Statistical analysis and modeling of the opening and closing auctions of financial markets

Mohammed Salek

université PARIS-SACLAY

Statistical analysis and modeling of opening and closing auctions in financial markets

Analyse statistique et modélisation des enchères d'ouverture et de clôture dans les marchés financiers

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 573 : Interfaces : matériaux, systèmes, usages
Spécialité de doctorat : Mathématiques appliquées
Graduate School : Sciences de l'ingénierie et des systèmes.
Référent : CentraleSupélec

Thèse préparée dans l'unité de recherche **Mathématiques et Informatique pour la Complexité et les Systèmes (Université Paris-Saclay, CentraleSupélec)**
sous la direction de **Damien CHALLET**, Professeur,
et la co-direction de **Ioane MUNI TOKE**, Maître de conférences.

Thèse soutenue à Paris-Saclay, le 04 juin 2024, par

Mohammed SALEK

Composition du jury

Membres du jury avec voix délibérative

Sophie MOINAS Professeure, Université Toulouse Capitole	Présidente
Fabrizio LILLO Professeur, University of Bologna	Rapporteur & Examinateur
Huyên PHAM Professeur, Université Paris Cité	Rapporteur & Examinateur
Eduardo ABI JABER Assistant professor, Ecole Polytechnique	Examinateur

THESE DE DOCTORAT
NNT : 2024UPAST045

université PARIS-SACLAY | ÉCOLE DOCTORALE Interfaces: matériaux, systèmes, usages

Titre : Analyse statistique et modélisation des enchères d'ouverture et de clôture dans les marchés financiers.
Mots clés : Microstructure ; Enchères ; Finance quantitative ; Comportement des agents.

Résumé : Cette thèse est dédiée à l'étude des enchères d'ouverture et de clôture sur les marchés européens d'actions, plus particulièrement celles de la bourse de Paris. Les enchères sont un mécanisme essentiel permettant d'ouvrir et clôturer les journées de négociation de manière ordonnée. En particulier, les prix de clôture sont d'une importance cruciale pour les investisseurs et les régulateurs. Contrairement à la littérature abondante sur la phase de négociation continue, les travaux sur les enchères sont rares, en partie à cause de la difficulté d'acquérir des données de haute qualité. Cette thèse est basée sur un ensemble de données de haute qualité qui nous permet de fournir un aperçu nouveau sur les phases d'enchères, de reconstruire leur dynamique évènement par évènement et de proposer des modèles précis des phénomènes observés. Premièrement, nous examinons l'impact des ordres sur les prix dans les enchères. Nous fournissons un cadre mathématique pour les enchères, calculons la forme moyenne du carnet d'ordres au moment de l'enchère et sa répartition en fonction de la latence des agents et du compte utilisé. Nous étudions l'impact au moment de l'enchère ainsi que l'effet d'un temps de compensation aléatoire. Deuxièmement, nous adaptons un modèle continu en prix et en temps aux spécificités des enchères d'actions. Nous montrons que des solutions générales peuvent être obtenues en formules fermées. Nous résolvons numériquement les équations du modèle et nous le calibrons à la dynamique complète du carnet d'ordres moyen à la clôture. Nous concluons en étudiant les causes de la sous-diffusivité des prix indicatifs.

Title: Statistical analysis and modeling of opening and closing auctions in financial markets.
Keywords: Microstructure; Auctions; Quantitative finance; Agents behavior.

Abstract: This thesis is devoted to the study of opening and closing auctions in European markets with a specific emphasis on the Paris stock exchange. Auctions serve as an essential mechanism to open and close trading days in an orderly way. In particular, closing prices are of crucial importance for both investors and regulators. In contrast to the abundant literature on the continuous trading phase, work on auctions is scarce, partly due to the difficulty in acquiring high-quality data. This thesis is based on a high-quality dataset allowing us to provide novel insights into auction phases, reconstruct their tick-by-tick dynamics, and propose accurate models of the observed phenomena. First, we examine price impact in auctions. We provide a mathematical framework for auctions, compute the average shape of the limit order book at auction time, and its breakdown by agents' latency and account type. We investigate price impact at auction time and address the effect of random clearing times. Second, we adapt a continuous price-time model to the specifics of equity auctions. We show that time-dependent solutions can be obtained in closed-form formulas. We numerically solve the full equations and fit the average dynamics of the limit order book in closing auctions. We conclude by investigating the causes behind the indicative price sub-diffusivity.

Remerciements

Cette thèse a grandement bénéficié de l'engagement et du soutien de plusieurs personnes que je tiens à remercier chaleureusement. En premier lieu, mes plus sincères remerciements vont à mes directeurs de thèse, Damien Challet et Ioane Muni Toke, dont la patience, l'enthousiasme et la disponibilité ont été des atouts majeurs pour mener à bien ce travail tout au long de ces trois années. Leur intuition physique et habileté mathématique ne cessent de m'impressionner.

Je souhaite également exprimer ma sincère gratitude envers les rapporteurs, Fabrizio Lillo et Huyên Pham, pour avoir pris le temps de lire et de commenter ce manuscrit. Leur implication m'honore et je leur en suis profondément reconnaissant pour le temps précieux qu'ils y ont investi. Je tiens également à remercier chaleureusement Sophie Moinas et Eduardo Abi Jaber, de faire partie des membres du jury lors de la soutenance. Votre expertise est un pilier essentiel pour l'évaluation de ce travail.

J'adresse également des remerciements chaleureux à Christian Bongiorno et Gaoyue Guo, qui, avec Ioane et Damien, m'ont présenté des opportunités inestimables d'enseigner. Instruire des élèves talentueux fut un réel plaisir, et je les remercie également pour leur assiduité, sérieux, et bonne humeur.

Un grand merci à Charles-Albert Lehalle pour ces moments passés à parler d'enchères entre autres. Bien que le projet n'ait pas (encore) abouti, nos échanges ont été riches d'enseignements. Je suis également redevable aux membres de mon comité de suivi, Michael Benzaquen et Bence Toth, pour le temps qu'ils m'ont consacré, et pour leurs interrogations et remarques pertinentes.

Mes remerciements s'étendent à tous les membres du laboratoire MICS que j'ai eu le plaisir de côtoyer durant ces trois années, en particulier à tous les doctorants, post-doctorants et stagiraires et plus spécifiquement ceux du groupe FiQuant, tant les anciens que les nouveaux, Mouhamad, Jeremi, Vincent, Timothée, Lamia, Victor, Bastien, Youssef, Michele. Une mention spéciale à Vincent, Timothée et Lamia pour ces discussions-déjeuners passionnants. Je suis également reconnaissant envers mes collègues qui ont contribué, ou contribuent à l'ambiance conviviale

des open space maths et biomaths, Romain, Lukas, El Mehdi, Théo, Yassine, Jules, Charlotte, Gabriel, Gabriel, Thomas, Francesco, Nicolas, Pierre, Dimitra, Imane, Malek, Margherita, Inès, Félicie. Merci à Jules pour nos échanges stimulants. Merci à Fabienne pour son aide précieuse dans le dédale des procédures administratives et son énergie positive. Merci également à ceux que j'aurais involontairement omis de mentionner.

Enfin, je remercie chaleureusement mes amis et ma famille pour leur soutien indéfectible tout au long de ce parcours. Merci à Chaimae de sa patience et de son aide capitale. À mon père, j'espère que tu serais fier du travail accompli. À ma mère Latifa Ennissay, merci de ton soutien décisif et sans limite ; à ma sœur Hajar Salek, à mon frère Houssam Salek et à mon neveu Jad, merci d'être une source inépuisable de bonheur et de sérénité.

Résumé

Cette thèse est dédiée à l'étude des enchères d'ouverture et de clôture sur les marchés européens d'actions, en particulier sur les enchères de la bourse de Paris. Les enchères sont un mécanisme essentiel permettant d'ouvrir et clôturer les journées de négociation de manière ordonnée. En particulier, les prix de clôture sont d'une importance cruciale pour les investisseurs et les régulateurs.

Pendant la journée de négociation, les marchés électroniques fonctionnent selon le principe de la double enchère continue : à tout moment, les agents proposent un certain volume d'offres et de demandes, à différents prix ; à tout moment, un agent peut accepter n'importe laquelle de ces propositions et commercer immédiatement. Pendant les enchères d'ouverture et de clôture, aucune transaction ne peut avoir lieu. Les agents peuvent néanmoins soumettre, modifier et annuler des ordres. Au moment de l'enchère, les ordres compatibles se concluent au prix de l'enchère. Il s'agit du prix qui maximise le volume échangé et minimise le volume restant au prix de l'enchère.

Le mécanisme d'enchères gagne en importance. Les bourses de Londres et de Francfort ont introduit une enchères quotidienne à la mi-journée. En outre, elles s'appuient désormais sur des enchères de volatilité au lieu d'arrêter brusquement la négociation continue lorsque des limites de prix anormales sont déclenchées. Les marchés européens de l'électricité fonctionnent désormais selon un mécanisme de vente aux enchères à l'avance. Le Chicago Board Options Exchange a lancé un marché d'actions dédié aux enchères successives. Enfin, les volumes de clôture augmentent constamment, en particulier sur les marchés européens où le volume de clôture peut dépasser la moitié du volume quotidien échangé pendant les jours de rebalancement et d'expiration de dérivées.

Contrairement à la littérature abondante sur la phase de négociation continue, les travaux sur les enchères sont rares, en partie à cause de la difficulté d'acquérir des données de haute qualité. Précisément, cette thèse est basée sur un ensemble de données de haute qualité qui nous

permet de fournir un aperçu nouveau sur les phases d'enchères, de reconstruire leur dynamique évènement par évènement et de proposer des modèles précis pour les phénomènes observés.

Dans le Chapitre 1, nous fournissons une introduction complète aux principaux mécanismes de compensation sur les marchés. Nous effectuons une revue de la littérature sur les enchères d'actions, présentons l'impact des ordres sur le prix dans les marchés ouverts et soulignons les différences notables lorsqu'on le considère dans le cadre des enchères. Nous présentons des modèles de liquidité latente qui seront utiles pour la modélisation des enchères et concluons ce chapitre par un aperçu de nos principales contributions.

Le chapitre 2 est consacré à l'analyse de l'ensemble de données de haute qualité utilisé dans cette thèse. Il présente les données et les méthodologies utilisées pour reconstruire les états du carnet d'ordres et la séquence d'événements. Ce chapitre se termine par une description statistique préliminaire, incluant une analyse de la contribution des différentes catégories d'agents à l'enchère de clôture, ainsi qu'une mesure des caractéristiques moyennes de chaque type d'ordre.

Dans le chapitre 3, nous examinons l'impact des ordres sur les prix dans les enchères. Nous fournissons un cadre mathématique pour les enchères, calculons la forme moyenne du carnet d'ordres au moment de l'enchère et sa répartition en fonction de la latence des agents et du compte utilisé. Spécifiquement, nous étudions l'impact au moment de l'enchère, en supposant qu'un opérateur plus rapide que les autres agit en dernier dans l'enchère. Nous étudions l'effet d'un temps de compensation aléatoire et ce que cela implique pour la négociation au moment de l'enchère. Nous concluons en calculant les fonctions de réponse du prix indicatif.

Dans le chapitre 4, nous adaptons un modèle continu en prix et en temps aux spécificités des enchères d'actions. Nous calibrons des solutions stationnaires aux données de carnets d'ordres au moment de l'enchère et montrons que des solutions dynamiques peuvent être obtenues en formules fermées. Nous mesurons les soumissions, les annulations et les changements de prix au cours de l'enchère, en les reliant aux paramètres de notre modèle. Nous résolvons numériquement les équations du modèle et nous le calibrons à la dynamique complète du carnet d'ordres moyen à la clôture. Nous concluons en étudiant les causes de la sous-diffusivité des prix indicatifs.

Dans le chapitre 5, nous discutons nos principaux résultats et nous proposons quelques extensions des travaux présentés dans cette thèse.

L'annexe A tente de répondre à la question suivante : comment un investisseur peut-il répartir de manière optimale un grand ordre entre la phase de négociation continue et l'enchère de clôture tout en minimisant les coûts liés à l'impact ?

Abstract

This thesis is dedicated to the study of opening and closing auctions in European markets with a specific emphasis on the Paris stock exchange. Auctions serve as an essential mechanism to open and close trading days in an orderly way. In particular, closing prices are of crucial importance for both investors and regulators.

During the trading day, electronic markets operate on the continuous double auction: at any time, agents offer a certain volume of bids and asks, at various prices; and at any time, an agent can accept any of these proposals and trade immediately. However, during opening and closing auctions, transactions can not occur. Still, agents can submit, modify, and cancel orders. At the auction time, eligible orders are cleared at the auction price: it is the price that maximizes the exchanged volume and minimizes the remaining volume at the auction price.

The auction mechanism is gaining importance and attracting increasing attention. The London and Frankfurt stock exchanges have introduced a daily midday auction. In addition, they now rely on volatility auctions instead of abrupt trading halts when abnormal price limits are triggered. European electricity markets now operate on a day-ahead auction mechanism. The Chicago Board Options Exchange has launched a dedicated equity market operating on successive auctions instead of the continuous double auction. Ultimately, closing auction volumes are consistently increasing, especially in European markets where the closing volume can exceed half the daily exchanged volume on rebalancing and expiry days.

In contrast with the abundant literature on the continuous trading phase, work on auctions is scarce, partly due to the difficulty in acquiring high-quality data. Precisely, this thesis is based on a high-quality dataset allowing us to provide novel insights into auction phases, reconstruct their tick-by-tick dynamics, and propose accurate models of the observed phenomena.

In Chapter 1, we provide a comprehensive introduction to the main clearing mechanisms in equity markets. We conduct a literature review on equity auctions, present price impact in open markets, and highlight notable differences when considered in auctions. We introduce latent

liquidity models that will be beneficial for auction modeling and conclude this Chapter with an outline of our main contributions.

Chapter 2 is dedicated to the analysis of the high-quality dataset used in this thesis. It presents the data and the methodologies for reconstructing order book snapshots and tick-by-tick events. This chapter concludes with direct observations of the refined data, including an analysis of the contribution of various agent categories in the closing auction, as well as a description of the average dynamics by order type.

In Chapter 3, we examine price impact in auctions. We provide a mathematical framework for auctions, compute the average shape of the limit order book at auction time, and its breakdown by agents' latency and account type. We investigate price impact at auction time, assuming a low latency trader acts last in the auction. We address the effect of random clearing times and what it implies for trading at auction time. We conclude by calculating response functions of the indicative price.

In Chapter 4, we adapt a continuous price-time model to the specifics of equity auctions. We use stationary solutions to fit order books at auction time and show that, in some cases, dynamic solutions can be obtained in closed-form formulas. We provide measurements of auction submissions, cancellations, and price updates, linking them to our model parameters. We numerically solve the full equations and fit the average dynamics of the limit order book during the closing auction. We conclude by investigating the causes behind the indicative price sub-diffusivity.

In Chapter 5, we discuss our core results and elaborate on possible extensions and future outlooks.

Appendix A tries to provide meaningful answers to the following question: How can an investor optimally split a large order between the continuous trading phase and the closing auction while minimizing impact-related costs?

Contents

Chapter 1

General introduction

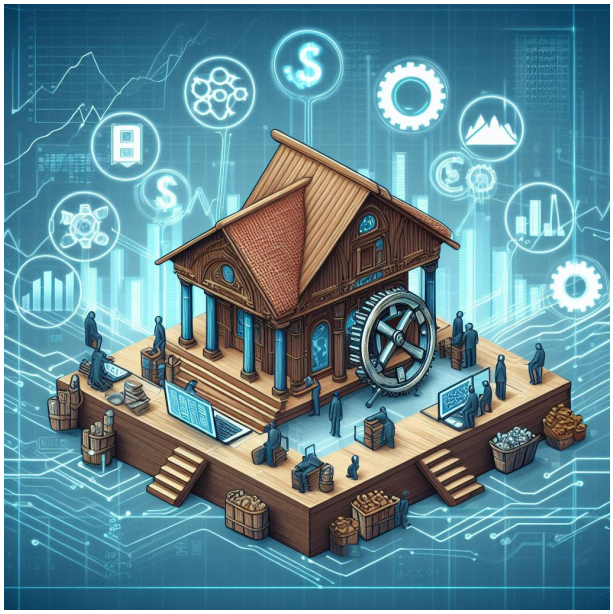

Figure 1.1 – AI-Generated image by DALL-E 3 using the prompt: A brief history of markets and mechanisms of market clearing.

1.1 A brief history of markets

1.1.1 Ancient markets

Humans used bartering, donating, and lending, to acquire and share the resources they needed to survive, collaborate, and prosper. To complete a bartering transaction, both parties had to agree on the quantity, quality, and nature of the items they were willing to offer and receive. To overcome this difficulty, humans invented (or evolved to use) currency (Svizzero and Tisdell, 2019). It is a medium of exchange that represents a standard unit of value. It should be stable, durable, divisible, portable, and widely accepted. It can be issued by a central authority, such as a government or a bank nowadays, or have intrinsic value, such as gold or silver. Currency has a long history, dating back to when people used shells, beads, and grains as money. Over time, currency evolved to include coins, paper notes, and digital forms. With currency, people who wanted to exchange items and goods only had to determine the fair price in the currency unit.

The common trust in the currency's value enabled the trade of commodities to flourish through price negotiation.

In parallel, the concept of a company's share or stock can be traced back to 1250 in Toulouse with what is likely the first joint-stock company, "Les moulins de Bazacle" (Sicard, 1953). Approximately a hundred shares, representing a fraction of ownership in this company, were traded, and their value fluctuated based on the company's profitability. However, the modern understanding of a joint-stock company, with freely transferable shares and limited liability for shareholders, was more formally instituted with the Dutch East India Company in the 17th century, whose shares were exchanged on the Amsterdam Bourse.

The establishment of the Amsterdam stock market (1602) laid the groundwork for the development of stock exchanges in other countries. Inspired by the Amsterdam model, London (1698) and Paris (1724) founded their respective stock exchanges. These new financial institutions adopted the principles of freely transferable shares and limited liability for shareholders, which had proven successful in Amsterdam. The foundation of these stock exchanges marked the beginning of a new era in the global financial market. They facilitated the growth of companies and contributed to economic development. Over time, these exchanges have evolved and adapted to changes in technology and market dynamics. Still, their core function remains the same: to provide a marketplace for buying and selling securities.

1.1.2 Modern markets

Financial markets have undergone structural changes in the past decades due to the massive electronization of exchanges, allowing quasi-instantaneous transactions anywhere in the world. Electronic stock exchanges operate digital platforms where market participants can execute transactions. They typically charge trading fees for liquidity takers and offer rebates for liquidity providers. Their role is similar to that of a poker house, collecting fees with each played hand. With the advent of MiFID regulations in Europe, alternative markets known as multilateral trading facilities (MTFs) have emerged. Examples include BATS and Chi-X, merged into BATS-Chi-X and acquired by CBOE Europe, Turquoise, acquired by the LSE, and other MTFs launched by large investment banks. These platforms compete with primary exchanges such as Euronext in Paris, LSE in London, and Xetra in Frankfurt. The rise of alternative exchanges results in market fragmentation, in contrast with market centralization around a central primary exchange.

Gaining direct access to financial markets is subject to strict regulatory requirements. How-

ever, thanks to the large number of possibilities offered by brokerage firms, accessing financial markets is becoming simpler for investors nowadays. A variety of agents can trade a vast array of asset classes. Although highly heterogeneous, market participants can be categorized into broadly defined groups that may sometimes overlap:

1. Market makers animate the market. They are consistently present and offer to sell and buy from impatient agents. They are usually high-frequency actors. Their business model is to collect the difference between the ask (the price at which they sell) and the bid (the price at which they buy), assuming prices are stable. Retail exchange offices and gold jewelers operate with the same market-making mechanism. However, market makers carry the inventory risk. For instance, when they buy a substantial amount of a security and hold it for a certain period awaiting sellers, the ask price may plummet far below the bid at which they bought it. This results in losses if they were to liquidate this position immediately, i.e., clear their inventory. This is why they continuously adjust their quotes in order to manage their inventory and need to be profitable on average.
2. Brokers, often large financial institutions, provide access to the market in return for fees to retail investors and other financial institutions who cannot operate directly in the markets. They bear the risk of default of their clients. Additionally, they are subject to several compliance and regulatory obligations.
3. Investment banks offer a wide range of financial services to their clients. They can function as market makers, brokers, and provide their clients with customized financial instruments such as complex derivative products or hybrid proprietary indices. Given their crucial role in the financial network, they are subject to numerous regulatory mandates. Indeed, the collapse of a major investment bank poses a systemic risk to the financial system and potentially the global economy. For example, the bankruptcy of the bank Lehman Brothers in 2008 was a triggering event of the great financial crisis that followed.
4. Institutional investors are significant players in the market due to the size of their assets, such as pension funds, sovereign funds, and insurance companies. They are usually mandated to manage and invest in a large volume of financial instruments. Depending on their objectives, mandates, and the maturity of claims, they usually engage in acceptable risks to maximize returns.
5. Portfolio managers, investors, or speculators. Examples include high-frequency trading agents that employ sophisticated technological infrastructures to conduct market making and profit from price discrepancies. Medium frequency agents such as hedge funds employ quantitative strategies to detect investment opportunities or statistical arbitrage. Low-

frequency agents such as mutual funds often rely on more traditional fundamental analysis to build portfolio positions on markets.

1.2 Mechanisms of market clearing

To monitor financial markets efficiently, regulators need to achieve a deep understanding of the market microstructure and the price formation process. Market microstructure is a branch of finance that studies how trading rules affect price determination and how information and behavior drive the price. Natural questions arise: what is the fair price for a given security at a given time? Does it even exist? Should prices fluctuate? How do they fluctuate? How can the exchange of securities be regulated in order to ensure fairness, competitiveness, and financial stability?

1.2.1 Auctions

One way to determine a price for a given security is to hold an auction. The word "auction" (late 16th century) is derived from the Latin word augere, then auctio, meaning to increase. Auctions are a market mechanism that allows one or multiple buyers to transact with one or multiple sellers. The most popular auction mechanism is the English auction (McAfee and McMillan, 1987; Krishna, 2009), where multiple buyers try to outbid each other for one asset. The highest bid wins. Throughout the auction, the current highest bid is known by bidders. Different variations of the English auction, also called an ascending-bid auction, are possible: bids may be announced by the bidder themselves, by the auctioneer, by proxies, and nowadays submitted electronically.

Other auction types may be driven by bidders, time, or a combination of both. In the auction by candle, another type of ascending-bid auctions, the highest bid right before a candle goes out wins. Conversely, in descending-bid auctions such as the Dutch auction, the announced prices by the auctioneer decrease and the first bid wins. Japanese auctions work in reverse as the announced price continuously increases and bidders abandon when the price becomes too large for them. The last standing bidder wins at the current price.

Auctions can be "sealed-bid" where bids are submitted simultaneously, and no bidder knows how much the others bid. Sealed-bid auctions can be first-price, where the highest bid wins and pays his bid, or second-price (Vickrey auction), where the winner pays the best second bid.

All the stated above are forward auctions where multiple buyers compete to buy an asset. In

reverse auctions, however, multiple sellers compete to sell to one buyer. The lowest ask price, the equivalent of the bid in reverse auctions, wins. Lastly, double auctions involve multiple buyers and multiple sellers.

Auctions may comprehend additional features that change the structure of bids, such as a bidding fee, a minimum bid increment, or the visibility of the bids. The determination of the winner and the paid price can also be subject to specific rules. The winners' curse emphasizes that the winner of an ascending-price auction often overestimates the asset's value and ends up paying more (Thaler, 1988). This is the case, for instance, when the distribution of bids is broadly distributed around the fair asset value.

Paul Milgrom and Robert B. Wilson won the 2020 Nobel in Economy for their significant contributions to auction theory (NobelPrize.org, 2020). For more about auction theory, see for instance Milgrom (2004); Krishna (2009).

1.2.2 Walrasian auctions: price *tâtonnement*

Economists addressed the price determination problem within the broader context of an economy, e.g., with Leon Walras formulating a general equilibrium theory of the economy in the 19th century (Walras, 1900). During this period, several theories were proposed regarding the origin of an asset's value, suggesting it could be derived from labor, utility, or scarcity. Walras' work is widely acknowledged today for its substantial contributions to the mathematization of economics.

The marketplace had identified the concepts of supply and demand long before they were mathematically theorized and documented. Supply is the total available volume to be sold, and demand is the total volume to be acquired, at a given time, at a given price. Supply often increases as the price increases and vice-versa for demand. Figure 1.2 shows the cumulative quantity to be sold in blue (supply) and the cumulative quantity to be bought in red (demand) as a function of the price. In this case, supply $p \to S(p)$ and demand $p \to D(p)$ intersect at a unique equilibrium price $p = p_a$ that maximizes the exchanged quantity $Q_a = S(p_a) = D(p_a)$.

When S and D are not fully revealed, the equilibrium price can be determined using the dynamic process known as Walras' tâtonnement (trial and error). In essence, the equilibrium price is established by an auctioneer (crieur) who gathers all potential buyers and sellers in a single place and then proceeds by tâtonnement. The auctioneer announces the first price and then collects, anonymously from each agent, the quantity to be sold or bought at the announced price. Then, the total quantity to be sold (supply) and bought (demand) are compared. If

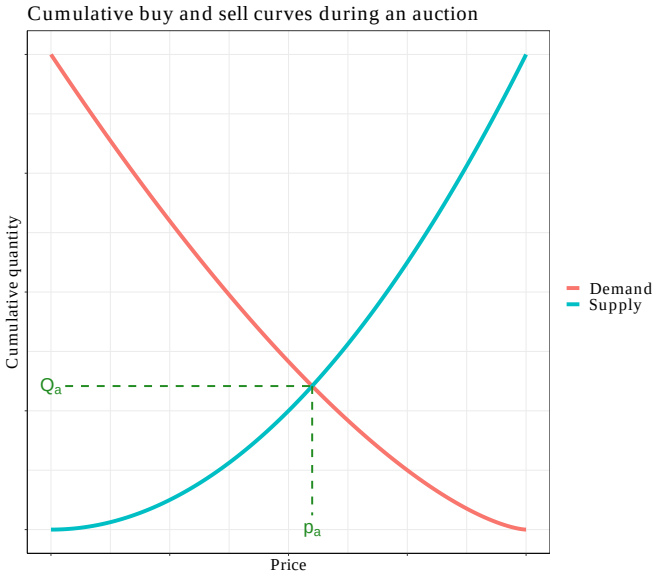

Figure 1.2 – Theoretical supply and demand curves during a Walrasian auction. The equilibrium price p_a is the price that maximizes the total exchanged quantity Q_a between buyers and sellers.

supply exceeds demand, the auctioneer understands that the announced price is larger than the equilibrium price, and vice-versa. He then iterates by announcing an adjusted price, collects the updated buy and sell intentions, and compares the total updated supply and demand. Iterations are carried on until converging to an equilibrium price where supply equals demand. Importantly, no transaction is allowed to happen during this procedure. Once an equilibrium price is reached, willing buyers and sellers transact at the auction price. Theoretically, the auction price p_a maximizes the possible volume that satisfied buyers and sellers can exchange, which is the auction volume Q_a (see Fig. 1.2). The convergence relies on strong assumptions such as perfect information and competition that are not always verified in practice.

Walras' auction allows buyers and sellers to exchange large quantities at a unique price and time. While the transaction price is determined by $S = P$, the time of the clearing is determined by the exchange.

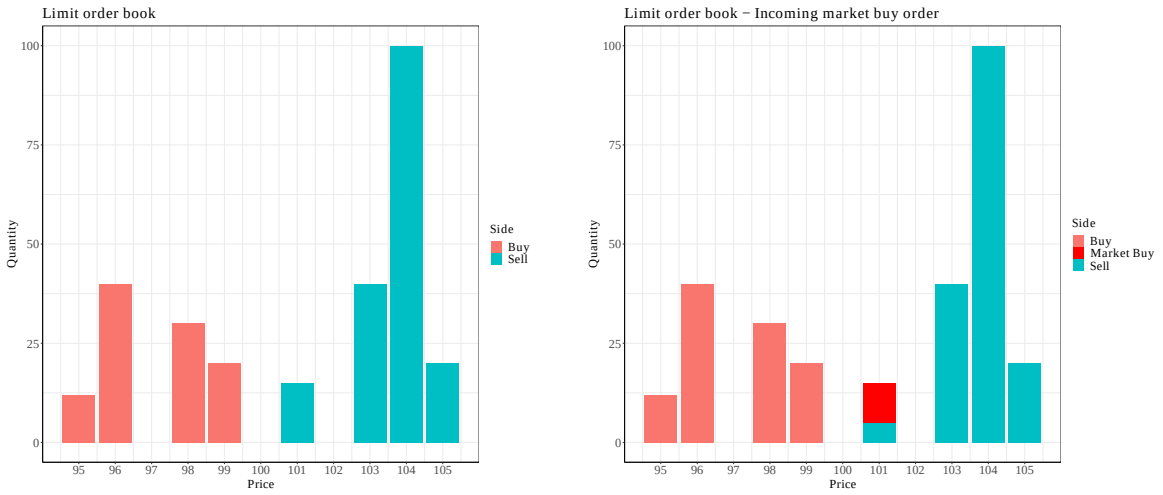

Figure 1.3 – Schematic view of a limit order book (left panel) and an incoming buy market order (right panel).

1.2.3 Continuous double auctions

Continuous double auctions enable impatient buyers and sellers to trade anytime they choose ("continuous"). The buy price and the sell price are determined by two distinct auctions that run in parallel, hence the name "double auctions". Patient sellers publicly display the quantities they are willing to sell and the corresponding prices, and likewise for patient buyers. These fixed-price, fixed-quantity orders are known as limit orders and result in an order book that resembles Fig. 1.3.

In this figure (left panel), patient sellers are offering the quantities in blue bars to be sold at prices 101, 103, 104, and 105. Conversely, patient buyers offer to buy the quantities in light red bars at prices 99, 98, 96, 95. Subsequently, impatient buyers can purchase the blue quantities offered in the order book, and impatient sellers can sell the red quantities to patient bidders using market orders that are instantaneously executed. The lowest offered price is the ask (here, 101), and the highest price a buyer is willing to pay is the bid (here, 99). The spread is the difference between the best ask and the best bid (here, 2); it is an important measure of the market liquidity. The larger the spread, the less liquid the asset. Market orders consume the available limit volume at the opposite side of the book, either fully or partially, across one or several price levels. The prevailing price of the asset can be defined as the last transaction price, the mid-price (the average of the bid and ask), or the weighted mid-price (the volume-weighted average of the bid and ask).

During the day, most stock exchanges, such as Euronext in France, Xetra in Germany, or

the LSE in the UK, operate using electronic order books, also known as continuous limit order books. In addition to limit and market orders, most exchanges provide multiple order types to market participants. For example, in Euronext, one finds:

1. Stop orders that can be limit or market orders, and activate only when the price crosses a specified threshold. The price threshold for buy orders must be higher than the current price and vice-versa for sell orders, hence their other name "stop loss orders".
2. Primary pegged orders are limit orders that track the best bid when on the buy side and the best ask when on the sell side, with possibly an offset from the best bid/ask.
3. Iceberg orders can be large limit orders with only a small disclosed quantity in the order book; generally, the hidden quantities acquire queue priority only when revealed.

