{mn}(l-1) + \sum_{T_{l-1}^m\leq T_k^n<T_l^m} e^{-\beta_j^{mn}(T_l^m - T_k^n)} & \text{if } m \neq n, \\ e^{-\beta_j^{mn}(T_l^m - T_{l-1}^m)}\left(1 + R_j^{mn}(l-1)\right) & \text{if } m = n. \end{cases} \end{aligned} \tag{4.12}$$

2. The special non-linear case considered by Bremaud and Massoulie [1996] is the positive Lipschitz transform of the usual linear intensity.

The final expression of the partial log-likelihood may thus be written :

$$\begin{aligned}\ln \mathcal{L}^m(\{N^m(t)\}_{t\leq T}) = T &- \int_0^T \lambda_0^m(s)ds \\ &- \sum_{i:T_i\leq T}\sum_{n=1}^{M}\sum_{j=1}^{P}\frac{\alpha_j^{mn}}{\beta_j^{mn}}\left(1-e^{-\beta_j^{mn}(T-T_i)}\right) \\ &+ \sum_{l:T_l^m\leq T}\ln\left[\lambda_0^m(T_l^m)+\sum_{n=1}^{M}\sum_{j=1}^{P}\alpha_j^{mn}R_j^{mn}(l)\right],\end{aligned} \tag{4.13}$$

where $R_j^{mn}(l)$ is defined with equation (4.12) and $R_j^{mn}(0)=0$.

4.3.4 Testing the calibration

A general result on point processes theory states that a given non-Poisson process can be transformed into a homogeneous Poisson process by a stochastic time change. A standard monovariate version of this result and its proof can be found in Bremaud [1981, Chapter II, Theorem T16]. Bowsher [2007] has shown that this can be generalized in a multidimensional setting, which provides specification tests for multidimensional Hawkes models. We reproduce here its result, with slightly modified notations to accommodate the notations chosen here.

Theorem 4.3.1 (Bowsher [2007, Theorem 4.1]) *Let N be a M-variate point process on $\mathbb{R}_+^*$ with internal history[3] $\{\mathcal{F}_t^N\}_{t\in\mathbb{R}_+}$, and $M\geq 1$. Also let $\{\mathcal{F}_t\}_{t\in\mathbb{R}_+}$ be a history of N (that is, $\mathcal{F}_t^N\subseteq\mathcal{F}_t, \forall t\geq 0$), and suppose, for each m, that N^m has the $\mathcal{F}_t$-intensity λ^m where λ^m satisfies $\int_0^\infty \lambda^m(s)ds=\infty$ almost surely. Define for each m and all $t\geq 0$ the $\mathcal{F}_t$-stopping time $\tau^m(t)$ as the (unique) solution to*

$$\int_0^{\tau^m(t)}\lambda^m(u)du = t. \tag{4.14}$$

Then the M point processes $\{\tilde{N}^m\}_{m=1,\dots,M}$ defined by $\tilde{N}^m(t)=N^m(\tau^m(t))$, $\forall t\geq 0$, are independent Poisson processes with unit intensity. Furthermore, the durations of each Poisson process $\tilde{N}^m$ are given by

$$\Lambda^m(T_{i-1}^m,T_i^m)=\int_{T_{i-1}^m}^{T_i^m}\lambda^m(s)ds,\ \ \forall i\geq 2. \tag{4.15}$$

Let us compute this integrated intensity of the m-th coordinate of a multidimensional Hawkes process between two consecutive events T_{i-1}^m and T_i^m of type m :

$$\begin{aligned}\Lambda^m(T_{i-1}^m,T_i^m) = &\int_{T_{i-1}^m}^{T_i^m}\lambda_0^m(s)ds \\ &+\sum_{n=1}^{M}\sum_{j=1}^{P}\sum_{T_k^n<T_{i-1}^m}\frac{\alpha_j^{mn}}{\beta_j^{mn}}\left[e^{-\beta_j^{mn}(T_{i-1}^m-T_k^n)}-e^{-\beta_j^{mn}(T_i^m-T_k^n)}\right] \\ &+\sum_{n=1}^{M}\sum_{j=1}^{P}\sum_{T_{i-1}^m\leq T_k^n<T_i^m}\frac{\alpha_j^{mn}}{\beta_j^{mn}}\left[1-e^{-\beta_j^{mn}(T_i^m-T_k^n)}\right].\end{aligned} \tag{4.16}$$

3. Which we previously called *natural filtration*.

As in the log-likelihood computation, following Ozaki [1979], we observe that :

$$\begin{aligned} A_j^{mn}(i-1) &= \sum_{T_k^n < T_{i-1}^m} e^{-\beta_j^{mn}(T_{i-1}^m - T_k^n)} \\ &= e^{-\beta_j^{mn}(T_{i-1}^m - T_{i-2}^m)} A_j^{mn}(i-2) + \sum_{T_{i-2}^m \leq T_k^n < T_{i-1}^m} e^{-\beta_j^{mn}(T_{i-1}^m - T_k^n)}, \end{aligned} \tag{4.17}$$

so that the integrated density can be written for all $i \geq 2$:

$$\begin{aligned} \Lambda^m(T_{i-1}^m, T_i^m) = \int_{T_{i-1}^m}^{T_i^m} \lambda_0^m(s)ds + \sum_{n=1}^{M}\sum_{j=1}^{P} \frac{\alpha_j^{mn}}{\beta_j^{mn}} \Bigg[A_j^{mn}(i-1) \\ \times \left(1 - e^{-\beta_j^{mn}(T_i^m - T_{i-1}^m)}\right) + \sum_{T_{i-1}^m \leq T_k^n < T_i^m} \left(1 - e^{-\beta_j^{mn}(T_i^m - T_k^n)}\right) \Bigg], \end{aligned} \tag{4.18}$$

where A_j^{mn} is defined as in equation (4.17) with for all $j = 1, \ldots, P, A_j^{mn}(0) = 0$.

Hence, simply following the method in Bowsher [2007], we can easily define tests to check the goodness-of-fit of a Hawkes model to our empirical data. Since the integrated intensity $\Lambda^m(T_{i-1}^m, T_i^m)$ is a time interval of a homogeneous Poisson Process, we can test for each $m = 1, \ldots, M$: (i) whether the variables $\left(\Lambda^m(T_{i-1}^m, T_i^m)\right)_{i\geq 2}$ are exponentially distributed ; (ii) whether the variables $\left((\Lambda^m(T_{i-1}^m, T_i^m)\right)_{i\geq 2}$ are independent. In Section 4.4.3, the independence test will be carried out with a Ljung-Box test up to the twentieth term, and we will use a standard Kolmogorov-Smirnov test the empirical data against the exponential distribution.

In this section, we have recalled results with detailed computations that we will now use to build and estimate models for trades-through. Several reasons make this approach attractive. Studying point processes taking the stochastic intensity as the main tool allows to define straightforwardly multivariate models, while multidimensional generalizations of duration-based models for example may be cumbersome, and at least quite artificial (see Daley and Vere-Jones [2003, p.41]). Furthermore, as for the special case of linear Hawkes processes with exponentially-decaying kernel, the results recalled in this section show that interpretation of the few parameters (base intensity as background market activity, jumps and exponential decay of the influence of past events), statistical estimation and specification tests are straightforward. Note also that simulation algorithms of Hawkes processes are easily available (see Ogata [1981], Moller and Rasmussen [2005] and Appendix A), allowing further developments for the use of the estimated models. These practical properties make Hawkes processes good candidates for the study of the clustering of events empirically observed in order book models (see Hautsch [2004], Bowsher [2007], Large [2007] and Bacry et al. [2012]).

We now turn to the modelling of trades-through in an order book model.

4.4 A simple Hawkes model for trades-through

4.4.1 Model

Since empirical evidence shows that trades-through obviously occur in a clustered way, it makes sense to try to model them with self-exciting Hawkes processes. We thus define our basic model as follows. Let $\left(T_i^A\right)_{i\geq 1}$ be the point process of trades-through occurring on the ask side of the limit order book, and $\left(T_i^B\right)_{i\geq 1}$ be the point process of trades-through occurring on the bid side. Let N^A and N^B denote the associated counting processes. These two processes are assumed to form a two-dimensional Hawkes process with intensities λ^A and λ^B defined with parameters $(\alpha_{ij}, \beta_{ij})_{(i,j)\in\{A,B\}^2}$ as follows :

	Average	Median	Min	Max	Stdev
λ_0^A	1.01E-02	8.42E-03	6.62E-06	3.52E-02	6.27E-03
α_{AA}	4.13E+00	6.45E-01	3.53E-02	3.09E+01	6.03E+00
α_{AB}	4.33E-01	1.05E-01	1.00E-10	4.78E+00	8.41E-01
β_{AA}	3.70E+01	4.78E+00	1.84E-01	2.34E+02	5.21E+01
β_{AB}	2.48E+01	1.33E+00	1.00E-10	1.48E+03	1.44E+02

TABLE 4.3 – Statistics summary for the maximum-likelihood estimates of the ask side of model (4.19).

	Average	Median	Min	Max	Stdev
λ_0^B	1.09E-02	9.08E-03	2.46E-03	3.98E-02	6.35E-03
α_{BA}	3.68E-01	7.56E-02	3.83E-13	4.46E+00	6.77E-01
α_{BB}	4.81E+00	3.04E+00	2.08E-02	4.62E+01	7.00E+00
β_{BA}	9.61E+00	1.46E+00	1.00E-10	1.00E+02	1.92E+01
β_{BB}	3.98E+01	2.00E+01	3.71E-02	3.75E+02	5.52E+01

TABLE 4.4 – Statistics summary for the maximum-likelihood estimates of the bid side of model (4.19).

$$\begin{aligned}
\lambda^A(t) &= \lambda_0^A(t) + \int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A + \int_0^t \alpha_{AB} e^{-\beta_{AB}(t-s)} dN_s^B, \\
\lambda^B(t) &= \lambda_0^B(t) + \int_0^t \alpha_{BA} e^{-\beta_{BA}(t-s)} dN_s^A + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B.
\end{aligned} \tag{4.19}$$

This is a standard bivariate Hawkes model of Section 4.3 with $P = 1$.

4.4.2 Calibration

We compute the maximum-likelihood estimates for the parameters of our model, using computations presented in Section 4.3, and the empirical data previoulsy described in Section 4.2.3. In a first step, we make the assumptions that base intensities λ_0^A and λ_0^B are constants. Taking into account the huge variations of trading activity during the day, we restrict our empirical observations to a two-hour interval, from 9 :30 am to 11 :30 am. Considering Figure 4.3, this may make a bit more realistic the stationarity assumption discussed in Section 4.3.2. [4]

During these five months of trading, we count in average 2737 trades each day during this time interval, 206 of which are trades-through (100 on the ask side and 106 on the bid side). Thus roughly 8% of the recorded transactions are trades-through.

Tables 4.3 and 4.4 summarize the statistics on the estimated values on the ask and bid sides.

These tables show that the median half-lives associated to the kernels AA, AB, BA and BB are respectively 145, 521, 474 and 35 milliseconds. It appears that we observe very large variations in the results of the numerical maximization of the likelihood. However, whatever the absolute size of the parameters, it is clear that the cross-excitation effect, i.e. the excitation of trades-through of a given side by the occurrence of trades-through on the opposite side, is much weaker than the

4. Actually, one can verify ex-post that the calibrated parameters obtained in the sequel for the constant parameters models satisfy the stationarity condition (4.8) on average (for the mean and median parameters), for example.

	Average	Median	Min	Max	Stdev
λ_0^A	1.18E-02	9.77E-03	3.47E-03	3.66E-02	6.53E-03
α_{AA}	6.02E+00	4.85E+00	1.00E-10	3.09E+01	6.14E+00
β_{AA}	4.76E+01	3.63E+01	3.92E-02	2.34E+02	4.87E+01
λ_0^B	1.26E-02	1.13E-02	3.80E-03	4.25E-02	7.00E-03
α_{BB}	8.05E+00	5.61E+00	1.28E-02	4.79E+01	8.81E+00
β_{BB}	6.64E+01	4.72E+01	2.64E-02	3.91E+02	7.22E+01

TABLE 4.5 – Statistics summary for the maximum-likelihood estimates of the simplified model (4.20) with λ_0^A and λ_0^B constant.

self-excitation effect, which translates the clustering of trades-through on a given side. The average value of α_{AB} is 9.5 times smaller that the average value of α_{AA}, while at the same time the associated exponential decay β_{AB} is only 1.5 times smaller than the average β_{AA}. The instantaneous effect is thus much smaller while its half-life is not significantly longer. This observation is also valid for the average α_{BA} which is 13 times smaller than the average α_{BB}, while the average exponentials decays differ only by a factor 4.

In other words, the ratio $\frac{\alpha}{\beta}$, which is equal to the total integrated intensity of an exponential kernel $\int_0^{+\infty} \alpha e^{-\beta u} du$, is much weaker in the cross-excitation cases (taking the average values, $\frac{\alpha_{AB}}{\beta_{AB}} = 0.017$, $\frac{\alpha_{BA}}{\beta_{BA}} = 0.038$) than in the self-excitation cases (still using the average values, $\frac{\alpha_{AA}}{\beta_{AA}} = 0.111$, $\frac{\alpha_{BB}}{\beta_{BB}} = 0.120$). Therefore, we can focus on the calibration and use of a simpler model, where trades-through are modelled by two one-dimensional Hawkes processes, with no cross-excitation :

$$\begin{aligned} \lambda^A(t) &= \lambda_0^A(t) + \int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A, \\ \lambda^B(t) &= \lambda_0^B(t) + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B. \end{aligned} \tag{4.20}$$

Table 4.5 summarizes the statistics of the estimated values of this simplified model with the assumption λ_0^A and λ_0^B constant. Values are similar to the previous case, confirming that the cross-effects were negligible. The effect of this simplification will be further discussed with the goodness-of-fit tests.

Finally, in an attempt to grasp small variations of activity independent of the clustering of the trades-through, following ideas presented in Bowsher [2007], we test a third version of the model by getting rid of the assumptions stating that λ_0^A and λ_0^B are constants. In this version of the simplified model (4.20), base intensities $\lambda_0^A(t)$ and $\lambda_0^B(t)$ are piecewise-linear continuous functions on the subdivision $(9:30 < 10:00 < 10:30 < 11:00 < 11:30)$ of the time interval $[9:30 \text{ am}; 11:30 \text{ am}]$. Note that this assumption implies that the process is not stationary anymore. Tables 4.6 and 4.7 summarize the statistics on the estimated values on the ask and bid sides.

Let us now discuss the goodness-of-fit of these three calibrations.

	Average	Median	Min	Max	Stdev
$\lambda_0^A(9:30)$	1.94E-02	1.65E-02	1.00E-20	5.40E-02	1.29E-02
$\lambda_0^A(10:00)$	1.13E-02	9.81E-03	1.00E-20	3.72E-02	9.73E-03
$\lambda_0^A(10:30)$	1.33E-02	1.25E-02	1.00E-20	5.01E-02	9.30E-03
$\lambda_0^A(11:00)$	7.67E-03	4.97E-03	1.00E-20	5.45E-02	9.93E-03
$\lambda_0^A(11:30)$	1.32E-02	1.03E-02	1.00E-20	1.54E-01	1.68E-02
α_{AA}	6.62E+00	5.10E+00	3.64E-13	3.09E+01	6.25E+00
β_{AA}	5.64E+01	4.61E+01	1.00E-20	2.34E+02	5.14E+01

TABLE 4.6 – Statistics summary for the maximum-likelihood estimates of the ask side of model (4.20) with λ_0^A and λ_0^B piecewise-linear continuous functions.

	Average	Median	Min	Max	Stdev
$\lambda_0^B(9:30)$	1.99E-02	1.67E-02	1.35E-03	6.51E-02	1.31E-02
$\lambda_0^B(10:00)$	1.25E-02	1.06E-02	1.00E-20	5.38E-02	1.00E-02
$\lambda_0^B(10:30)$	1.26E-02	1.14E-02	1.00E-20	5.15E-02	9.18E-03
$\lambda_0^B(11:00)$	9.32E-03	7.77E-03	1.00E-20	5.65E-02	9.65E-03
$\lambda_0^B(11:30)$	1.33E-02	9.16E-03	7.06E-14	1.26E-01	1.53E-02
α_{BB}	8.20E+00	5.60E+00	8.70E-04	4.79E+01	8.68E+00
β_{BB}	6.82E+01	5.15E+01	1.25E-03	3.91E+02	6.98E+01

TABLE 4.7 – Statistics summary for the maximum-likelihood estimates of the bid side of model (4.20) with λ_0^A and λ_0^B piecewise-linear continuous functions.

Model	Performance	2.5%		1%	
Full Model (4.19) with λ_0 constant	4 passed	70	(64.2)	83	(76.1)
	3 passed	29	(26.6)	26	(23.9)
	2 or less	10	(9.2)	0	(0.0)
No Cross (4.20) with λ_0 constant	4 passed	59	(54.1)	77	(70.6)
	3 passed	35	(32.1)	25	(22.9)
	2 or less	15	(13.8)	7	(6.4)
Full Model (4.19) with λ_0 piecewise-linear	4 passed	84	(77.1)	94	(86.2)
	3 passed	20	(18.3)	13	(11.9)
	2 or less	5	(4.6)	2	(1.8)
No Cross (4.20) with λ_0 piecewise-linear	4 passed	83	(76.1)	95	(87.2)
	3 passed	20	(18.3)	14	(12.8)
	2 or less	6	(5.5)	0	(0.0)

TABLE 4.8 – Performance of the calibration of the Hawkes models. For each model, this table gives the number of trading days (out of 109) where 4, 3, or 2 or less tests out of for where successfully passed. The four tests are two independence Ljung-Box tests and two Kolmogorov-Smirnov tests for the exponential distribution. Values in parentheses are percentages.

4.4.3 Goodness-of-fit

For each trading days, we have extracted the time series $\left(t_i^A\right)_{i\geq 1}$ and $\left(t_i^B\right)_{i\geq 1}$. For each of the three models discussed above and for each trading day, we can compute the integrated intensities $\left(\Lambda^A(t_i^A, t_{i+1}^A)\right)_{i\geq 1}$ and $\left(\Lambda^B(t_i^B, t_{i+1}^B)\right)_{i\geq 1}$ defined as in (4.18) and perform the four tests of goodness-of-fit described in Section 4.3. This gives us four tests per model and per trading day. Table 4.8 shows the results of the tests for a risk of the first kind equal to 1% and 2.5%.

These results confirm that the cross-excitation of trades-through of one side of the book on the other side is weak. In the case where λ_0 is constant, the percentage of trading days where the model passes all 4 statistical tests is 76% in the full specification case, and stays at 71% when cross-excitation is not taken into account. And in the case where λ_0 is allowed to vary as a piecewise-linear continuous function, these two percentages are even equal : in this latter case, we don't have any statistical improvement by including the cross-excitation effect.

Moreover, these results show that adding more flexibility in the modelling of λ_0 using piecewise-linear continuous functions helps the model to grasp the dynamics of trades-through : all tests are passed in more that 87% of the trading days tested in both cases.

4.5 Toy-trading strategy based on trades-through autocorrelated arrival times

4.5.1 Hawkes signal computation

We want to empirically test the quality of the trades-through model we developed in (4.20). So, we build a signal from the fitted parameters and we use it in a toy-trading strategy. Many different toy-trading signals are possible : for example, Carlsson et al. [2007] looked at a simple buy-sell signal based on intensity ratio of general trades. Here, our approach is different as we are not comparing the relative strengths of buy and sell signals, modelled by the two intensities λ^A and λ^B. We want to use the particular structure of the autocorrelated arrival of trades-through on the same side of the order book and their stronger market-impact (cf 3.4). So, we will use a 'momentum' toy-trading strategy by looking at each trade-trade-through, taking position in the same direction

of it (at the midprice) and closing the position (also at the midprice) after a given duration. Our empirical validation of the model will be based on the comparison of a reference strategy in which we invest one unity of the stock after each trade-through, and a weighted strategy in which we invest the normalized value of the signal after each trade-through. If the average PnL is higher in the latter case, we can consider it as an empirical validation of our modelling : our signal (based on fitted parameters) included information about the trades-through particular structure.

Assume the intensity is given by the model with no cross-excitation (4.20) and λ_0 constant, which was previously statistically validated. The signal $\phi(\delta)$ will be the conditional expectation of the intensity process at time $t+\delta$ given the process history $\mathcal{F}_t^N$. By taking the conditional expectation of the equation (4.20), ϕ is solution of the following integral equation :

$$\phi(\delta) = E(\lambda(t+\delta)|\mathcal{F}_t^N) = \lambda_0 + e^{-\beta\delta}\int_0^t \alpha e^{-\beta(t-u)}dN_u + \int_0^\delta \alpha e^{-\beta(\delta-u)}\phi(u)du \tag{4.21}$$

Using the Laplace transform technique (as Hawkes [1971]), we transform the integral equation in a simple equation of the Laplace transform of $\phi(\delta)$, denoted hereafter $\mathcal{L}\phi(p)$:

$$\mathcal{L}\phi(p) = \frac{\lambda_0}{p} + \frac{1}{p+\beta}\int_0^t \alpha e^{-\beta(t-u)}dN_u + \frac{\alpha}{p+\beta}\mathcal{L}\phi(p) \tag{4.22}$$

At this point, we obtain the Laplace transform of $\phi(\delta)$:

$$\mathcal{L}\phi(p) = \frac{\lambda_0}{p}\left(\frac{p+\beta}{p+\beta-\alpha}\right) + \frac{1}{p+\beta-\alpha}\int_0^t \alpha e^{-\beta(t-u)}dN_u \tag{4.23}$$

Using the inverse Laplace transform, we obtain the final expression for the conditional intensity :

$$\phi(\delta) = E(\lambda(t+\delta)|\mathcal{F}_t^N) = \lambda_0\frac{\beta-\alpha e^{-(\beta-\alpha)\delta}}{\beta-\alpha} + e^{-(\beta-\alpha)\delta}\int_0^t \alpha e^{-\beta(t-u)}dN_u \tag{4.24}$$

4.5.2 Results

We are not interested in doing a complete quantitative parameters optimization and design of this toy-trading strategy. But, we would like to verify if there is a qualitative improvement of the mean PnL results when using the Hawkes fitted parameters in the input signal. The results obtained with this toy-trading strategy on the French stock BNP-Paribas and on the European DAX future in the figure 4.4 show that this is the case.

It is worth noticing that in order to build a profitable strategy with trades-through based on Hawkes calibrated parameters, one has to predict the current parameters of the process with the previous parameters obtained the days before. As already said, the parameters of the calibration may vary from day to day. But the statistical analysis we carried out showed that a prediction of the next calibrated parameters with the average of the parameters calibrated on the two weeks before gives satisfactory results, allowing one to obtain most of the excess return of the mid-to-mid strategy without knowing beforehand the realized values of the parameters.

We also want to mention that, to take further advantage of the structure of trades-through in a trading strategy, it would be interesting to look at a more developped strategy based on limit orders placement, and not only to look at 'midprice' signals. It would also be interesting to consider strategies done with market orders only, but with adaptive entry and exit times based on trades-through. We let this interesting development for future research.

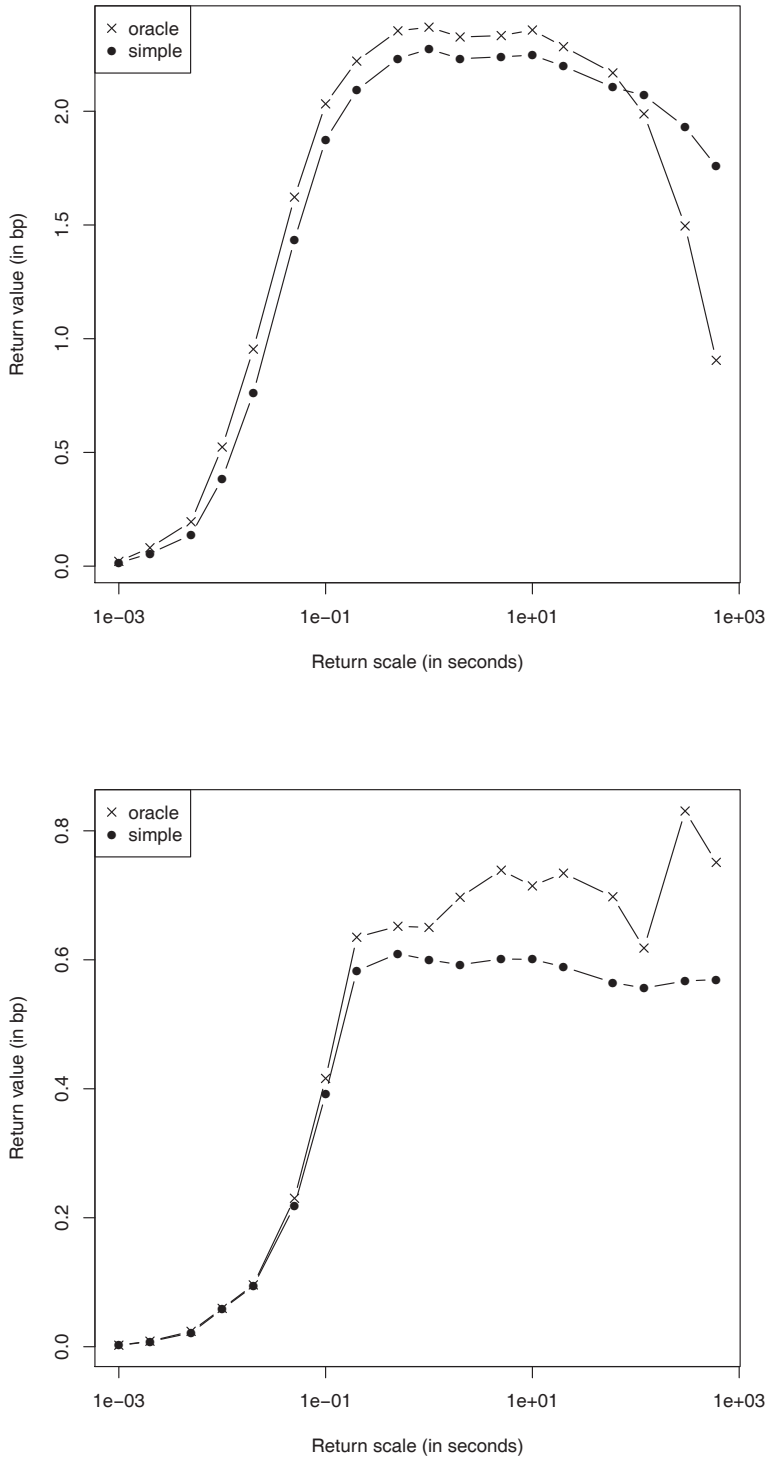

FIGURE 4.4 – "Mid-to-mid" strategy using trades-through on DAX Future FDX@ *(down)* and BNPP stock *(up)*. Benchmark (simple) : invest 1 after an observed trade-through. Hawkes (oracle) : invest $\phi(\delta)$ after an observed trade-through.

4.6 Searching for inhibitory effects in trades-through

4.6.1 Models

In the previous sections of this chapter, we looked at modelling the intensity of trades-through with linear intensity Hawkes models. As we saw, those models are statistically adapted to our purpose. But the linear intensity Hawkes models have some restriction in the range of parameters available for the calibration. In fact, in order for the intensity to be positive in this family of models, we have to limit ourselves to positive base intensities λ_0, positive jump sizes α and positive exponential decay parameters β. In this type of model, this implies that the intensity can never be less than λ_0 and can only temporarily increase after the arrival of an event.

As mentioned in Bowsher [2007], such constraints may appear to limit the applicability of a Hawkes approach to situations where the arrival of events can only increase the future arrival rate of events. So, it may be interesting to at least allow for negative jumps ($\alpha < 0$) in the modelling of the intensity, as the arrival of events of a certain type may inhibit the arrival of events of some other type and, as a consequence, may temporarily decrease their intensity, for example.

In the following, we will use two different models where the constraints $\alpha \geq 0$ have been removed. The first model is simply obtained by taking the exponential of the previous intensity model (4.19), as an example of a general log-linear intensity :

$$\begin{aligned}
\lambda^A(t) &= \exp\left(\lambda_0^A + \int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A + \int_0^t \alpha_{AB} e^{-\beta_{AB}(t-s)} dN_s^B \right), \\
\lambda^B(t) &= \exp\left(\lambda_0^B + \int_0^t \alpha_{BA} e^{-\beta_{BA}(t-s)} dN_s^A + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B \right), \\
\text{with} &: (\beta_{AA}, \beta_{AB}, \beta_{BA}, \beta_{BB}) \in {\mathbb{R}_+^*}^4 \text{ and } (\lambda_0^A, \lambda_0^B, \alpha_{AA}, \alpha_{AB}, \alpha_{BA}, \alpha_{BB}) \in \mathbb{R}^6.
\end{aligned} \tag{4.25}$$

One can also consider the following threshold effect model :

$$\begin{aligned}
\lambda^A(t) &= \lambda_0^A + \max\left(\int_0^t \alpha_{AA} e^{-\beta_{AA}(t-s)} dN_s^A + \int_0^t \alpha_{AB} e^{-\beta_{AB}(t-s)} dN_s^B, \epsilon^A - \lambda_0^A \right), \\
\lambda^B(t) &= \lambda_0^B + \max\left(\int_0^t \alpha_{BA} e^{-\beta_{BA}(t-s)} dN_s^A + \int_0^t \alpha_{BB} e^{-\beta_{BB}(t-s)} dN_s^B, \epsilon^B - \lambda_0^B \right), \\
\text{with} &: (\lambda_0^A, \lambda_0^B, \beta_{AA}, \beta_{AB}, \beta_{BA}, \beta_{BB}) \in {\mathbb{R}_+^*}^6 \text{ and } (\alpha_{AA}, \alpha_{AB}, \alpha_{BA}, \alpha_{BB}) \in \mathbb{R}^4.
\end{aligned} \tag{4.26}$$

It simply reflects the fact that once the financial market intensity for type m events has fallen to the threshold level ϵ^m (typically a very small strictly positive value), the occurrence of otherwise inhibitory events has no more (immediate) impact.