Orders can be submitted with a validity parameter or execution instructions:

1. Good-for-day orders are canceled at the end of the trading day if not executed.
2. Good-till-cancel orders persist in the order book until completion.
3. Valid-for-auction or Valid-for-closing orders can be submitted throughout the trading day and become active only at the start of the relevant auction phase.
4. Fill-or-kill orders are executed either totally or canceled.
5. Immediate-or-cancel orders are executed either totally or partially, and the remaining volume is automatically canceled.

When a limit order crosses the opposite best quote, it becomes a marketable order: it acts as a market order until it depletes the available volume up to its limit price. If a quantity remains unexecuted, it transforms into a limit order at a new best (bid or ask). Marketable orders are far more used than market orders during continuous trading sessions as they offer more control over the trading price. Market orders represented 3.7%, on average, of the orders that triggered a transaction during the continuous trading phase for the TotalEnergies stock on Euronext between 2013 and 2017.

There are many other types and possible parametrizations for order submissions that may depend on the considered exchange. We refer the reader to the complete rulebooks, e.g., Euronext (2023a,b) for further information on general trading rules, order types, and parametrizations on Euronext.

1.2.4 Equity auctions

Limit order books allow market participants to trade instantaneously when needed and to engage in market-making activities. However, most stock exchanges only operate during the day, unlike the FX market for instance, which operates on a 24/7 basis. Before the introduction of call auctions, opening and closing prices were subject to a significant amount of manipulation. When the closing price was simply the last transaction price, prior to the introduction of the closing auction on the Paris stock exchange, Hillion and Suominen (2004) found that spread and volatility experienced a significant increase in the last minute of the trading day, along with a sharp increase in the bid/ask spread and the number of hidden orders. Using an equilibrium model, they argued that the observed patterns are likely due to the manipulation of the closing price.

The closing price is an important metric for stocks as it facilitates the benchmarking of daily returns, portfolios, and net asset values of funds, among others. In an effort to curb the manipulation, primary exchanges adapted their opening and closing mechanisms during the last decades. Some switched to mechanisms similar to a volume-weighted price in the last minutes of the trading day to determine the closing price. For instance, in 1993, the Hong Kong stock exchange started to use the median of five prices in the last minute of the trading day as the closing price (Park et al., 2020). Some others, such as the Paris stock exchange, adopted similar mechanisms to Walras' auction. Currently, most electronic exchanges, including primary exchanges in European countries, use auctions to start and end trading days on liquid stocks. Unlike US equity auctions where continuous trading runs in parallel, continuous trading is halted in European equity auctions.

Equity auctions start with an accumulation period and end with a clearing process. During the accumulation period, participants submit orders (quantity, price, side, order type, ...) to the exchange. Modifications and cancellations are allowed, but transactions cannot occur. Figure 1.4 (left panel) shows an example of a limit order book during an equity auction. Buy and sell limit orders overlap in the order book. When supply and demand overlap during the accumulation period, there is no best bid, no best ask, and no spread. At all times, an indicative price p_a can be computed: it maximizes the matched volume (indicative volume) Q_a and minimizes the unmatched volume at the indicative price. The unmatched volume at the indicative price is known as the imbalance or the surplus. Finding the indicative price and volume boils down to finding the intersection of the supply and demand curves, as shown in the right panel of Fig 1.4. At auction time, buy orders whose prices are larger than the auction price and sell orders whose prices are smaller than the auction price are executed at the auction price. Limit orders whose

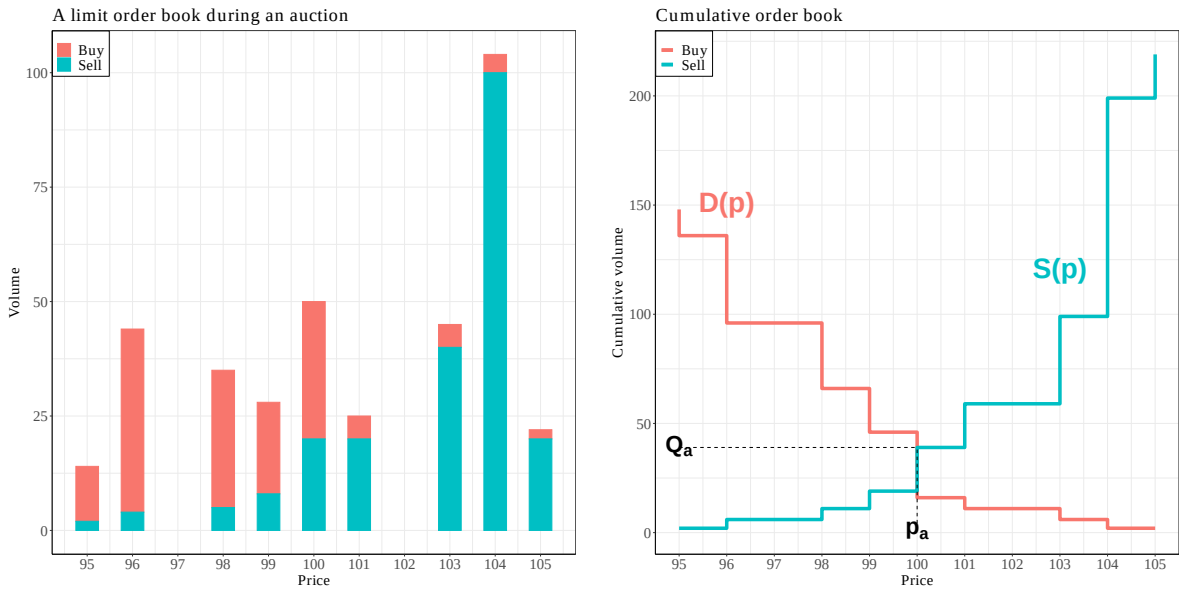

Figure 1.4 – Schematic view of a limit order book during an equity auction (left panel) and its cumulative counterpart (right panel).

limit price equals the auction price may be matched or remain in the order book depending on a time priority principle.

Note that the auction price is often not uniquely defined by the maximization of the exchanged volume alone; this is why exchanges implement a complementary set of rules such that p_a is always well defined. In the case of the Euronext markets, when multiple prices maximize the exchanged volume, the chosen p_a is the one with the smallest imbalance. Then, if multiple prices with the highest executable volume and the smallest imbalance still exist, the auction price is the one closest to the reference price (last traded price).

In the case of Euronext Paris, the accumulation period of the opening auction starts as soon as 07:15:00, and the clearing occurs randomly in a thirty-second window starting at 09:00:00. Likewise, the accumulation period of the closing auction starts at 17:30:00 and lasts for five minutes. The close clearing occurs randomly between 17:35:00 and 17:35:30. Following the closing auction clearing and until 17:40:00, the exchange operates a period of trading-at-last in which agents can further trade at the closing price. Figure 1.5 shows a sketch of a typical trading day in the Paris stock exchange for a liquid stock.

The closing auction volume has seen consistent growth, particularly in European markets, where it can surpass half the daily volume on days of index rebalancing and derivative expiry (Raillon, 2020). To enhance liquidity during midday, typically the period of lowest liquidity, and

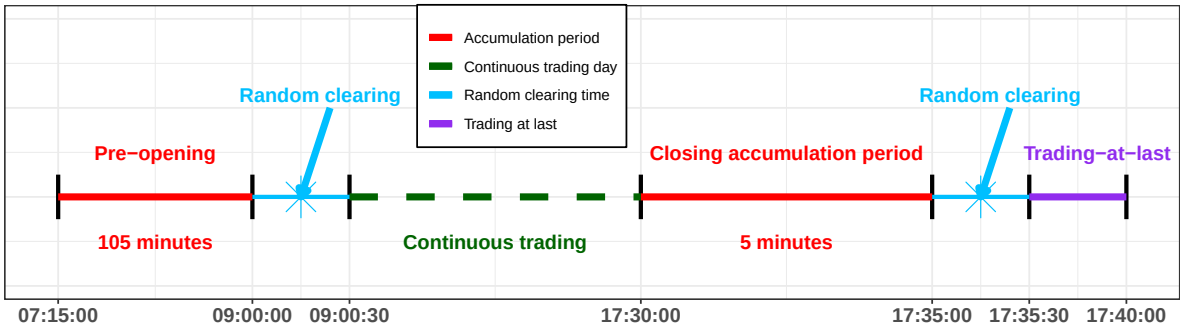

Figure 1.5 – Schematic view of a typical trading day on Euronext Paris.

to mitigate the exposure to high-frequency market makers which dominate the total turnover during day trading (AMF, 2017), some primary exchanges, such as Xetra and LSE, also introduced a daily intraday auction. For less liquid stocks in Paris, Euronext (2019) implements the double-fixing trading method, i.e., two auctions per day with periods of trading at last after the uncrossings, and no continuous double auction.

Financial markets are highly complex systems where numerous human and algorithmic agents interact at different timescales. Identifying the fundamental laws that govern these systems, ranging from micro- to macro-level, is a considerable challenge. However, with the advent of increasingly precise data, the significance of market microstructure research has greatly increased. Many stylized facts have been uncovered, and theoretical models can now be tested and confronted with real-world data, leading to a robust body of knowledge in this field (Mantegna and Stanley, 1999; Chakraborti et al., 2011a,b; Abergel et al., 2016; Bouchaud et al., 2018a; Lehalle and Laruelle, 2018).

Despite their importance, equity auctions have been the subject of very few studies in comparison with open markets, partly due to the difficulty of obtaining suitable granular data. This thesis is based on high-quality data, which will enable us to reconstruct the tick-by-tick dynamics of accumulation periods and propose adequate modeling of the observed phenomena.

1.3 Empirical studies of call auctions

1.3.1 Stylized facts

In the following, call auctions refer to the auction market mechanism, pre-opening refers to the accumulation period of the opening auction, and frequent batch auctions refer to a market mechanism where call auctions are sequentially held one after the other. They are also named periodic auctions.

Early studies of the pre-opening

The Paris stock exchange used to operate on a (unique) daily call auction before 1986; then it gradually switched to a computerized continuous limit order book mechanism (Biais et al., 1995). This change allowed for continuous trading from 10:00 to 17:00 each day, preceded by a pre-opening period, i.e. an opening auction. In 1998, a closing auction mechanism was introduced for the most actively traded stocks. The continuous trading phase in Euronext Paris currently spans from 09:00 to 17:30.

Generally, most primary stock exchanges introduced call auctions in order to determine opening and closing prices in the late 90s' and at the start of the 21st century.

Biais et al. (1999) study the pre-opening period on the Paris bourse, highlighting its high degree of similarity with the Walrasian tâtonnement process. They test whether the indicative price reflects learning vs noise and find that early in the pre-opening period, the noise hypothesis is not rejected. However, as the opening time gets closer, the informational content and efficiency of prices increase, and the learning hypothesis can no longer be rejected.

The effect of introducing closing call auctions on market quality

The seminal work of Pagano and Schwartz (2003) study the effect of this introduction on market quality. They find a substantial reduction in transaction costs after the introduction of the closing auction and enhanced accuracy of the price determination for the overall market.

Similarly, Aitken et al. (2005) analyze the effect of the introduction of a closing call auction on the Australian stock exchange in 1997. They document a shift of activity from the last hour of continuous trading to the closing auction. In addition, they find that the closing auction consolidates liquidity without having an adverse effect on the cost of trading in the continuous session.

Lastly, they document higher activity and exchanged volumes during Fridays and quarter-end days.

Kandel et al. (2012) conduct similar research on the Borsa Italiana and report a volume shift from the last minutes of the continuous phase to the closing auction. Additionally, they report reduced spreads, volatility, and average trade sizes in the last minutes of continuous trading. Interestingly, they evidence an increased average order size in the closing auction one year after its introduction; they link this finding to increased confidence in agents' ability to execute large trades at the close.

Park et al. (2020) examine the effect of introducing a closing call auction on price manipulation in the Hong Kong stock exchange. The latter introduced a closing call auction in 2016 after it removed it in 2008, only 10 months after its adoption, due to suspicions of widespread price manipulation. The authors indicate that the call auction mechanism is prone to price manipulation unless it is coupled with other precautionary mechanisms. These include the establishment of price thresholds, volatility interruptions, randomization of the clearing, and price stabilization mechanisms such as the order-balancing mechanism in NYSE and NASDAQ.

Likewise, Li et al. (2021) investigate the impact of the introduction of a closing auction in the Shanghai stock exchange in 2018. They use the Shenzhen stock exchange as a control group where closing auctions have been held since 2006. Their findings indicate significant shifts in the trading volume from closing auctions to preceding continuous trading. In addition, they observe a significant decrease in the closing price deviation. The authors argue that the introduction of a closing call auction could reduce manipulation and liquidity noise in the closing price.

Alexakis et al. (2021) use tick-by-tick data from the Athens stock exchange to investigate the effect of introducing a closing call auction on price manipulation. They find that the closing call auction manages to lower manipulation compared with a value-weighted average price (VWAP) mechanism. In addition, they indicate that the closing auction does not fully eliminate manipulation, even when precautionary mechanisms are included in the auction. Lastly, they suggest continuous monitoring from regulators with a specific emphasis on the reference price.

Pre-opening, high-frequency trading, and fragmentation

Boussetta et al. (2017) investigate the quality of price discovery in the pre-opening of Euronext Paris within the context of fragmented markets and the rise of high-frequency trading. They report an early activity of slow brokers and a later activity of fast traders (HFTs) and dedicated liquidity providers (market makers) in the last half-hour. Additionally, they find that the pre-opening activity of slow brokers is strongly related to the price discovery process across trading venues.

In the same vein, Bellia et al. (2017) examine the strategic behavior of high-frequency traders in Euronext Paris pre-opening. They document that HFTs neither harm nor improve the quality of the price discovery process and liquidity. Additionally, they report increased profitability for later submissions, whether stemming from HFTs or slow agents in comparison with earlier submissions. HFTs lead the price discovery process using their own accounts in contrast with market-making accounts, and the presence of liquidity providers in the pre-opening is marginal. They argue that the latter is the result of Euronext's policy that rewards market makers only during continuous trading; raising questions about this policy's relevance for the overall market quality.

Note that both Boussetta et al. (2017) and Bellia et al. (2017) employ the BEDOFIH database (European High-Frequency Financial Database; Base Européenne de Données Financières à Haute Fréquence), a high-quality dataset developed and maintained by EUROFIDAI (The European Financial Data Institute), which is used in this thesis. A detailed presentation of the BEDOFIH database, the data processing, and preliminary results on the ecology of closing auctions in Euronext Paris are provided in Chapter 2.

Statistical regularities in equity auctions

Gu et al. (2010) is one of the earliest works analyzing empirical regularities of equity auctions using high-frequency data from the Shenzhen stock exchange in 2003. They find that the distribution of limit prices is a skewed bell shape around the reference price. The maximum of the distribution of buy limit prices is reached below the reference price and vice versa for sell limit prices with asymmetric buys and sells. Interestingly, they evidence a large peak at the reference price, i.e., yesterday's closing price. Additionally, they study the distribution of order sizes, highlighting number preferences known as order size clustering. Lastly, they show that the average order book, i.e., the available limit volume as a function of the distance to the bid (for buys) and the ask (for sells), is exponentially decreasing.

Raillon (2020) documents the increasing share of the closing auction in the total traded volume, especially in European markets such as France where 41% of the CAC40 volume was exchanged on the close in June 2019. The author proposes several explanations for this trend such as the growth in passive investing, MiFID II regulations (best execution, transaction cost analysis), lower exposure to HFTs, and an amplification by automated execution algorithms. In addition, the study describes cases of indicative price and volume formation during the accumulation period. It highlights that indicative volume and price converge rapidly to their terminal values. 90% of the auction volume and 0.3% distance from the auction price are, on average, reached in the first two minutes of the closing accumulation period.

Challet and Gourianov (2018) conduct a thorough study of the dynamical regularities in US equity auctions, highlighting the same trend of increasing closing volumes. They find that the indicative price is mean-reverting with Hurst exponent $H < 1/2$, particularly in the closing auction. In addition, the authors examine response functions in auctions, i.e., the response of the auction price conditional on an order submission or cancellation at a given time during the accumulation period. They find that earlier actions have more impact than later ones, and that conditional on improving or worsening the imbalance, response functions can be markedly different. For instance, at the close, submissions that improve imbalance have the opposite effect on the auction price than those that worsen it, In contrast with the open, where imbalance-improving submissions have barely the same effect to no effect at all compared with imbalance-worsening submissions.

Lehalle and Laruelle (2018) provides a comprehensive overview of call auctions in its second chapter. The authors highlight the importance of closing auctions for primary exchanges in the context of market fragmentation. In addition, microstructural details of the indicative price determination across numerous primary exchanges are discussed. The authors report the temporal shape of the indicative volume and discuss informational peaks during the accumulation period. In parallel, the first chapter of Bouchaud et al. (2018a) formalizes the mechanism of price determination in Walrasian auctions. Features of opening and closing auctions are discussed in its third chapter.

Challet (2019) examines statistical regularities of the opening and closing auctions of French equities with a focus on the diffusive properties of the indicative auction price. The author finds that as the auction goes by, the typical price change decreases, favoring the sub-diffusion of the indicative price. Conversely, the author shows that the rate of events increases as the clearing approaches, favoring super-diffusion. The author concludes with the need for an additional mechanism —strategic behavior— to produce nearly diffusive prices.

Frauendorfer and Müller (2020) analyze closing auctions of Swiss equities and find that closing prices are highly sensitive to the removal of both small limit and market orders. For instance, removing 25% of the executed volume from one side results in the dislocation of 20 bps (basis points) in the closing price. In addition, the authors highlight that most of the price discovery process occurs within the first two minutes of the auction, with the rest of the period having no major influence on price discovery.

Asset managers Bank (2020); Blackrock (2020) discuss the key role of closing auctions in day trading. They both highlight the increasing importance of the closing auction and attribute the upward trend of the share of closing volumes to the rise in passive investing. They argue that additional factors are responsible for this trend, such as the large liquidity at the close and the desire to avoid being exposed to high-frequency traders. The authors view the closing auction as a funnel of liquidity and accurate price discovery. They argue that fears around rising market-on-close activity are unlikely to materialize.

Comerton-Forde and Rindi (2022) use data from European primary exchanges with a focus on UK stocks to investigate the trading activity at the close. They document an increase in market-on-close activity and highlight that the rise in passive investing does not fully explain this trend. They report that the auction returns increase with activity, especially during index rebalancing days, and that these returns are permanent in liquid stocks. Aramian and Comerton-Forde (2023) show that despite the large panel of competitive alternative closing mechanisms, closing auctions of primary exchanges still capture the lion's share (more than 80% in 2023) in European equity markets. The authors show that alternative trading mechanisms attract a higher market share than usual in index rebalancing days (40%) and end-of-months (30%). Conversely, the market share of alternative trading mechanisms decreases significantly on volatile and less liquid days. They discuss the divergent perspectives on the impact of fragmentation at the close and raise the need for further research on the subject.

Besson and Fernandez (2021) leverage a dataset from Euronext to examine the price impact of order submission at the close. They argue that the lower cost of trading at the close is a key factor in explaining the increase of the closing volume share alongside its large liquidity. They report that price impact at the close is two to three times lower than its continuous counterpart and use linear functions to fit the price impact at the close. In addition, they report several statistical regularities such as the price distribution of limit orders, order types, and temporal pattern of the indicative price and volume. Lastly, the authors discuss the potential adverse consequences of the internalization of market orders, which is the leading alternative closing mechanism, on the auction's volatility.

1.3.2 Auctions and market design

Reboredo (2012) study the effect of volatility auctions. These are five-minute auctions triggered during continuous trading whenever the price abnormally moves beyond predetermined thresholds. Volatility auctions are intended to have an overall positive impact on market quality in comparison with trading halts. The author uses intraday data from the Madrid stock exchange on a period that includes the switch from trading halts to volatility auctions. Reboredo (2012) finds that although the trading volume and intensity peak around these auctions, they systematically return to pre-event levels. Additionally, the author reports a decrease in post-events volatility in contrast with trading halts events where the post-events volatility remains high.

Since the mid-2000s, high-frequency trading firms have engaged in a technological arms race, i.e., large investments in their technology platforms to minimize the travel time of their signals (latency). The main goal of these firms is to be the fastest to exploit arbitrage opportunities in the markets. Amid the race, markets witnessed anomalies that are very likely to be attributed to high-frequency trading. A striking example is the NYSE flash crash on May 6th, 2010. This event triggered an ongoing debate among the research community on whether high-frequency trading is beneficial or harmful. Budish et al. (2015) argue that the arms race is the result of the flawed design of continuous limit order books and the sequential nature of order processing on the stock exchange. They employ a simple model that reproduces their empirical observations. The authors advocate for the use of frequent batch auctions, which is also motivated by their model. They find that a lower bound of 100 ms on the batch period totally eliminates the arms race problem.

Building on the frequent batch auction design, Paul et al. (2021) use a stochastic order book model to compute the optimal batch duration. They assume the case of naive market takers and that of rational ones looking to minimize their transaction costs. The optimal auction duration is derived from the minimization of a quadratic tracking error between the clearing prices and an ad-hoc fair price. The model's parameters are related to observable quantities for 77 stocks on Euronext markets. The authors find that continuous limit order books are usually sub-optimal, with optimal auction durations from 2 to 10 minutes. Derchu et al. (2020) use a similar framework to compare continuous limit order books, frequent batch auctions, and the AHEAD mechanism (Ad-Hoc Electronic Auction Design). The latter is a market design where agents trade at a fixed price and can launch auction phases whenever they are unsatisfied with the current price. Once the auction clears, trading resumes at the (fixed) auction price, and so on. The authors argue that the AHEAD mechanism significantly improves the market microstructure from the investors' viewpoint.

CBOE (2020) implemented a periodic auction book for European equities in 2015 with a randomized auction duration that can not exceed 100ms. Besson et al. (2019) find that there is essentially only one market order involved in each auction batch. Thus, a small batch period does not allow buyers and sellers to net their market orders. Aquilina et al. (2022) review the recent developments in the HFT arms race. They find that arbitrage opportunity races are very frequent, extremely fast, and account for more than 20% of the total traded volume. They argue that latency arbitrage is tantamount to a half-basis point tax on trading and that five billion dollars a year are at stake in global markets equity alone. Budish et al. (2023) propose a comprehensive market design combining frequent batch auctions with the possibility of constructing a portfolio using single order submission.

Donier and Bouchaud (2016) argue that the inter-auction time should be large enough for liquidity to accumulate. However, as the inter-auction time grows, waiting costs increase, which may result in fragmentation in favor of secondary continuous markets.

1.3.3 Auction modeling

The seminal work of Mendelson (1982) is one of the earliest in auction modeling. The author models a clearing house with buyers and sellers submitting orders of unit size on a discrete price grid. When buyers and sellers have the same intensity, the distribution of the exchanged volume is derived along with its first moments. When there is an imbalance between buyers and sellers, weaker results are presented. Another seminal work is that of Madhavan (1992) comparing continuous trading with batch auctions. The author uses a rational expectation framework to show that periodic auctions offer greater price efficiency, albeit agents must sacrifice trading continuity and bear informational and waiting costs. In light of these findings, Madhavan (1992) suggests a switch from trading halts to volatility auctions when the continuous trading mechanism fails, i.e., when large movements of the price occur.

Muni Toke (2015b) shows that the call auction problem is analytically tractable even in cases where Mendelson (1982) claims such forms did not exist. The author derives the exact distribution of the auction volume and the lower and upper clearing prices. Weak limits of these distributions are shown to be asymptotically normal and independent of the distribution of limit prices. The model is extended to allow for cancellations. Derksen et al. (2020) presents a similar model, allowing for possibly different distributions of limit order placement. The authors present the asymptotic limit of the distribution of the clearing price and study the influence of the order flow parameters. Building on the previous model, Derksen et al. (2022a) theoretically relate the

tail exponent to the order flow parameters and examine the empirical tails of the closing auction return distribution. They argue that limit orders are likely the main cause of the observed heavy tails in the return distribution. Empirical fits on the empirical distribution tails are conducted to support the theoretical results.

Lastly, Donier and Bouchaud (2016) introduce a continuous (time and price) limit order book model incorporating mechanisms for order deposition, cancellation, and price updates without involving transactions. Doing so, they recover the Walrasian auction mechanism and extend the framework to allow for successive batch auctions (see Section 1.5.1).

1.4 Price impact

Price impact, sometimes coined as market impact, refers to the fact that a buy pressure pushes the price up whereas a sell pressure pushes it down. In continuous double auctions, every transaction is an agreement between a buyer and a seller to exchange a given quantity at an agreed price. Thus, it is not obvious how a buyer or a seller will impact the price. The key to this paradox is the heterogeneity in the agents' opinions. In particular, a transaction can be viewed as a disagreement: the buyer can only hope that his freshly acquired asset will appreciate in the future. If the seller had a similar opinion, he would not have sold in the first place.

Milgrom and Stokey (1982) prove the "no-trade theorem", stating that under some constraints (efficient market structure, no noise trading or irrational actions, and accessible public information), no trader can profit from a piece of private information. The underlying idea behind the proof is the key difference between patient and impatient agents. The initiator of a transaction reveals his buying/selling intention to other agents. If there is no noise and all agents are rational, no one should accept that trading proposition. In practice, patient agents adjust their reservation price once they digest the information surrounding impatient agents'. This is particularly well illustrated by the following dialogue taken from Bouchaud et al. (2018a) (translated from Laumonier (2014))

Buyer: How much is it?

Seller: £1.50.

Buyer: OK, I'll take it.

Seller: It's £1.60.

Buyer: What? You just said £1.50.

Seller: That was before I knew you wanted it.

Buyer: You cannot do that!

Seller: It's my stuff.

Buyer: But I need a hundred of those!

Seller: A hundred? It's £1.70 a piece.

Buyer: This is insane!

Seller: It's the law of supply and demand, buddy. Do you want it or not?

1.4.1 Price impact in open markets

In open markets, price impact is an important source of transaction costs for agents with sizable portfolios. To minimize the price impact of large orders, agents often resort to optimal execution schemes (Almgren and Chriss, 2001; Alfonsi et al., 2010; Obizhaeva and Wang, 2013; Abi Jaber and Neuman, 2022; Hey et al., 2023). Studying comprehensively price impact is important for regulators to enhance effective policies that promote market stability.

Price impact of single market orders

Bouchaud et al. (2018a) measure an unconditional price impact of individual trades. It is the correlation between the side of an incoming market order and the subsequent price change.

$$R = \langle \varepsilon_t \cdot (m_{t+1} - m_t) \rangle, \tag{1.1}$$

where m_t is the mid-price at time t, $\varepsilon_t = +1$ for a buyer-initiated transaction and -1 for a seller-initiated transaction, and $\langle \cdot \rangle$ denotes the average over time. The time t increases by one at the arrival of every trade. R is found to be strictly positive with high statistical significance for 20 US stocks using one year (2015) worth of data. Actually, R is the linear regression slope estimate of $(m_{t+1} - m_t)$ on ε_t given zero-intercept. Using 2017 data for TotalEnergies, we also find that $R = 0.47$ bps is strictly positive with high statistical significance. This proves that buy market(able) orders are, on average, followed by an upward price move and vice versa for sell orders.

Conditioning on the size of the traded volume, we can define the price impact of a single

market order of volume v (until the arrival of the next market order)

$$R^1(v) = \langle \varepsilon_t \cdot (m_{t+1} - m_t) | V_t = v \rangle \,. \tag{1.2}$$

R^1 is the lag-1 response function, where V_t is the transaction size at time t. Empirically, R^1 is found to be weakly dependent on v and is often described as a concave function with exponent $\simeq 0.3$ (Lillo et al., 2003; Lillo and Farmer, 2004) and sometimes logarithmic (Potters and Bouchaud, 2003; Bouchaud et al., 2003). Actually, this strong sub-linearity results from selective liquidity taking (Bouchaud et al., 2009). This refers to agents usually considering the volume at the opposite best to size their orders as they often do not consume more than one price level. Farmer et al. (2004) argue that large price fluctuations are not driven by the size of market orders but rather by the distribution of gaps in the order book. They highlight the important role of liquidity fluctuation in defining price impact (Lillo and Farmer, 2005).

General lag response functions can be defined as

$$R(l, v) = \langle \varepsilon_t \cdot (m_{t+l} - m_t) | V_t = v \rangle, \tag{1.3}$$

where the average is now taken over time. Bouchaud et al. (2003) find that the general response function is separable $R(l, v) \approx R(l) \times R^1(v)$, where $R(l)$ is an unconditioned response function that measures the impact of an initiated trade at time t, l lags further.

Another relevant measure of price impact relates to volume imbalance $\Delta V(t, t+T) = \sum_{t<t'<t+T} \varepsilon_{t'} V_{t'}$ over a given time horizon T. An aggregate impact over a time horizon T is defined as

$$\mathrm{R}(\Delta V, T) = \mathbb{E}[m_{t+T} - m_t | \quad \Delta V(t, t+T) = \Delta V], \tag{1.4}$$

with the expectation taken over periods of length T that can span from minutes to a whole trading day. When T is the average inter-arrival times of market orders such that $\Delta V(t, t+T)$ contains on average one market order, R and R^1 coincide. Similarly, R is found to be a concave function of ΔV when T is small while becoming linear as T increases (Plerou et al., 2002).

Virtual price impact

If the order book is visible, one may want to measure price impact from the shape of the order book. Averaging impact values from the observed volumes and quotes would yield a virtual impact function. Weber and Rosenow (2005) calculate a virtual impact function on order book data and find that it is four times stronger than the actual impact function $\mathrm{R}(\Delta V, T)$. Bouchaud et al. (2009) conduct similar calculations and find that the virtual impact significantly differs from the actual price impact for actual orders. By definition, price impact is the average price move of actual orders conditioned on their volumes, whereas virtual impact is the average order book depth at a given volume. Until the arrival of the next order that triggers a transaction, the market reacts and digests some of the recently sent market orders. Additionally, agents can adjust quotes, cancel orders, and submit new limit orders between the arrivals of two market orders. The virtual impact does not allow for the measurement of such interactions and market reactions.

The long memory of market orders and propagators

Lastly, an accurate measure of price impact in continuous markets relates to the concept of meta-orders. These refer to large orders that are executed incrementally, as the order book liquidity does not permit a one-shot execution without resulting in a large price change. Agents slice and dice their large-size orders into several child orders that are executed over a given time horizon (typically one trading day). Meta-orders are responsible for the long memory of order flow (Lillo and Farmer, 2004): It refers to the slow decay of the autocorrelation function of market order signs (ε_t). Using agents' labeled data, Toth et al. (2015) break down the autocorrelation function of (ε_t) into two contributions: splitting (slice and dice) and herding. They find that the persistence of order flow is overwhelmingly due to splitting rather than herding.