Both models have been considered in the study of Bowsher [2007], where stationarity conditions are even presented in some restricted cases for which the general result of Bremaud and Massoulie [1996, Theorem 7] is valid [5].

4.6.2 Results

The calibration of the models is still done by maximum-likelihood estimation, using equations (4.9) and (4.10). As the kernel is not linear with an exponential term, the recursive computation technique used before (see (4.12)) is not possible anymore. This forces us to a more intensive

5. Basically, the stationarity of those processes is still valid when one considers the above models with a cut-off maximal value (see Bowsher [2007, Theorem 2.4]).

computation with the numerical evaluation of the integral present in the log-likelihood expression (4.10).

After calibrating the two previous models on the same data than before, we looked at the number of parameters $\alpha < 0$, which indicates empirical inhibitory effects. We only found 21 such parameters for the first model (4.26) and 10 for the second model (4.27), out of the 436 parameters α calibrated during the whole period (which represents less than 5% in both cases). This result confirms our initial choice of Hawkes linear intensity models (where all α parameters were forced to be positive), as there is no statistically significant sign of inhibitory effects in trades-through arrival times data.

4.7 Calibration of the empirical decay kernel of trades-through arrival times by a non-parametric method

In this section, we want to estimate the empirical shape of the trades-through decay kernel, without assuming any exponential-linked forms. Exponential-based Hawkes kernels are very convenient because they allow to fully compute the log-likelihood of the process when doing maximum likelihood estimations. They also greatly increase the computation time of kernel-based values, thanks to the recursive computation technique 4.12.

4.7.1 A reminder on nonparametric Hawkes process calibration

We use the MISD ("Model Independent Stochastic Declustering") algorithm of Marsan and Lengline [2008] for calibrating general Hawkes processes. In our case, there is no space or earthquake magnitude in the intensity model we consider, so we can use the simplified framework review of Lewis and Mohler [2011]. The conditional intensity λ, which represents the rate of events, is modelled by :

$$\lambda(t) = \mu(t) + \sum_{t_i<t} g(t - t_i) \tag{4.27}$$

We already know (cf 4.4.2) that the background rate $\mu(t)$ may be considered constant in time or be time-varying. For the sake of simplicity, we will restrict ourselves in the sequel to the case of a constant background rate μ. We present the complete calibration method in the Appendix B and briefly discuss how empirical results are modified at the end of the following subsection.

The intensity $\lambda(t)$ depends on past events t_i through the triggering kernel $g(t)$. As we want to investigate the real empirical form of the kernel, there is no exponential-form assumptions on $g(t)$.

Let p_{ij} be the probability that event j triggers event i. There are tow different kinds of events : the background events ($i = j$) which are due to the background rate μ of the intensity, and the offspring events ($i > j$) which are the responses of previous events ($j = 1, ..., i - 1$) through the kernel $g(t)$.

We use the notation : $A_m = \{(i, j) \in [1, n]^2 | (t_i - t_j) \in [m\delta t, m\delta t + \delta t]\}$. The numerical discretization of the triggering kernel function g is $g_m = g(m\delta t)$, where δt is the minimal time-variation of the kernel. The upper script ^{k} is used to indicate a value obtained after the k-th step of the algorithm.

The MISD algorithm is an example of the very general EM ("Expectation Maximization") class of statistical calibration algorithms (see Bishop [2007], Hastie et al. [2009] and Duda et al. [2000], for an introductory presentation of the EM algorithm). After guessing a starting point, it works by iterating the following two steps, until convergence is reached.

The expectation step updates the offspring events rates p_{ij} (with $i > j$) and the background events rates p_{ii} by :

$$
\begin{aligned}
p_{ij}^k &= \frac{g^k(t_i - t_j)}{\mu^k + \sum_{j=1}^{i-1} g^k(t_i - t_j)} = \frac{g^k(t_i - t_j)}{\lambda^k(t_i)}, \\
p_{ii}^k &= \frac{\mu^k}{\mu^k + \sum_{j=1}^{i-1} g^k(t_i - t_j)} = \frac{\mu^k}{\lambda^k(t_i)}.
\end{aligned} \tag{4.28}
$$

Then, the maximization step estimates the new parameter values of the triggering weight g and the background weight μ with :

$$
\begin{aligned}
g_m^{k+1} &= \frac{1}{n\delta t} \sum_{(i,j)\in A_m} p_{ij}^k, \\
\mu^{k+1} &= \frac{1}{T} \sum_{i=1}^{n} p_{ii}^k, \text{ where } T \text{ is the duration of the time series.}
\end{aligned} \tag{4.29}
$$

Finally, we want to mention that Lewis and Mohler [2011] developed a new algorithm (less intuitive, but faster than MISD on big data) for calibrating general linear intensity based models of the form 4.27 which is based on a Maximum Penalized Likelihood Estimation (MPLE). We benchmarked the following empirical results based on the MISD algorithm with the MPLE algorithm and we found no statistical difference.

4.7.2 Application to trades-through arrival times

Empirical kernels are more peaked than exponential ones. The shape of the curve indicates that a power-law decay kernel might be more appropriate to model trades-through kernel decay. Indeed, a power-law fit of the empirical kernel gives satisfactorily results with R^2 statistics about 90% ($R^2_{bid} = 94\%$ and $R^2_{ask} = 87\%$) with empirical slopes of -1.2 ($slope_{bid} = -1.16$ and $slope_{ask} = -1.31$).

There is no significant difference when generalizing our calibration with a time-varying background rate (detailed in the Appendix B).

4.8 Conclusion

We have studied in this paper a model for trades-through based on Hawkes processes. We have shown that the clustering properties of trades-through can be well modelled with such self-exciting processes. Although calibration results may vary a lot from trading day to trading day, general patterns remain, such as the weak cross-excitation effects.

The clustering of trades-through highlighted here may not be a surprise for a market practitioner, but our results may help describing periods of high liquidity consumption in order book models. Our observations are related to the dynamics of aggressive market orders described e.g. in Large [2007]. Note though that the set of aggressive market orders, usually defined as the set of market orders that move the best price, contains, but is not equal to, our set of trades-through defined in Section 4.2.2. The link between order flows and price changes is of prime interest when modelling order books, and aggressive market orders are essential to a good understanding of the market impact (see e.g. Eisler et al. [2012]). Our empirical study of the clustering of trades-through hopefully contribute to that line of work. Both the nature of our observations and the simplicity of the models could lead to future practical work on trades-through-based trading strategies and order

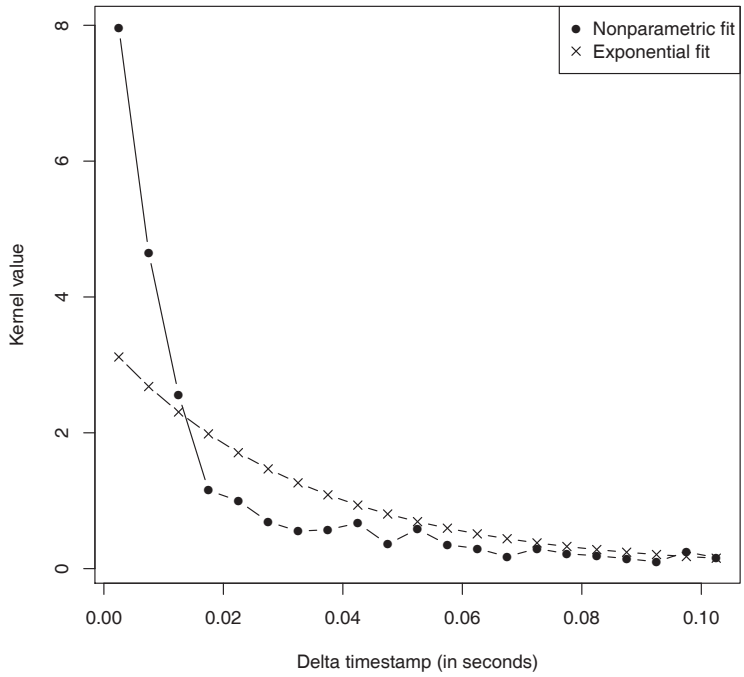

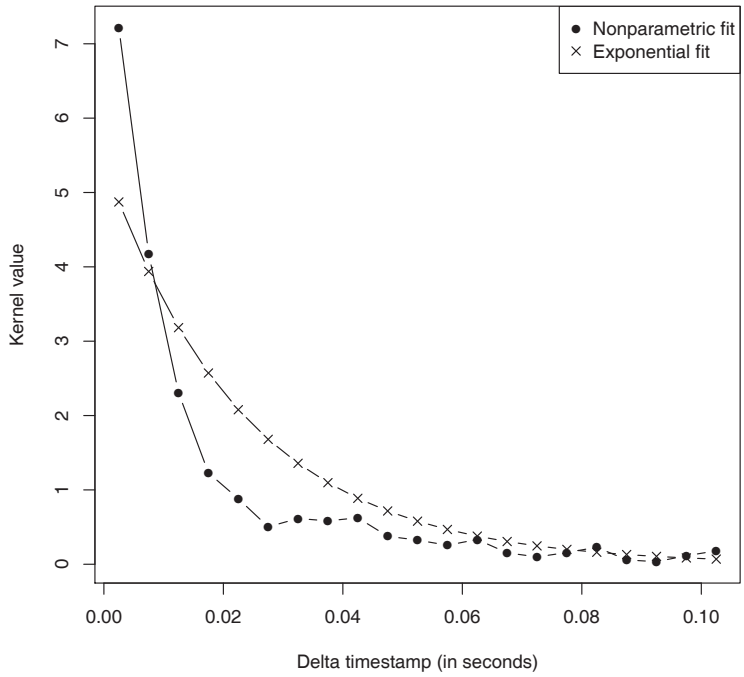

FIGURE 4.5 – Comparison of trades-through kernel calibration (nonparametric with 20 bins vs exponential) for the stock BNP Paribas on June–October 2010 (10 a.m. to 2 p.m.) on the bid side *(up)* and the ask side *(down)* of the limit order book.

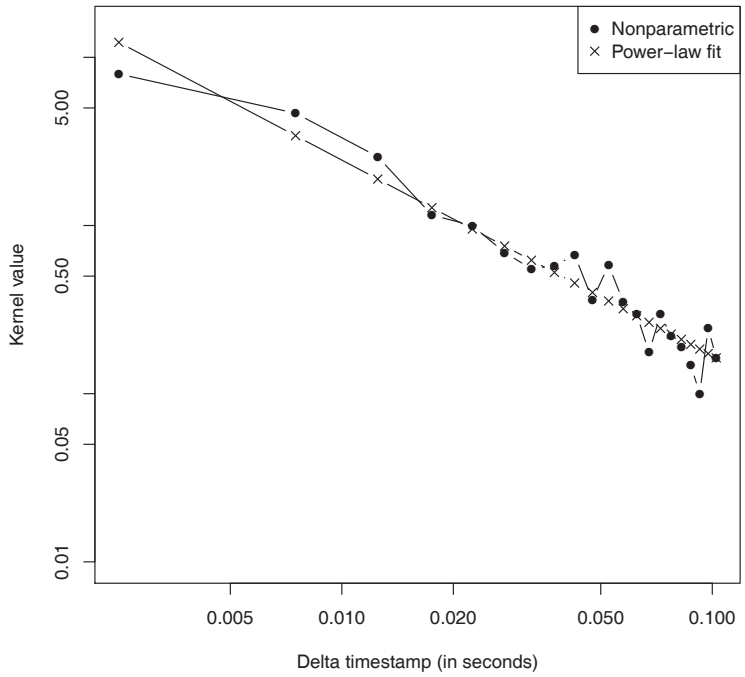

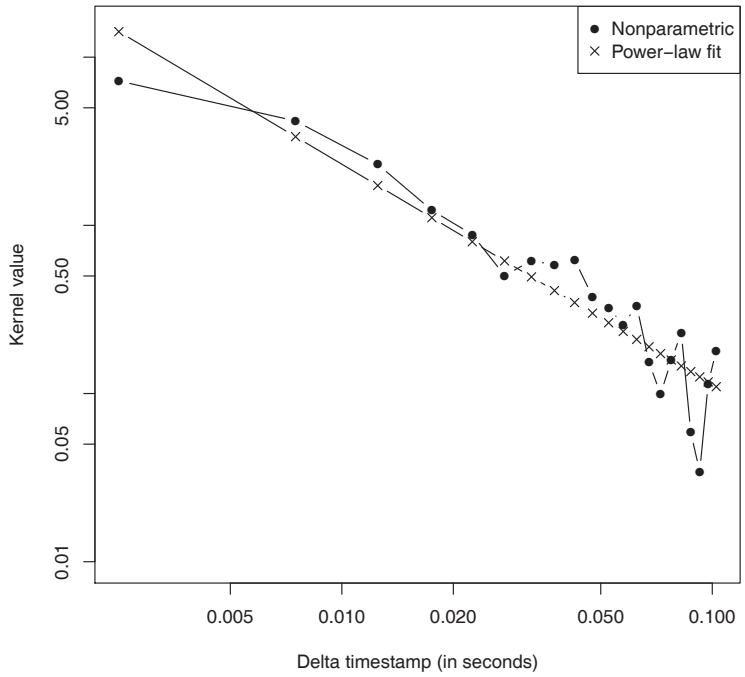

FIGURE 4.6 – Power-law fit of the empirical nonparametric trades-through kernel for the stock BNP Paribas on June–October 2010 (10 a.m. to 2 p.m.) on the bid side *(up)* and the ask side *(down)* of the limit order book (on a log-log scale).

Δ_{TS} (in seconds)		Average	Median	Min	Max	Stdev
0.0025	λ_1	7.96	7.57	0	21.0	4.14
0.0075	λ_2	4.65	4.75	0	14.7	3.04
0.0125	λ_3	2.56	2.40	0	9.99	2.16
0.0175	λ_4	1.15	3.33E-03	0	5.65	1.48
0.0225	λ_5	1.00	0	0	10.3	1.56
0.0275	λ_6	6.91E-01	0	0	9.87	1.35
0.0325	λ_7	5.52E-01	0	0	4.25	9.73E-01
0.0375	λ_8	5.73E-01	0	0	3.73	8.85E-01
0.0425	λ_9	6.70E-01	0	0	3.89	9.88E-01
0.0475	λ_{10}	3.60E-01	0	0	3.63	7.38E-01
0.0525	λ_{11}	5.83E-01	0	0	4.16	9.50E-01
0.0575	λ_{12}	3.52E-01	0	0	3.38	7.35E-01
0.0625	λ_{13}	2.96E-01	0	0	3.41	7.08E-01
0.0675	λ_{14}	1.76E-01	0	0	3.65	5.53E-01
0.0725	λ_{15}	2.96E-01	0	0	3.01	6.54E-01
0.0775	λ_{16}	2.20E-01	0	0	2.71	5.77E-01
0.0825	λ_{17}	1.90E-01	0	0	3.14	5.63E-01
0.0875	λ_{18}	1.47E-01	0	0	2.12	4.48E-01
0.0925	λ_{19}	1.00E-01	0	0	2.19	3.87E-01
0.0975	λ_{20}	2.47E-01	0	0	5.55	7.67E-01

TABLE 4.9 – Statistics summary for the expectation-maximization estimates of the bid side of model (4.20) with 20 bins.

Δ_{TS} (in seconds)		Average	Median	Min	Max	Stdev
0.0025	λ_1	7.21	6.89	0	26.4	4.93
0.0075	λ_2	4.17	3.96	0	19.7	3.21
0.0125	λ_3	2.31	2.15	0	8.73	1.94
0.0175	λ_4	1.23	1.72E-01	0	6.78	1.62
0.0225	λ_5	8.81E-01	0	0	5.07	1.30
0.0275	λ_6	5.03E-01	0	0	7.14	1.07
0.0325	λ_7	6.15E-01	0	0	7.53	1.28
0.0375	λ_8	5.83E-01	0	0	8.16	1.25
0.0425	λ_9	6.25E-01	0	0	5.71	1.23
0.0475	λ_{10}	3.76E-01	0	0	5.63	8.19E-01
0.0525	λ_{11}	3.25E-01	0	0	2.81	7.06E-01
0.0575	λ_{12}	2.56E-01	0	0	3.44	6.07E-01
0.0625	λ_{13}	3.32E-01	0	0	2.77	7.12E-01
0.0675	λ_{14}	1.53E-01	0	0	2.35	4.57E-01
0.0725	λ_{15}	9.89E-02	0	0	3.44	4.55E-01
0.0775	λ_{16}	1.57E-01	0	0	3.77	5.52E-01
0.0825	λ_{17}	2.29E-01	0	0	3.70	6.77E-01
0.0875	λ_{18}	5.87E-02	0	0	1.59	2.19E-01
0.0925	λ_{19}	3.43E-02	0	0	1.66	1.89E-01
0.0975	λ_{20}	1.15E-01	0	0	2.19	3.91E-01

TABLE 4.10 – Statistics summary for the expectation-maximization estimates of the ask side of model (4.20) with 20 bins.

book modelling.

We have improved a toy-trading strategy based on the autocorrelation of trades-through arrival times, thanks to our Hawkes model of trades-through. We have also shown the absence of empirical inhibitory effects. Finally, we have obtained the complete empirical profile of trades-through decay kernel using a nonparametric statistical method, and we found its decay to be faster than exponential, and closer to a power-law.

It would be interesting to generalize the indicator used in our toy-trading strategy to this power-law decay kernel of trades-through, for example. As the power-law decay kernel is also more adapted to describe trades-through time decay, one can also expect to increase the stability of the calibrated parameters across time, and so contribute to increase the predictability of the model when using it in a trading strategy. We let this interesting development for further research.

It would also be interesting to compare the statistical calibration results obtained here with trades-through with those of other types of events, all the trades or only the aggressive trades (those who move the midprice as defined in Large [2007]) for example. This would confirm the different nature of those different type of events.

Appendix A – Simulation of Hawkes processes

Ogata [1981] proposes an algorithm for the simulation of Hawkes processes, based on a general procedure called thinning. Taking a different point of view, a more recent work by Moller and Rasmussen [2005] claims to provide better quality simulations, without edge effects. In this appendix however, we describe the simple thinning algorithm in a multidimensional setting, because of its intuitive progression and its ability to be straightforwardly generalized to any point process described with its conditional intensity.

Let $\mathcal{U}_{[0,1]}$ denote the uniform distribution on the interval $[0,1]$ and $[0,T]$ the time interval on which the Hawkes process defined by equation (4.4) is to be simulated. We define $I^K(t) = \sum_{n=1}^{K} \lambda^n(t)$ the sum of the intensities of the first K components of the multivariate process. The algorithm is written as follows.

1. **Initialization** : Set $i \leftarrow 1$, $i^1 \leftarrow 1, \ldots, i^M \leftarrow 1$ and $I^* \leftarrow I^M(0) = \sum_{n=1}^{M} \lambda_0^n$.
2. **First event** : Generate $U \rightsquigarrow \mathcal{U}_{[0,1]}$ and set $s \leftarrow -\frac{1}{I^*} \ln U$ (s is exponentially distributed with parameter I^*).
   (a) **If** $s > T$ **Then** go to step 4.
   (b) **Attribution Test** : Generate $D \rightsquigarrow \mathcal{U}_{[0,1]}$ and set $t_1^{n_0} \leftarrow s$ where n_0 is such that $\frac{I^{n_0-1}(0)}{I^*} < D \leq \frac{I^{n_0}(0)}{I^*}$.
   (c) Set $t_1 \leftarrow t_1^{n_0}$.
3. **General routine** : Set $i^{n_0} \leftarrow i^{n_0} + 1$ and $i \leftarrow i + 1$.
   (a) **Update maximum intensity** : Set $I^* \leftarrow I^M(t_{i-1}) + \sum_{n=1}^{M} \sum_{j=1}^{P} \alpha_j^{nn_0}$.
   I^* exhibits a jump of size $\sum_{n=1}^{M} \sum_{j=1}^{P} \alpha_j^{nn_0}$ as an event of type n_0 has just occurred.
   (b) **New event** : Generate $U \rightsquigarrow \mathcal{U}_{[0,1]}$ and set $s \leftarrow s - \frac{1}{I^*} \ln U$.
   If $s > T$,
   Then go to step 4.
   (c) **Attribution-Rejection test** : Generate $D \rightsquigarrow \mathcal{U}_{[0,1]}$.
   If $D \leq \frac{I^M(s)}{I^*}$,
   Then set $t_{i^{n_0}}^{n_0} \leftarrow s$ where n_0 is such that $\frac{I^{n_0-1}(s)}{I^*} < D \leq \frac{I^{n_0}(s)}{I^*}$, and $t_i \leftarrow t_{i^{n_0}}^{n_0}$ and go through the general routine again,
   Else update $I^* \leftarrow I^M(s)$ and try a new date at step (b) of the general routine.
4. **Output** : Retrieve the simulated process $\{\{t_i^n\}_i\}_{n=1,\ldots,M}$.

Of course, still using using notations introduced in Section 4.3, one can simulate the double sequence $\{t_i, z_i\}_i$ instead of the M sequences $\{t_i^n\}_i, n = 1, \ldots, M$ with very minor and straightforward modifications.

Appendix B – Nonparametric calibration of Hawkes processes with time-varying background intensity by MISD algorithm

We consider the general calibration problem of a Hawkes process with time-varying intensity $\mu(t)$ whose conditional intensity is given by : $\lambda(t) = \mu(t) + \sum_{t_i<t} g(t - t_i)$

There are two different time scales useful in the sequel : the one of the numerical discretization of the kernel δt, and the one of the numerical discretization of the background rate Δt. As we make the assumption that the jumps components are 'well-separated', which means that on average the characteristic time linked to the triggering kernel is far bigger than the characteristic time linked to the background rate, we have the relationship : $\Delta t >> \delta t$.

We first define two useful functions to map event indices and time sets :

$$\delta^{-1}(i,j) = m \text{ such that } (t_i - t_j) \in [m\delta t, m\delta t + \delta t]$$

$$\text{and } \Delta^{-1}(i) = m \text{ such that } t_i \in [m\Delta t, m\Delta t + \Delta t]$$

We also define two useful sets :

$$A_m = \{(i,j) \in [1,n]^2 | (t_i - t_j) \in [m\delta t, m\delta t + \delta t]\}$$

$$\text{and } B_m = \{i \in [1,n] | t_i \in [m\delta t, m\delta t + \delta t]\}$$

We have two numerical discretizations, one for the background rate $\mu_m = \mu(m\Delta t)$ and one for the triggering kernel $g_m = g(m\delta t)$. The upper script ^{k} is used to indicate a given value after the k-th step of the algorithm.

In this case, the MISD algorithm generalizes to iterating the following two steps until convergence is reached.

Expectation step (for $t_i > t_j$) :

$$\begin{aligned} p_{ij}^k &= \frac{g_{\delta^{-1}(i,j)}^k (t_i - t_j)}{\mu_{\Delta^{-1}(i)}^k + \sum_{j=1}^{i-1} g_{\delta^{-1}(i,j)}^k (t_i - t_j)} = \frac{g_{\delta^{-1}(i,j)}^k (t_i - t_j)}{\lambda^k(t_i)}, \\ p_{ii}^k &= \frac{\mu_{\Delta^{-1}(i)}^k}{\mu_{\Delta^{-1}(i)}^k + \sum_{j=1}^{i-1} g_{\delta^{-1}(i,j)}^k (t_i - t_j)} = \frac{\mu_{\Delta^{-1}(i)}^k}{\lambda^k(t_i)}. \end{aligned} \tag{4.30}$$

Maximization step :

$$\begin{aligned} g_m^{k+1} &= \frac{1}{n\delta t} \sum_{(i,j)\in A_m} p_{ij}^k, \\ \mu_m^{k+1} &= \frac{1}{\Delta t} \sum_{i \in B_m} p_{ii}^k. \end{aligned} \tag{4.31}$$

Chapitre 5

Comparison of trades-through microstructural jumps with macroscopic jumps (detected by a statistical test robust to microstructure noise)

Résumé

We investigate to what extent the microscopic jumps measured by trades-through (see Pomponio and Abergel [2010]) aggregate at the macroscopic level to compose large jumps. To this end, we compare the moments of the day when the proportion of trades-through is high with those linked to extrema and big variations of a statistical test used to detect jumps in high-frequency data with market microstructure noise (see Aït-Sahalia et al. [2012]). We first generalize their empirical results concerning the daily trajectories jumps detection, obtained on 2008 DJIA's data of the TAQ database, with Thomson-Reuters 2011 data on the DAX and Footsie. We empirically study how this jump detection method behaves at higher sampling frequency with intraday returns trajectories and we obtain the empirical intraday proportion of detected jumps. The maximum proportion of detected jumps and its biggest upwards move happenned at 14 :30 and 16 :00 (Paris time), corresponding to the two high peaks in the intraday distribution of trades-through. We found that the upwards regime-switch in the mean number of trades-through, from 15 :30 to the end of the day, is linked with a smaller proportion of detected jumps. We explain it using the global assets moves decomposition assumed by the statistical test, as a sum of continuous-volatility moves and pure-jumps moves. This upwards switch of trades-through probably causes an increase of the volatility factor of the continuous part, thus allowing more large assets moves to be considered as acceptable moves of the continuous part.

Contents

5.1 Introduction 86
5.2 Testing for jumps in a discretely observed noisy process *a la* Aït-Sahalia and Jacod 87
5.2.1 The case with no microstructure noise 87
5.2.2 The case with microstructure noise 88
5.3 Empirical comparison of high occurrence of jumps, detected by trades-through and by a statistical test method 90
5.3.1 Empirical data and parameters used 90

5.3.2 Testing for jumps in daily trajectories 91
5.3.3 Testing for jumps in intraday trajectories 91
5.3.4 Comparing the intraday distributions of jumps obtained with trades-through and with the statistical test method 92
5.4 Conclusion 93

5.1 Introduction

Several problems coming from the financial industry, ranging from modelling and hedging financial assets to their risk management, are strongly dependent on the presence of jumps in the assets considered. This probably explains why an important question arising in the quantitative finance literature has been to determine, based on empirical observations, whether those data are compatible with a continuous underlying process, or if a jump component has to be considered in the modelling of the asset. Of course, in the reality, only discrete observations (completely composed of jumps by construction) are available and one has to say if those empirical jumps are due to the sampling of a continuous underlying process, or if they reflect the presence of true jumps in the asset.

In this chapter, we want to compare an approach available in the literature to find jumps in prices/returns series (see Aït-Sahalia et al. [2012]) with the microstructural method of jump detection based on tick-by-tick data, that we earlier called 'trades-through'. What are the similarities and differences in those two jump detection methods, when tested on empirical data ? More precisely, can we compare the periods when the proportion of trades-through is high during the day with those linked to the extrema of this statistical test used to detect jumps in high-frequency data with market microstructure noise ?

We first review some of the main approaches to detect jumps in a discretely observed process, by closely following the introductions done in Aït-Sahalia and Jacod [2009] and Aït-Sahalia et al. [2012]. Aït-Sahalia [2002] tested for a continuity inequality based on the transition function of the process when assuming a diffusion hypothesis. Carr and Wu [2003] test is based on short-dated options : when time to maturity decays to zero, the option price time value also goes to zero at a rate depending on the structure of the underlying process (presence of jumps or not) and of the moneyness (at-the-money or out-the-money). Barndorff-Nielsen and Shephard [2004] initiated a new series of tests based on generalizations of the realized variance, called bipower variations and more generally multipower variations. Jiang and Oomen [2008] invented a test based on the cumulative replication error of a discretely hedged variance swap. Lee and Mykland [2008], as well as Lee and Hannig [2010], computed a test ratio based on the comparison between returns and instantaneous volatilities (estimated with realized bipower variations for the former and with realized truncated quadratic variation for the latter). Finally, Aït-Sahalia and Jacod [2009] built a test based on the ratio of power variations sampled at different frequencies.

Usually in statistical estimation problems, the higher the sampling frequency, the better the convergence of the estimators. But, when using financial data and increasing the frequency up to the high-frequency domain, the presence of the microstructure noise (due to bid-ask bounces and discreteness of price changes, for example) completely changes the problem. In this case, increasing the frequency comes with higher measurement errors and the previous tests generally lose their statistical power for finding jumps.

This microstructure noise problem has been studied in the literature (see Aït-Sahalia et al. [2012] and the references therein for a more detailed review) using two main models, the additive error and the rounding error. Robustifications of the previous tests have been designed to tackle

this issue, mainly by pre-averaging or post-averaging the data and/or the statistical measurements. Fortunately enough, there exists today a very general test developed in Aït-Sahalia and Jacod [2009] and robustified to deal with the market microstructure noise in Aït-Sahalia et al. [2012] by a pre-averaging technique (introduced in Jacod et al. [2009]), and this is the one we will focus on in our study.

In a first part of this chapter, we describe how this statistical test works for finding jumps in discretely observed high-frequency prices/returns data with microstructure noise. We use it with empirical data to detect jumps on daily trajectories for major European (Footsee and DAX) and US (DJIA) futures and their compositions in 2011, comparing our results to those of Aït-Sahalia et al. [2012] on DJIA future's composition in 2008. Then, we empirically see how the test distribution evolves when trying to find jumps in intraday trajectories. Finally, we compare the intraday distribution of jumps using this statistical test with the intraday distribution of trades-through. In particular, we compare the moments of the day when the proportion of trades-through is high with those linked to extrema of this statistical test. Using trades-through study, we finally suggest a possible explanation for those results, based on the global decomposition of the assets moves assumed by the statistical test.

5.2 Testing for jumps in a discretely observed noisy process *a la* Aït-Sahalia and Jacod

5.2.1 The case with no microstructure noise

For the sake of simplicity, we first review the main properties of the statistical test built in Aït-Sahalia and Jacod [2009] to determine if jumps are present in discretely sampled asset returns when there is no microstructure noise. This will greatly help us to gain some understanding of the far more complicated expressions necessary to deal with the microstructure noise case. We will not rigorously review all the different technical assumptions needed in the demonstrations (and we directly refer the reader to the original paper for more details), but basically almost any Itô semimartingale X on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ with sufficient regularity conditions and with a volatility σ_t also being an Itô semimartingale is included in this study.