The predictability in market order signs may push one to question the price predictability. For example, if each trade moves the current mid-price permanently with impact R such that

$$r_t = R\,\varepsilon_t + \xi_t, \tag{1.5}$$

where ε_t is the sign of the incoming market order at t and ξ_t an idiosyncratic noise, then

$$\mathbb{E}[r_t\, r_{t+l}] = (R)^2 \mathbb{E}[\varepsilon_t\, \varepsilon_{t+l}]. \tag{1.6}$$

Since $C(l) = \mathbb{E}[\varepsilon_t\, \varepsilon_{t+l}]$ slowly decays as power law of parameter γ, Eq. (1.6) is clearly in-

compatible with the unpredictability of returns. A way around (Bouchaud et al., 2003, 2006, 2009, 2018a) is to consider that impact is not permanent and decays with time with a kernel G. Then, the mid-price is $m_t = m_0 + \sum_{n<t} G(t-n)\varepsilon_n + \sum_{n<t} \xi_t$ and the return is $r_t = G(1)\varepsilon_1 + \sum_{n<t}(G(t-n+1) - G(t-n)) \cdot \varepsilon_n + \xi_t$. The first term is the immediate impact of the incoming market order at time t and the second one describes the decay of previous market orders. The propagator G can be calibrated on empirical data using response functions. Interestingly, the absence of correlation in returns implies that G should decay as a power law of parameter $\beta = (1-\gamma)/2$. This suggests that open markets operate near criticality.

Another way around (Lillo-Mike-Farmer model) is to relate the exponent of the distribution of metaorders with γ (Lillo et al., 2005; Sato and Kanazawa, 2023).

Price impact of metaorders

In continuous markets, the price impact of meta-orders is consistently found to be of square root type w.r.t. the order size; independently of the market, the asset, the considered period, and the execution style (Almgren et al., 2005; Tóth et al., 2011; Donier and Bonart, 2015; Zarinelli et al., 2015; Mastromatteo et al., 2014; Tóth et al., 2016). The impact of a meta-order of size Q is

$$I(Q) = Y\,\sigma_d \sqrt{\frac{Q}{V_d}}, \tag{1.7}$$

where σ_d is the daily volatility, V_d is the daily traded volume, and Y is a constant of order unity. The square root law in open markets is very likely a universal phenomenon. However, it only holds when Q/V_d is small, up to a few percent. Furthermore, it is only valid when the execution horizon is larger than the typical time needed for the limit order book to refill its liquidity (order of minutes). Otherwise, if execution is fast, the available liquidity is rapidly consumed, and price impact turns out to be convex. Lastly, the execution horizon should be smaller than the typical memory time of the latent (hidden) liquidity (order of days).

Price impact of order book events

Other studies investigated the price impact of order book events and not only marketable orders. Said et al. (2017) conduct a large-scale investigation on the market impact of (meta) limit orders, distinguishing between aggressive (in this context, marketable) and passive limit orders. The authors confirm the square root law on their proprietary dataset with lesser impact for passive limit orders. Cont et al. (2014) introduce the order flow imbalance that measures

the changes in supply and demand at the best bid/ask, thereby including market orders, limit orders, and cancellations. They conduct linear regression analysis on a fixed time scale (10 seconds) and find a positive linear correlation between returns and the order flow imbalance. The authors argue that order flow imbalance has more predictive power than volume imbalance. They suggest that the square root law may be a post-aggregation statistical artifact related to the trade's duration along with Capponi and Cont (2019). However, Bucci et al. (2019b) indicate that "price impact should not be misconstrued as volatility", i.e., the average directional price move conditioned on an individual trading decision (price impact of a meta-orders) has nothing to do with the average standard deviation of the price over that trade's duration (price diffusion and volatility).

Eisler et al. (2012) extend the propagator model to account for limit orders and cancellations. Whereas events happening on the same side of the order book (bid or ask) are long-range correlated, signed events of the same side are short-ranged: these observations corroborate price unpredictability even when market order signs are predictable. Although similar to the Hasbrouck VAR model (Hasbrouck, 1991), the generalized propagator model is microscopically interpretable in contrast with the ad-hoc econometric assumptions of the former.

1.4.2 Price impact in auctions

Research on auction impact is scarce. Donier and Bouchaud (2016) show that under sufficient regularity constraints of the supply and demand curves, price impact is linear in Walrasian auctions. Besson and Lasnier (2022) use linear price impact models to best fit empirical data from Euronext closing auctions. They find a smaller instantaneous impact for later submissions, a larger impact decay for earlier submissions, with a mitigated overall impact when distinguishing early and late submissions. Derksen et al. (2020) assert a concave market impact for market orders, treating them as a liquidity surplus that shifts the clearing price distribution.

Auctions differ in a number of ways from open markets. The auction time frame can be very short. For instance, five minutes at the closing auction is all it takes to exchange considerable amounts of volume. As the clearing approaches, the perspective of slicing a large order is narrower. Additionally, a large amount of liquidity builds up toward the end of the auction; the order book becomes resilient to large submissions, and the typical price change decreases (Challet, 2019), favoring one-shot executions. Importantly, when supply and demand overlap, there is no bid/ask. In continuous markets, trades must account for a half-spread cost when the benchmark is at the mid-price. In auctions instead, there is no half-spread to be paid. More

surprisingly, large market orders may result in zero impact on the indicative price at the auction time (see Chapter 3).

Although response functions of the indicative price yield the shape of market impact during auctions, they can not characterize the actual cost of trading. The transaction price of a given submission is not an impacted indicative price but rather the final auction price (assuming that the corresponding order is matchable at auction time). Therefore, the actual cost of trading for an order that is sent during the accumulation period stems from the impact on the auction price. Response functions of the auction price measured by Challet and Gourianov (2018) are conditioned on the time of the submission/cancellation and unconditional on the order size. These are consistent with measuring the cost of trading during equity auctions.

If the auction clearing time is known with certainty, some (low-latency) agent can —in theory— act last in the auction and send an order just before the clearing. In this case, the market does not have enough time to react as it clears right away. Transaction costs originate fully from the instantaneous impact caused by this agent's submission, who can then infer the incurred cost by calculating a virtual/instantaneous impact function from the order book. Thus, an apparent benefit of the clearing time randomization is to prevent fast agents from having an edge over the other participants, i.e., act last knowing what every other agent has submitted.

One way to examine auction impact is to investigate the statistical regularities of price impact at the auction clearing time, i.e., measuring the order book response at auction time. Although closely related to the concept of virtual impact, price impact at auction time is a faithful measure of the actual costs incurred by agents who act very late in the accumulation period. In parallel, calculating 1-lag response functions of the indicative price will provide us with useful information about the shape of market impact in equity auctions. We develop these contributions on price impact in equity auctions in Chapter 3.

1.5 Latent liquidity models

To describe numerous stylized facts in financial markets, various models have been developed ranging from simple phenomenological models to complex micro-founded ones. In contrast with standard economic models, where a representative rational agent maximizes a utility function, zero-intelligence agent-based models assume that agents do not learn. The resulting dynamics are analyzed from a statistical mechanics standpoint. Despite their simple hypotheses, zero-intelligence models are powerful predictive tools (Farmer et al., 2005; Smith et al., 2003).

In continuous markets, the available liquidity in the whole order book at a given time is, on

average, 10^{-5} of the market capitalization, and the average daily volume rarely exceeds 0.5% of the market capitalization (Tóth et al., 2011). Therefore, visible liquidity in LOBs is the result of a dynamic process, and most liquidity latent, i.e., not revealed yet. Starting from the skewed bell shape of the average order book, Bouchaud et al. (2018a) argues that there is no reason for the true supply and demand of a given asset to be sharply localized around the current price. Thus, the revealed liquidity at a given time in the order book can only be "the tip of the iceberg". By introducing the latent order book concept, Tóth et al. (2011) succeed in explaining the shape of market impact with minimal ingredients, i.e., using purely statistical considerations without resorting to fair price implications or expected utility maximization. The main hypothesis is the existence of a latent limit order book containing all buy and sell intentions that are partially revealed in the observable limit order book.

1.5.1 Walras' auctioneer

Donier and Bouchaud (2016) build on the latent order book model and provide a fully consistent framework for the free evolution of the marginal supply and demand curves ρ_S and ρ_B defined as

$$\begin{aligned} \rho_S(p,t) &= \partial_p S(p,t); \\ \rho_B(p,t) &= \partial_p D(p,t), \end{aligned} \tag{1.8}$$

where S and D are the supply and demand curves. Under the Walrasian mechanism, i.e., no transaction occurs before the auction final time, the free evolution of marginal supply and demand curves is governed by the following partial differential equations

$$\begin{aligned} \partial_t \rho_S(x,t) &= D\partial_x \rho_S(x,t) - \nu_S(x)\rho_S(x,t) + \lambda_S(x); \\ \partial_t \rho_B(x,t) &= D\partial_x \rho_B(x,t) - \nu_B(x)\rho_B(x,t) + \lambda_B(x), \end{aligned} \tag{1.9}$$

where $x = p - p_t$, is the centered log price p around the current log price p_t. The first term on the right-hand side (RHS) of (1.9) represents price updates with diffusion coefficient D. Originally, the PDEs are derived with respect to p and contain a drift term $-S_t\partial_p\rho(p,t)$. The underlying idea is that between t and $t+\delta t$, each agent updates its reservation price p to $p+\beta \mathrm{d}p_t + \eta$, where η is an idiosyncratic noise with zero mean and variance Σ^2, and β is a prefactor encoding agents' reaction to price movements. Agents over-react when $\beta > 1$ and under-react when $\beta < 1$. When derived in the reference frame of the current price p_t with $dp_t = dS_t$, the drift term disappears. When the price is deterministic, there is no additional contribution to the diffusion coefficient that consists of idiosyncratic price updates $D = \Sigma^2/2$. When p_t is a Brownian motion with

volatility σ, an additional contribution to the diffusion coefficient comes from agents' reaction to price movements $\mathrm{Var}(\beta)\,\sigma^2/2$. Hence

$$D = (\Sigma^2 + \mathrm{Var}(\beta)\,\sigma^2)/2. \tag{1.10}$$

If $\beta = 1$ for every agent, this volatility term disappears as well. The second term of the RHS of (1.9) corresponds to cancellations with rate ν; usually, cancellations are symmetric $\nu_S = \nu_B$. The third term of the RHS of (1.9) refers to the deposition of new orders with rate λ. A constant cancellation rate $\nu(x) = \nu$ allows the retrieval of a closed-form formula for the buy and sell latent order books

$$\rho(x,t) = \frac{1}{\sqrt{4\pi D t}} \cdot \int_{\mathbb{R}} \rho(y,0) e^{-\frac{(y-x)^2}{4Dt} - \nu t} \mathrm{d}y + \int_0^t \frac{1}{\sqrt{4\pi D(t-t')}} \int_{\mathbb{R}} \lambda(y) e^{-\frac{(x-y)^2}{4D(t-t')} - \nu(t-t')} \mathrm{d}y \mathrm{d}t. \tag{1.11}$$

Note that although market orders (corresponding to $|x| \to \infty$), buy limit orders for $x > 0$, and sell limit orders for $x < 0$ directly affect the equilibrium point of the supply and demand curves, this effect is not accounted for in the model.

Donier and Bouchaud (2016) allow for successive auctions with inter-auction time τ. When the inter-auction time is large $\tau \to \infty$, they evidence a linear price impact proportional to the sum of the marginal supply and demand around the auction price. When the inter-auction time is very short $\tau \to 0$, they show that the order book becomes locally linear around the current price, which leads to a square root impact. The theory's prediction is confirmed using Bitcoin data from 2015: the authors document an average order book that is locally linear and vanishing around the current price.

In open markets, the model is slightly modified to account for the transaction mechanism by adding a reaction term $-\kappa\rho_B\rho_S$ to both equations of system (1.9) and setting $\kappa \to +\infty$. In this case, marginal supply and demand do not overlap (Donier et al., 2015). Introducing a buy meta-order with a trading rate m_t adds a terms $m_t\delta(x - x_t)$ to the sell equations where $x_t = p_t - \hat{p}_t$ is the impacted price and $\hat{p}_t$ is the un-impacted price assuming the absence of the meta-order. Donier et al. (2015) study the behavior of the impact price x_t and find square root impact regimes in the infinite memory limit (very large timescales for the cancellation and the deposition rates, $\lambda, \nu \to 0$, with $\lambda/\sqrt{D\nu}$ constant). In addition, this model is proven to be arbitrage-free and consistent with diffusive prices. Benzaquen and Bouchaud (2018a) extends the study on the impacted price x_t when the order book's memory is finite ($\lambda, \nu \neq 0$) and show that the square root law is valid in multiple execution regimes by including market participants

with heterogeneous timescales. Importantly, the multi-time scale framework reconciles the latent liquidity model with the diffusivity puzzle, i.e., the slow decay of the autocorrelation of market order signs using a fine-tuning argument analogously to propagator models (Bouchaud et al., 2003).

1.5.2 From latent liquidity to revealed liquidity

Latent liquidity models are the subject of extensive research. Starting from a discrete stochastic order book model (Cont et al., 2010), Gao and Deng (2018) show that the average order book converges in the hydrodynamic limit (zero tick size, infinite rates of order arrivals, zero-order size relative to queue size, continuous-time), to a deterministic shape and verifies the partial differential equations of Equation (1.9) without the diffusion term. Lemhadri (2019) adds a mean-reverting agent behavior and provides existence results for the impacted price x_t.

Recently, Dall'Amico et al. (2019) introduced a conversion mechanism between the latent and the revealed order books during the continuous trading phase and fitted stationary solutions of their model to order book snapshots convincingly. In summary, orders materialize in the revealed order book $\rho^{(r)}$ coming from the latent order book $\rho^{(l)}$ with a rate $\nu_r\Gamma_r$ and disappear from the revealed order book all the way to the latent order book with a rate $\nu_l\Gamma_l$. See Fig. 1.6 for an illustration. In the infinite memory limit, the sell side equations read

$$\begin{aligned}\partial_t\rho_S^{(r)} &= D_r\partial_{xx}\rho_S^{(r)} + \nu_r\Gamma_r(x)\rho_S^{(l)} - \nu_l\Gamma_l(x)\rho_S^{(r)} - \kappa\rho_S^{(r)}\rho_B^{(r)};\\ \partial_t\rho_S^{(l)} &= D_l\partial_{xx}\rho_S^{(l)} - \nu_r\Gamma_r(x)\rho_S^{(l)} + \nu_l\Gamma_l(x)\rho_S^{(r)},\end{aligned} \tag{1.12}$$

and the buy side equation equations read

$$\begin{aligned}\partial_t\rho_B^{(r)} &= D_r\partial_{xx}\rho_B^{(r)} + \nu_r\Gamma_r(-x)\rho_B^{(l)} - \nu_l\Gamma_l(-x)\rho_B^{(r)} - \kappa\rho_B^{(r)}\rho_S^{(r)};\\ \partial_t\rho_B^{(l)} &= D_l\partial_{xx}\rho_B^{(l)} - \nu_r\Gamma_r(-x)\rho_B^{(l)} + \nu_l\Gamma_l(-x)\rho_B^{(r)},\end{aligned} \tag{1.13}$$

where ν_l, ν_r are constant rates and Γ_r, Γ_l are probability functions for revealing and unrevealing orders. Dall'Amico et al. (2019) take $\nu_r = \nu_l$ and $\Gamma_r = 1 - \Gamma_l \in [0, 1]$. Using $\phi = \rho_S^{(r)} - \rho_B^{(r)}$, closed form stationary solutions are obtained, for instance, when $D_r = 0$ and $D_r = D_l \neq 0$. When $D_r \neq D_l$, numerical schemes are used to obtain the revealed order books and these are calibrated to empirical order book data.

The stationary latent books are found to be linear with a decreasing exponential correction. The stationary revealed order books are found to have a skewed bell shape around the current price. When $D_r = 0$, the limit volumes at the current price are found to be strictly positive,

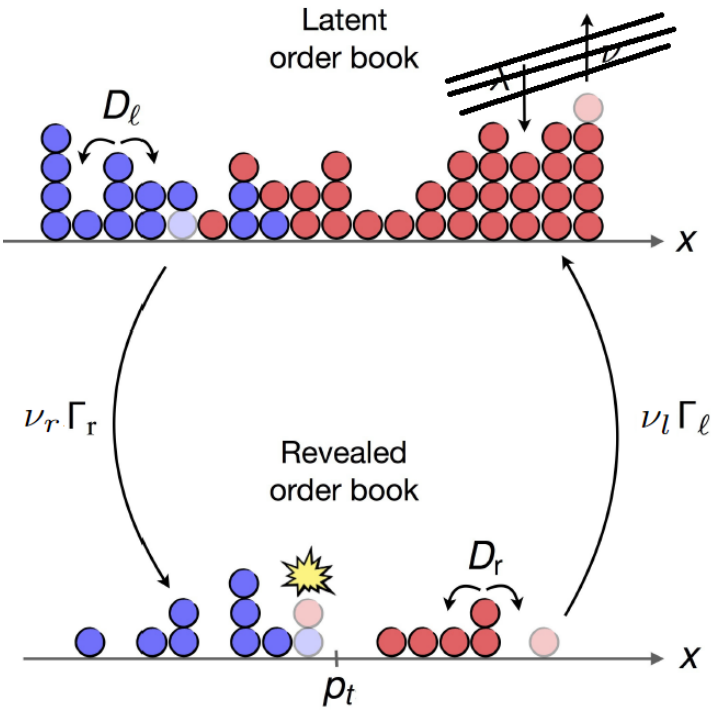

Figure 1.6 – Sketch of the conversion mechanism between the latent and the revealed order books. Adapted from Dall'Amico et al. (2019).

and ϕ experiences a discontinuity at $x = 0$. When $D_r \neq 0$, ϕ becomes continuous at $x = 0$. Importantly, empirical order book data can be calibrated using the latent/revealed framework, thereby providing an estimate of the involved latent liquidity. Dall'Amico et al. (2019), use the calibrated parameters of their model to sort the studied assets on a stability map. Calibrating assets during periods of high volatility yields indeed positions in the map that are near critical regions. This may be a useful tool for regulators who monitor market stability.

Due to the considerable variability in the temporal dynamics of open markets, only stationary solutions are calibrated to successive snapshots of a given asset in a given period. Auctions, however, offer a suitable avenue for deriving and calibrating time-dependent solutions thanks to their distinctive temporal patterns. A good starting point would be to measure the time and price dependencies of depositions, cancellations, and diffusions (price updates) during the accumulation period. This is precisely the adopted approach in Chapter 4.

1.6 Objectives, outline, and main findings

1.6.1 Order book reconstruction

To investigate the statistical regularities of equity auctions and their dynamics, this thesis leverages high-quality data from the European high-frequency database (BEDOFIH) with microsecond resolution. Naturally, a substantial part of the research work involved handling and mining this database. An upfront effort of pre-processing and reconstruction of snapshots as well as tick-by-tick events was needed.

In Chapter 2, we provide a detailed description of the data, the pre-processing, and the reconstruction of snapshots and tick-by-tick events. Moreover, we provide a comprehensive overview of the closing auction ecology, i.e., the types of agents participating in the closing auction. In addition, we map the contribution of each category to the total events of the accumulation period as well as the contribution to the closing volume. Importantly, we dissect the temporal patterns of the auction events (submissions, cancellations, and updates), first by looking through the lens of market orders and then that of limit orders.

1.6.2 Price impact in equity auctions

In Chapter 3, we formalize mathematically the relevant quantities in equity auctions. We show that market orders can result in zero impact on the indicative price up to a large fraction of the auction volume.

Indeed, if the size of a buy market order is smaller than $V_S^R(p_a) + V_B^M(p_a)$, the indicative price does not change. Here, $V_S^R(p_a)$ is the sell limit volume that remains at the auction price after the clearing, and $V_B^M(p_a)$ is the matched buy limit volume that was posted at the auction price. We prove a similar assumption for the sell side. Essentially, for a market order to move the price in a given direction, its size should not only be larger than the remaining limit volume at the auction price of the opposite side (i.e., the opposite side's imbalance) but to the sum of that imbalance and the matched limit volume at the auction price of the market order side.

The intuition behind this result is that for a market order to be executed at the clearing, it should not only deplete the available liquidity on the opposite side (imbalance) but also gain priority over matched orders with the lowest priority — specifically, matched limit orders at the auction price. Fig. 1.7 provides an illustration of how the indicative price first changes under the influence of market orders.

In open markets, a market order can also have a zero instantaneous impact on the mid-price provided it does not deplete the opposite best. However, the cost of trading remains strictly positive, amounting to a half-spread (when benchmarked at mid-price). This cost of immediacy represents the market maker's gain. Importantly, even though the market order's immediate impact is zero, the average impact until the next market order is strictly positive (albeit lower than that of a market order that crosses the opposite best). In fact, the market is continuously reacting to incoming market orders through cancellations of opposite-side limit orders, price updates, and limit order submissions.

Taking into account these remarks, we investigate the (instantaneous) price impact at auction time. The reason for this is twofold. First, if we assume that an agent acts last during the accumulation period, he can exactly infer the impact on the auction price by looking at the order book. In this case, impact costs stem fully from this instantaneous impact. Most primary exchanges recently implemented a clearing time randomization, in part to avoid price manipulation by low-latency actors. We address the effect of this randomization as well by comparing the impact at auction time with that at the last deterministic time of the auction.

Second, we argue that agents have a strong incentive to reveal their intentions at auction time, especially around the indicative price and particularly at the end of the trading day, hence

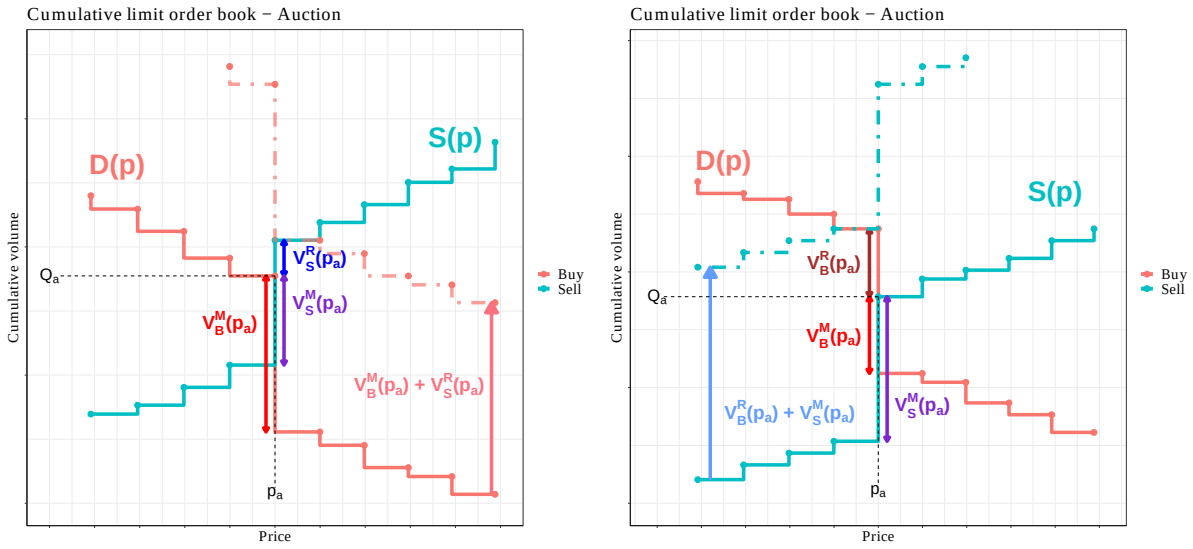

Figure 1.7 – Cumulative buy (red curves) and sell curves (blue curves) during hypothetical auctions. Left panel: the buy volume is totally matched at the auction price; right panel: the sell volume at the auction price is totally matched at the auction price. Dash-dotted lines: effect of an addition buy market order (left plot) and sell market order (right plot): the auction price can change only when the market order is larger than the matched volume plus the imbalance, which explains why zero impact is prevalent.

the focus on the closing auction. This allows us to relate the average limit order book at auction time with the latent order book that determines the shape of market impact.

Scaling the zero impact volume introduced above by the auction volume Q_a, we define

$$\begin{aligned} \omega_B^{(0)} &= \frac{V_S^R(p_a) + V_B^M(p_a)}{Q_a}; \\ \omega_S^{(0)} &= \frac{V_S^M(p_a) + V_B^R(p_a)}{Q_a}. \end{aligned} \tag{1.14}$$

We show that $\omega_B^{(0)}$ and $\omega_S^{(0)}$ can be simultaneously large. This means that sending a large market order either to buy or to sell will result in zero impact on the indicative price. Additionally, we provide the daily distribution of the zero-impact volumes and compare them with subsequent incremental volumes that are needed to change the tick price further. We find that sending a market order of size $1\%\, Q_a$ (which is already a large volume compared with the typical size of a market order during day trading) at auction time and not impacting the price occurs in 46% to 74% of the days depending on the studied stock.

We show that zero impact volumes have bursty dynamics where they can be very large at a given time and be totally depleted in the next few events. Conversely, they can be null at a given time and abnormally high right after. In the latter case, large limit volumes are posted at the auction price and result in price pinning.

Removing the zero impact part, we show that price impact at auction time is linear whenever the sum of buy and sell orders sum up to a constant in the direction of the indicative price

$$I(\omega^{(i)}) - I(\omega^{(0)}) = \frac{1}{p^{(1)}\widetilde{\mathscr{L}}}\left(\omega^{(i)} - \omega^{(0)}\right), \tag{1.15}$$

where $\widetilde{\mathscr{L}} = \tilde{\rho}_S(x) + \tilde{\rho}_B(x)$ is a constant, $\tilde{\rho}_S$, $\tilde{\rho}_B$ are the scaled order densities. We show that the $\tilde{\rho}_S$ and $\tilde{\rho}_B$ sum up to constant day by day with variable price range Δ. We determine this interval using a change detection criterion and show that the estimated slope is an accurate proxy for price impact at auction time. We emphasize that the investigated price impact at auction time is examined on a day-to-day basis rather than being averaged across days.

In addition, we investigate the influence of derivatives expiry days on the resilience of the order book and find that the impact slope is often lower than during other days of the week or the month. This means that the auction order book is prepared to absorb larger shocks during these days. We examine the evolution of the liquidity $p_a Q_a \tilde{\mathscr{L}}$, which links with the inverse of the impact slope during the accumulation period. We find that it has a distinctive temporal

pattern, first linear, then concave, and finally convex w.r.t. time as the auction time approaches.

After the introduction of random clearing times, it is no longer possible for (low-latency) agents to act last during the auction. Nevertheless, we show that acting the latest deterministic time of the accumulation period T, i.e., right before 09:00:00 for the opening auction and 17:35:00 for the closing auction, yields a similar outcome as acting at auction time. In fact, the order book stabilizes after T with a 77% reduction in the indicative price volatility in comparison with the last five seconds of the deterministic closing auction time. In addition, the impact slope is stable during this random time window: its changes in absolute value are capped by 12% in more than 90% of the days. Lastly, as pointed out earlier, the dynamics of zero impact are bursty and can drastically vary in a very short interval. Therefore, agents can act conservatively at 17:35:00, and their impact can be decomposed into a stable and predictable linear component plus a random zero impact bonus that can decrease the incurred cost.

Finally, we compute 1-lag response functions of the indicative price for aggressive orders (market orders, buy limit orders with limit price larger than the indicative price, and sell limit orders with limit price smaller than the indicative price) and find that it is linear as well. This result proves that the nature of price impact is the same during the accumulation time and at the auction time, i.e., mostly mechanical, and contrasts with results for open markets, where selective liquidity taking causes very different shapes between the average virtual impact (using the instantaneous shape of the book) and market impact of actual trades.

1.6.3 Modeling the auction dynamics

In Chapter 4, we adapt the latent/revealed liquidity framework to equity auctions. Our starting point is the coupled partial differential equations of Dall'Amico et al. (2019). In that sense, our model is very similar to Eqs. (1.12), (1.13). For the sell side, they read

$$\begin{cases} \partial_t \rho_S^{(r)} &= D_r \partial_{xx} \rho_S^{(r)} + (\nu_r \Gamma_r)(x,t) \rho_S^{(l)} - (\nu_l \Gamma_l)(x,t) \rho_S^{(r)}; \\ \partial_t \rho_S^{(l)} &= D_l \partial_{xx} \rho_S^{(l)} - (\nu_r \Gamma_r)(x,t) \rho_S^{(l)} + (\nu_l \Gamma_l)(x,t) \rho_S^{(r)}. \end{cases} \tag{1.16}$$

Eqs (1.16) are complemented with the following boundary conditions

$$\begin{cases} \partial_x \rho_S^{(l)} \underset{x \to +\infty}{\longrightarrow} a > 0, \quad \rho_S^{(l)} \underset{x \to -\infty}{\longrightarrow} b \geq 0; \\ \rho_S^{(r)} \text{ does not diverge when } |x| \to +\infty. \end{cases} \tag{1.17}$$

The main innovation is that the reveal and unrevealed rates $\nu_r \Gamma_r$, $\nu_l \Gamma_l$ depend on both price

and time without a guarantee of variable separation at first. As latent exogenous depositions and cancellations are omitted, revelations are equivalent to submissions and unrevelations are equivalent to cancellations. For symmetry reasons, we focus on the sell side and omit the subscript S from $\rho^{(l)}, \rho^{(r)}$. When the revealed diffusion coefficient is zero, we obtain a simple stationary solution

$$\begin{aligned} \rho^{(r)}(x) &= \frac{\nu_r}{\nu_l} \frac{\Gamma_r(x)}{\Gamma_l(x)} \rho^{(l)}(x); \\ \rho^{(l)}(x) &= \max(ax + b, b). \end{aligned}$$

By assuming the simplest price functions for the reveal and unreveal rates, i.e., a constant $\Gamma_l = 1$ and an exponentially decreasing $\Gamma_r \propto e^{-|x|}$, we are able to fit the average order book at auction time. This yields estimates of the involved latent liquidity (a, b).