In our case, X is the log-price, defined on a given time interval $[0, T]$ (typically one day), and perfectly observed at times $i\Delta_n$ for a discretization step $\Delta_n = \frac{T}{n}$. The discrete increments of X (the log-returns) are defined by : $\Delta_i^n X = X_{i\Delta_n} - X_{(i-1)\Delta_n}$. We also define the useful estimator $\hat{B}$ for $p > 0$, based on p-variations of increments at timescale Δ_n : $\hat{B}(p, \Delta_n)_t = \sum_{i=1}^{[\frac{t}{\Delta_n}]} |\Delta_i^n X|^p$. We need the family of constants for $p > 0$ of the p-th absolute moment of a standard centered normal distributed variable U : $m_p = E(|U|^p) = \frac{2^{p/2}}{\sqrt{\pi}}\Gamma(\frac{p+1}{2})$. [1]

The test statistics used to discriminate between the presence and the absence of jumps is obtained for $p > 2$ and $k \geq 2$ by computing the following ratio of power variations sampled at different frequencies : $\hat{S}(p, k, \Delta_n)_t = \frac{\hat{B}(p, k\Delta_n)_t}{\hat{B}(p, \Delta_n)_t}$. We have the two different convergences in probability when $\Delta_n \to 0$: $\hat{S}(p, k, \Delta_n)_t \to 1$ if X has jumps and $\hat{S}(p, k, \Delta_n)_t \to k^{p/2-1}$ if X is continuous.

1. where we denote by Γ the usual function defined for any complex number z with strictly positive real part by $\Gamma(z) = \int_0^{+\infty} t^{z-1}e^{-t}dt$.

We will now detail the estimator of the asymptotic variance of this test, under the hypothesis of continuity. To this purpose, we first need the following estimator of the quantity $\int_0^t |\sigma_s|^p ds$, based on truncations of big returns and defined for any constants $\alpha > 0$ and $\bar{\omega} \in (0, \frac{1}{2})$ by : $\hat{A}(p, \Delta_n)_t = \frac{\Delta_n^{1-p/2}}{m_p} \sum_{i=1}^{[\frac{t}{\Delta_n}]} |\Delta_i^n X|^p 1_{\{|\Delta_i^n X| \leq \alpha \Delta_n^{\bar{\omega}}\}}$. We know that, if X is continuous, $\hat{S}(p, k, \Delta_n)_t \to k^{p/2-1}$ at the speed $\sqrt{\hat{V}^c_{n,t}}$ with $\hat{V}^c_{n,t} = \frac{\Delta_n M(p,k) \hat{A}(2p, \Delta_n)_t}{\hat{A}(p, \Delta_n)_t^2}$. More precisely, we have the following central limit theorem stating that, if X is continuous, and $p \geq 2$ the variables $\frac{\hat{S}(p, k, \Delta_n)_t - k^{p/2-1}}{(\hat{V}^c_{n,t})^{1/2}}$ converge in law towards a standard centered normal variable, conditionally on $\mathcal{F}$.

Finally, we have the following form of critical rejection region of the continuous hypothesis : $C^c_{n,t} = \{\hat{S}(p, k, \Delta_n)_t < c^c_{n,t}\}$. By denoting, for every level $\alpha \in (0, 1)$, z_α the usual α-quantile of the standard gaussian variable [2], we have for $p > 3$: $c^c_{n,t} = k^{p/2-1} - z_\alpha \sqrt{\hat{V}^c_{n,t}}$ with $\hat{V}^c_{n,t} = \frac{\Delta_n M(p,k) \hat{A}(2p, \Delta_n)_t}{\hat{A}(p, \Delta_n)_t^2}$, which gives us a α-level statistical test. In the particular cases of $p = 4$ and $k = 2$, we get : $c^c_{n,t} = 2 - z_\alpha \sqrt{\frac{160 \Delta_n \hat{A}(8, \Delta_n)_t}{3 \hat{A}(4, \Delta_n)_t^2}}$.

5.2.2 The case with microstructure noise

Because of the market microstructure noise, the previous statistic ratio test $\hat{S}(p, k, \Delta_n)_t$ may converge to the same value for both the continuous case and the jump case (to $\frac{1}{k}$ for the additive noise case and to $\frac{1}{k^2}$ for the rounding noise case), which prevents us from using it with empirical noisy data. So, this statistical test was robustified by pre-averaging the data with a kernel in Aït-Sahalia et al. [2012].

The statistical test ratio

We begin with the generalization of the previous statistics ratio $\hat{S}(p, k, \Delta_n)_t$ in the presence of microstructure noise. X is still the log-price of the process and is still assumed to be an Itô semi-martingale, with its volatility also following an Itô semi-martingale, on some filtered probability space. But this time, because of the noise, the process observed at times $i\Delta_n$ is no longer the true process X but a noisy version of it that we will denote by $Z : Z = X + \epsilon$, where the observation error ϵ is a mean zero and mutually independent process, conditionally on X. Note that this conditional white noise model includes both the usual additive and rounding noise cases. Aït-Sahalia et al. [2012] precise that this noise model should be valid for financial data up to the medium scale frequency (with a sampling frequency of one data every 5 seconds, approximately).

We still denote the discrete increments of X by : $\Delta_i^n X = X_{i\Delta_n} - X_{(i-1)\Delta_n}$. But this time, in order to get rid of the microstructure noise, we will not directly use any given increment, but we will replace it by its local average with a given weight function g (or kernel). To this purpose, we need an averaging window of the form $k_n \Delta_n$, with $k_n = [\frac{\theta}{\sqrt{\Delta_n}}]$ for some $\theta > 0$. We also need the moments of both the weight function and its derivative : $\bar{g}(p) = \int |g(s)|^p ds$ and $\bar{g'}(p) = \int |g'(s)|^p ds$. The discretization of the kernel is $g_i^n = g(\frac{i}{k_n})$ and its increments are $g'^n_i = g_i^n - g_{i-1}^n$.

The increments of X are now replaced by the locally averaged increments of Z : $\bar{Z}(g)_i^n =$

2. with $P(U > z_\alpha) = \alpha$

$\sum_{j=1}^{k_n-1} g_j^n \Delta_{i+j}^n Z$. We also need a local estimator of the variance of the noise : $\hat{Z}(g)_i^n = \sum_{j=1}^{k_n} (g'^n_j \Delta_{i+j}^n Z)^2$.

By combining the previous elements, one obtains : $V(Z,g,q,r)_t^n = \sum_{i=0}^{[\frac{t}{\Delta_n}]-k_n} |\bar{Z}(g)_i^n|^q |\hat{Z}(g)_i^n|^r$.

For any even integer $p \geq 4$, the robustified versions of the p-th power variation estimators are now defined by : $\bar{V}(Z,g,p)_t^n = \sum_{l=0}^{p/2} \rho(p)_l V(Z,g,p-2l,l)_t^n$, where the $(\rho(p)_j)_{j=0,...,p/2}$ are the unique numbers solving the triangular system : $\rho(p)_0 = 1$ and $\sum_{l=0}^{j} 2^l m_{2j-2l} \binom{p-2l}{p-2j} \rho(p)_l = 0$ for $j = 1, 2, ..., p/2$. [3].

We can now generalize the statistics ratio test to its robustified version. For two weight functions g and h, we need the moments ratios : $\gamma = \frac{\bar{g}(2)}{\bar{h}(2)}$, $\gamma' = \frac{\bar{g}(p)}{\bar{h}(p)}$ and $\gamma'' = \frac{\gamma^{p/2}}{\gamma'}$. The robustified test statistic is : $S_{RJ}(g,h,p)_n = \frac{\bar{V}(Z,g,p)_T^n}{\gamma' \bar{V}(Z,h,p)_T^n}$ We have the following convergences in probability : $S_{RJ}(g,h,p)_n \to 1$ if X has jumps or $\to \gamma''$ if X is continuous. By choosing specific weight functions of the form $h(s) = g(sk)$ for some $k > 1$, we even obtained the same limit values for the statistical test than in the no microstructure noise case : $S_{RJ}(g,h,p)_n \to 1$ if X has jumps or $\to k^{p/2-1}$ if X is continuous.

The asymptotic variance of the statistical test, when testing for absence of jumps

We now detail the variance of the statistical test when testing the continuity hypothesis. [4] Consider two functions ϕ and ψ and any integers $w \geq 1$ and $w' \in \{0, ..., 2w\}$, we first define the following variables :

$$
\begin{aligned}
a(\phi,\psi)_t &= \int_{1\vee t}^{1+1\wedge t} \phi(u-1)\psi(u-t)du \\
a'(\phi,\psi;w,w')_t &= \sum_{r=0}^{[w'/2]} \binom{w'}{2r} m_{2r} m_{2w-2r} a(\phi,\phi)_1^{w-w'} a(\phi,\psi)_t^{w'-2r} (a(\phi,\phi)_1 a(\psi,\psi)_1 - a(\phi,\psi)_t^2)^r
\end{aligned}
$$

$$
\begin{aligned}
A(g,h;w)_t &= \sum_{\substack{l,l' \in \{0,...,p/2\} \\ l+l' \leq p-w}} [\sum_{w'=(2w-p+l)^+}^{(2w)\wedge(p-2l')} \rho(p)_l \rho(p)_{l'} \binom{2w-w'}{p-2l} \binom{w'}{p-2l'} \\
&\quad (2\bar{g}'(2))^l (2\bar{h}'(2))^l a'(g,h;w,w')_t a'(g',h';p-l-l'-w,p-2l'-w')_t] \\
A'(g,h;w) &= \int_0^2 A(g,h;w)_t dt - 2m_p^2 \bar{g}(2)^{p/2} \bar{h}(2)^{p/2} 1_{\{w=p\}}
\end{aligned}
$$

Then, we estimate the asymptotic variance $\Sigma_{R,J,n}^c$ using truncating levels : $u_n = \alpha \Delta_n^{\bar{\omega}}$ for some

3. We already defined the (m_r) constants family at 5.2.1

4. The following expressions are somewhat similar to those of the case with no microstructure noise, in which we replaced the values by their local averaged versions, even if the different constants involved hereafter are not easily explainable.

$\alpha > 0$ and $\frac{1}{12} < \bar{\omega} < \frac{1}{4}$:

$$
\begin{aligned}
V^*(Z,g,q,r)_t^n &= \sum_{i=0}^{[\frac{t}{\Delta_n}]-k_n} |\bar{Z}(g)_i^n|^q 1_{\{|\bar{Z}(g)_i^n| \leq u_n\}} |\hat{Z}(g)_i^n|^r \\
M^*(g,h,\phi)_t^n &= \Delta_n^{1-p/2} \sum_{w=0}^{p} \frac{\theta A'(g,h;w)}{m_{2w} 2^{p-w} \bar{\phi}(2)^w \bar{\phi}'(2)^{p-w}} \sum_{l=0}^{w} \rho(2w)_l V^*(Z,\phi,2w-2l,p+l-w)_t^n \\
\Sigma_{R,J,n}^c &= \frac{M^*(g,g,\phi;p)_T^n - 2\gamma^{p/2} M^*(g,h,\phi,p)_T^n + \gamma^p M^*(h,h,\phi;p)_T^n}{(\Delta_n^{1-p/4} \bar{V}(Z,g,p)_T^n/\gamma'')^2}
\end{aligned}
$$

Finally, we have a central limit theorem under the hypothesis of continuity of the process X, with the convergence in law towards a standard centered gaussian variable (conditionally on $\mathcal{F}$) for the standardized statistical test $t : t = \frac{S_{RJ}(g,h,p)_n - \gamma''}{(\Delta_n)^{1/4}\sqrt{\Sigma_{R,J,n}^c}}$.

The critical rejection region of the no-jumps null hypothesis for the level $\alpha \in (0,1)$ is now obtained with : $C_n^c = \{S_{RJ}(g,h,p)_n < \gamma'' - z_\alpha \Delta_n^{1/4}\sqrt{\Sigma_{R,J,n}^c}\}$.

Jump detection test summary

The validity of the continuity hypothesis will be characterized by a standardized statistical test t distribution close to a standard normal distribution, whereas a distribution shifted to the left with negative means will indicate the presence of jumps. Empirically, the more the distribution of the standardized statistical test is shifted to the left, the more jumps are present.

5.3 Empirical comparison of high occurrence of jumps, detected by trades-through and by a statistical test method

5.3.1 Empirical data and parameters used

We use Thomson-Reuters [5] tick-by-tick data from almost all 2011, as we excluded from the statistics all the days when Daylight Saving Time were present in the US and not in Europe (which correspond to two weeks in March and one week in November 2011). We focus on European equity futures data, from the Eurex limit order book for the future DAX (FDX) and from the London Stock Exchange limit order book for the future Footsie (FTSE), and also on US equity future data, from the Chicago Mercantile Exchange limit order book for the future E-mini DJIA (YM). For a given day, we always focus on the most liquid traded future, which is the one maturing first. For each of those 3 major futures, we added its 30 most liquid stocks (which means all its composition for the DAX and the DJIA futures, and only 30 out of the 100 stocks for the FTSE future) as traded on their principal exchange (XETRA for the DAX stocks, LSE for the FTSE stocks, NYSE and NASDAQ for the DJIA stocks). As the composition of each index is changing over time, we choose to consider for the study the fixed composition of February 1st, 2011.

Concerning the parameters used to compute the standardized statistical test t values, we use the power $p = 4$, the truncation power $\bar{\omega} = 0.249$ and the kernel $g(x) = (0.5 - |x - 0.5|)^+$ (plotted on the figure 5.1). We also choose $k = 2$, ie $h(x) = g(2x)$ and $\phi = g$, so $\gamma'' = 2$. Finally, we take a pre-averaging window $k_n = 100$ and the truncation variance level is fixed at $u_n = C(\bar{V}(Z,g,2)_T^n/T)^{1/2}\Delta_n^{0.249}$, with $C = 5$.

5. See 3.2.1 for more detail on TRTH data.

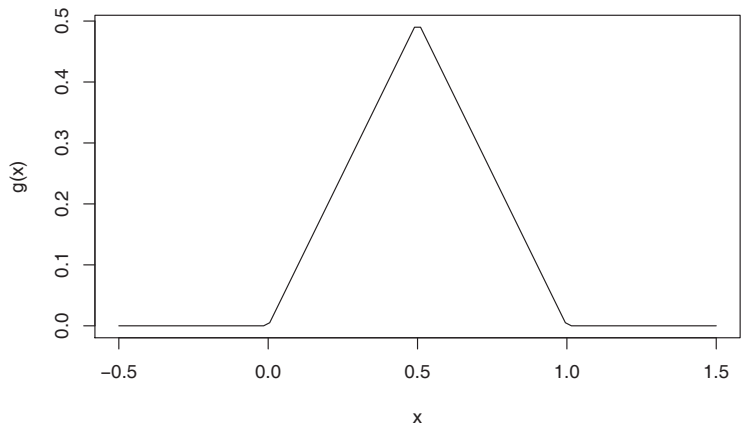

FIGURE 5.1 – Triangular kernel g chosen for the robustification of the statistical test in presence of microstructure noise by the pre-averaging method

5.3.2 Testing for jumps in daily trajectories

We focus on daily trajectories with $T = 1$ day. We first qualitatively compare our results on DJIA in 2011 to those of Aït-Sahalia et al. [2012] on the same index in 2008 from the TAQ database. To this purpose, we take the same number of points in our time series, $N = 5000$. We can see in the upper graph of the following figure 5.2 that the distribution we obtained for the standardized statistical test t (under the continuity hypothesis) for the DJIA's stocks in 2011 is very similar to the one they obtained, with a normal-shaped distribution with approximately a -1 mean, a peak value at 0.35 and two abscissa x at -2 and 0 for the 0.2 y level-line, characterizing a similar standard deviation.

For the other indices, we increase the daily number of points of each time series to $N_{\text{daily}} = 10000$. This doesn't change the results for the daily trajectories we are currently looking at, but this will allow us to do intraday higher-frequency jump detection in the next section. We can see in the two following graphs of the same figure 5.2 that the global shapes of the distributions for the DAX and the FTSE indices are similar to the one of the DJIA index, with empirical means about -1 and a slightly smaller (respectively bigger) variance for the DAX (respectively FTSE) than for the DJIA.

5.3.3 Testing for jumps in intraday trajectories

We want to understand how the jump detection statistical test evolves when looking at intraday trajectories with $T = 15$ minutes. We keep the same sampling frequency for the returns time series as before, with a total daily number of points $N_{\text{daily}} = 10000$. This time, it gives us 15 minutes time series of approximately 400 points for the DJIA and 300 points for the DAX and FTSE, corresponding to a sampling frequency of approximately one point every 3 seconds. There is a trade-off between a higher number of points (higher sampling frequency) in the returns time series to get good convergence results of the test t and a smaller sampling frequency to stay within the test validity domain (regarding the noise model assumptions, for example), and we empirically find that those values give satisfactorily results.

We compute the standardized statistical test t (under the continuity hypothesis) for the intraday

time series and plot their distributions in the following graphs 5.3. The intraday test distribution is clearly less shifted to the left than the daily test distribution, as confirmed by the difference in the means (hereafter estimated with the median, which is a more robust statistic), which is the sign of a smaller proportion of detected jumps in the intraday case. We can see that a fraction of the negative part of the daily density test distribution is transfered to positive values in the intraday case, revealing that a smaller proportion of jumps is detected in the intraday case. This is clearly understandable, as a given jump detected in a given daily trajectory should only be detected in one of the intraday trajectories composing this full daily trajectory, leading to a smaller proportion of detected jumps.

5.3.4 Comparing the intraday distributions of jumps obtained with trades-through and with the statistical test method

We want to measure the intraday variations of the proportion of detected jumps, using the standardized statistical test t. We know that a possible indicator of the proportion of detected jumps is to look at how much the t distribution is shifted to the left, with respect to a standard normal distribution. To this end, we plot the intraday distribution of the median of the standardized statistical test t and we look at how negative this median was across the day. We use non-overlapping 15 minutes intervals, spanned across the whole trading session, and centered at each quarters of hours.

In the figure 5.4, we reported the two intraday distributions of the normalized statistic median, for each 15 minutes interval, and averaged on 2011 for the DAX and FTSE futures data. We tested some stability features of these two curves by changing the sampling frequency from 3 seconds to bigger values (5 and 10 seconds) with an adapted pre-averaging window $k_n = 30$. We found no statistical difference for the intraday curve shape of the FTSE normalized statistic median, unlike the DAX for which the results were less stable with parameters changes. As a consequence, we will focus our study on the FTSE results which are more reliable, and we only report the DAX results for indicative purpose. Moreover, when we consider the boxplots of the standardized statistical test t for the FTSE (see figure 5.5), the range of intraday variations of the median of the t test is about one third of its interquartile range. This explains why there is no clear intraday pattern for the median standardized statistical test t. As a consequence, we will mainly focus our analysis on the most extreme values (minimum and maximum) and on the biggest intraday variations of the median curves. The minimum of the intraday distribution of the t median is obtained at 14 :30 [6] (corresponding to the highest proportion of jumps detected during the day) and the maximum at 15 :30 (corresponding to the lowest proportion of jumps detected during the day). The two biggest intraday moves happened at 15 :30 (a global minimum of the jumps proportion) and at 16 :00 (a local maximum of the jumps proportion).

We will now throw some light on those results by comparing them with the intraday distributions of microscopic jumps, as defined by trades-through. So, we represented in the figure 5.6 the mean intraday histograms of trades-through timestamps, averaged over the 30 most liquid components of the FTSE and DAX indices on 2011 [7], using one-minute bins and 15-minutes bins [8]. We first used the same timescale of 15 minutes as the one used to detect jumps with the statistical test, but even if the 15-minutes bins histograms reveal two local maxima at 14 :30 and 16 :00

6. All hours are given at the Paris time reference.

7. We continue to exclude from the statistics all the days when Daylight Saving Time was present in the US and not in Europe (which correspond to two weeks in March and one week in November 2011).

8. The zero value obtained in the DAX index components histogram is due to an intraday auction occurring at 13 :00.

(linked to US major macro news release), it doesn't seem to reveal anything clear happening at 15 :30. Taking advantage of the flexibility of the trades-through approach, we then zoomed in at the 1-minute timescale level to look at the finer structure of trades-through intraday histograms. The 1-minute timescale intraday histograms of trades-through have two very sharp intraday peaks (at 14 :30 and 16 :00) and we also see that a switch of regime happens at 15 :30, when the mean number of trades-through is increased by approximately 50% for the DAX and 100% for the Footsee.

Remember that, as detailed in Aït-Sahalia and Jacod [2009] and Aït-Sahalia et al. [2012], assuming that the asset is part of the Itô semi-martingales class, we can decompose the asset returns as the sum of a continuous part (basically a Brownian motion with a stochastic volatility factor σ_t) plus a pure jumps part. In their framework, the authors assume that the volatility component σ_t can also have jumps. As a consequence, the global moves of the asset returns are also the sum of two parts : the continuous-volatility moves, which come from a Brownian motion weighted by σ_t factor, and the pure-jumps moves.

With this general decomposition of the assets moves in mind, we can now link the two intraday distributions of jumps (detected by the t test and by the trades-through) at the specific hours of 14 :30, 15 :30 and 16 :00.

First, at 15 :30, we have a switch of regime in the mean number of trades-through, due to the open of the US market (NYSE and NASDAQ), with mean values 100% higher for the FTSE. This higher level of the mean number of trades-through goes on until the end of the day. This reflects into a higher σ_t volatility factor of the continuous part after 15 :30, allowing to integrate larger moves in the asset returns as acceptable moves of the continuous part, and as a consequence leading to a decrease in the proportion of detected jumps. Indeed, the statistical test mean is shifted to the right after 15 :30 (excluding the particular case of the 16 :00 peak) with a mean value of -0.61 for the FTSE, whereas the mean test values were more negative (corresponding to a higher jumps proportion) strictly before 15 :30 (also excluding the particular case of the 14 :30 peak) at -0.73.

At 14 :30, there is a sharp peak in the mean number of trades-through per 1-minute interval, with peak values 3 to 4 times bigger than the previous values and a fast return to values of the previous order of magnitude in less than 5 minutes. This peak in the mean number of trades-through cannot be included in the continuous part of the asset decomposition, as it would need two jumps of the volatility factor σ_t (one upwards then one downwards) very close in time, which is improbable. So, this peak is directly reflected in an increase of the pure-jumps component. It leads to a large increase in the proportion of detected jumps by the statistical test (a more negative mean of the test), which in fact reaches its intraday minimum.

Finally, at 16 :00, we have a peak very similar to the one of 14 :30, with a fast return to mean level values of the after-15 :30 regime. But, as we just said before, the σ_t volatility factor of the continuous part is higher after 15 :30, allowing more large asset moves to be considered as acceptable moves of the continuous part (whereas they would have been considered as jumps before). This explains why the 16 :00 peak corresponds to a decrease in the statistical test mean value (corresponding to a higher jumps proportion), but less important than the one of 14 :30, even if the peak intensities are very similar.

5.4 Conclusion

We introduced and empirically tested the jump detection method, robust to market microstructure noise, introduced in Aït-Sahalia et al. [2012]. We validated their empirical results concerning the distribution of the standardized statistical t test on daily trajectories from 2008 DJIA's future data of the TAQ database with our Thomson-Reuters 2011 data, and extended it to the European futures DAX and FTSE on the same period of time. We empirically studied the modification of the

distribution of the standardized statistical test t, when increasing the sampling frequency to look at intraday trajectories with $T = 15$ minutes.

We obtained the intraday median values of the standardized statistical t test distribution for the DAX and FTSE futures composition. Maximum and minimum proportions of detected jumps were revealed at precise hours (14 :30 and 15 :30). Moreover, the biggest intraday variations occurred at 15 :30 (global minimum of jumps proportion) and at 16 :30 (local maximum of jumps proportion).

We compared and explained those results by using the microscopic jump definition of trades-through. The intraday histograms of the mean number of trades-through reveal the existence of peaks at 14 :30 and 16 :00, and of a regime-switch at 15 :30 to higher mean values. This trades-through approach helped us to better understand the extrema in the intraday proportion of detected jumps by the t test. The maximum detected jumps proportion at 14 :30 corresponds to a sharp peak in the trades-through distribution and this directly reflects in the pure-jumps component of the asset moves. The jumps detection decrease after 15 :30 is linked to the trades-through upwards regime-switch, causing an increase of the volatility factor of the continuous part of the asset moves and, subsequently, diminishing the proportion of asset moves considered as jumps by the t test. Finally, the 16 :00 peak of trades-through is reflected as a local maximum of the proportion of detected jumps, less important than the one of 14 :30 because of a higher volatility factor of the continuous part after 15 :30.

To conclude, we tried to understand to what extent the microscopic jumps measured by trades-through emerge at the macroscopic level to compose large jumps. We obviously found that large marcroscopic jumps cannot be explained by trades-through only. In other words the nature of macroscopic jumps is not in general a simple addition of microscopic jumps. But trades-through do help to better understand some features of the macroscopic jumps, and especially extreme values obtained for the intraday distribution of the proportion of jumps detected.

In the future, it would be interesting to further study this relationship between trades-through and jumps detected by the statistical t test. Moreover, it would be interesting in order to get a more complete picture of the situation to include a liquidity factor in this analysis, and to relate those jumps and the liquidity fluctuations in the order book.

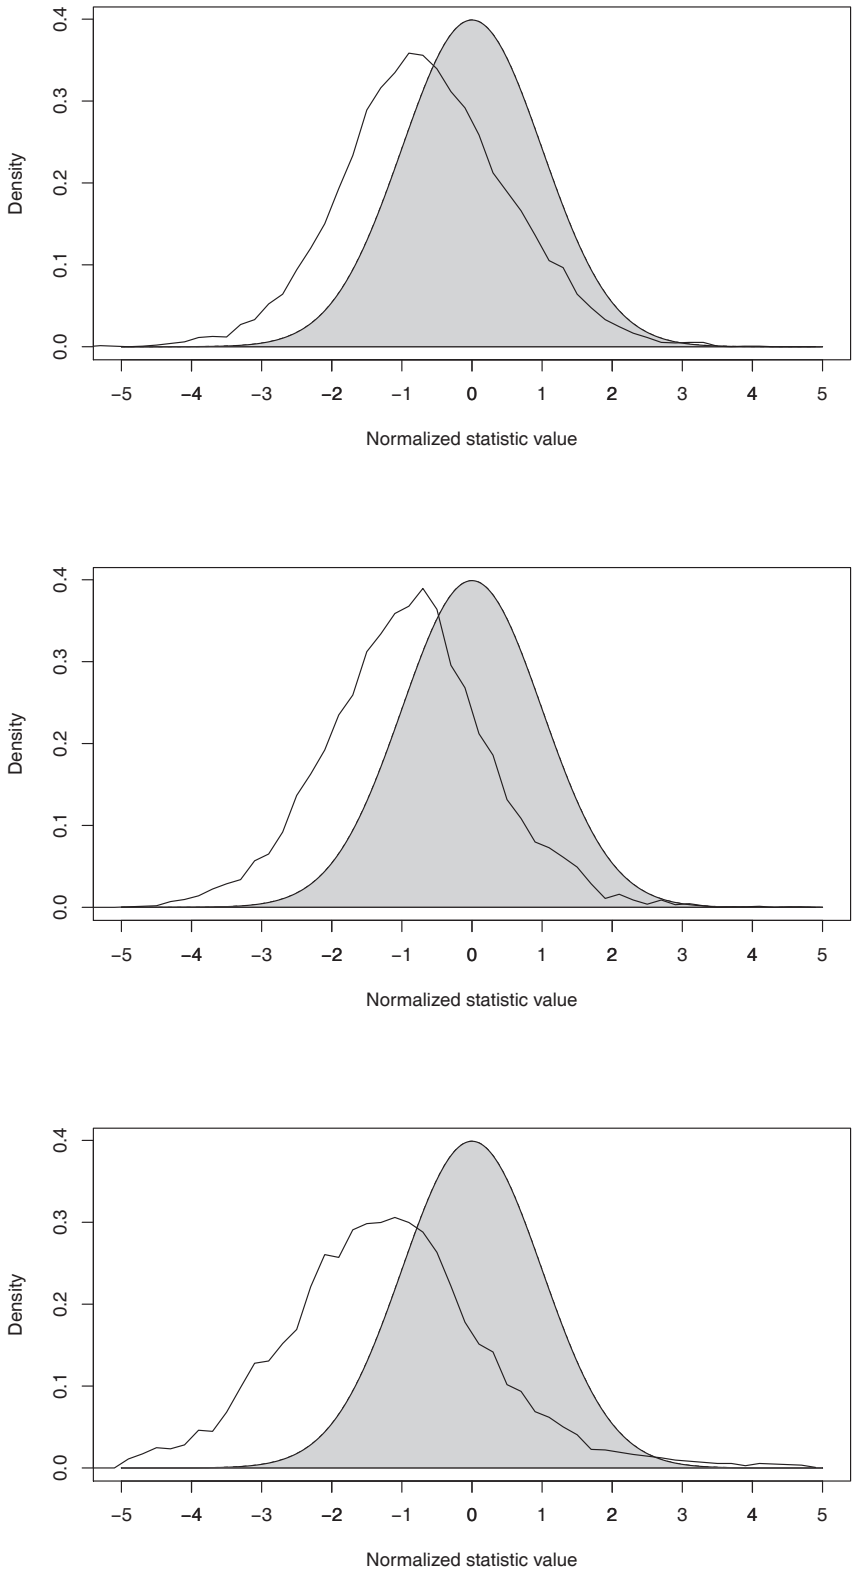

FIGURE 5.2 – Daily normalized statistic distribution in 2011 (solid line) compared to standard normal distribution (shaded area).
(up) DJIA index (30 stocks composition) with $N_{\text{daily}} = 5000$;
(middle) DAX index (future and 30 stocks composition) with $N_{\text{daily}} = 10000$;
(down) Footsee index (future and 30 most liquid stocks composition) with $N_{\text{daily}} = 10000$.