Finding closed-form general solutions of Eqs (1.16) is challenging. However, taking $D_r = D_l = 0$ allows tractable calculations. In addition, it is later shown that diffusion has a lesser influence on the order book than submissions and cancellations during the accumulation period. In this new setup, Eqs (1.16) can be decoupled, and each price level x can be treated independently

$$\begin{aligned} \partial_t \rho^{(r)} + (\nu_r \Gamma_r + \nu_l \Gamma_l)\ \rho^{(r)} &= \nu_r \Gamma_r \rho^{\Sigma}; \\ \partial_t \rho^{(l)} + (\nu_r \Gamma_r + \nu_l \Gamma_l)\ \rho^{(l)} &= \nu_l \Gamma_l \rho^{\Sigma}; \\ \rho^{\Sigma}(x) &= \max(ax + b, b). \end{aligned} \tag{1.18}$$

Simplifying further by assuming stationary rates $(\nu\Gamma)(x,t) = \nu \times \Gamma(x)$, we obtain the following general solution

$$\begin{aligned} \rho^{(r)}(x,t) &= \rho_\infty \left[1 - e^{-[\nu_r \Gamma_r(x) + \nu_l \Gamma_l(x)]\, t}\right]; \\ \rho_\infty(x) &= \frac{\nu_r\, \Gamma_r(x)\, \rho^{\Sigma}(x)}{\nu_r\, \Gamma_r(x) + \nu_l\, \Gamma_l(x)}. \end{aligned} \tag{1.19}$$

The revealed order book of Eq. (1.19) converges to ρ_∞ in the long run $t \to +\infty$. This solution fails to reproduce the accelerating auction dynamics around the indicative price as the auction deadline approaches $t \to T$. Introducing time-dependent rates, we assume that they evolve as the inverse of the remaining time to the clearing starting from a certain time threshold $t^{(0)}$

$$(\nu\Gamma)(x,t) = \frac{C}{\gamma + T - t}\, \Gamma(x), \tag{1.20}$$

where γ is an offset to the deadline. When $\gamma_l = \gamma_r$ and $t_l^{(0)} = t_r^{(0)}$, we obtain a closed-form

formula for the revealed order book

$$\begin{aligned}\rho^{(r)}(x,t) &= \rho_T - (\rho_T - \rho_0)\left(\frac{\gamma + T - t}{\gamma + T - t^{(0)}}\right)^{C_r\Gamma_r(x)+C_l\Gamma_l(x)};\\ \rho_T(x) &= \frac{C_r\,\Gamma_r(x)\,\rho^{\Sigma}(x)}{C_r\,\Gamma_r(x) + C_l\,\Gamma_l(x)},\end{aligned} \tag{1.21}$$

and ρ_0 is the initial condition at $t^{(0)}$. The revealed order book of Eq. (1.21) is convex w.r.t. time whenever the exponent $C_r\Gamma_r(x) + C_l\Gamma_l(x)$ is below 1. Empirically, the fitted exponents are not always below one. However, we find $\gamma_l > \gamma_r$, which yields a convex time behavior when $t \to T$. In addition, when the cancellation rate is constant, and the submission rate evolves as the inverse of the remaining time to the deadline, we obtain convex time solutions as the auction time approaches.

Before solving Eqs (1.16) numerically in the general case, we provide thorough measurements of the submission and cancellation rates as well as the revealed diffusion coefficient using tick-by-tick data. We find that the submission rate is exponentially decreasing $\propto e^{-|x|/x_r}$ for positive prices $x \geq 0$ with an apparently constant price scale x_r in time. For prices below a negative price threshold $x < -x_0$, submissions become noisy and constant. For $-x_c < x < 0$, submissions are exponentially decreasing with a decreasing price scale w.r.t. time. Additionally, at the beginning of the auction, we find that submissions are constant w.r.t. time (price-wise), then increasing as the auction time approaches. We verify that submissions satisfy the scaling relationship $\nu_r\Gamma_r \propto (\gamma_r + T - t)^{\alpha}$ with an optimal exponent $\alpha = 1$ around the indicative price. We show that submissions have the same magnitude and shape whether for high frequency agents (HFTs) or non high frequency agents (non-HFTs) with a smaller price scale for HFTs.

We conduct a similar analysis on the cancellation rate, showing a complex overall price dependence. The time dependence of cancellation is first decreasing (price-wise), then increasing only for limit prices around the indicative price. We verify that cancellations satisfy the scaling relationship $\nu_l\Gamma_l \propto (\gamma_l + T - t)^{\alpha}$ with an optimal exponent $\alpha = 1$ in a narrow region around the indicative price. Importantly, we show that cancellations are mostly driven by HFTs, when the cancellation rate for non-HFTs is significantly smaller, noisy and may be assumed constant to a good approximation.

We break down the revealed diffusion coefficient into two contributions. The first contribution is linked with the indicative price volatility. We estimate the latter to be of order 10^{-8}. The second contribution represents idiosyncratic price updates. We find that it is negligible compared with the volatility contribution. In addition, we highlight a time-dependent pattern of the indicative price volatility $\sigma \propto t^{-1/2}$. Thus, $D_r \propto 1/t$.

To solve the coupled PDEs, we take $\nu_l \Gamma_l = \nu$ to be constant and model the submission rate as the sum of two exponential terms with different price scales. We allow for the fast exponential term (smaller price scale) to evolve as the inverse of the remaining time to the clearing. We study three diffusion settings: zero diffusion $D_r = D_l = 0$, constant diffusion where we consider (D_r, D_l) as free parameters, and time diffusion where we add an additional free parameter by considering $D_r \propto 1/t$. The calibration to auction data shows that the model is able to reproduce the full price and time dynamics of the average order book density during the accumulation period. The obtained order of magnitude of D_r is similar to that estimated previously. We find that diffusion regularizes the discontinuities of the first derivative of the revealed order book while improving the fit accuracy, particularly around the indicative price.

In a separate section, we study the reasons that are responsible for the sub-diffusivity of the indicative price, i.e., $\langle \log(p_{t+\tau}/p_t)^2 \rangle_t \sim \tau^{2H}$, with $H < 1/2$. We conduct a similar analysis to Chen et al. (2017) that show that H can be decomposed as $H = J + L + M - 1$. Each of the exponents J, L, M is respectively associated with the failure of one condition in the central limit theorem. We divide the closing auctions into separate regimes and fit average exponents H, J, L, M on 1-second auction data for five active stocks. We find $H \approx 0.3$ (sub-diffusivity), $L \approx 0.7$ (heavy-tailed increments), $M \approx 0$ (time decreasing increments), $J \approx 1/2$ (absence of long-term memory). The latter indicates the indicative price process is efficient (unpredictable) in the sense of Chen et al. (2017).

Chapter 2

Analyzing a high-quality dataset

Figure 2.1 – BEDOFIH's Logo.

Contents of this Chapter

2.1 Dataset description

BEDOFIH is a large financial database managed by the EUROFIDAI Institute. In addition to the high-frequency database, BEDOFIH also comprises a refined daily database and an ESG (Environmental, Social, and Governance) database. The high-frequency database contains detailed market data (Orders, trades, references, flags) with a millisecond or microsecond accuracy, depending on the market. It covers a wide range of instruments, including stocks, bonds, structured products, and exchange-traded funds (ETFs), among others, across four main exchanges: three primary exchanges (LSE, Xetra, and Euronext), and a secondary exchange, BATS Chi-X. For this thesis, we extracted data from the 34 most active stocks on Euronext Paris under the regulation of the AMF (Autorité de Marché Financiers) between 2013 and 2017. In this case, the timestamp resolution is one microsecond.

For each stock and each trading day, information is available in four files:

— a history orders file that contains all the orders that remained in the central limit order book from the previous trading day;

— a current orders file that contains all submissions, modifications, and cancellations for the current trading day;

— a trades' file that lists all the transactions that took place during the current trading day;

— an events file that lists special market events, if any, such as a delayed opening, a halt in trading, etc.

Each order is uniquely determined by its order ID. When an order is first submitted, a new line is inserted in the current orders file with a characteristic ID equalling one. The characteristic ID tracks the changes of the order during its lifetime and increases by one each time the order is modified; a new line in the orders file is inserted then. In addition to the order ID and its characteristic ID, precise times of the order submission and its time of validity in the central limit order book are displayed. Interestingly, if an order is, a posteriori, modified, we know about the time of modification in the corresponding line. We also know, ex-ante, the time when the order is released from the order book, i.e., if it is canceled or totally executed. Similarly, we have at our disposal the state variable of an order (partially/totally filled, canceled, expired, rejected by the trading system, etc), its validity (Good for the day, good till cancel, valid for auction, valid for closing, good until a specified date, fill or kill, etc), and its type (market, limit, stop order, etc).

For every order, detailed information is available, including the side (buy/sell), the limit price (zero for market orders), the initial submitted quantity, the price threshold if it is a stop order, the remaining quantity of an order—which can be useful when there are partial executions between modifications. More importantly, two flags are of interest: the HFT flag corresponding to the AMF's classification of market participants (pure HFTs, MIXED HFTs, and NON-HFTs) [1], and the user account flag (Own account, client account, parent company account, market maker, retail market organization, or retail liquidity provider). An agent can operate using multiple user accounts. Meanwhile, its HFT flag remains the same.

1. "A participant is considered a high-frequency trader (HFT) if he meets one of the two following conditions:

— The average lifetime of its canceled orders is less than the average lifetime of all orders in the book, and it has canceled at least 100,000 orders during the year.

— The participant must have canceled at least 500,000 orders with a lifetime of fewer than 0.1 seconds, and the top percentile of the lifetime of its canceled orders must be less than 500 microseconds.

An investment bank meeting one of these conditions is described as mixed-HFT (MIX). If a participant does not meet any of the above conditions, it is a non-HFT (NON)."(AMF, 2017), p. 33.

The trades file contains information about every transaction during the trading day. Each transaction comprises the buyer and seller order IDs in synchronization with the orders file, the transaction price, the exchanged quantity, and the transaction time. A unique order ID may generate one or multiple transactions. In the latter case, the generated transactions are processed sequentially in time with a resolution of a few microseconds. This is also the case for the generated transactions of opening and closing auctions. A session flag allows us to distinguish the opening auction transactions from those of the rest of the session. For the continuous trading session, each transaction comprises a flag indicating whether it is buyer or seller initiated. More on the files' description can be found in EUROFIDAI (2020).

We extract 2 to 5 years' worth of data for each stock, amounting to $N = 34,977$ stock days.

2.2 Reconstructing order book snapshots

In order to reconstruct the exact state of the limit order book at any point in time, including at auction time, we combine the information from the four different files for each stock and each trading day to create a snapshot. The idea is to keep track of the last update for orders whose validity time is below the snapshot time and release time is above the snapshot time. After combining the history file with the current file, we encountered the "partial execution problem". Some history orders may have been partially executed in the past. Without a modification from the order sender, no update regarding the remaining quantity is available in the order files. Therefore, one should look at all previous trade files to see if they involve partial executions of the history orders. When constructing snapshots for the closing auction, one should take into account the possible partial executions that occur during the day.

Very few of these stock-days result in errors or mismatches (e.g., dataset errors, non-crossing supply and demand for the opening auction, or half-day trading/halted trading before 17:30 for the closing auction). After removing these invalid snapshots, we are left with $N_o = 34,971$ valid snapshots at the opening auction time and $N_c = 34,820$ valid snapshots at the closing auction time.

Using these reconstructed snapshots just before the auction time, we compute reconstructed prices and volumes as per Euronext rules, i.e., by maximizing the exchanged volume and minimizing the imbalance. This boils down to finding the intersection of the reconstructed supply and demand curves. Table 2.1 reports the percentage of snapshots for which the reconstructed price (resp. volume) matches the actual auction price (resp. volume) among valid snapshots. The actual auction price and volume are easily accessible in the trades' file.

Table 2.1 – Percentages of auction snapshots with accurate reconstruction.

	Opening auction	Closing auction
Number of valid snapshots	34,971	34,820
% snapshots matching the auction price	99.6%	99.9%
% snapshots matching the auction volume	99.0%	99.7%
% snapshots matching both	98.9%	99.6%

The remaining discrepancies are likely a result of using simplified rules to account for stop orders. For these few unmatched snapshots, we note that the discrepancies between computed and actual quantities are small: less than 1 basis point on the absolute average difference from the auction price and 0.2% on the absolute average distance from the auction volume. These few unmatched auctions are discarded from the sample in the subsequent analysis, though they would not alter the outcome of our experiments.

Stop orders that are sent during the accumulation period are not taken into account in the calculation of the executable volume at the clearing, even if they are triggerable during that phase. However, they may generate additional execution right after the auction clearing. On the contrary, stop orders that were sent before the start of the accumulation period are taken into account in the calculation of the auction quantities if the threshold price was crossed prior to the start of the auction phase. Thus, for each stop order, one should investigate the whole price history (since the submission of the corresponding stop order) to assess if it is active before the start of an accumulation period.

Note that there are many other order types that might used in the accumulation period such as pegged orders and iceberg orders. Although these does not represent a significant fraction of order types used in the auction, they might however induce small discrepancies in the matching of the indicative volume. For iceberg orders, the whole quantity of the order is used to calculate the auction price and volume.

By computing 1-second snapshots, we build an auction replayer that allows us to visualize the full dynamics for a given auction. Figure 2.2 provides an overview of the built dashboard paused at four moments: 17:32:00 for the top left panel, 17:33:00 for the top right panel, 17:34:55 for the bottom left, and 17:35:29 at the auction closing time for the bottom right panel. We can track, for instance, the evolution of the cumulative limit order book as shown in the top left panel of a given snapshot of the dashboard. The intersection of the cumulative curves (supply and demand) points to the indicative price and volume, that are shown using black dashed curves. The upper right panel of the replayer depicts the usual limit order book. The

bottom left panel shows the evolution of the total amount of market orders (buy, sell, and their difference). Finally, the evolution of the indicative price is shown in the bottom right panel. The replayer can be fully customized to account for other or new quantities of interest (e.g., the indicative volume, the matched volume of HFTs, etc.) and provide us with useful insights if careful enough when watching. For instance, we can notice that a large sell limit order has been inserted during that closing auction between 17:34:55 and 17:35:29 when inspecting the cumulative LOB evolution. This order provides a large resistance against upward price moves while being mostly unmatchable. This may suggest strategic price pinning.

2.3 Reconstructing tick-by-tick events

To have a comprehensive overview of the auction dynamics, we reconstructed tick-by-tick events during the accumulation period in order to monitor every change in price limits (queues), indicative prices, and volumes. The preliminary idea is to compute a first snapshot, 07:15:00 for the opening auction (history orders) and 17:30:00 for the closing auction. Next, we can derive the differences in quantity by order ID for each order that is present during the accumulation period. Then, the total volume at a limit price is obtained by proceeding with a cumulative sum of the previous quantity differences for each limit price.

However, when an order is canceled, no new line is present in the data to announce it. Thus, the corresponding cancellation events are created, concatenated, then sorted chronologically with the rest of the events. Importantly, we amend the corresponding limit queues whenever there is an update in an order volume, its price, or both. Once a sequential description of the events and the queues is obtained, we are able to compute an indicative price and volume for each event.

We conducted the tick-by-tick reconstruction of the opening and closing auction for the five actively traded stocks on Euronext Paris (BNP Paribas, LVMH, Sanofi, Société Générale, and TotalEnergies) between 2013 and 2017. The obtained accuracy of the final reconstructed indicative price and volumes is larger than 97%, similar to the results in Table 2.1.

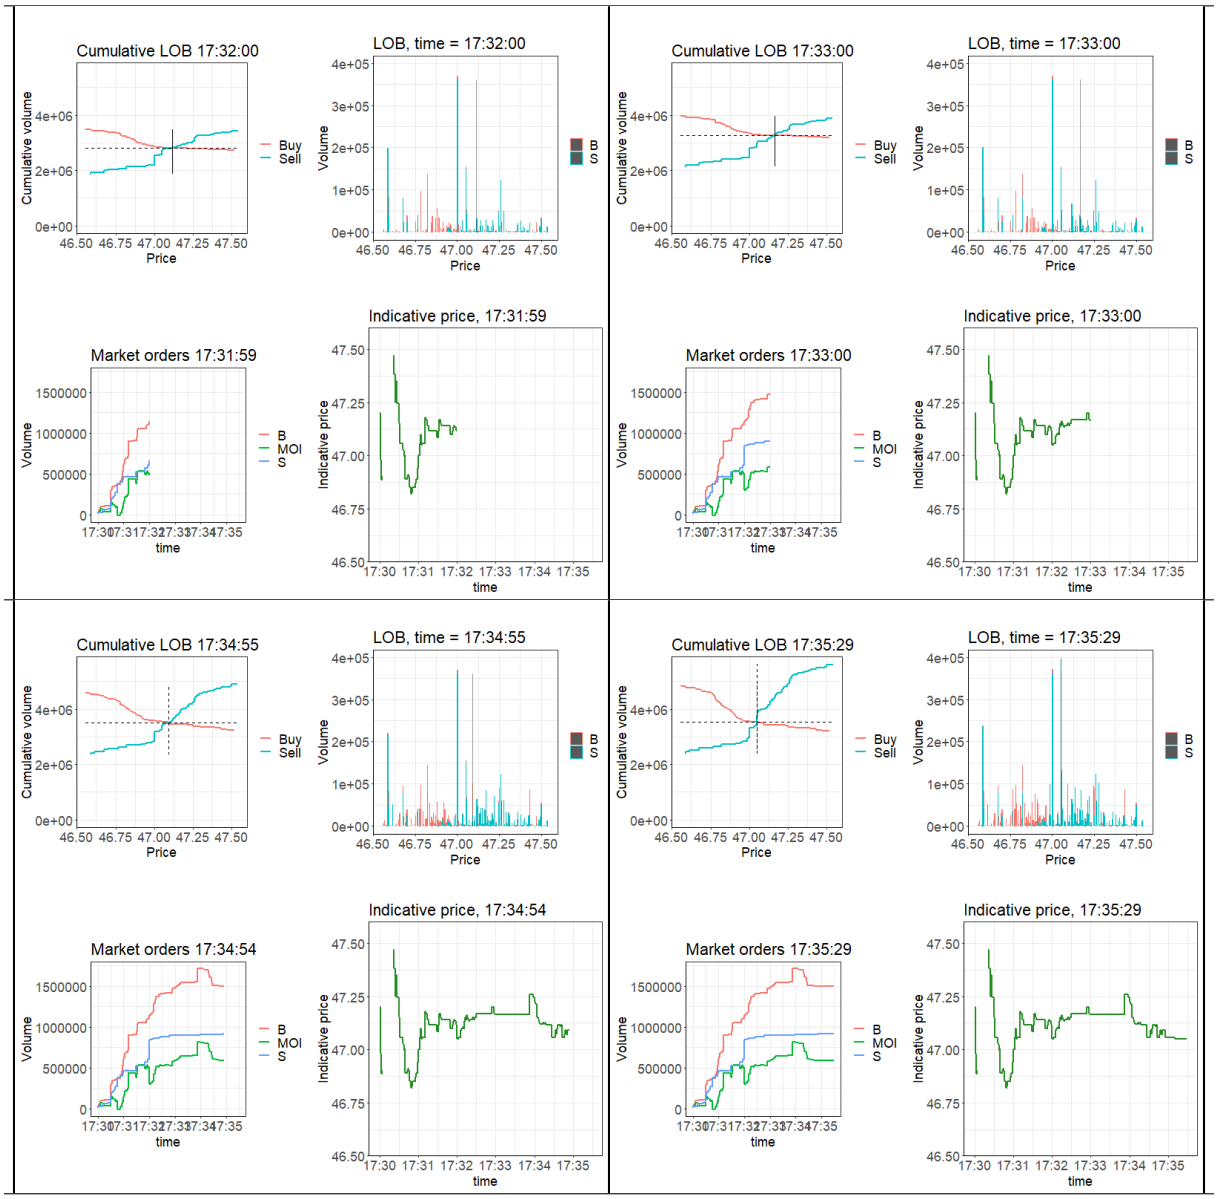

Figure 2.2 – Snapshots of the *auction replayer* dashboard for the closing auction of the TotalEnergies stock on February 28th, 2017.

2.4 Closing auction ecology

Using the reconstructed tick-by-tick events, we provide an overview of the closing auction ecology for the TotalEnergies stock between 2013 and 2017.

2.4.1 Dissecting the contribution of agent categories

We report in Table 2.2 the average participation in the closing auction volume broken down by agent type and latency. MIX agents (investment banks) hold the lion's share in the closing auction volume (80% of the closing volume on average), followed by slow agents (15.5%). The market share of HFTs in the closing volume is only 4.5% on average. However, HFTs engage in more than 40% of the price-changing events of the closing auction (see Table 2.3), more than 80% of all events, and more than 55% of all events using their market-making account. Thus, high-frequency market makers display a large activity during auctions when they hold a very low share of the closing volume (1%). In addition, market orders represent less than 3% of the auction events and less than 0.002% of HFT-related events.

Table 2.2 – In bold: average daily participation in the closing auction volume. In parenthesis: daily standard deviation. Buy and sell sides are nearly symmetrical. TotalEnergies 2013∼2017.

	Client account	Own account	Market maker	Parent company
HFT	**2.5** *(3.8)*	**1** *(1.4)*	**1** *(1.4)*	**0** *(-)*
MIX	**18** *(10.9)*	**50** *(12.1)*	**3** *(3.3)*	**9** *(6.3)*
NON	**7.5** *(6.2)*	**8** *(7.1)*	**0** *(0.4)*	**0** *(0.6)*

Table 2.3 – Participation in % of price changing events for a total of 560, 816 events. TotalEnergies 2013∼2017.

	Client account	Own account	Market maker	Parent company
HFT	0.65	14.98	24.87	0
MIX	6.83	37.35	3.12	3.27
NON	2.86	5.91	0.02	0.09

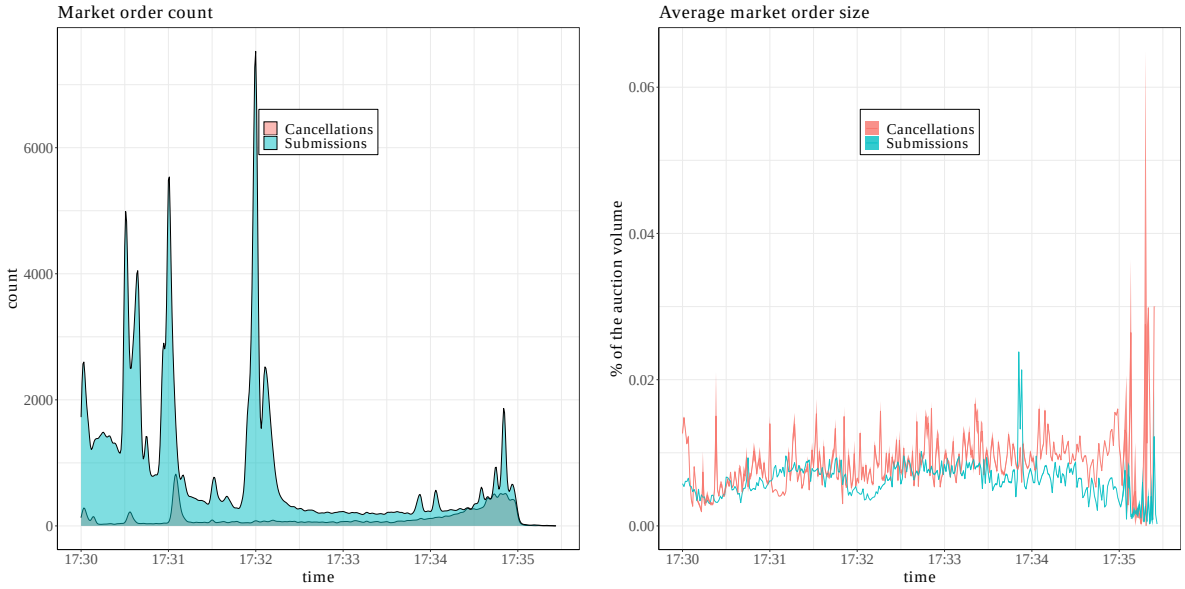

Figure 2.3 – Market orders during the closing auction. Left panel: market order count as a function of time. Right panel: average market order size using 1-second time bins. Buy and sell are aggregated as they are symmetrical. A breakdown by market order submission (blue) and cancellation (red) is shown.

2.4.2 An overview on order dynamics

Market orders

Market orders represent 312,785 out of 10,234,512 closing auction events on TotalEnergies between 2013 and 2017. Figure 2.3 shows that the market order count displays significant peaks at round minutes (17:31, 17:32, 17:34) and multiples of 30 seconds (17:30:00, 17:31:30). Some peaks are present right after round times and may indicate reactions to round time actions. Most market order activity takes place before 17:32:30. However, we observe an increase in activity as the auction time approaches. The average size of market order submission or cancellation is roughly constant throughout the auction.

Regarding the average dynamics of the indicative price and volume, Challet (2019); Raillon (2020) find that 80% of the final auction volume is matched in the first two minutes of the closing auction. Besson and Lasnier (2022) report that the indicative price overreacts in the first minute of the accumulation period before mean reverting in the last minute.

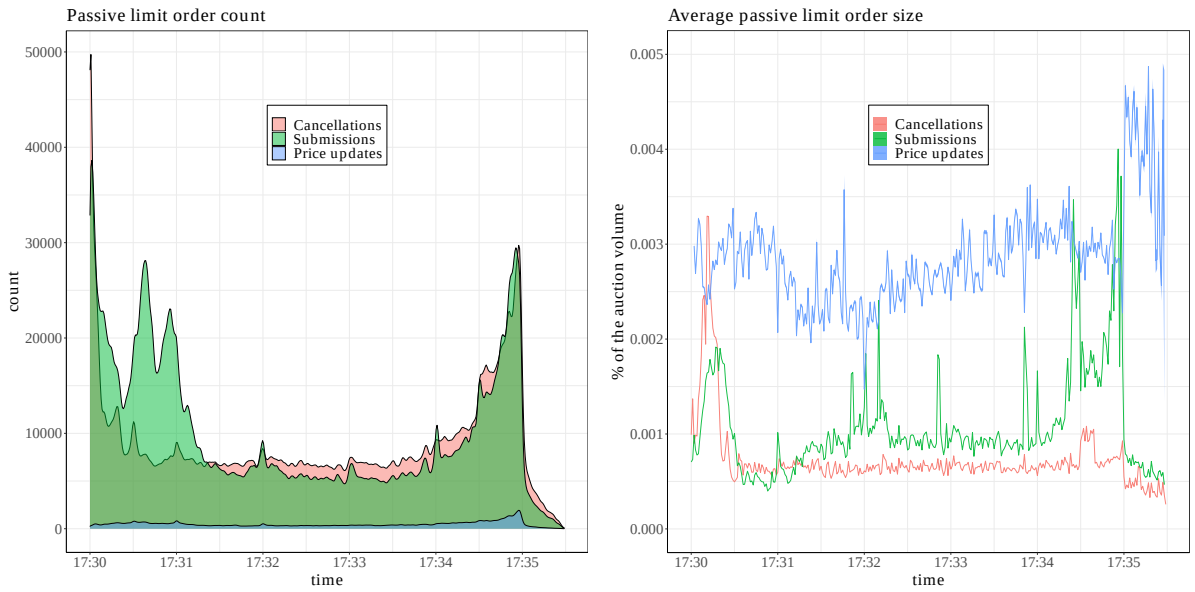

Figure 2.4 – Passive limit orders during the closing auction. Left panel: passive order count as a function of time. Right panel: average passive limit order size using 1-second time bins. Buy and sell are aggregated as they are symmetrical. A breakdown by market order submission (green), cancellation (red), and price update (blue) is shown.

Passive limit orders

Passive limit orders (buy limit orders whose limit price is strictly below the indicative price and sell limit orders with limit price strictly above the indicative price at the time they were submitted/canceled/updated) represent $6,548,086$ out of the total events. Figure 2.4 shows that the count of the passive limit orders displays an overall U shape and increases as the auction time approaches. The events count is moderately distributed across the two halves of the accumulation period. The average cancellation size is stable after the first thirty seconds of the accumulation period. However, the average submission rate displays large peaks and increases during the last minute of the auction. Price update events represent less than 2.5% of passive limit events.

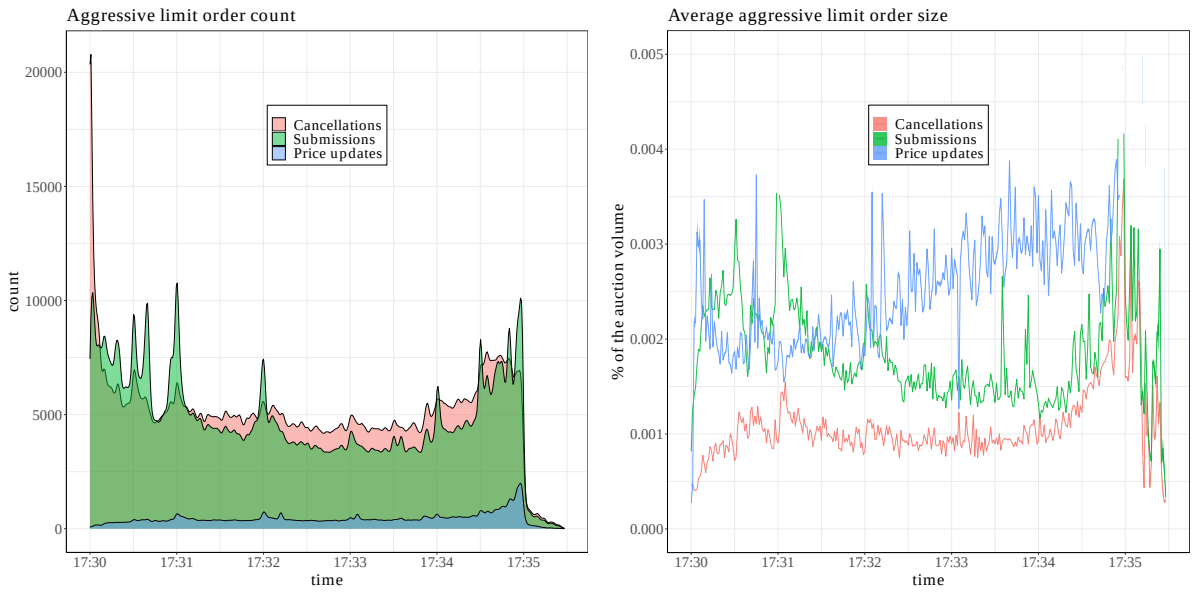

Figure 2.5 – Aggressive limit orders during the closing auction. Left panel: Aggressive order count as a function of time. Right panel: average aggressive limit order size using 1-second time bins. Buy and sell are aggregated as they are symmetrical. A breakdown by market order submission (green), cancellation (red), and price update (blue) is shown.

Aggressive limit orders

Aggressive limit orders (buy limit orders whose limit price is above the indicative price and sell limit orders with limit price below the indicative price at the time they were submitted/canceled/updated) represent $3,373,641$ out of the total events. Figure 2.5 shows that the count of the passive limit orders displays an overall U shape as well, although with a less pronounced increase than passive orders at the auction time. The average submission and cancellation rates are 50% to twice larger than those of passive limit orders. Both aggressive rates (submission and cancellation) increase and double during the last minute of the auction. Price update events represent less than 4.5% of aggressive limit events. Note that $3,054,425$ out of $3,373,641 + 312,785$ (Aggressive limit orders + market orders) do not induce a change in the indicative price. These zero-impact events have similar dynamics to the ones shown in Fig. 2.5 with lower baselines of the average order size.

To sum up, the comprehensive nature of the BEDOFIH dataset enables us to thoroughly examine the statistical regularities of the accumulation period phase while distinguishing the behavior of different market participant categories. Additionally, it provides us with the opportunity to verify the accuracy of our models with a high degree of precision.