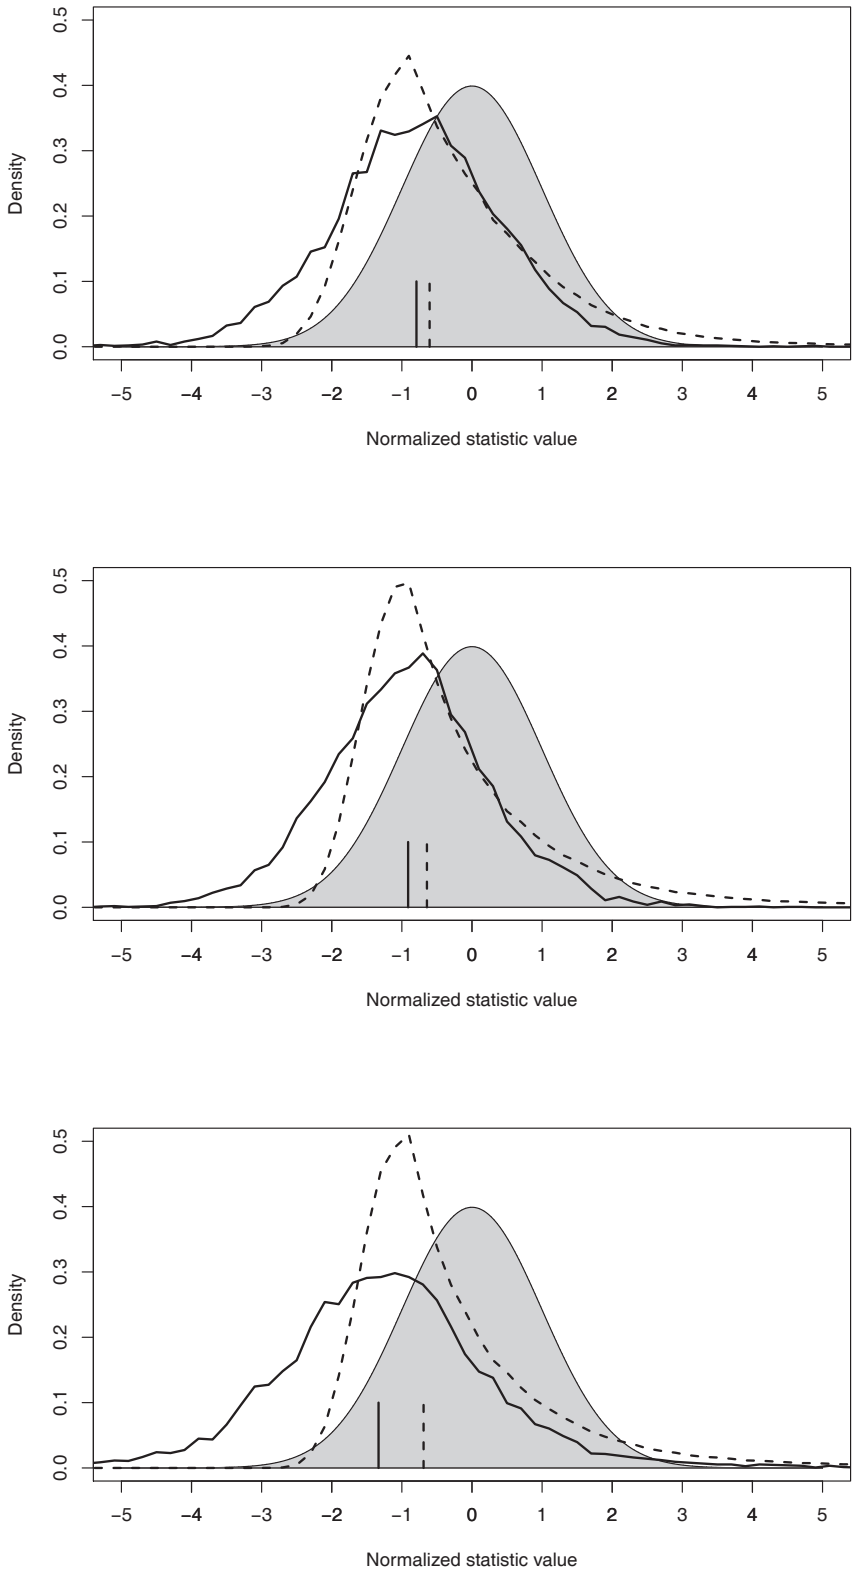

FIGURE 5.3 – Normalized statistic distribution comparison in 2011 with $N_{\text{daily}} = 10000$, between $T = 1$ day (solid line) and $T = 15$ min (dotted line), with standard normal distribution (shaded area). A vertical line indicates the median of the distribution.
(up) DJIA index (30 stocks composition) ; *(middle)* DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition).

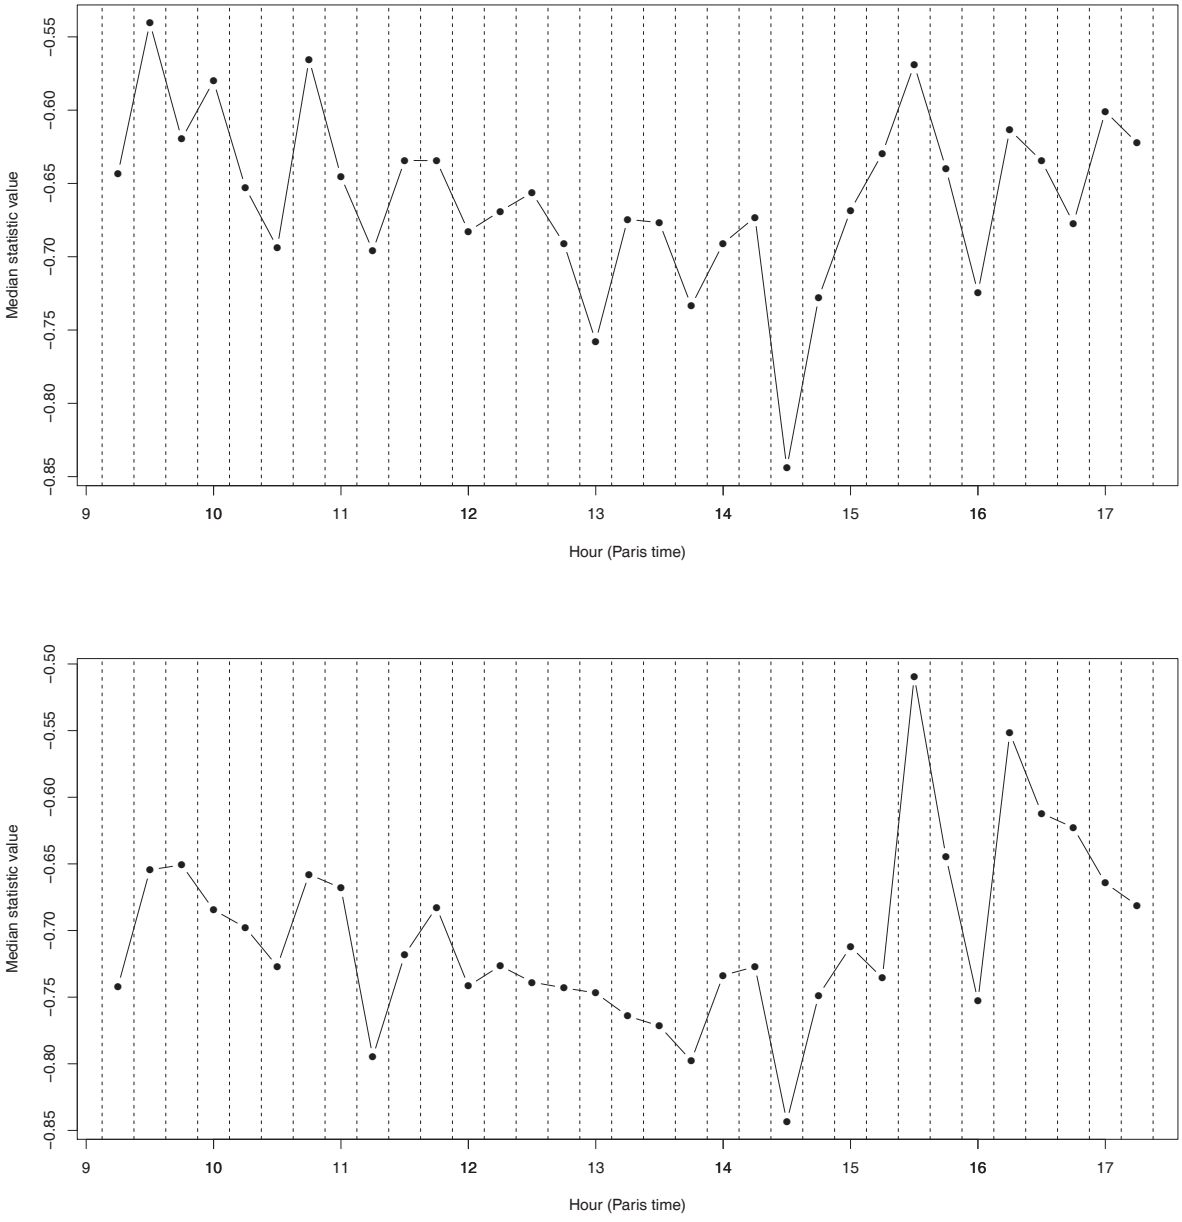

FIGURE 5.4 – Intraday distribution of the normalized statistic median in 2011 with $N_{\text{daily}} = 10000$ and $T = 15$ min.
(up) DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition).

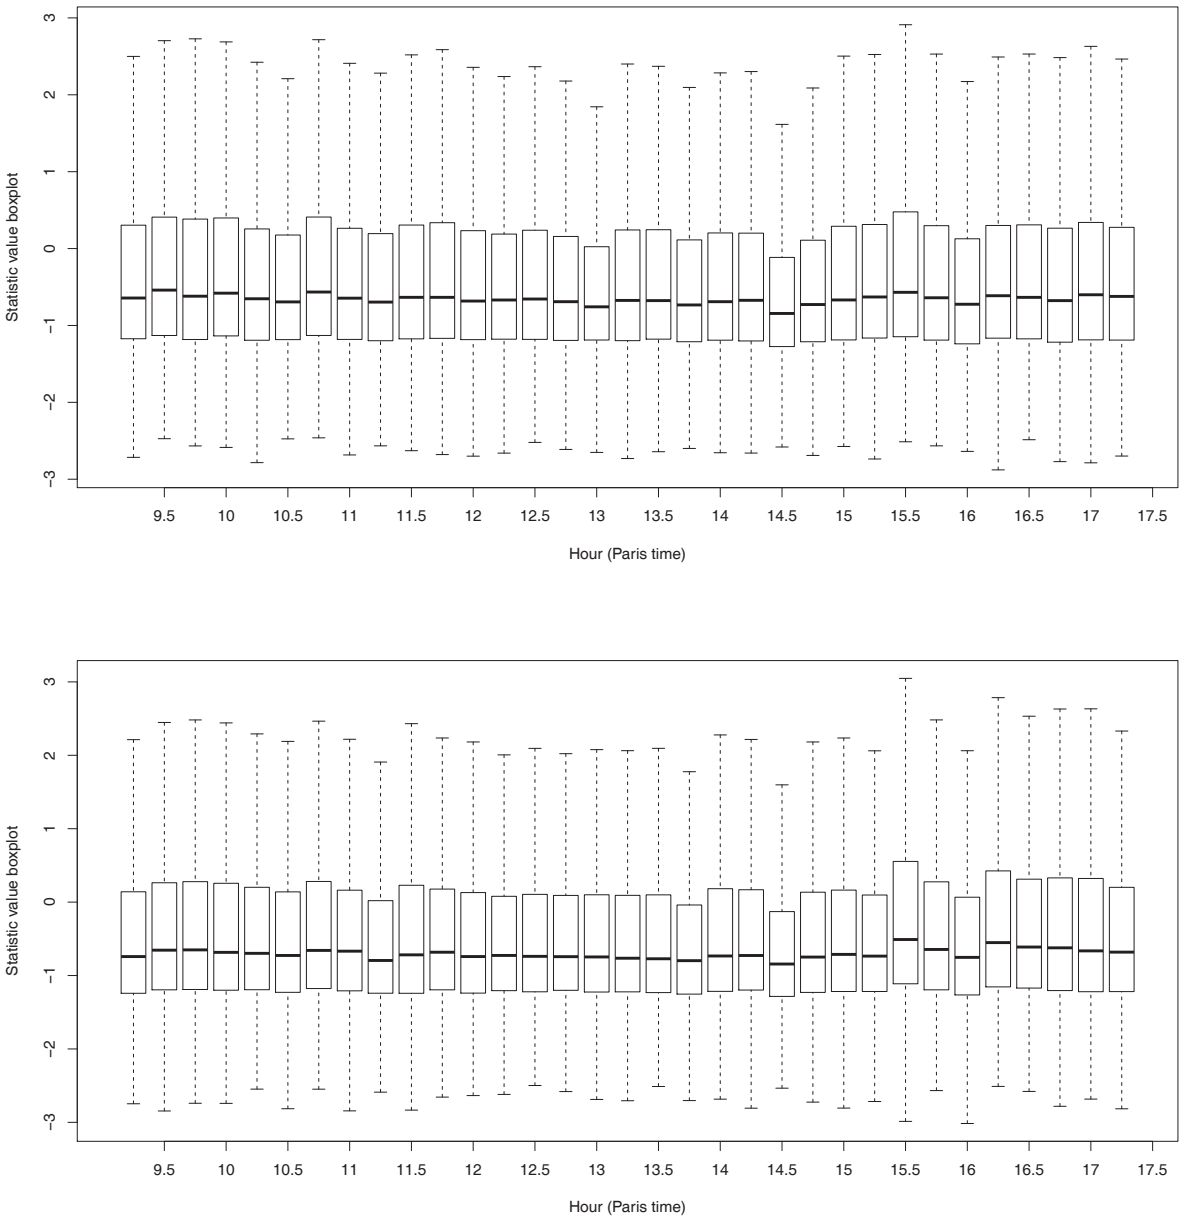

FIGURE 5.5 – Intraday distribution of the normalized statistic boxplot in 2011 with $N_{\text{daily}} = 10000$ and $T = 15$ min.
(up) DAX index (future and 30 stocks composition) ; *(down)* Footsee index (future and 30 most liquid stocks composition).

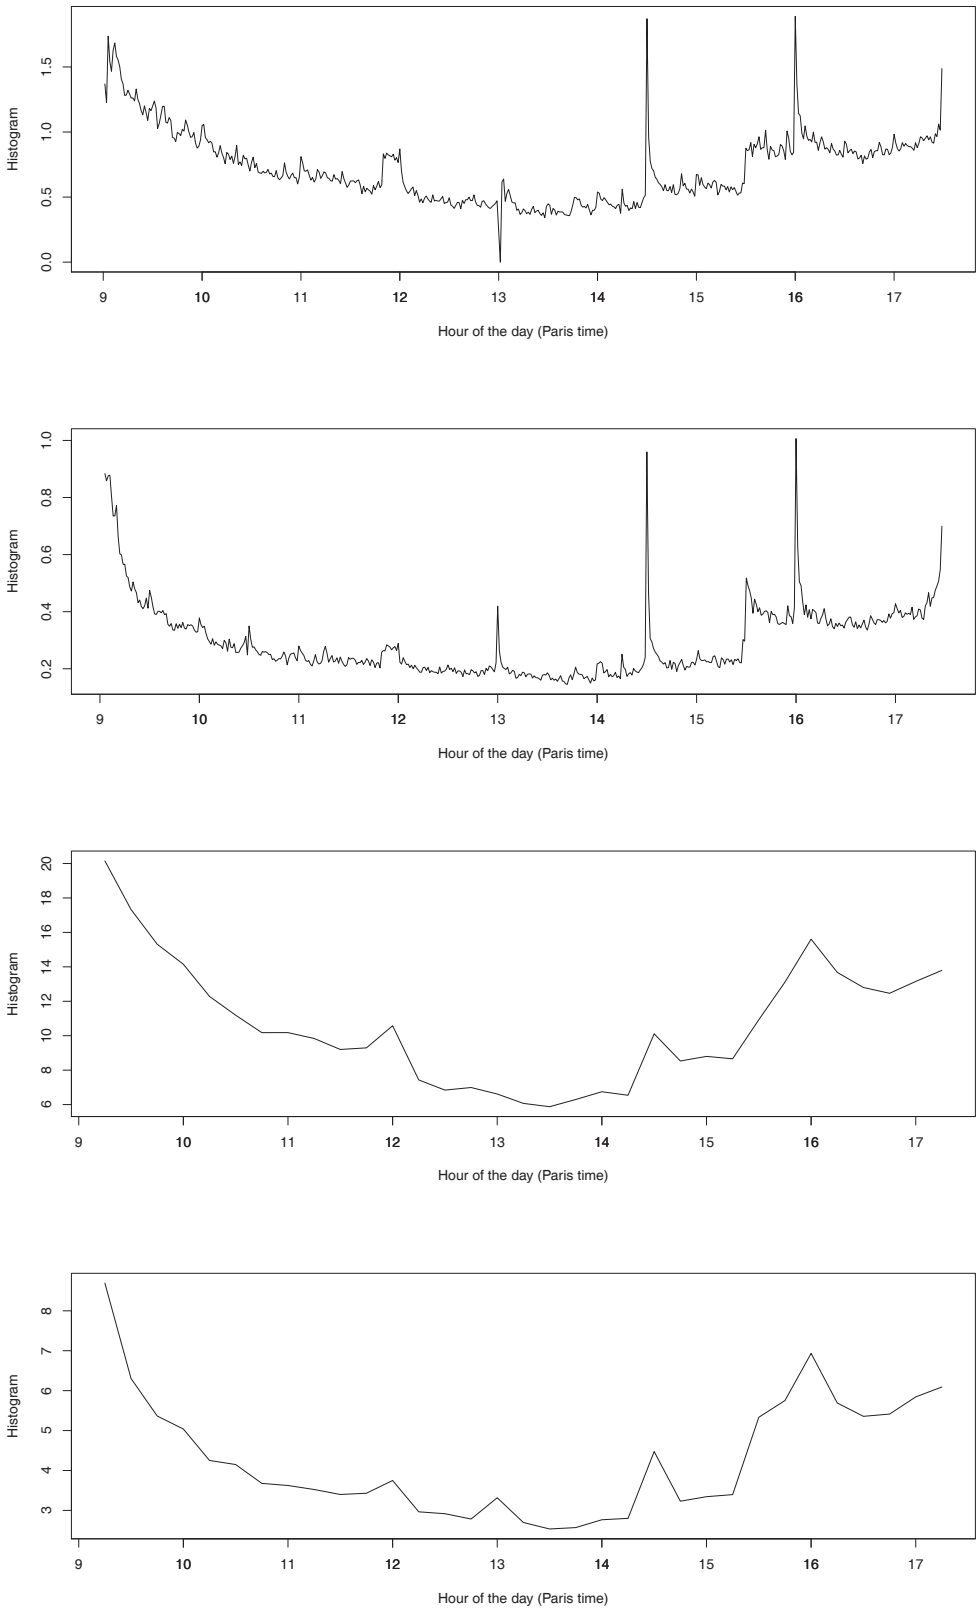

FIGURE 5.6 – Mean intraday histograms of timestamps of trades-through for the 30 most liquid components of FDX and FFI indices on 2011, using one-minute bins (*2 graphs up*) and 15-minutes bins (*2 graphs down*)

Conclusion

We empirically studied trades-through, we modelled them and used them in applications. Let us summarize the results we attained :

- We studied trades-through empirical properties : frequency, volume, intraday distribution, market impact, clustering and spread relaxation.
- A new method for the measurement of lead-lag parameters between assets, sectors or markets was presented.
- The trades-through intensity signal was measured and used in a simple trading strategy based on lead-lag.
- Trades-through clustered arrival times were modelled using Hawkes processes.
- We improved a toy-trading strategy based on the autocorrelation of trades-through arrival times by the use of our Hawkes model of trades-through.
- We obtained the empirical profile of trades-through decay kernel and showed it was well-fitted by a power-law.
- We generalized empirical results concerning the daily trajectories jumps detection on new financial indices.
- The empirical intraday proportion of detected jumps was given and linked to the intraday distribution of trades-through.
- We gave a possible explanation of this empirical proportion, using trades-through and intraday changes in the assets moves decomposition (between the continuous-volatility moves and pure-jumps moves).

In the future, it would be interesting to include optimal placement of limit orders in the different strategies based on trades-though we looked at, in order to take further advantage of their particular structure and definition. It would also be interesting to generalize our results on the strategy indicator based on our trades-through Hawkes model, with the empirical power-law decay kernel of trades-through. Finally, a broader study of the jumps structure across the different scales in financial time series, based on trades-through, liquidity indicators and statistical tests, would also be of great interest.

Bibliographie

F. Abergel and N. Huth. High frequency lead/lag relationship-empirical facts. 2011. URL `http://arxiv.org/abs/1111.7103`.

K. Al Dayri, E. Bacry, and J.F. Muzy. Non-parametric kernel estimation for symmetric Hawkes processes. Application to high frequency financial data. *Euro. Phys. Journal B*, 2011. To appear.

R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. 2005.

Y. Aït-Sahalia. Telling from discrete data whether the underlying continuous-time model is a diffusion. *The Journal of Finance*, 57 :2075–2112, 2002.

Y. Aït-Sahalia and J. Jacod. Testing for jumps in a discretely observed process. *Annals of statistics*, 37 :184–222, 2009.

Y. Aït-Sahalia, J. Jacod, and L. Jia. Testing for jumps in noisy high frequency data. *Journal of econometrics*, 168 :207–222, 2012.

E. Bacry, S. Delattre, M. Hoffmann, and J.F. Muzy. Modeling microstructure noise with mutually exciting point processes. *Quantitative Finance*, 2012. URL `http://www.tandfonline.com/doi/abs/10.1080/14697688.2011.647054`. Forthcoming.

O. Barndorff-Nielsen and N. Shephard. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2 :1–37, 2004.

C.M. Bishop. *Pattern recognition and machine learning*. Springer, 2007.

J.-P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing : from statistical physics to risk management*. Cambridge university press, 2003.

J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets : the subtle nature of 'random' price changes. *Quantitative Finance*, 4(2) :176–190, 2004.

J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. In Thorsten Hens and Klaus Reiner Schenk-Hoppe, editors, *Handbook of financial markets : dynamics and evolution*, pages 57–160. Elsevier, North-Holland, 2009.

C.G. Bowsher. Modelling security market events in continuous time : Intensity based, multivariate point process models. *Journal of Econometrics*, 141(2) :876–912, December 2007. URL `http://ideas.repec.org/a/eee/econom/v141y2007i2p876-912.html`.

P. Bremaud. *Point processes and queues, martingale dynamics*. Springer-Verlag New York Inc., October 1981.

P. Bremaud and L. Massoulie. Stability of nonlinear Hawkes processes. *The Annals of Probability*, 24(3) :1563–1588, 1996. ISSN 00911798. URL `http://www.jstor.org/stable/2244985`.

J. Carlsson, M. Foo, H. Lee, and H. Shek. High frequency trade prediction with bivariate Hawkes process. Stanford University, 2007.

P. Carr and L. Wu. What type of process underlies options ? a simple robust test. *The Journal of Finance*, 58 :2581–2610, 2003.

A. Chakraborti, I. Muni Toke, M. Patriarca, and F. Abergel. Econophysics review I : Empirical facts. *Quantitative Finance*, 7 :991–1012, 2011. URL `http://www.mendeley.com/research/econophysics-review-i-empirical-facts/`. In press.

R. Cont, A. Kukanov, and S. Stoikov. The price impact of order book events. 2011. URL `http://arxiv.org/pdf/1011.6402.pdf`.

D.J. Daley and D. Vere-Jones. *An introduction to the theory of point processes*, volume I : Elementary Theory and Methods. Springer, 2nd edition, 2003.

F. De Jong and M.W.M. Donders. Intraday lead-lag relationships between the futures-, options and stock market. *European Finance Review*, 1 :337–359, 1998.

F. De Jong and T. Nijman. High frequency analysis of lead-lag relationships between financial markets. *Journal of Empirical Finance*, 4 :259–277, 1997.

R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification*. Wiley-Interscience, 2000.

Z. Eisler, J.-P. Bouchaud, and J. Kockelkoren. The price impact of order book events : market orders, limit orders and cancellations. *Quantitative Finance*, 2012. URL `http://arxiv.org/abs/0904.0900`. Forthcoming.

J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and Sen A. What really causes large price changes ? *Quantitative Finance*, 4(4) :383–397, 2004.

A. Gerig. *A Theory for Market Impact : How Order Flow Affects Stock Price*. PhD thesis, University of Illinois, 2007.

T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning : Data Mining, Inference, and Prediction*. Springer, 2009.

N. Hautsch. *Modelling irregularly spaced financial data : Theory & practice of dynamic duration models*. Springer-Verlag Berlin and Heidelberg, April 2004.

N. Hautsch and R. Huang. The market impact of a limit order. *Journal of Economic Dynamics and Control*, 36(4) :501–522, 2012.

A.G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58 (1) :83 –90, April 1971. URL `http://www.jstor.org/stable/2334319`.

A.G. Hawkes and D. Oakes. A cluster process representation of a self-exciting process. *Journal of Applied Probability*, 11(3) :493–503, 1974. URL `http://www.jstor.org/stable/10.2307/3212693`.

M. Hoffmann, M. Rosenbaum, and N. Yoshida. Estimation of the lead-lag effect from nonsynchronous data. *Bernoulli*, 2011. To appear.

J. Jacod, Y. Li, P.A. Mykland, M. Podolskij, and M. Vetter. Microstructure noise in the continuous case : The pre-averaging approach. *Stochastic Processes and their Applications*, 119 :2249–2276, 2009.

G. Jiang and R. Oomen. Testing for jumps when asset prices are observed with noise–a "swap variance" approach. *Journal of Econometrics*, 144 :352–370, 2008.

A. Joulin, A. Lefevre, D. Grunberg, and J.-P. Bouchaud. Stock price jumps : news and volume play a minor role. *Wilmott Magazine*, Sep/Oct :1–7, 2008.

L. Kullmann, J. Kertesz, and K. Kaski. Time-dependent cross-correlations between different stock returns : A directed network of influence. *Physical Review E*, 66, 2002.

J. Large. Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets*, 10(1) :1–25, February 2007. URL `http://ideas.repec.org/a/eee/finmar/v10y2007i1p1-25.html`.

R. LeBaron and R Yamamoto. Long-memory in an order-driven market. *Physica A : Statistical Mechanics and its Applications*, 383(1) :85–89, 2007.

S. Lee and J. Hannig. Detecting jumps from Lévy jump diffusion processes. *Journal of Financial Economics*, 96 :271–290, 2010.

Suzanne S. Lee and Per A. Mykland. Jumps in financial markets : A new nonparametric test and jump dynamics. *The Review of Financial Studies*, 21 :2535–2563, 2008.

E. Lewis and G. Mohler. A nonparametric EM algorithm for multiscale Hawkes processes. 2011.

F. Lillo and J.D. Farmer. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics*, 8(3) :1, 2004. URL `http://ideas.repec.org/a/bpj/sndecm/v8y2004i3n1.html`.

F. Lillo, S. Mike, and J.D. Farmer. Theory for long memory in supply and demand. *Physical Review*, 71(6), 2005.

A.W. Lo and A.C. MacKinlay. When are contrarian profits due to stock market overreaction ? *The review of financial studies*, 3(2) :175–205, 1990.

D. Marsan and O. Lengline. Extending earthquakes' reach through cascading. *Science*, 319 : 1076–1079, 2008.

J. Moller and J.G. Rasmussen. Perfect simulation of Hawkes processes. *Advances in Applied Probability*, 37(3) :629–646, 2005. ISSN 00018678. URL `http://www.jstor.org/stable/30037347`.

I. Muni Toke. "Market making" in an order book model and its impact on the bid-ask spread. In F. Abergel, B.K. Chakrabarti, A. Chakraborti, and M. Mitra, editors, *Econophysics of Order-Driven Markets*, New Economic Windows, pages 49–64. Springer-Verlag Milan, 2011. URL `http://www.springerlink.com/content/x811111w104v5hk3/`.

Y. Ogata. On Lewis' simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1) :23–31, January 1981. URL `http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=01056305`.

T. Ozaki. Maximum likelihood estimation of Hawkes' self-exciting point processes. *Annals of the Institute of Statistical Mathematics*, 31(1) :145–155, 1979. URL `http://www.springerlink.com/content/hr3q7667x3522235/`.

F. Pomponio and F. Abergel. Trade-throughs : Empirical facts – application to lead-lag measures. October 2010. URL `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1694103`.

A. Ponzi, F. Lillo, and R. N. Mantegna. Market reaction to a bid-ask spread change : A power-law relaxation dynamics. *Physical Review*, 80(1), 2009.

M. Potters and J.-P. Bouchaud. More statistical properties of the stock order books and price impact. *Physica A : Statistical Mechanics and its Applications*, 324(1-2) :133–140, 2003.

B. Toth, J. Kertész, and J.D. Farmer. Studies of the limit order book around large price changes. *The European Physical Journal B - Condensed Matter and Complex Systems*, 71(4) :499–510, 2009.

B. Toth, I. Palit, F. Lillo, and J. Doyne Farmer. Why is order flow so persistent ? 2011. URL `http://arxiv.org/abs/1108.1632`.

P. Weber and B. Rosenow. Order book approach to price impact. *Quantitative Finance*, 5(4) : 357–364, 2005.

P. Weber and B. Rosenow. Large stock price changes : volume or liquidity ? *Quantitative Finance*, 6(1) :7–14, 2006.

A.G. Zawadowski, J. Kertész, and G. Andor. Large price changes on small scales. *Physica A : Statistical Mechanics and its Applications*, 334(1–2) :221–226, December 2004.

A.G. Zawadowski, J. Kertész, and G. Andora. Short-term market reaction after extreme price changes of liquid stocks. *Quantitative Finance*, 6(4) :283–295, 2006.