Chapter 3

Price impact in equity auctions: zero, then linear

Based on Salek et al. (2023);
by Mohammed Salek, Damien Challet, and Ioane Muni Toke;
To appear in Market Microstructure and Liquidity;
Poster presentation in the market microstructure Workshop, CFM/Imperial College/HSBC, London, December 11-12, 2023;
Presented in the 17th Financial Risks International Forum, Paris, March 18-19, 2024.

Abstract

Using high-quality data, we report several statistical regularities of equity auctions in the Paris stock exchange. First, the average order book density is linear around the auction price at the time of auction clearing and has a large peak at the auction price. While the peak is due to slow traders, the order density shape is the result of subtle dynamics. The impact of a new market order or cancellation at the auction time can be decomposed into three parts as a function of the size of the additional order: (1) zero impact, caused by the discrete nature of prices, sometimes up to a surprisingly large additional volume relative to the auction volume (2) linear impact for additional orders up to a large fraction of the auction volume (3) for even larger orders price impact is non-linear, frequently super-linear.

Contents of this Chapter

3.1 Introduction

Most electronic markets rely on auctions to start and end trading days in an orderly way. Because the volume involved during auctions is larger than the liquidity available at a given time in a typical open-market limit order book, auctions reduce price impact and fluctuations. The share of the closing auction in the total exchanged volume has significantly increased over the years (Blackrock, 2020), especially in European markets (Raillon, 2020). This increase highlights the importance of the auction mechanism in the price formation process.

In contrast to the abundant literature about open-market dynamics, work on auctions is scarce. On the theoretical side, Muni Toke (2015b) derives the distribution of the exchanged volume and the auction price using a stochastic order flow model during a standard call auction. In the same vein, Derksen et al. (2020) propose a stochastic model for call auctions which produces a concave price impact function of market orders; in addition, Derksen et al. (2022a) build on the previous model to demonstrate the heavy-tailed nature of price and volume in closing auctions. Besides, Donier and Bouchaud (2016) show that under sufficient regularity conditions (continuous price and time) and using a first-order Taylor expansion of supply and demand curves, price impact in Walrasian auctions is linear in the vicinity of the auction price.

Empirically, Pagano and Schwartz (2003) find that introducing opening and closing call auctions improves market quality and lowers execution costs in the Paris stock exchange. Boussetta et al. (2017) add that although opening volumes are decreasing and the market is fragmenting, the opening auction still improves market quality on Euronext Paris. They also report that slow brokers submit orders early, whereas high-frequency traders tend to act moments before

the clearing. Challet and Gourianov (2018) analyze US equities data and compute the auction price response functions conditional on the addition, and cancellation of an order. In addition, Challet (2019) demonstrates that a strategic behavior of agents is needed to explain the antagonistic effects of activity acceleration and indicative price volatility decrease as the auction end approaches.

More recently, Jegadeesh and Wu (2022) assess the robustness of closing auctions by comparing the price impact between NASDAQ and NYSE exchanges and find that the cost of trading during closing auctions is generally smaller than during trading hours. They also find that closing auctions mainly attract uninformed and passive investors, while informed traders prefer to act during continuous market hours. In the same spirit, Besson and Fernandez (2021) analyze the closing auction in European markets and use a linear function to fit the impact of market orders; they report a smaller instantaneous impact for later submissions, and an overall cost of trading on close two to three times smaller than during trading hours.

Here, we characterize in detail the empirical properties of liquidity and price impact in equity auctions. At auction time, price impact is fully determined by the state of the order book, and we focus on the instantaneous impact caused by an order if sent just before the clearing. We do not find a straightforward linear impact: while adding or canceling a market (or marketable) order at the auction time has a linear component, the discreteness of the limit order book mechanically leads to zero price impact for small enough orders. These free-of-cost volumes can represent a fairly large fraction of the total matched volume. Before auction time, the order book shape yields a virtual/instantaneous price impact that can differ from that of actual submissions/cancellations. However, we find that the average impact of actual orders is of the same nature, i.e., linear.

This paper is organized as follows: first, we introduce a discrete-price auction mathematical framework (Section 3.2) suitable to derive the conditions under which price impact is zero or linear. Next, we present the high-quality data used in this work: a large dataset from the European high-frequency financial (BEDOFIH) database (Section 3.3). The main part consists in a detailed study of several statistical regularities of auctions, focusing on limit order book shapes and price impact during the auctions (Sections 3.4 and 3.5). Our main results are as follows:

1. the average limit order book of buy (sell) orders has a skewed bell shape whose maximum is attained below (above) the auction price. Both distributions roughly mirror each other and can be considered linear in the vicinity of the auction price;
2. there is an often large peak of volume at the auction price that builds up towards the end

of the auction;

3. breaking down the average limit order book densities by the agent latency (HFT, MIXED, NON) and their account type (own account, client account, market maker, parent company, retail market organization ...) makes it clear that each category has a different behavior; the peak is not due to HFTs but to slower traders, and some traders post buy and sell orders asymmetrically;
4. at any time during the auction, instantaneous price impact is zero for small enough volumes for both buy and sell orders simultaneously because of the discreteness of prices. The presence of a peak for both buy and sell limit order densities increases the importance of zero impact in auctions;
5. for large enough volumes, instantaneous price impact is linear for most of the days and not only on average. This holds when the sum of the buy and sell order densities is constant as a function of the price around the indicative/auction price, which happens on most days. Using a change point detection algorithm, we characterize the linear impact price region day by day and asset by asset at the auction time;
6. the average price impact of actual submissions/cancellations during the accumulation period is linear as well. This contrasts with open markets where the dependence of the average impact on the order size is much weaker. In some exchanges, limit order books are not disseminated during auctions and selective liquidity taking is not possible;
7. price impact at auction time is smaller during option expiry dates.

3.2 A mathematical framework for auctions

In Euronext markets, equity auctions start with an accumulation period and end with a clearing process. During the accumulation period, participants send their orders (quantity, price, side, order type, ...) to the exchange. Types of orders include market orders, limit orders, activated stop orders, and valid for auction orders. Modifications and cancellations are allowed, but transactions cannot occur. At any time during the accumulation process and at the end of the auction, the price that maximizes the matched volume and minimizes the imbalance is computed. At the auction time, buy (resp. sell) orders whose prices are larger (resp. smaller) than the auction price are executed, while limit orders whose price equals the auction price may be matched or remain in the order book after the auction.

Definition 1 (**Supply and demand**)**.** For an auction $\mathscr{A} = (a, d)$, where a is the auction type (open, close, ...) at date d, we define the available supply $S(p, t)$ and demand $D(p, t)$ at a price

p and time t as, dropping the (a, d) for the sake of clarity,

$$\begin{aligned} S(p,t) &= \sum_{p' \leq p} V_S(p', t), \\ D(p,t) &= \sum_{p' \geq p} V_B(p', t), \end{aligned} \tag{3.1}$$

where $V_S(p', t)$ (resp. $V_B(p', t)$) is the available sell volume (resp. buy volume) at a price p' and time t.

Limit orders can only be submitted on a discrete price grid. Therefore, at any time t, $p \mapsto S(p, t)$ is a non-decreasing right-continuous step function, and $p \mapsto D(p, t)$ is a non-increasing left-continuous step function.

Definition 2 (**Auction price and volume**)**.** For an auction $\mathscr{A} = (a, d)$, the auction volume Q_a^d noted Q_a is the one maximizing the exchanged quantity between buyers and sellers at the time of the clearing T_a^d noted T_a. For a given price p at time t, buyers and sellers can exchange a volume equal to $\min\{S(p, t), D(p, t)\}$ at most. Thus

$$Q_a = \max_p \min \left\{ S(p, T_a), D(p, T_a) \right\}.$$

The auction price p_a^d noted p_a is the price that maximizes the exchanged quantity. As it may not be unique, we have

$$p_a \in \left\{ p \mid Q_a = \min \left\{ S(p, T_a), D(p, T_a) \right\} \right\}.$$

In this work we will always assume that supply $S(p, T_a)$ and demand $D(p, T_a)$ intersect, so that Q_a always exists and is unique. Note however that p_a is often not uniquely defined by the maximization of the exchanged volume alone; this is why exchanges implement a complementary set of rules such that p_a is always well defined. In the case of the Euronext markets used in this work, when multiple prices maximize the exchanged volume, the chosen p_a is the one with the smallest imbalance. Then, if multiple prices with the highest executable volume and the smallest imbalance coexist, the auction price is the one closest to the reference price (last traded price).

Definition 3 (**Indicative price and volume**)**.** For an auction $\mathscr{A}$, the indicative price p_t^{ind} and the indicative volume Q_t^{ind} at time $t \leq T_a$ are the hypothetical auction price and the total matched volume if the clearing took place at time t.

Obviously, we have $p_a = p_{T_a}^{\text{ind}}$ and $Q_a = Q_{T_a}^{\text{ind}}$. From now on, the time notation will be omitted

when we work at time $t = T_a$ (e.g., $S(p)$ stands for $S(p, T_a)$). Note however that subsequent definitions and results can be stated for any time $t \leq T_a$ using time-dependent notations and substituting p_a with p_t^{ind} and Q_a with Q_t^{ind}.

Definition 4 (**Buy and sell densities**)**.** For an auction $\mathscr{A} = (a, d)$, we define the buy (resp. sell) density ρ_B^d (resp. ρ_S^d) at a price p as

$$\rho_\bullet^d(p) = \frac{V_\bullet(p)}{\delta p}, \quad \bullet \in \{B, S\}, \tag{3.2}$$

where δp is the difference between the price p and the next non-empty tick price when $\bullet = B$, and δp is the difference between p and the previous non-empty tick price when $\bullet = S$.

To define a meaningful average density over a large number of days, volumes can be scaled by the auction volume Q_a^d at day d, and prices can be substituted with log-price differences from the auction price $p \leftarrow \log(p/p_a)$.

Definition 5 (**Scaled buy and sell densities**)**.** For an auction $\mathscr{A} = (a, d)$, we define the scaled buy and sell densities as

$$\tilde{\rho}_\bullet^d(x) = \frac{\rho_\bullet^d(p_a e^x)}{Q_a^d}, \quad \bullet \in \{B, S\}, \tag{3.3}$$

where $x = \log\left(\frac{p}{p_a}\right)$. Furthermore, if we substitute δp by a constant δx, we can compute for a given stock the average scaled density as

$$\langle \tilde{\rho}_\bullet(x) \rangle = \left\langle \frac{V_\bullet(p_a e^x)}{Q_a^d \times \delta x} \right\rangle, \quad \bullet \in \{B, S\}, \tag{3.4}$$

where $\langle \cdot \rangle_d$ denotes the average across days of the computed quantity at time $t = T_a$.

Observe that this quantity is a discrete version of the continuous marginal supply and demand curves defined in Donier and Bouchaud (2016), where $\rho_B(p) = -\partial_p D$ and $\rho_S(p) = \partial_p S$.

Definition 6 (**Matched and remaining volumes**)**.** For an auction $\mathscr{A}$, we define $V_\bullet^M(p)$ as the matched (executed) volume at a price p and side $\bullet \in \{B, S\}$, and $V_\bullet^R(p)$ as the remaining (non-executed) volume at a price p and side $\bullet$. Hence, any limit volume $V_\bullet(p)$ at price p is the sum of the matched and remaining volumes

$$V_\bullet(p) = V_\bullet^M(p) + V_\bullet^R(p), \quad \bullet \in \{B, S\}. \tag{3.5}$$

Obviously, for any price $p > p_a$, all the buy volume is matched and all the sell volume remains. Thus $V_B^M(p) = V_B(p)$, $V_S^M(p) = 0$, $V_B^R(p) = 0$, and $V_S^R(p) = V_S(p)$. Symmetrically, for any, price $p < p_a$, we have $V_B^M(p) = 0$, $V_B^R(p) = V_B(p)$, $V_S^M(p) = V_S(p)$, and $V_S^R(p) = 0$. Consequently, $V_\bullet^M(p) \times V_\bullet^R(p)$ can be non-zero only if $p = p_a$.

Proposition 1. Let $\mathscr{A}$ be an auction with an auction price p_a and an auction volume Q_a. The following equalities stand:

(a) $Q_a = S(p_a) - V_S^R(p_a) = D(p_a) - V_B^R(p_a)$;

(b) $V_S^R(p_a) \times V_B^R(p_a) = 0$.

Proof. (a): as the auction volume Q_a is the sum of all matched volumes, we have

$$\begin{aligned}
Q_a &= \sum_p V_B^M(p) = \sum_p V_S^M(p), \\
&= V_B^M(p_a) + \sum_{p>p_a} V_B^M(p) = V_S^M(p_a) + \sum_{p<p_a} V_S^M(p_a), \\
&= V_B(p_a) - V_B^R(p_a) + \sum_{p>p_a} V_B(p) = V_S(p_a) - V_S^M(p_a) + \sum_{p<p_a} V_S(p_a), \\
&= D(p_a) - V_B^R(p_a) = S(p_a) - V_S^R(p_a).
\end{aligned}$$

(b) is proved by contradiction: if $V_S^R(p_a) \times V_B^R(p_a) \neq 0$, then $\left(V_S^R(p_a), V_B^R(p_a)\right) \neq (0,0)$. This implies that a residual volume $\delta V = \min\left(V_S^R(p_a), V_B^R(p_a)\right) > 0$ can be matched between buyers and sellers at the auction price and thus contradicts the fact that Q_a is maximizing the exchanged volume during the auction. □

Let us now introduce volumes scaled by the auction volume: given an integer volume of shares $q \in \mathbb{N}$, we define the scaled volume $\omega = q/Q_a$.

Definition 7 (**Price impact**)**.** For an auction $\mathscr{A}$, for any $\omega > 0$, we define the price impact before the auction clearing of a buy (resp. sell) market order $I_B(\omega)$ (resp. $I_S(\omega)$) as the absolute change in the auction log-price immediately after submitting a buy (resp. sell) market order of size $q = \omega \times Q_a$

$$I_\bullet(\omega) = \left|\log\left(\frac{p_\omega}{p_a}\right)\right|, \quad \bullet \in \{B, S\}, \tag{3.6}$$

where p_ω is the new auction price after injecting the market order.

Note that $I_\bullet$ refers to the instantaneous impact of an order submission at auction time $t = T_a$, i.e., assuming a market order is sent just before the clearing. In this case, the market can not react

to this submission as the clearing happens right away, and no relaxation can occur. However, if a submission/cancellation is sent to the exchange way before the clearing, the corresponding price impact $I_\bullet$ at $t < T_a$ with $p_a \leftarrow p_t^{\text{ind}}$ and $Q_a \leftarrow Q_t^{\text{ind}}$ refers to a virtual/instantaneous price impact that may differ from the price impact of an actual submission/cancellation since the market can still react to it.

Proposition 2. Let $\mathscr{A}$ be an auction with an auction price p_a and an auction volume Q_a. We inject a market order of size $q = \omega Q_a$ before the auction clearing. The new auction price is p_ω. We have:

(a) The function $I_\bullet : \omega \mapsto \left|\log\left(\frac{p_\omega}{p_a}\right)\right|$, for $\bullet \in \{B, S\}$ and $\omega > 0$, is a non-decreasing and right-continuous step function.

(b) Let $(\omega_B^{(i)})_{i\geq 0}$ be the ordered points of discontinuity of I_B. Then

$$\begin{aligned} \omega_B^{(0)} &= \frac{V_S^R(p_a) + V_B^M(p_a)}{Q_a}, \\ \omega_B^{(i)} &= \omega_B^{(i-1)} + \frac{V_S(p_B^{(i)}) + V_B(p_B^{(i)})}{Q_a} \quad , \quad i \geq 1, \end{aligned} \tag{3.7}$$

where $p_B^{(i)} > p_a$ is the i^{th} non-empty price tick strictly greater than the auction price.

(c) Let $(\omega_S^{(i)})_{i\geq 0}$ be the ordered points of discontinuity of I_S. Then

$$\begin{aligned} \omega_S^{(0)} &= \frac{V_S^M(p_a) + V_B^R(p_a)}{Q_a}, \\ \omega_S^{(i)} &= \omega_S^{(i-1)} + \frac{V_S(p_S^{(i)}) + V_B(p_S^{(i)})}{Q_a} \quad , \quad i \geq 1, \end{aligned} \tag{3.8}$$

where $p_S^{(i)} < p_a$ is the i^{th} non-empty price tick strictly lower than the auction price.

Obviously $I_\bullet(\omega_\bullet^{(i)}) = \left|\log(p_\bullet^{(i+1)}/p_a)\right|$. Also, remark that if all price ticks contain non null volume ($V_B + V_S > 0$), then $p_\bullet^{(i)} = p_a \pm i\theta$, where θ is the tick size. The proof of Proposition 2 is given in Appendix 3.A. Proposition 2 allows us to compute the impact function at any time of a given auction, including during the accumulation period. In addition, the price impact of a new order is zero if its size is smaller than $\omega_\bullet^{(0)} Q_a$. Figure 3.1 provides a graphical explanation of $\omega_\bullet^{(0)}$ formulas. On the left panel for example, the buy volume at the auction price is totally matched ($V_B(p_a) = V_B^M(p_a)$ and $V_B^R(p_a) = 0$). In this case, in order to shift the price, a buyer would need to execute a market buy order of minimal volume $V_S^R(p_a) + V_B(p_a)$. Alternatively, a seller would need to execute a market sell order of minimal volume $V_S^M(p_a)$. The right panel of Figure 3.1

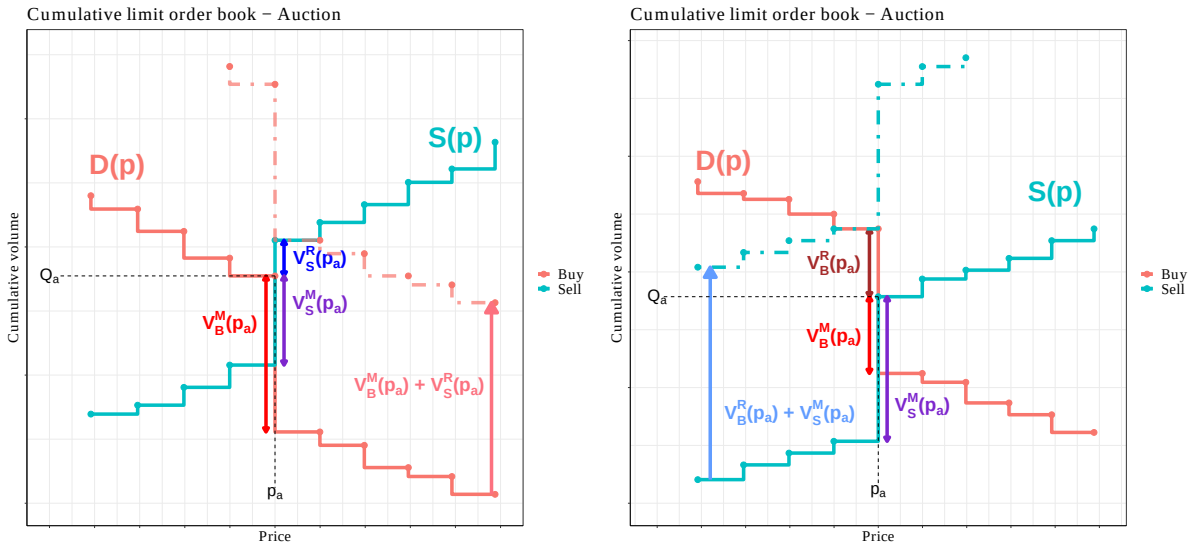

Figure 3.1 – Cumulative buy (red curves) and sell curves (blue curves) during hypothetical auctions. Left panel: the buy volume is totally matched at the auction price; right panel: the sell volume at the auction price is totally matched at the auction price. Dash-dotted lines: effect of an addition buy market order (left plot) and sell market order (right plot): the auction price can change only when the market order is larger than the matched volume plus the imbalance, which explains why zero impact is prevalent.

illustrates the symmetric case in which the sell volume at the auction price is totally matched ($V_S(p_a) = V_S^M(p_a)$ and $V_S^R(p_a) = 0$). Moreover, observe that if a trader sends a market order of exact size $q = \omega_\bullet^{(0)} \times Q_a \in \mathbb{N}$, then both p_a and $p_\bullet^{(1)}$ maximize the auction volume. As explained above, the new auction price would be the one with the smallest imbalance, i.e. equal remaining volumes. If p_a and $p_\bullet^{(1)}$ have equal imbalances, then the new auction price is the closest to the reference price. Here, we assumed that whenever $q = \omega_\bullet^{(0)} \times Q_a \in \mathbb{N}$, the price automatically shifts to $p_\bullet^{(1)}$.

Also, by Proposition 2, $V_S(p_\bullet^{(i)}) + V_B(p_\bullet^{(i)}) = Q_a \times (\omega_\bullet^{(i)} - \omega_\bullet^{(i-1)})$ for $i \geq 1$ is the necessary volume to take the price from $p_\bullet^{(i)}$ to $p_\bullet^{(i+1)}$. We therefore define $\delta\omega_\bullet^{(i)} = \omega_\bullet^{(i)} - \omega_\bullet^{(i-1)}$ for $i \geq 1$ to denote this scaled incremental volume, with the convention that $\delta\omega_\bullet^{(0)} = \omega_\bullet^{(0)}$. Finally, notice that a cancellation of a buy market order of size q affects the price in the same way as submitting a sell market order of the same size: in both cases the new price p_ω is a solution of $S(p_\omega) + q = D(p_\omega)$. Similarly, cancelling a sell market order has the same effect as submitting a buy market order. Consequently, we only focus on the price impact of market order submissions in the following.

3.3 Data

The dataset used in this work is part of the BEDOFIH database (Base Européenne de Données Financières à Haute-fréquence) built by the European Financial Data Institute (EUROFIDAI). The dataset provides detailed order data for all stocks traded on Euronext Paris between 2013 and 2017. For each stock and each trading day, information is provided in four files:

— a history orders file that contains all the orders that remained in the central limit order book from the previous trading day ;
— a current orders file that contains all submissions, modifications, and cancellations for the current trading day ;
— a trades file that lists all the transactions that took place during the current trading day ;
— an events file that lists special market events, if any, such as a delayed opening, a halt in trading, etc.

In addition to standard information such as time with microsecond precision, price, side (buy/sell), quantity, and price threshold for stop orders, we have access to additional order details in these files, some of which are computed ex-post. These include the order type and its temporal validity (market, limit, valid-for-auction, valid-for-closing, etc.), the high-frequency status of the market participant (HFT, NON-HFT, or MIXED), and the account type (own account, client account,

Table 3.1 – Percentages of auction snapshots with accurate reconstruction.

	Opening auction	Closing auction
Number of valid snapshots	34,971	34,820
% snapshots matching the auction price	99.6%	99.9%
% snapshots matching the auction volume	99.0%	99.7%
% snapshots matching both	98.9%	99.6%

market maker, parent company, retail liquidity provider, retail market organization).

In order to reconstruct the exact state of the limit order book (LOB) at any point during the auction, we combine the information from the four different files for each stock and each trading day to create a snapshot. We select the 34 most traded stocks on Euronext Paris between 2013 and 2017 and analyze 2 to 5 years worth of data for each stock, totaling $N = 34,977$ stock-days. A small number of these stock-days result in errors or mismatches (e.g., dataset errors, non-crossing supply and demand for the opening auction, or half-day trading/halted trading before 17:30 for the closing auction). After removing these invalid snapshots, we are left with $N_o = 34,971$ valid snapshots at the opening auction time and $N_c = 34,820$ valid snapshots at the closing auction time.

Using these reconstructed snapshots just before the auction time, we compute reconstructed prices and volumes as per Euronext rules, i.e., by maximizing the exchanged volume and minimizing the imbalance. This boils down to finding the intersection of the reconstructed supply and demand curves. Table 3.1 reports the percentage of snapshots for which the reconstructed price (resp. volume) matches the actual auction price (resp. volume) among valid snapshots. The remaining discrepancies may be a result of using simplified rules to account for stop orders and occasional contradictions between recorded data in the orders file and the trades file. For these few unmatched snapshots, we note that the discrepancies between computed and actual quantities are small: less than 1 basis point on the absolute average difference from the auction price and 0.2% on the absolute average distance from the auction volume. These few unmatched auctions are discarded from the sample in the subsequent analysis, though they would not alter the outcome of our experiments.

3.4 Average shape of the auction limit order book

This section investigates the typical shape of the limit order book at auction time T_a, what it implies for post-clearing price impact, and how the average LOB shape can be broken down by latency and account type of market participants.

3.4.1 Pre-clearing vs. post-clearing LOB shape

For each stock of the dataset, we compute the buy and sell average empirical densities $\langle \tilde{\rho}_\bullet \rangle$ (see Definition 5) as a function of the log-price difference $x = \log\left(\frac{p}{p_a}\right)$. Figure 3.2 shows the average LOB density for the most traded stock in our dataset (ISIN FR0000120271, TTE.PA, TotalEnergies). We distinguish the orders that are cleared by the auction process (dotted lines) from the ones that remain in the LOB after the end of the auction (full lines). Average LOB densities are very similar across all the studied stocks.

Figure 3.2 shows the average LOB densities at the closing auction: the buy and sell densities have a skewed bell-shaped curve around the auction price. Opening and closing auctions have clearly different LOB densities. As expected, the average LOB density is noisier at the opening auction than at the closing auction which reflects the typical liquidity available at either auction Challet (2019). However, the following remarks hold for both auctions:

— there is a peak at the auction price, i.e. $\langle \tilde{\rho}_\bullet \rangle (0)$ is larger than typical values taken near 0. This translates an accumulation of orders on $p = p_a$ on average at the time of the clearing;

— $\langle \tilde{\rho}_\bullet \rangle$ is linear around $x = 0$, i.e. $p = p_a$.

As shown by Fig. 3.2, all buy orders with $p > p_a$ are cleared, and all buy orders with $p < p_a$ remain in the LOB after the auction as long as their temporal validity extends beyond the clearing; similarly, all sell orders with $p < p_a$ are cleared and all sell orders with $p > p_a$ remain in the LOB after the auction. For $p = p_a$, some orders are matched, some are not. This explains why the peaks of buy and sell volumes at p_a are reduced after the clearing. Finally, the auction-only orders are removed from the LOB after the clearing if they are not executed.

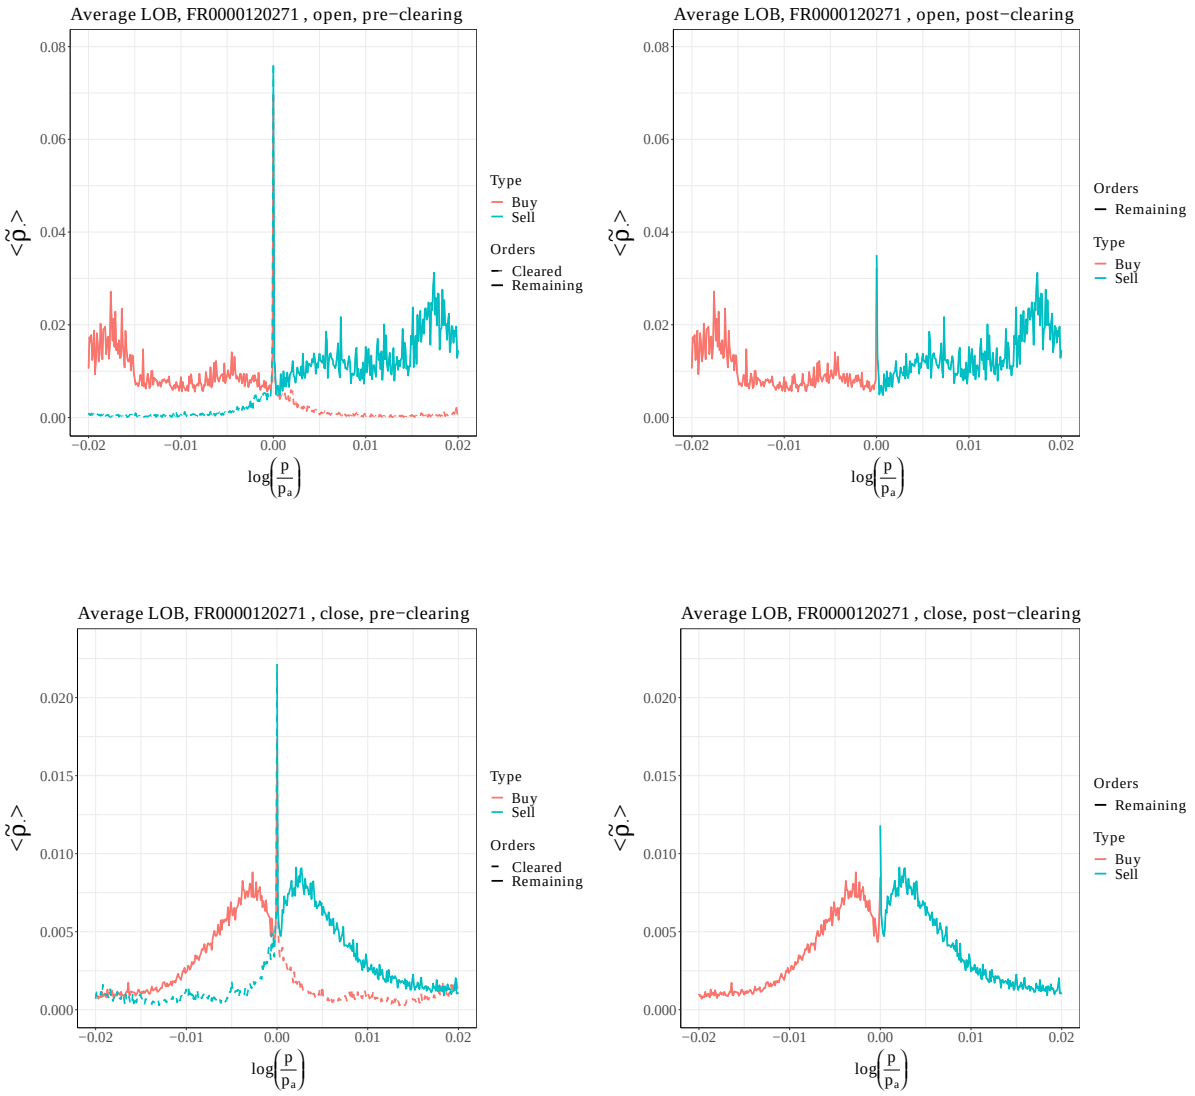

Figure 3.2 – Average density of the limit order book $\langle \tilde{\rho}_\bullet \rangle$ as a function of the log difference from the auction price p_a at the opening auction (top) and at the closing auction (bottom); left plots: pre-clearing, right plots: post-clearing (right). TTE.PA (TotalEnergies) between 2013 and 2017. In all panels, the mean density is computed on price intervals of size $\delta x = 1$bp over $N = 1266$ days.

3.4.2 Post-clearing instantaneous price impact

Let us briefly discuss the instantaneous post-clearing price impact during a continuous trading phase just after an auction. The following remarks are valid whenever there is a continuous trading phase right after the auction clearing (that is, after the open auction here). Consider the case of a trader sending a buy market order during the continuous trading phase just after auction clearing. This trader can expect to match up to all the remaining sell orders at p_a without impacting the price. Once the liquidity at p_a is consumed, sending an additional buy volume $q > 0$ will result in a sub-linear price impact. Indeed, since $\langle \tilde{\rho}_S \rangle$ has been observed to be linear around 0 (peak excluded), we may write $\langle \tilde{\rho}_S \rangle (x) = a_1 + b_1 x$ on this neighborhood so that we have on average

$$\int_0^x \langle \tilde{\rho}_S \rangle (u) \mathrm{d}u = q, \tag{3.9}$$

which implies

$$\frac{b_1}{2} x^2 + a_1 x - q = 0. \tag{3.10}$$

Hence, the post-clearing instantaneous price impact x is sub-linear and ranges between a square root limit when $q \gg \frac{a_1^2}{2b_1}$ and a linear impact limit $q \ll \frac{a_1^2}{2b_1}$. This reproduces in a stylized way the crossover between linear and square-root market impact observed in continuous double auctions (Bucci et al., 2019a). The latter can be explained for example by assuming the existence of a hidden, latent LOB (Tóth et al., 2011), which is only partially revealed but whose shape largely determines that of market impact. At auction times instead, market participants are forced to reveal their intentions at least in the vicinity of p_a, and one can relate the auction LOB with the latent LOB.

3.4.3 Breakdown by market participant latency

Figure 3.3 displays a breakdown of the average empirical densities $\langle \tilde{\rho}_\bullet \rangle$ at the closing auctions by the speed of market participants. We used the latency flag in our data which specifies the HFT category of the order sender as per the AMF definition. Let us make three remarks regarding Figure 3.3. First, we notice that the MIX LOB has the same order of magnitude and shape as the total LOB (Fig. 3.2, bottom). This indicates that the contribution of traders flagged as fast (HFT) and slow (NON) to the liquidity provision of the closing auction (limit orders in the neighbourhood of the auction price) is smaller than the contribution of investment banks (flagged MIX). Second, the HFT LOB does not display an outstanding peak of volumes at the auction price. This suggests that this peak is actually caused by slow traders and may result

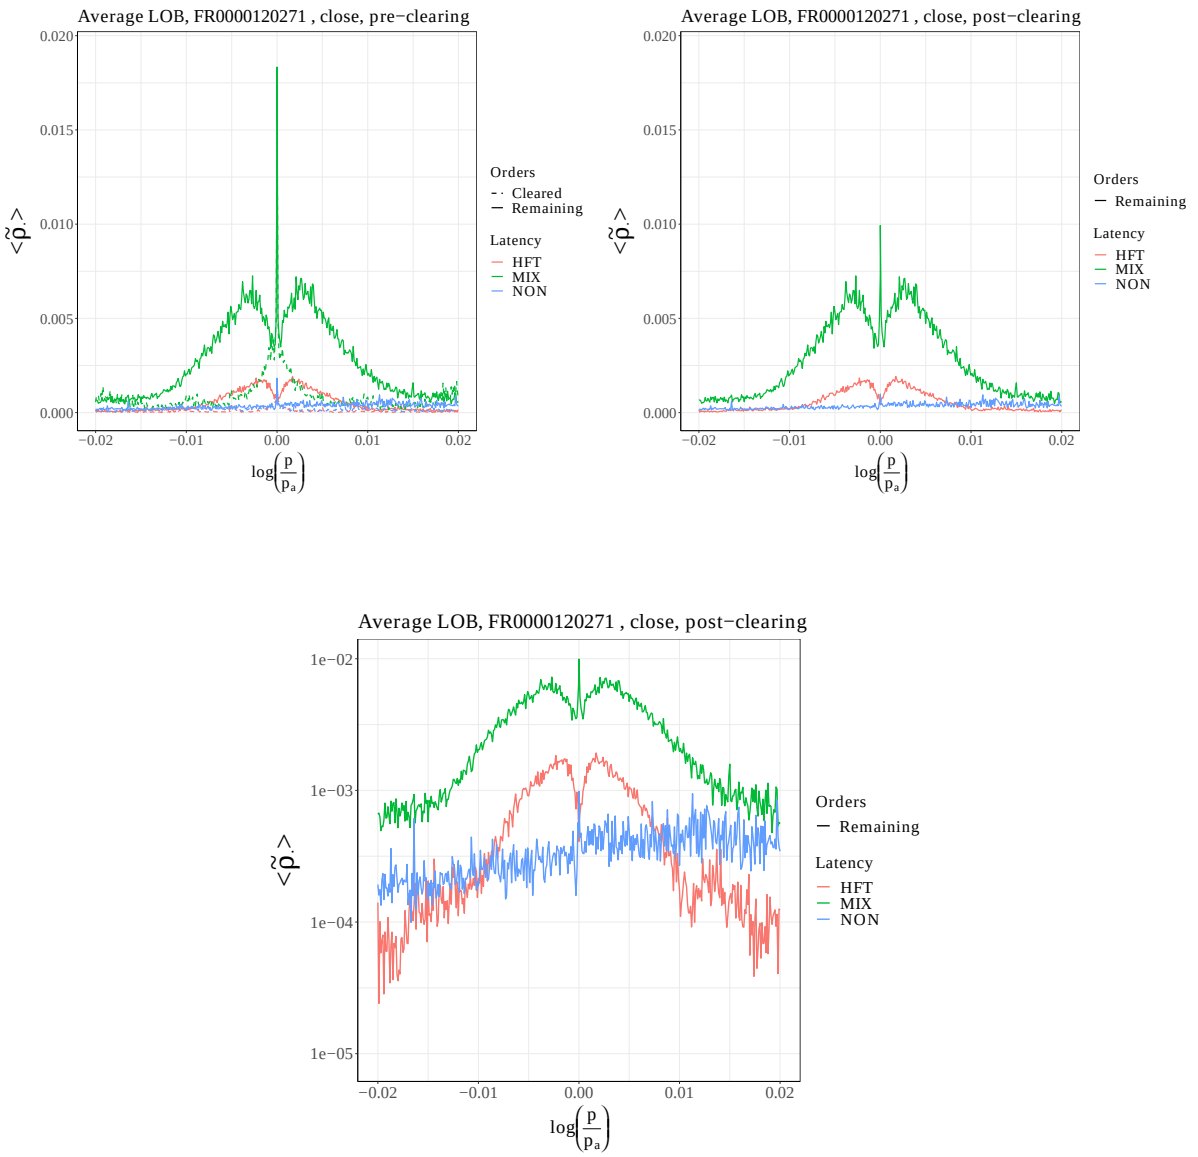

Figure 3.3 – Average density of the limit order book $\langle \tilde{\rho}_\bullet \rangle$ as a function of the log difference from the auction price p_a: breakdown by user latency of the average LOB during the closing auction just before the clearing (left), right after the clearing (right and bottom), with Y-axis in a log-scale (bottom) for TTE.PA between 2013 and 2017. The HFT flag denotes pure high frequency traders, MIX denotes investment banks with high frequency trading activities, and NON denotes traders without HFT activities.

in auction price pinning. Third, Figure 3.3 deals with the most liquid stock of the sample, but some stocks have a very small HFT-flagged LOB with the same order of magnitude as the low frequency LOB: HFT-flagged traders do not place sizeable limit orders in the closing auction of all stocks.

As stated in AMF (2017); Benzaquen and Bouchaud (2018a), open markets are dominated by fast trading algorithms, which suggests considering the HFT LOB only (up to a multiplicative constant) when relating the auction LOB with the latent continuous-auction LOB. In this setting, the post-clearing price impact is much closer to a square root because of the sharp linear shape of the HFT LOB that vanishes around the current price.

3.4.4 Breakdown by account type

Figure 3.4 shows a breakdown of the average empirical densities $\langle \tilde{\rho}_\bullet \rangle$ at the closing auction by the account type. This particular flag tells on whose behalf an order was sent: client account, market maker, own account, parent company account, retail market organization (RMO), and retail liquidity provider (RLP). We notice that traders operating on behalf of their own account, which includes a significant fraction of investment bank activities, and market makers provide most of the liquidity in the vicinity of the auction price. In addition, the density of orders sent on behalf of clients and slow traders have the same shape (see Fig. 3.3). This decomposition will be valuable in designing realistic agent-based models in addition to incorporating multi-time scale liquidity [1].

3.5 Price impact

This section investigates a set of statistical regularities of price impact in equity auctions focusing on closing auctions. In the first part, we study price impact at the auction time, which was fixed at 17:35:00 before the 28$^{\text{th}}$ of September 2015, and then randomly between 17:35:00 and 17:35:30: we assume that a trader wishes to know by how much the auction price would have moved if she had sent a market order right before the clearing, supposing that she could know the clearing time in advance. In the second part, we study the behavior of price impact before the auction time. To this end, we examine the evolution of the virtual/instantaneous price impact throughout the accumulation period. Then, we relate the price impact at auction time with that

1. There are only 126 authorized participants on the cash market (that includes equities) of Euronext Paris. See: https://live.euronext.com/en/resources/members-list.

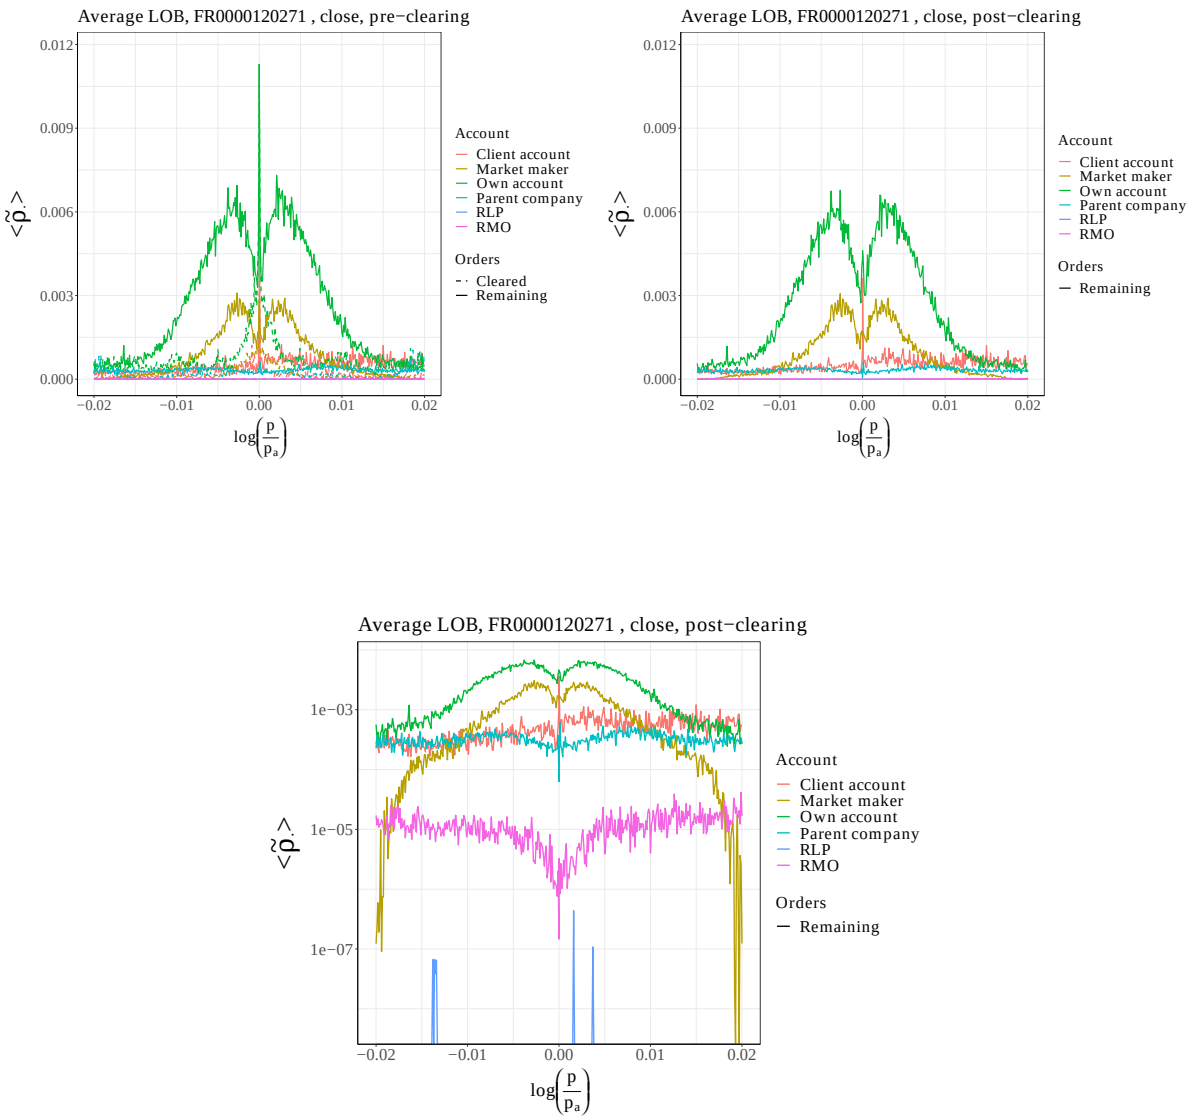

Figure 3.4 – Average density of the limit order book $\langle \tilde{\rho}_\bullet \rangle$ as a function of the log difference from the auction price p_a: breakdown of the average LOB during the closing auction by user account type just before the clearing (left), right after the clearing (right and bottom), with Y-axis in a log-scale (bottom) for TTE.PA between 2013 and 2017. Colors represent orders executed on the behalf of: a client account, a market maker, an own account, a parent company account, a retail market organization (RMO), and a retail liquidity provider (RLP).

at 17:35:00. Finally, we compute the average impact of actual submissions/cancellations during auctions and discuss why it is markedly different from that of open markets.

3.5.1 At the auction time

In this first part, we investigate the impact of a market order submitted (or canceled) to the exchange just before the clearing. We explicitly assume that the trader would have been able to insert or cancel her order just before the clearing process. In this setting, we highlight the existence of a significant zero impact volume below which the auction price would not have changed and explain why this zero impact is purely mechanical. We then show that any additional volume has a linear price impact over a volume range that we determine, not only on average but for most stocks and days. We also derive a simple formula for the impact slope that we validate empirically using a simplified optimization routine. Finally, we examine the influence of derivative expiry days on closing auctions.

Zero impact: $\omega < \omega_\bullet^{(0)}$

When inspecting the price impact function over several days and auctions, we observe that the minimal volume necessary to change the auction price ($Q_a \times \omega_\bullet^{(0)}$ using the notations of Proposition 2), can be much larger than the typical volumes needed to impact the price further ($Q_a \times \delta\omega_\bullet^{(i)}$, $i \geq 1$). A compelling example is given by Figure 3.5, which shows the price impact function for TTE.PA at the closing auction of May 5, 2017, with the following quantities: $p_a = 48.00$€, $Q_a = 2,246,617$, $\omega_B^{(0)} = 27.45\%$, and $\omega_S^{(0)} = 9.61\%$. Hence, if sent just before T_a, a buy order of a cash volume lower than $Q_a \times \omega_B^{(0)} \times p_a = 29.6$ million€ would not have resulted in an auction price change. Similarly a sell order of a cash volume lower than $Q_a \times \omega_S^{(0)} \times p_a = 10.3$ million€ would have had zero impact.

In our sample, zero price impact is present in more than 98% of the total processed days and sides. This means that in more than 98% of the time, sending one share, either on the buy or the sell side, will not change the auction price. In addition, and maybe more surprisingly, zero impact on both sides simultaneously is by far the most common situation. This comes from the fact that the prices are discrete and thus the cumulative buy and sell volumes $D(p)$ and $S(p)$ are step functions. At the auction price, these steps only overlap partially. To change the auction price, one needs to shift vertically either D or S in such a way that the overlap at the auction price disappears (see Figure 3.1 for an illustration). Thus, zero price impact only disappears when both $V_B(p_a)$ and $V_S(p_a)$ only have one share at most at p_a.

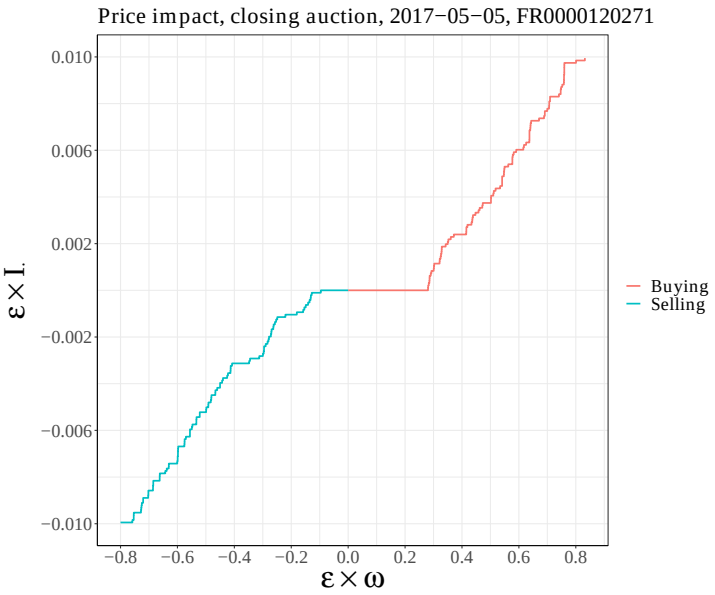

Figure 3.5 – Virtual price impact $\varepsilon \cdot I$ as a function of the (scaled) added signed volume $\varepsilon \cdot \omega$ at the closing auction of TTE.PA on 2017-05-05. $\varepsilon = +1$ for a buy market order and $\varepsilon = -1$ for sell market order.

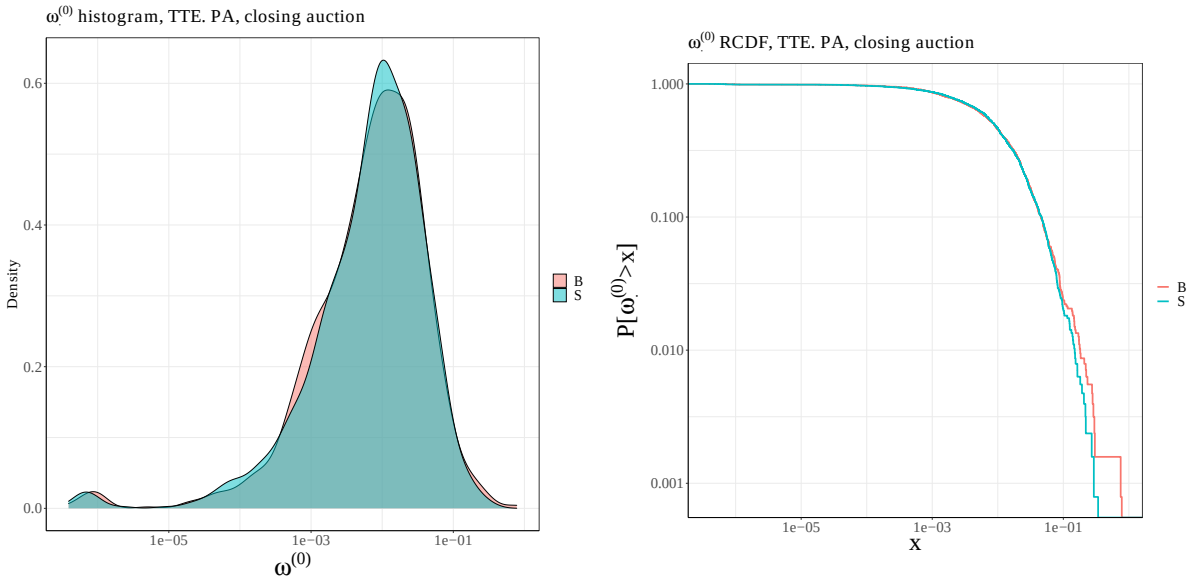

Figure 3.6 – Left panel: smoothed histograms of zero impact volumes (buy and sell) $\omega_\bullet^{(0)}$; right panel: empirical reverse cumulative distribution function (RCDF) of zero impact volumes $\omega_\bullet^{(0)}$.

Price impact can be zero for relatively large orders because of the peak of volume at p_a: recall that the (scaled) zero-impact volume $\omega_\bullet^{(0)}$ is the minimal volume needed to change the auction price; by the definition of $\omega_\bullet^{(0)}$ in Proposition 2, having large matched buy and sell volumes at the auction price leads to large zero impact volumes on both sides (see Figure 3.1). This is confirmed empirically: we report in Table 3.2 the probability $\mathbb{P}_{1\%}$ to send a market order of size $q = 1\% \times Q_a$ just before the clearing without moving the closing price. For the stocks in our sample, this probability ranges from 46% to 74%. The randomization of the clearing time prevents fast agents from using their low latency to size their trades so as to have zero impact.

We also report several statistical observations on $\omega_B^{(0)}$ and $\omega_S^{(0)}$. First, their statistical distribution can not be distinguished (as shown in Figure 3.6). This is confirmed by a Kolmogorov-Smirnov test reported in Table 3.2, which also reports the empirical Spearman correlation between these two quantities: quite surprisingly, given the observation above, the correlation between $\omega_B^{(0)}$ and $\omega_S^{(0)}$ is rather weak, -0.15 on average, and is non-significant for some very liquid stocks (e.g., TTE.PA the most traded stock in our dataset). This confirms that zero-impact is mostly a mechanical effect, not a strategic one.

Let us finally compare $\delta\omega_\bullet^{(0)} = \omega_\bullet^{(0)}$, the minimal scaled volume needed to move the auction price, to $\delta\omega_\bullet^{(i)} = \omega_\bullet^{(i)} - \omega_\bullet^{(i-1)}$, $i \geq 1$, the minimal scaled volumes needed to take the price from $p_\bullet^{(i)}$ to $p_\bullet^{(i+1)}$ (see Proposition 2). Table 3.3 presents results for the stock TTE.PA of pairwise

Table 3.2 – Spearman correlation and Kolmogorov-Smirnov test statistics for $\omega_B^{(0)}$ and $\omega_S^{(0)}$, as well as the probability $\mathbb{P}_{1\%}$ to send a market order of size $q = 1\% \times Q_a$ without impacting the auction price just before the clearing across the stocks in our sample.

ISIN	$\text{cor}(\omega_B^{(0)}, \omega_S^{(0)})$	KS statistic$(\omega_B^{(0)}, \omega_S^{(0)})$	$\mathbb{P}_{1\%}$	Observations
CH0012214059	-0.559***	0.088*	64%	510
FR0000031122	-0.267***	0.056	67%	1009
FR0000045072	-0.321***	0.035	65%	1014
FR0000073272	-0.102**	0.049	66%	1015
FR0000120073	-0.210***	0.025	64%	1014
FR0000120172	-0.156***	0.045	66%	1268
FR0000120271	-0.048	0.022	46%	1266
FR0000120354	-0.152***	0.088***	70%	1014
FR0000120404	-0.089**	0.052	69%	1015
FR0000120537	-0.004	0.044	70%	504
FR0000120578	-0.139***	0.054*	48%	1268
FR0000120628	-0.263***	0.043	64%	1261
FR0000120644	-0.166***	0.027	60%	1012
FR0000120685	-0.209***	0.044	68%	1014
FR0000121014	-0.272***	0.021	68%	1014
FR0000121147	-0.015	0.027	73%	1013
FR0000121261	-0.225***	0.027	65%	1013
FR0000121501	-0.311***	0.053	68%	1014
FR0000121667	-0.338***	0.021	69%	1012
FR0000121972	-0.095***	0.035	58%	1264
FR0000124141	-0.288***	0.026	73%	1012
FR0000125007	-0.14***	0.036	57%	1265
FR0000125338	-0.172***	0.042	68%	1012
FR0000125486	-0.132***	0.038	59%	1013
FR0000127771	-0.248***	0.04	68%	1015
FR0000130338	-0.098*	0.033	74%	613
FR0000130809	-0.061*	0.032	56%	1259
FR0000131104	-0.128***	0.049	53%	1264
FR0000131708	-0.118**	0.042	68%	771
FR0000131906	-0.075**	0.032	62%	1269
FR0000133308	-0.242***	0.043	67%	1012
FR0010208488	-0.271***	0.025	68%	1010
FR0013176526	-0.349***	0.065	57%	401
NL0000235190	-0.096***	0.035	61%	1264

The symbols ***,**, and * indicate significance at the 0.1%, 1%, and 5% level, respectively.

Table 3.3 – Kolmogorov-Smirnov statistics for pairs of rescaled incremental volumes $\delta\omega_\bullet^{(i)}$ and $\delta\omega_\bullet^{(j)}$ for the stock TTE.PA.

	$\delta\omega^{(0)}$	$\delta\omega^{(1)}$	$\delta\omega^{(2)}$	$\delta\omega^{(3)}$	$\delta\omega^{(4)}$	$\delta\omega^{(5)}$	$\delta\omega^{(6)}$	$\delta\omega^{(7)}$	$\delta\omega^{(8)}$	$\delta\omega^{(9)}$	$\delta\omega^{(10)}$
$\delta\omega^{(0)}$											
$\delta\omega^{(1)}$	0.091***										
$\delta\omega^{(2)}$	0.047**	0.08***									
$\delta\omega^{(3)}$	0.063***	0.103***	0.034								
$\delta\omega^{(4)}$	0.057***	0.105***	0.033	0.023							
$\delta\omega^{(5)}$	0.053**	0.089***	0.02	0.026	0.028						
$\delta\omega^{(6)}$	0.055**	0.112***	0.04*	0.021	0.028	0.032					
$\delta\omega^{(7)}$	0.068***	0.117***	0.044*	0.019	0.026	0.031	0.024				
$\delta\omega^{(8)}$	0.043*	0.086***	0.02	0.03	0.03	0.018	0.037.	0.037			
$\delta\omega^{(9)}$	0.047**	0.091***	0.019	0.025	0.021	0.02	0.032	0.036	0.016		
$\delta\omega^{(10)}$	0.042*	0.081***	0.022	0.04*	0.037	0.022	0.043*	0.041*	0.016	0.023	

The symbols ***, **, and * indicate significance at the 0.1%, 1%, and 5% level, respectively.

Kolmogorov-Smirnov tests on the empirical distribution functions of $\delta\omega_\bullet^{(i)}$ and $\delta\omega_\bullet^{(j)}$. For the sake of brevity, results are presented for $i, j \leq 10$, but the statistical testing has actually been conducted up to $i, j = 40$. We clearly observe that $\delta\omega_\bullet^{(0)}$ and $\delta\omega_\bullet^{(1)}$ have specific statistical properties, while the distributions of the incremental volumes $\delta\omega_\bullet^{(i)}$ for $2 \leq i \leq 32$ could hardly be distinguished as the null hypothesis could not be rejected at the 1% significance level. Figure 3.7 shows smoothed histograms and empirical reverse cumulative distribution function for $\delta\omega_\bullet^{(i)}$, $0 \leq i \leq 5$. This observation is not easily generalized to all stocks since additional factors come into play: small tick vs. large tick stocks and the randomization of the clearing time. These factors have a non-negligible influence on the distribution of $\delta\omega_\bullet^{(i)}$s across different stocks and over the years.

Linear impact: $\omega_\bullet^{(0)} < \omega < \omega_\bullet^{(\max)}$

According to Donier and Bouchaud (2016), in a Walrasian auction with continuous prices, average volumes around the auction price are non-null, which leads to a linear impact (in a first-order expansion), while in a continuous double auction, average volumes vanish around the current price and lead to a square root impact.

It is useful to first assume that price is continuous in order to derive a simple condition for the price impact to be strictly linear. If we send a buy market order of size $\omega \times Q_a$ before the auction clearing, and assuming we work in a log-price frame of reference $x = \log(p/p_a)$ in a

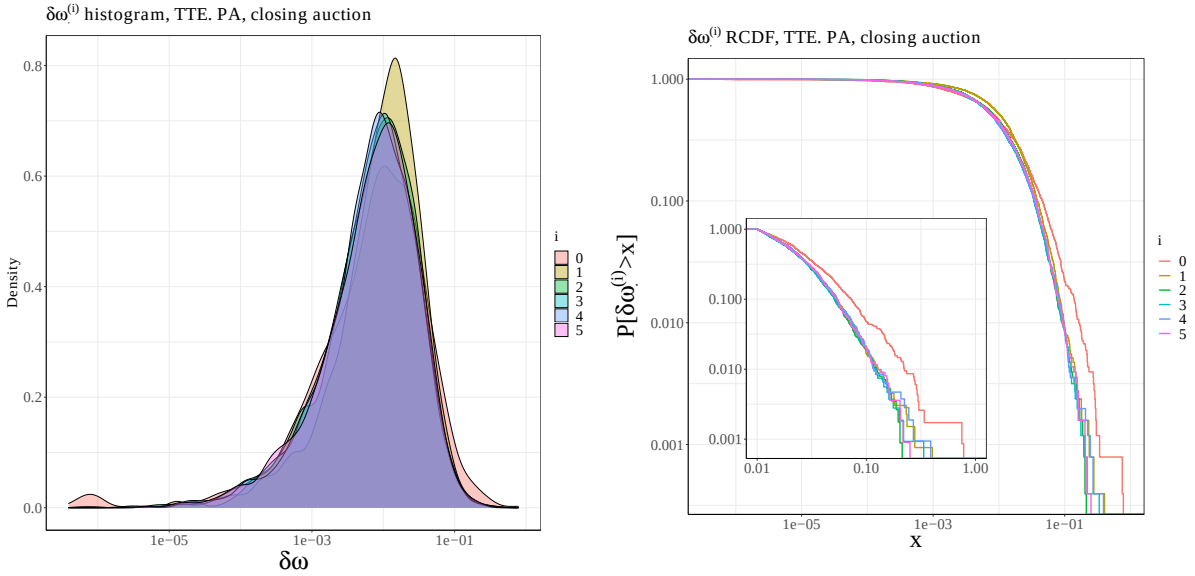

Figure 3.7 – Left panel: smoothed histograms of scaled incremental volumes $\delta\omega_\bullet^{(i)}$ for $i = 0 \cdots 5$; right panel: empirical reverse cumulative distribution function (RCDF) of scaled incremental volumes $\delta\omega_\bullet^{(i)}$ for $i = 0 \cdots 5$.

continuous price setting, we have

$$\begin{cases} S(0) = D(0), \\ S\left(I_B(\omega)\right) = D\left(I_B(\omega)\right) + \omega Q_a, \end{cases} \tag{3.11}$$

hence,

$$S\left(I_B(\omega)\right) - S(0) = D\left(I_B(\omega)\right) - D(0) + \omega Q_a. \tag{3.12}$$

Donier and Bouchaud (2016) perform a first-order expansion to write

$$\partial_x S(0) \times \left(I_B(\omega) - 0\right) = \partial_x D(0)(I_B(\omega) - 0) + \omega Q_a, \tag{3.13}$$

and approximate

$$I_B(\omega) = \frac{1}{\tilde{\rho}_S(0) + \tilde{\rho}_B(0)} \times \omega. \tag{3.14}$$

However, instead, we use equation (3.12) to find exactly

$$\int_0^{I_B(\omega)} (\tilde{\rho}_S + \tilde{\rho}_B)(x)\mathrm{d}x = \omega, \tag{3.15}$$

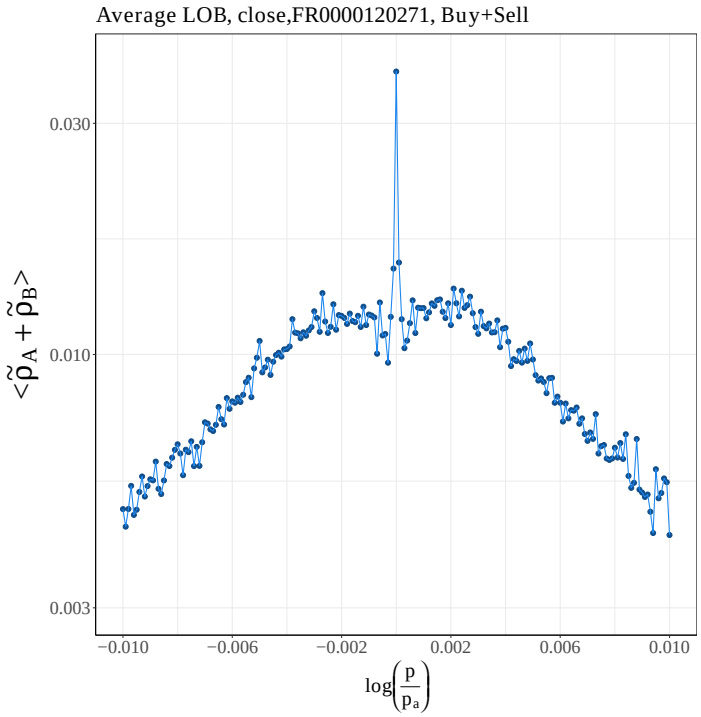

Figure 3.8 – Average empirical density of total (buy + sell) volumes for TTE.PA at the closing auction.

thus

$$\begin{aligned} I_B(\omega) &= F^{-1}(\omega), \\ F(x) &= \int_0^x (\tilde{\rho}_S + \tilde{\rho}_B)(u)\mathrm{d}u. \end{aligned} \tag{3.16}$$

Having a linear impact requires that F^{-1} and F are linear functions, therefore that $x \mapsto (\tilde{\rho}_S + \tilde{\rho}_B)(x)$ is constant.

Figure 3.8 shows the average empirical density $\langle \tilde{\rho}_S + \tilde{\rho}_B \rangle_d$ of the sum of buy and sell volumes for the most liquid stock in the sample. It strongly suggests the existence of a price interval, on each side of the auction price, in which the sum of buy and sell volumes can be well approximated by a constant.

We now include this observation in the discrete-price theoretical framework introduced Section 3.2 and we prove in Proposition 3 that if buy and sell densities sum up to a constant around the auction price p_a (removing the zero-impact part), price impact is linear.

Proposition 3. If $x \mapsto (\tilde{\rho}_S + \tilde{\rho}_B)(x)$ is constant on some intervals $]-\Delta_S, 0[$ and $]0, \Delta_B[$, then the price impact $I_\bullet$ is linear. More precisely, if $\tilde{\rho}_S(x) + \tilde{\rho}_B(x) = \widetilde{\mathscr{L}_B}$ positive constant for all $x \in]0, \Delta_B[$ and $\tilde{\rho}_S(x) + \tilde{\rho}_B(x) = \widetilde{\mathscr{L}_S}$ positive constant for all $x \in]-\Delta_S, 0[$, then for all i such

that $I(\omega^{(i)}) < \Delta$, we have

$$I(\omega^{(i)}) - I(\omega^{(0)}) = \frac{1}{p^{(1)}\widetilde{\mathscr{L}}}\left(\omega^{(i)} - \omega^{(0)}\right), \tag{3.17}$$

where we omitted the $\bullet \in \{B, S\}$ notation from I, ω, $p^{(1)}$ and $\widetilde{\mathscr{L}}$. Recall that $p_\bullet^{(1)}$ is the first non empty price tick after (resp. before) the auction price when $\bullet = B$ (resp. when $\bullet = S$), as in Proposition 2.

The proof of Proposition 3 is given in Appendix 3.B. Notice that $\widetilde{\mathscr{L}}$ represents a constant scaled liquidity around p_a. Also, since $\widetilde{\mathscr{L}}$ and ω are both scaled by Q_a, the price impact as written in the right-hand side of equation (3.17) does not depend on the auction volume Q_a. For large-tick stocks, if $V_B(p) + V_S(p) = V_c$ constant around p_a, then the scaled liquidity is given by $\widetilde{\mathscr{L}} = V_c/(Q_a\theta)$, where θ is the tick size. For small-tick stocks, one can obtain an approximation by substituting θ with a fraction of the average spread.

Following Proposition 3, we want to characterize the intervals in which $\tilde{\rho}_S + \tilde{\rho}_B$ can be considered constant. Therefore we need to find $\widetilde{\mathscr{L}_\bullet}$ and $\Delta_\bullet$ for $\bullet \in \{B, S\}$, such that

$$\begin{aligned} \tilde{\rho}_S(x) + \tilde{\rho}_B(x) &= \widetilde{\mathscr{L}_B} \text{ for all } x \in]0, \Delta_B[; \\ \tilde{\rho}_S(x) + \tilde{\rho}_B(x) &= \widetilde{\mathscr{L}_S} \text{ for all } x \in]-\Delta_S, 0[. \end{aligned} \tag{3.18}$$

For symmetry reasons, we focus on Δ_B and $\widetilde{\mathscr{L}_B}$: the problem is to find Δ_B and $\widetilde{\mathscr{L}_B}$ for a given day d by resorting to a simple change point detection algorithm. This method minimizes the residual sum of squared errors between $\log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right)$ and its mean $\eta(y)$ for $x \in]0, y]$ plus the residual sum of errors of a linear fit of $\log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right)$ for $x > y$. We choose to work with logarithms, since errors are multiplicative. The resulting cost function is

$$\begin{aligned} f(y) &= \sum_{0<x\leq y} \left|\log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right) - \eta(y)\right|^2 + \sum_{x>y} \left|\log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right) - \hat{\beta}(y)x - \hat{\alpha}(y)\right|^2; \\ \eta(y) &= \frac{1}{N_y} \sum_{0<x\leq y} \log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right), \end{aligned} \tag{3.19}$$

where $(\hat{\alpha}(y), \hat{\beta}(y))$ is the linear regression estimate of $\log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right)$ over x for $x > y$, and N_y is the number of non-null observations $(x, \log\left(\tilde{\rho}_S(x) + \tilde{\rho}_B(x)\right))$ for $x \in]0, y]$. We then define

$$\Delta_B = \arg\min_y f(y). \tag{3.20}$$

This definition means that for $x \leq \Delta_B$, the sum of the logarithm of the sum of scaled empirical

buy and sell densities is better approximated by its mean than by a non-constant (linear) fit, whereas for $x > \Delta_B$, the opposite holds. Then, we calculate $\widetilde{\mathscr{L}_B}$ as the mean of $(\tilde{\rho}_S + \tilde{\rho}_B)(x)$ for $0 < x \leq \Delta_B$ in order to avoid an underestimation due to the convexity of the exponential function. Finally, we define $\omega_\bullet^{(\max)}$ as the maximum scaled volume of a market order that would result in a null or linear impact, i.e.,

$$\omega_\bullet^{(\max)} = \omega_\bullet^{(0)} + \sum_{0<|x|\leq\Delta_\bullet} \frac{V_B(x) + V_S(x)}{Q_a}. \tag{3.21}$$

Figure 3.9 shows examples for Δ detection using the previous optimisation for two different days at the closing auction, and plots in each case the theoretical impact given by Proposition 3 with respect to the actually observed impact function. One sees that the estimated cut-off Δ, as well as the slope estimate $(p^{(1)}\widetilde{\mathscr{L}})^{-1} \approx (p_a\widetilde{\mathscr{L}})^{-1}$, fit very well the actual slope and domain of the linear price impact. This is actually the case of most days, as shown by Figure 3.10, where we plot the observed slope against the theoretical slope.

We also plot the smoothed histograms of Δ and $\omega^{(\max)}$ issued by our detection algorithm for the stock TTE.PA between 2013 and 2017 (1266 stock-days and two sides (buy and sell)) (see Fig. 3.11) . Note that we truncated the closing auction snapshots at a maximum log-price distance $x \leq 2\%$, which is twice the average impact of a market order of a size equal to the auction volume Q_a. In addition, only fits with a number of points ≥ 20 are kept, which happens in about $\approx 90\%$ of the days and sides: this shows that the price impact is linear for most of the days with an average value of Δ above 50 basis points. Finally, $\mathbb{P}\left[\omega^{(\max)} > 0.5\right] = 0.73$: this means that a trader has 73% chance to execute 50% of the total auction volume just before the close clearing and still result in zero or linear impact.

Appendix 3.C reports empirical properties of the impact slope at auction time computed for every asset, which may be of some use in transaction cost analysis.

Influence of derivatives expiry dates

When there is no derivatives expiry, the liquidity in currency units defined by $L^\$:= p_a \times Q_a \times \widetilde{\mathscr{L}}$ whether on Friday or other days of the week (Fig. 3.12, right panel) seem to be drawn from the same distribution, as we could not reject the null hypothesis associated with Kolmogorov-Smirnov tests for any pair of weekdays outside the third week of the month. However, on expiry days (third Fridays of the month), liquidity in currency units is typically larger than for other weekdays during the same week and seems to be drawn from a different shifted distribution to the right (Fig. 3.12, left panel). This finding is confirmed by one tailed Kolmogorov-Smirnov

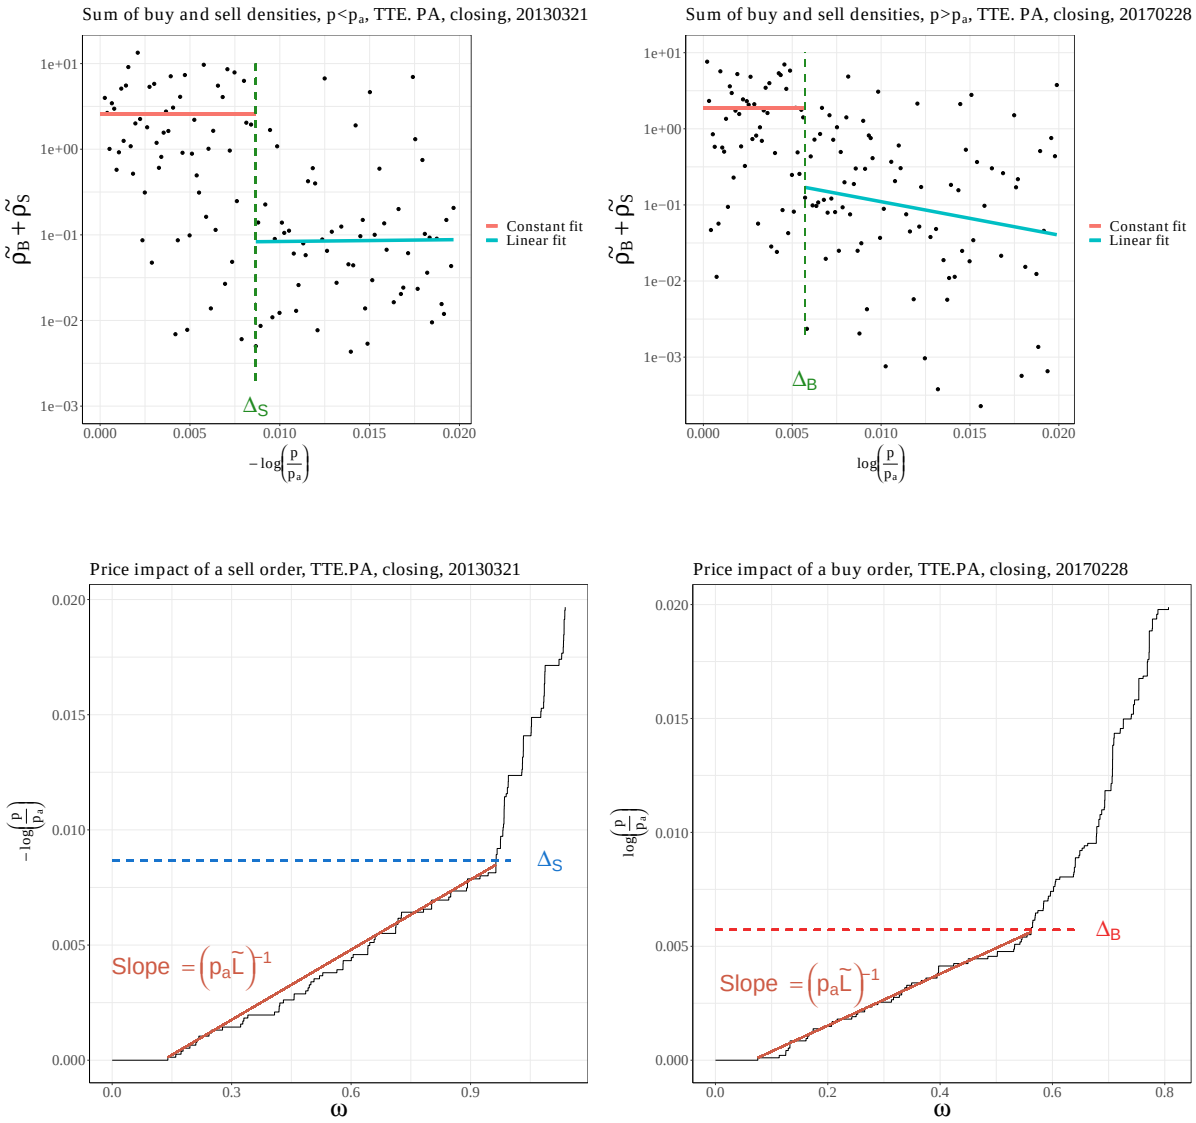

Figure 3.9 – Simplified change point detection algorithm applied on the buy side of the closing auction of TTE.PA at 2013-03-21 (left) and the sell side of the closing auction of TTE.PA at 2017-02-28 (right). The upper plots show the sum of the buy and sell empirical densities and estimated cut-off Δ with a green dashed line. Lower plots show the fit of the estimated impact slope $(p^{(1)}\widetilde{\mathscr{L}})^{-1} \approx (p_a\widetilde{\mathscr{L}})^{-1}$ on the corresponding impact functions.

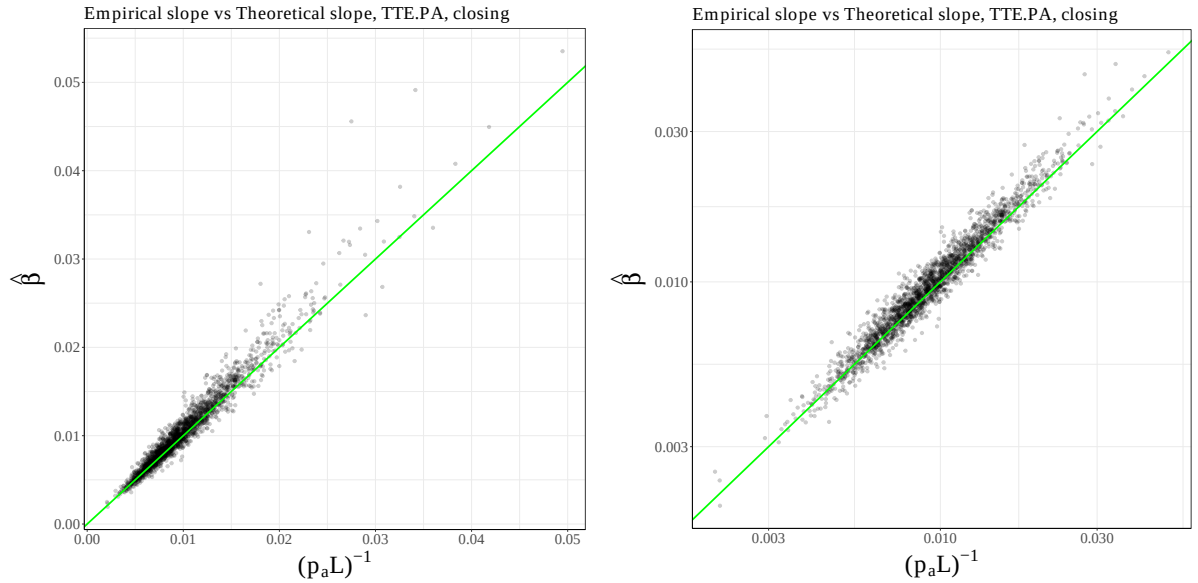

Figure 3.10 – Empirical slope $\hat{\beta}$ vs Theoretical slope $(p^{(1)}\widetilde{\mathscr{L}})^{-1} \approx (p_a\widetilde{\mathscr{L}})^{-1}$ given by equation (3.17) at the closing auction for TTE.PA. Left panel: normal scale. Right panel: log-log scale. A straight line with unit slope and null intercept is plotted in green for visual guidance.

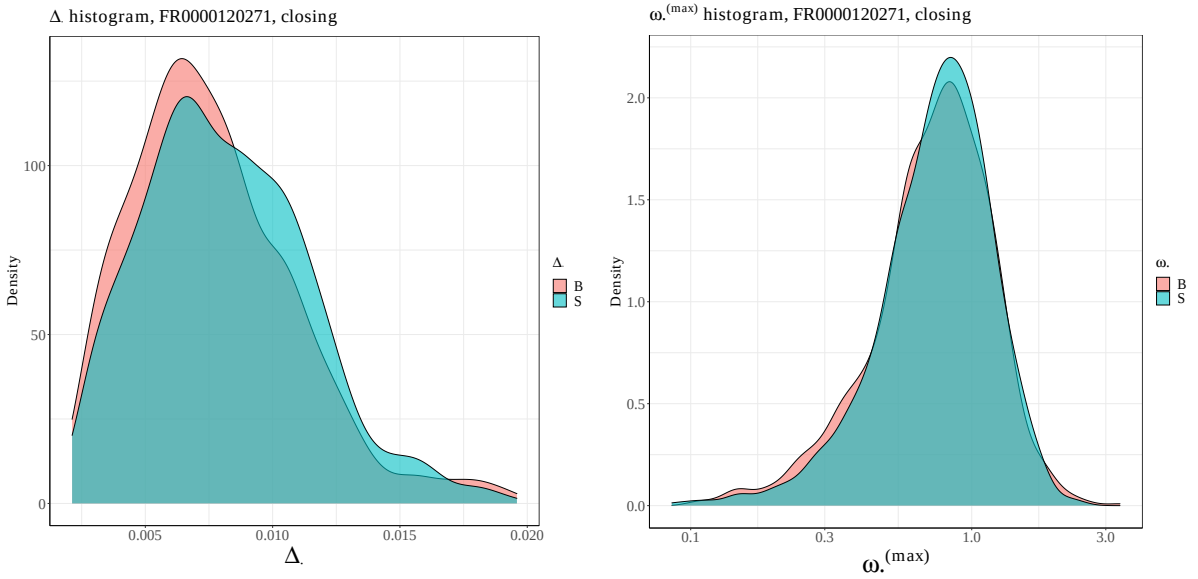

Figure 3.11 – Smoothed histograms of the maximum log-distance Δ over which the sum of the buy and sell densities can be considered constant (left) and the maximum scaled volume $\omega^{(\max)}$ that results in a null or linear impact. These are outputs of the optimization of equation (3.19) applied to closing auctions of TTE.PA between 2013 and 2017.

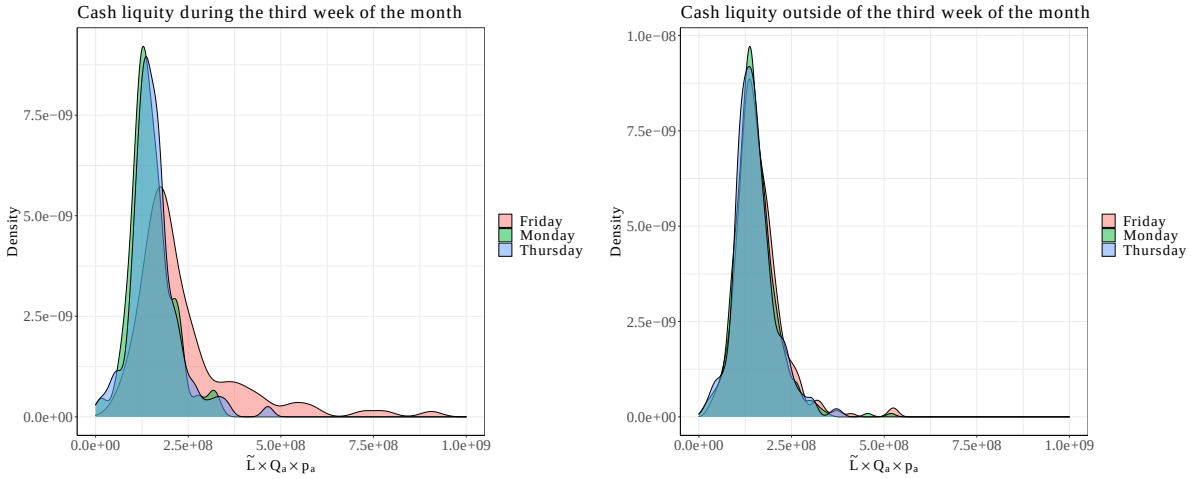

Figure 3.12 – Smoothed histograms of the closing auction estimated cash liquidity $L^{\$} := \widetilde{\mathscr{L}} \times Q_a \times p_a$ during the third week of the month (left) and outside of the third week of the month (right).

tests for Friday and any other weekday during third weeks of a month. Therefore, the impact slope is typically smaller during expiry days and the final auction order book is more resistant to price changes.

3.5.2 Before the auction time

In this second part, we study price impact before the auction time. First, we examine the evolution of virtual price impact throughout the accumulation period by looking at the evolution of liquidity as well as the maximum volume resulting in a linear impact. Second, we assume that traders have means to infer the impact slope at 17:35:00, which is the latest time that ensures not missing the clearing with certainty. We then relate zero impacts and the impact slope at 17:35:00 with those at the auction time. Finally, we study the average impact on the indicative price of actual submissions/cancellations between 17:30:30 and the auction time by means of response functions.

Price impact evolution

We investigate how the virtual price impact behaves throughout the accumulation period. We construct successive snapshots at 5-second intervals for TTE.PA at the closing auction. Then, we compute the (virtual/ instantaneous) price impact for $t \leq T_a$ with $p_a \leftarrow p_t^{\text{ind}}$ and $Q_a \leftarrow Q_t^{\text{ind}}$. We define the absolute liquidity $\mathscr{L}_t$ as the (constant) sum of buy and sell empirical densities at time t: $\mathscr{L}_t = (V_B + V_S)(t)/\delta p$ (Recall that the buy and sell densities sum up to a constant around the current indicative price). Similarly, we define $Q_t^{(\max)}$ as the maximum (absolute) volume that results in a null or linear impact time t. Figure 3.13 shows that averages of both the absolute liquidity (w.r.t. Q_a) $\mathscr{L}_t/Q_a$ and the fraction of the final liquidity $\mathscr{L}_t/\mathscr{L}_{T_a}$ follow the same pattern, i.e., a strong concave monotonicity at the start of the accumulation period followed by strong convex evolution as the clearing nears. Likewise, the average of the maximum linear volume with respect to the final volume $Q_t^{(\max)}/Q_a$ has the same shape. Nonetheless, the average mean of $Q_t^{(\max)}/Q_{T_a}^{(\max)}$ has a more complex pattern and suggests a strong effect of cancellations.

Impact at auction time vs. 17:35:00

Let us now relate the virtual market impact at 17:35:00 and at auction time after the introduction of the randomized clearing time. When the limit order book is not disseminated, traders have no direct way to estimate its shape or their virtual impact, at either time. However, sending a large market order and gradually cancelling it is a way around, and is observed at times.

The relationship between the two parts of price impact (zero, then linear) at both times is markedly different. The relative change of zero impact volumes is distributed over several orders of magnitude (see Fig. 3.14); agents do have an incentive to send zero-impact orders between 17:35:00 and the auction time. On the contrary, the slopes of the linear impact part are closely related: in 90% of the days, the relative change in the impact slope is smaller than 12% in absolute value (see Fig. 3.14). This means that the auction book stabilizes after 17:35:00 as one can expect since the clearing can occur at any time after 17:35:00. For TotalEnergies stock, the average absolute price change between 17:34:55 and 17:35:00 is 7 basis points. It is only 1.6 basis points between 17:35:00 and the auction time.

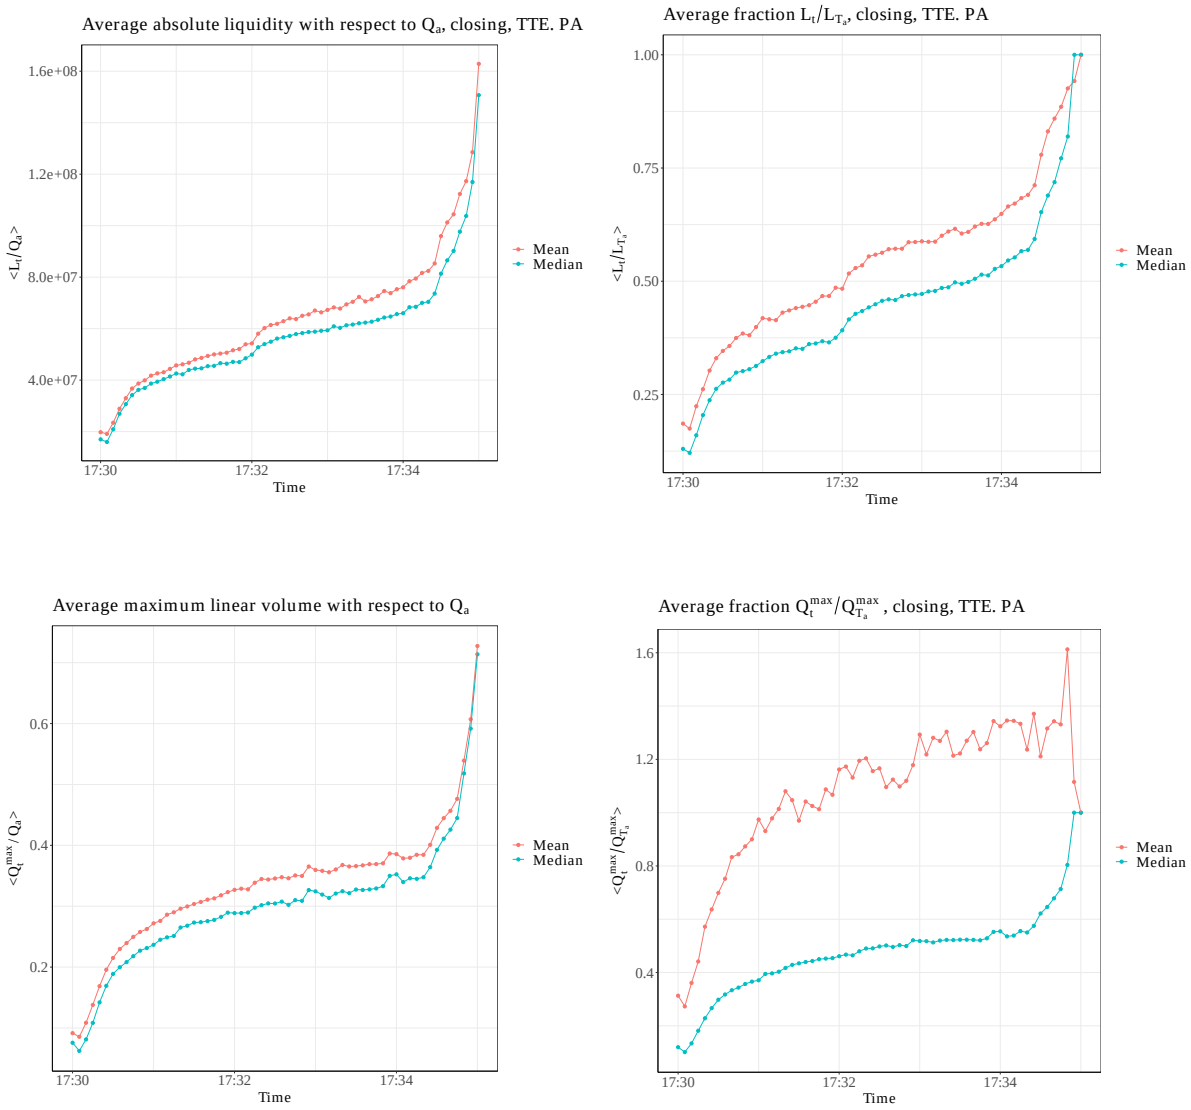

Figure 3.13 – Average absolute liquidity with respect to Q_a during the closing auction (upper left). Average fraction $\langle \mathscr{L}_t / \mathscr{L}_{T_a} \rangle$ of the absolute liquidity at time t $\mathscr{L}_t$ by final -at the clearing- liquidity $\mathscr{L}_{T_a}$ (upper right). Average (absolute) maximum linear-impact volume with respect to Q_a during the closing auction $\left\langle Q_t^{(\max)} / Q_a \right\rangle$ (lower left). Average fraction $\left\langle Q_t^{(\max)} / Q_{T_a}^{(\max)} \right\rangle$ of the absolute maximum linear-impact volume at time t $Q_t^{(\max)}$ by final -at the clearing- absolute maximum volume $Q_{T_a}^{(\max)}$ (lower right).

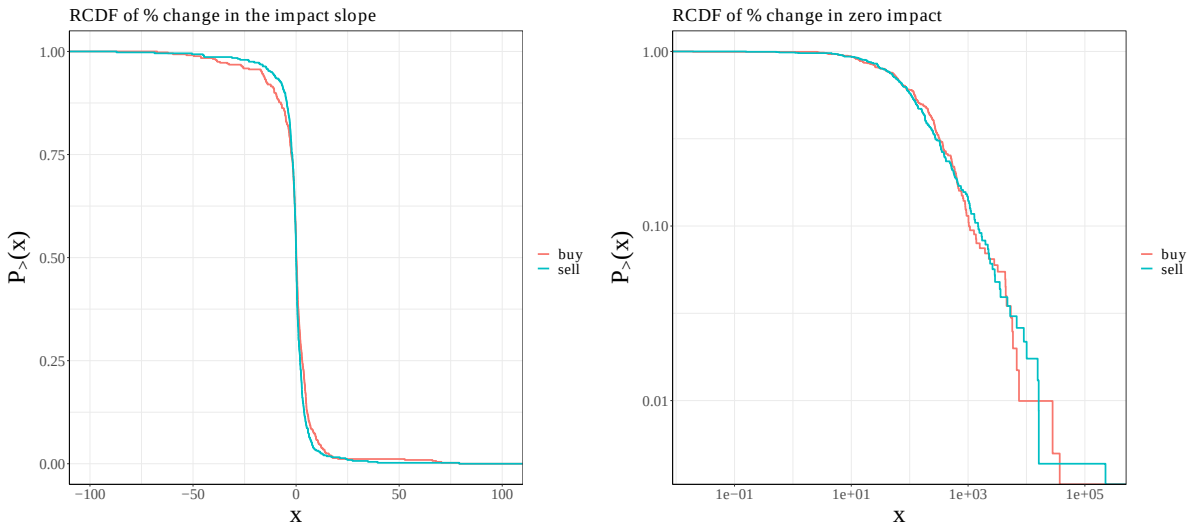

Figure 3.14 – Empirical reverse cumulative distribution of the % change in the linear impact slope (left panel), and the positive % changes in zero impact (right panel) between 17:35:00 and the clearing after the introduction of a random clearing window between 17:35:00 and 17:35:30 at closing auctions. Negative changes in % of zero impact are broadly distributed between 0% and -100%. Shown results are for TotalEnergies stock.

The linear impact of market order submission/cancellation before the auction

Finally, we evaluate the average impact of actual submissions/cancellations during the accumulation period. To this end, we compute the one lag response function R^1 for marketable orders (market orders and limit orders with an aggressive limit price) conditional on the order (scaled) size ω

$$R^1(\omega) = \langle \varepsilon_t \cdot (p_{t+1} - p_t) | \omega \rangle \,, \tag{3.22}$$

where, p_t is the indicative price just before the arrival of t^{th} marketable order submission, $\varepsilon = +1$ for a buy, -1 for a sell, and the time is incremented at each marketable order submission. Additionally, we compute the one lag mechanical response function R^M conditional on the order (scaled) size ω

$$R^M(\omega) = \left\langle \varepsilon_t \cdot (p_t^+ - p_t) | \omega \right\rangle \,, \tag{3.23}$$

where p_t^+ is the indicative price just after the marketable order arrival.

In contrast to open markets where R^1 is sub-linearly dependent on the volume (Lillo et al., 2003; Potters and Bouchaud, 2003; Bouchaud et al., 2018a), we observe in Fig. 3.15 that R^1 scales linearly with ω for marketable orders larger than a certain threshold $\omega^* \approx 3 \cdot 10^{-3}$. For

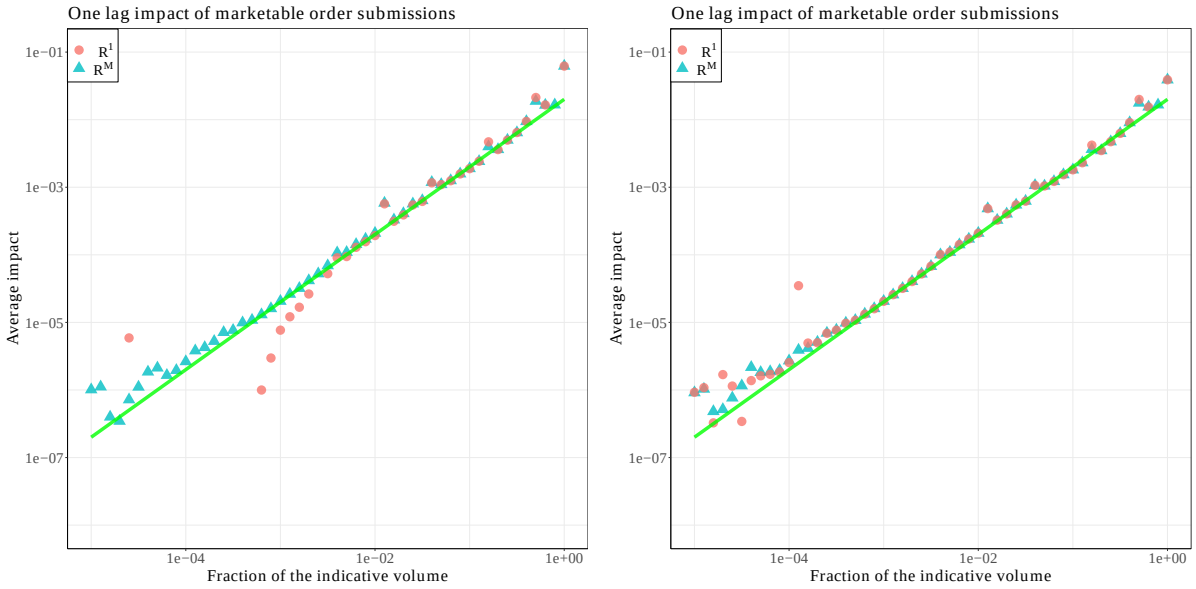

Figure 3.15 – Left panel: the average one lag response function R^1 and the mechanical impact R^M for marketable order submissions as a function of the scaled order size $\omega = V/Q_t^{\text{ind}}$; right panel: R^1 and R^M for submissions and cancellations. We used tick-by-tick closing auction data from BEDOFIH for TotalEnergies stock between 2013 and 2017. We discarded the first 30 seconds of each auction as it contains abnormal submissions related to the activation of VFA/VFC orders (Valid For Auction/Valid For Closing). The green line is the curve of $y = 0.02 \times x$.

$\omega < \omega^*$, values of R^1 can be negative indicating a strong mean reversion of the price, with values smaller than a tenth of a basis point in absolute value. For $\omega > \omega^*$, we have essentially $R^1 \approx R^M$ indicating that the price impact of individual orders is mostly mechanical and linear in ω. There is no selective liquidity taking as R^1 and R^M scale linearly with ω. Incorporating marketable order cancellations with $\varepsilon = +1$ for sell cancellations and -1 for buy cancellations yields $\omega^* \longrightarrow 0$, as we account for almost all price-changing events (right panel of Fig. 3.15). These results imply that the nature of price impact is the same during the accumulation time and at the auction time, and contrasts with results for open markets, where selective liquidity taking causes very different shapes between the average virtual impact (using the instantaneous shape of the book) and market impact of actual trades (Weber and Rosenow, 2005; Bouchaud et al., 2009).

3.6 Conclusion

The discrete nature of prices in limit order books mechanically causes the price impact at auction time to be zero at first, sometimes for quite a substantial fraction of the total exchanged volume. Surprisingly, zero price impact happens most of the time simultaneously on both sides of the auction book, for additional sell and buy market orders or equivalently for cancellations of buy or sell market orders. For volumes larger than zero-impact ones, price impact at auction time is linear in a limited price range around the auction price not only on average but for more than 90% of days. The theoretical work of Donier and Bouchaud (2016) shows the linearity of the auction impact locally around the auction price using a first-order expansion and under strong regularity assumptions of supply and demand in a continuous price setting. Here, we showed that the linearity of auction impact is due instead to the fact that the sum of buy and sell volumes around the auction price is constant.

While this work mainly describes the final result of the order accumulation process and characterizes the limit order book at the auction time, a more microscopic description of the dynamics of order submission, cancellation, and perhaps diffusion (price update) is needed. Even though market orders submitted during the accumulation period do not play a significant role in shaping the price response of the final limit order book, the action-reaction game between market orders and limit orders throughout the auction (Raillon, 2020; Besson and Fernandez, 2021) is probably a major driver of its dynamics. Similarly, the interplay between the various categories of agents (HFTs, market makers, agents trading on their behalf, or agents trading on behalf of their clients, ...) is clearly of great interest. For example, Boussetta et al. (2017) show that HFTs submit their orders in a markedly different way than slow traders. A good starting point would be a substantial modification of the model of Donier and Bouchaud (2016) in the spirit of the work done by Lemhadri (2019).

Acknowledgements

This publication stems from a partnership between CentraleSupélec and BNP Paribas.

The authors acknowledge the use of the EUROFIDAI BEDOFIH's database acquired through "Equipex PLADIFES ANR-21-ESRE-0036 (France 2030)".

Appendix

3.A Proof of Proposition 2

Proof. We only prove the proposition for an additional buy market order resulting in a price impact denoted by I_B. The case of a sell market order resulting in an impact I_S is symmetric. By definition, $\omega \mapsto p_\omega$ is a non-decreasing right-continuous step function; the same holds for $\omega \mapsto I(\omega)$. Obviously $I(0) = 0$ and $I(\omega) = 0$ if and only if $p_w = p_a$. Since $\omega^{(0)}$ denotes the first point of discontinuity of I, by monotonicity, the condition $p_w = p_a$ is equivalent to $\omega < \omega^{(0)}$. In the original auction $\mathscr{A}$ with auction price p_a and auction volume Q_a, we have

$$\begin{cases} S(p_a) - V_S^R(p_a) = D(p_a) - V_B^R(p_a) = Q_a, \\ V_S^R(p_a) \times V_B^R(p_a) = 0. \end{cases} \tag{3.24}$$

All these quantities are fixed by the original auction setting. If we add a buy market order of size $q = \omega \times Q_a$ in this setting, the new auction price p_ω satisfies

$$\begin{cases} S(p_\omega) - V_S^R(\omega) = D(p_\omega) - V_B^R(\omega) + q, \\ V_S^R(\omega) \times V_B^R(\omega) = 0, \end{cases} \tag{3.25}$$

where S and D are the original supply and demand functions, and $V_S^R(\omega)$ (resp. $V_B^R(\omega)$) is the remaining sell quantity (resp. buy quantity) *at price* p_ω in the new setting. These volumes depend clearly on ω.

Let us now determine the first point of discontinuity $\omega_B^{(0)}$. It is clear that the first price change due to the addition of a market order of size $q = \omega_B^{(0)} Q_a$ occurs when $V_S^R(\omega) = V_S(p_\omega)$, $V_B^R(\omega) = 0$, and the new auction price $p_\omega = p_B^{(1)}$ is the first non empty price tick after p_a in the sense of $V_S + V_B$, i.e., the first tick price strictly greater than the auction price which contains buy or sell shares. (see Figure 3.1 to build an intuition). Equation (3.25) yields

$$S(p_B^{(1)}) - V_S(p_B^{(1)}) = D(p_B^{(1)}) + q. \tag{3.26}$$

Using the fact that $S(p_B^{(1)}) = S(p_a) + V_S(p_B^{(1)})$ and $D(p_B^{(1)}) = D(p_a) - V_B(p_a)$ we obtain

$$S(p_a) = D(p_a) - V_B(p_a) + q, \tag{3.27}$$

hence, using equation (3.24), one finds

$$V_S^R(p_a) = V_B^R(p_a) - V_B(p_a) + q. \tag{3.28}$$

Using $V_B(p_a) = V_B^R(p_a) + V_B^M(p_a)$, we obtain

$$q = V_S^R(p_a) + V_B^M(p_a), \tag{3.29}$$

which yields

$$\omega_B^{(0)} = \frac{1}{Q_a}\left(V_S^R(p_a) + V_B^M(p_a)\right) \tag{3.30}$$

Let us now determine $\omega_B^{(i)}$, $i \geq 1$: which is the $(i+1)^{\text{th}}$ point of discontinuity of I_B. We proceed similarly: the $(i+1)^{\text{th}}$ price change due to the injection of a market order occurs when $V_S^R(\omega) = V_S(p_\omega)$, $V_B^R(\omega) = 0$, and $p_\omega = p_B^{(i+1)}$ is the $(i+1)^{\text{th}}$ non empty price tick greater than p_a (in the sense of $V_S + V_B$). Equation (3.25) yields

$$S(p_B^{(i+1)}) - V_S(p_B^{(i+1)}) = D(p_B^{(i+1)}) + q, \tag{3.31}$$

$$\sum_{p' < p_B^{(i+1)}} V_S(p') = \sum_{p' \geq p_B^{(i+1)}} V_B(p') + q, \tag{3.32}$$

$$S(p_a) + \sum_{p_a < p' < p_B^{(i+1)}} V_S(p') = D(p_a) - \sum_{p_a \leq p' < p_B^{(i+1)}} V_B(p') + q. \tag{3.33}$$

Using equation (3.24), we obtain

$$Q_a + V_S^R(p_a) + \sum_{p_a < p' < p_B^{(i+1)}} V_S(p') = Q_a + V_B^R(p_a) - (V_B(p_a) + \sum_{p_a < p' < p_B^{(i+1)}} V_B(p')) + q. \tag{3.34}$$

Finally,

$$\sum_{p_a < p' < p_B^{(i+1)}} (V_S + V_B)(p') = q - (V_S^R(p_a) + V_B^M(p_a)). \tag{3.35}$$

Thus,

$$\omega_B^{(i)} = \omega_B^{(0)} + \frac{1}{Q_a} \sum_{p_a < p' < p_B^{(i+1)}} (V_S + V_B)(p') \quad , \quad i \geq 1, \tag{3.36}$$

which leads to

$$\omega_B^{(i)} = \omega_B^{(i-1)} + \frac{V_S(p_B^{(i)}) + V_B(p_B^{(i)})}{Q_a} \quad , \quad i \geq 1. \tag{3.37}$$

□

3.B Proof of Proposition 3

Proof. Using Proposition 2 we have

$$\begin{aligned}
\omega_B^{(i)} - \omega_B^{(0)} &= \frac{1}{Q_a} \sum_{p_a < p' < p_B^{(i+1)}} V_S(p') + V_B(p') \\
&= \frac{1}{Q_a} \sum_{k=1}^{i} (V_S + V_B)(p_B^{(k)}) \\
&= \sum_{k=1}^{i} (p_B^{(k+1)} - p_B^{(k)}) \left(\tilde{\rho}_S + \tilde{\rho}_B\right)(p_B^{(k)}) \\
&= \widetilde{\mathscr{L}}_B \sum_{k=1}^{i} (p_B^{(k+1)} - p_B^{(k)}) \\
&= \widetilde{\mathscr{L}}_B (p_B^{(i+1)} - p_B^{(1)}) \\
&\approx \widetilde{\mathscr{L}}_B p_B^{(1)} \left[I_B\left(\omega_B^{(i)}\right) - I_B\left(\omega_B^{(0)}\right) \right],
\end{aligned} \tag{3.38}$$

where we used the approximation $I_B\left(\omega_B^{(i)}\right) - I_B\left(\omega_B^{(0)}\right) = \log(p_B^{(i+1)}/p_B^{(1)}) \approx p_B^{(i+1)}/p_B^{(1)} - 1$.

□

3.C Empirical properties of impact slopes at auction time

In this appendix, we report empirical observations on the impact slope at auction time on day d defined as

$$\widetilde{S}_d = (p^{(1)} \widetilde{\mathscr{L}})^{-1}. \tag{3.39}$$

Figure 3.C.1 plots $\widetilde{S}_d$ for TotalEnergies as a function of time. It oscillates around a typical value and has a positive autocorrelation over a few days. The distribution of $\widetilde{S}_d$ for the 34 stocks is reported in Fig. 3.C.2: while its shape is similar for all the assets, its parameters depend on each stock.

We also report the distribution of the absolute value log-changes of the slopes in Fig. 3.C.3, which clearly appear to be exponentially distributed. Its one-step autocorrelation is negative.

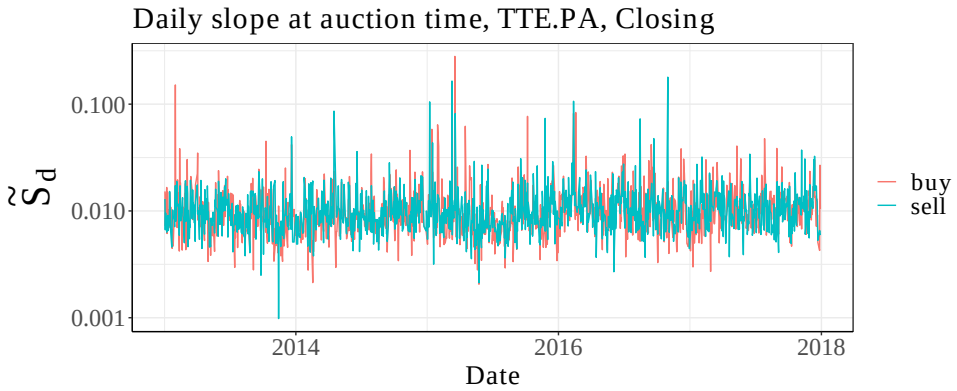

Figure 3.C.1 – Daily impact slope at the closing auction time for TotalEnergies stock.

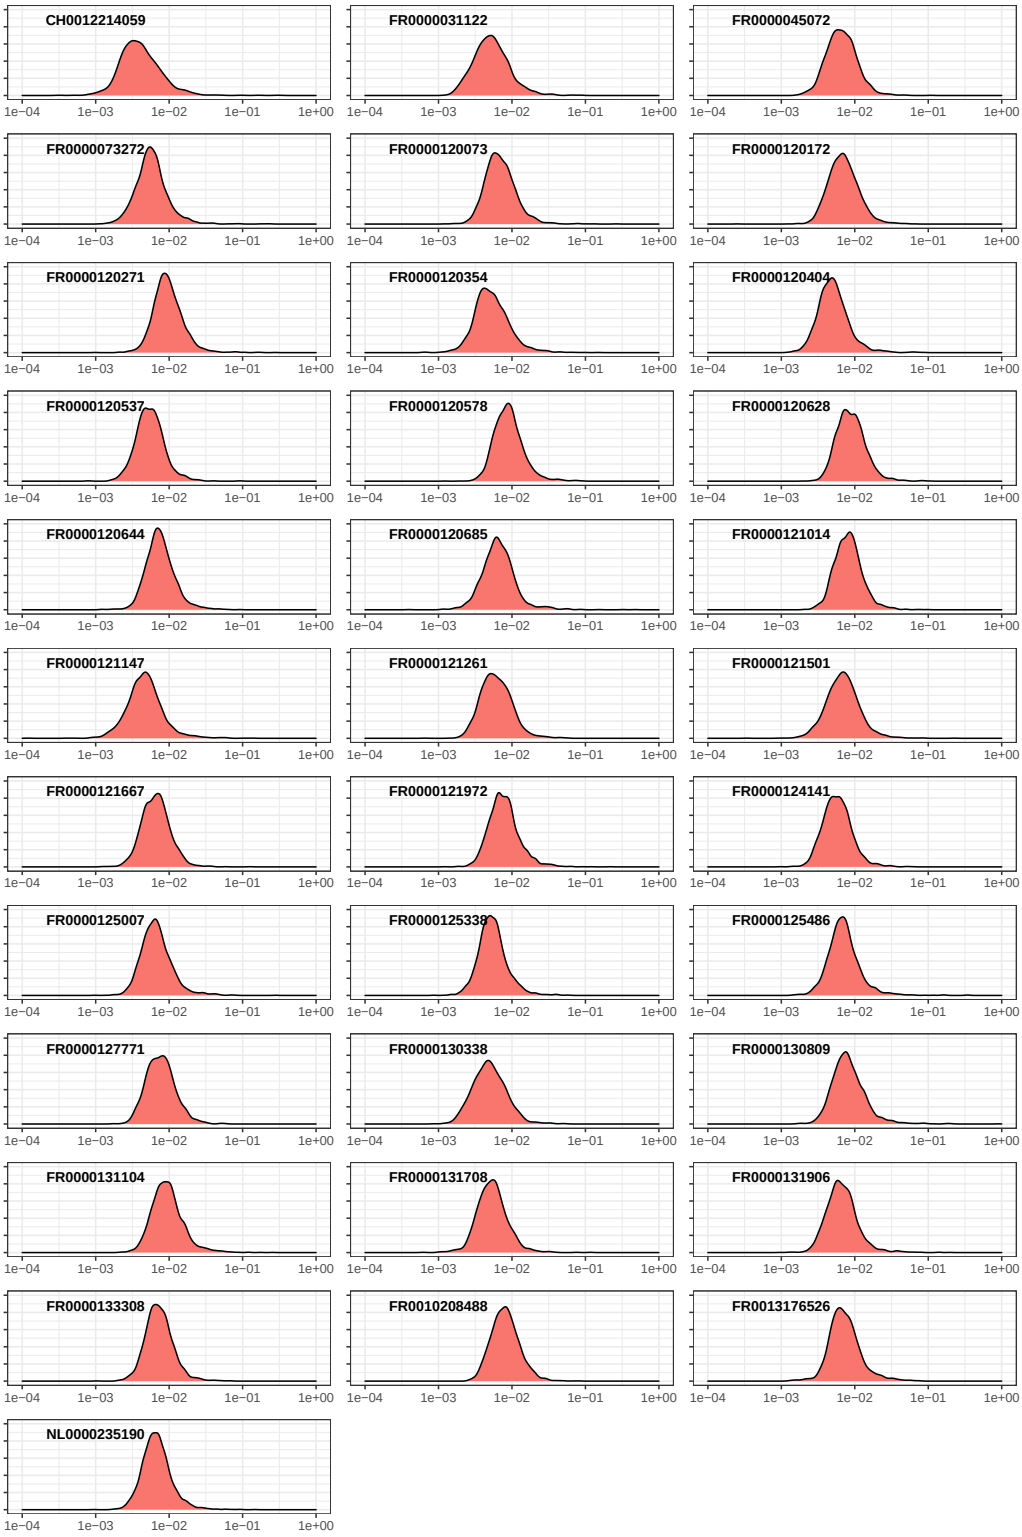

Figure 3.C.2 – Kernel density of the impact slope at the closing auction time for the 34 studied assets.

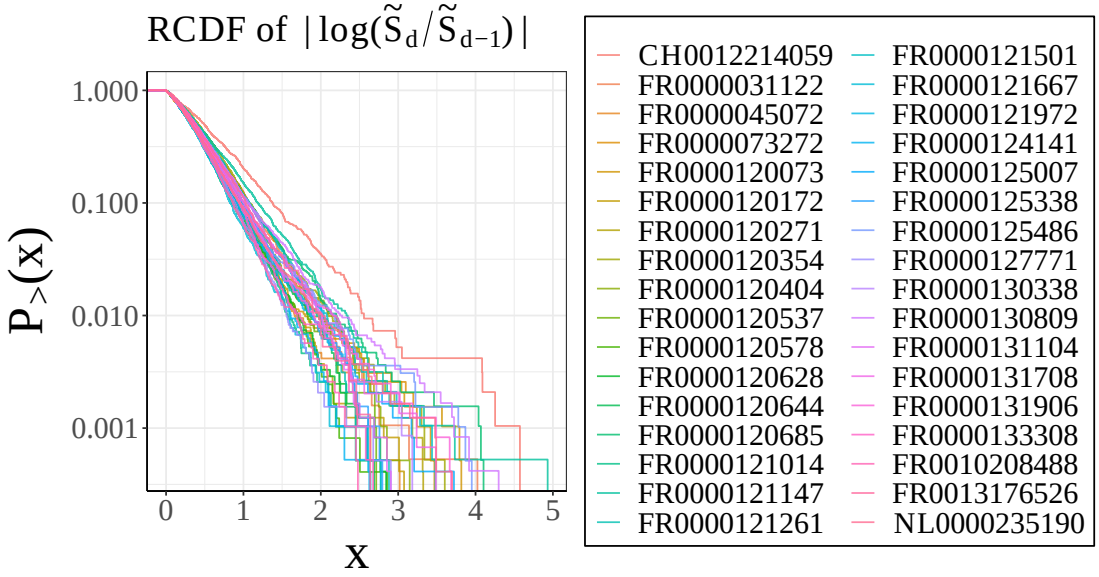

Figure 3.C.3 – Reverse cumulative distribution function of the absolute change in the logarithm of the slope.

Chapter 4

Equity auction dynamics: latent liquidity models with time acceleration

Based on Salek et al. (2024)
by Mohammed Salek, Damien Challet, and Ioane Muni Toke;
To appear in Quantitative Finance.

Abstract

Equity auctions display several distinctive characteristics in contrast to continuous trading. As the auction time approaches, the rate of events accelerates causing a substantial liquidity buildup around the indicative price. This, in turn, results in a reduced price impact and decreased volatility of the indicative price. In this study, we adapt the latent/revealed order book framework to the specifics of equity auctions. We provide precise measurements of the model parameters, including order submissions, cancellations, and diffusion rates. Our setup allows us to describe the full dynamics of the average order book during closing auctions in Euronext Paris. These findings support the relevance of the latent liquidity framework in describing limit order book dynamics. Lastly, we analyze the factors contributing to a sub-diffusive indicative price and demonstrate the absence of indicative price predictability.

Contents of this Chapter

4.1 Introduction

Auctions play an essential role in modern equity markets, facilitating the matching of large volumes at a single price. Major equity exchanges such as Euronext (Paris), Xetra (Germany), and LSE (UK) among others, start and end trading days with an auction to set opening and closing prices for liquid stocks. The matched volume in the closing auction is a significant portion of the daily exchanged volume. This closing volume has seen consistent growth, particularly in European markets, where it can surpass half the daily volume on days of index rebalancing and derivative expiry (Raillon, 2020). To enhance liquidity during midday, typically the period of lowest liquidity, and to mitigate the exposure to high-frequency market makers which dominate the total turnover during day trading (AMF, 2017), some primary exchanges such as Xetra and LSE, also introduced a daily intraday auction. For less liquid stocks in Paris, Euronext (2019) implements the double-fixing trading method, i.e., two auctions per day and no continuous double auction. Some researchers advocate for periodic batch auctions as an alternative market design to continuous trading (Paul et al., 2021; Derchu et al., 2020; Budish et al., 2015). In 2015, CBOE (2020) implemented periodic batch auctions for liquid European stocks.

During the so-called auction accumulation period, auction limit and market orders can be sent, modified, or canceled but no transaction occurs. At all times, an indicative price can be computed: it maximizes the matched volume (indicative volume) and minimizes the remaining order imbalance at the indicative price (surplus). Depending on the exchange, the order book data can be fully open: opening and closing auctions in the London Stock Exchange (2018), or partially opaque: the exchange only disseminates the indicative price, volume, the surplus, and its side as in Xetra (2021). At the auction time, the exchange clears all matched orders at the auction price. Unlike US equity auctions where continuous trading runs in parallel, continuous trading is halted during accumulation periods in European equity markets.

Auctions are found empirically to improve the overall price formation process (Pagano and Schwartz, 2003), even when the market is fragmented (Boussetta et al., 2017). Similarly, Besson and Fernandez (2021) find that the cost of trading during closing auctions in European markets is reduced by a factor two compared to open markets. Challet and Gourianov (2018) examine US opening and closing auctions and find markedly different auction response functions for each auction. Challet (2019) shows that as the auction clearing approaches, the indicative price volatility decreases while the rate of events accelerates. Salek et al. (2023) analyze equity auctions on Euronext Paris and find that price impact at the auction time is first zero due to large limit orders that are present at the auction price, then linear for most days.

While there is an extensive literature on the microstructure modeling of open markets (Chakraborti et al., 2011a,b; Bouchaud et al., 2018a; Lehalle and Laruelle, 2018), auction-specific models are relatively rare (Derksen et al., 2020; Mendelson, 1982; Muni Toke, 2015b). The zero-intelligence model of Donier and Bouchaud (2016) is a promising framework for Walrasian auctions; this model assumes the existence of a latent (hidden) limit order book containing all buy and sell intentions that may be partially revealed in the visible limit order book, building on Tóth et al. (2011). Remarkably, the latent order book model is able to reproduce the shape of market impact with minimal ingredients.

Latent order book models have been the subject of extensive research (Donier et al., 2015; Lemhadri, 2019; Benzaquen and Bouchaud, 2018a,b; Mastromatteo et al., 2014). Recently, Dall'Amico et al. (2019) introduced a conversion mechanism between the latent and revealed order books and fitted the resulting model to market data convincingly. In this work, we specifically adapt the framework of Dall'Amico et al. (2019) to equity auctions, allowing for more generic dependencies of the model parameters on price and time. In addition to stationary solutions, we demonstrate that general solutions can be obtained in closed form when diffusion is negligible. When diffusion is not negligible, numerical schemes are used to fit the full auction dynamics. Thanks to high-quality tick-by-tick data, we meticulously measure the price and time dependencies of the submission, cancellation, and diffusion rates within our framework. We show that the time acceleration near the clearing can be achieved using assumptions similar to those introduced by Alfi et al. (2009, 2007) which describe human behavior when faced with a deadline.

Our main findings are as follows:

1. Calibrating the average book at auction time under our framework requires only a few essential parameters. This provides estimates of the involved latent liquidity.
2. The submission rate is found to be an exponentially decreasing function of the distance to the indicative price, and inversely proportional to the remaining time to the auction

clearing, both close enough to the indicative price.

3. The cancellation rate is predominantly influenced by high-frequency agents. It is a decreasing time function at the start of the auction, then inversely proportional to the remaining time to the auction clearing around the indicative price.
4. The revealed diffusion coefficient is primarily driven by the indicative price volatility, with price reassessments being, on average, negligible.
5. Despite the indicative price being sub-diffusive, it is nevertheless efficient in the sense of Chen et al. (2017).

4.2 Modeling the auction book

4.2.1 Model description

Our starting point is the coupled reaction-diffusion equations for the revealed $\rho^{(r)}$ and the latent $\rho^{(l)}$ orders densities derived in Dall'Amico et al. (2019): they posit the existence of interactions between the latent order book $\rho^{(l)}$ and the revealed order book $\rho^{(r)}$. The reveal rate $\nu_r\Gamma_r$ is the rate at which latent trading intentions in the latent order book materialize into actual orders in the revealed order book and the unreveal rate $\nu_l\Gamma_l$ is the rate at which actual orders revert back to latent intentions. In addition, a diffusion mechanism is included in both the revealed and the latent order book corresponding to price updates in each order book with diffusion coefficient D_r and D_l, respectively.

Assuming the absence of exogenous depositions and cancellations in the latent order book, an order submission is equivalent to an order revelation, and similarly, an order cancellation is equivalent to an order unrevelation. By allowing for general price and time dependencies, the buy side equations read

$$\begin{cases} \partial_t \rho_B^{(r)} &= D_r \partial_{xx} \rho_B^{(r)} + (\nu_r \Gamma_r)(-x,t)\rho_B^{(l)} - (\nu_l \Gamma_l)(-x,t)\rho_B^{(r)}; \\ \partial_t \rho_B^{(l)} &= D_l \partial_{xx} \rho_B^{(l)} - (\nu_r \Gamma_r)(-x,t)\rho_B^{(l)} + (\nu_l \Gamma_l)(-x,t)\rho_B^{(r)}, \end{cases} \tag{4.1}$$

where $x = \log(p/p_t)$ is the centered log price around the log indicative price $\log(p_t)$. For the sell side, we have

$$\begin{cases} \partial_t \rho_S^{(r)} &= D_r \partial_{xx} \rho_S^{(r)} + (\nu_r \Gamma_r)(x,t)\rho_S^{(l)} - (\nu_l \Gamma_l)(x,t)\rho_S^{(r)}; \\ \partial_t \rho_S^{(l)} &= D_l \partial_{xx} \rho_S^{(l)} - (\nu_r \Gamma_r)(x,t)\rho_S^{(l)} + (\nu_l \Gamma_l)(x,t)\rho_S^{(r)}. \end{cases} \tag{4.2}$$

We complement Eqs (4.1) and (4.2) with a set of boundary conditions. Let us focus on the sell side for the time being. It makes sense to assume that the latent liquidity book $\rho_S^{(l)}$ is an increasing function of x, as more people are willing to sell assets at a higher price. This means that one can impose a boundary condition on the slope of the latent order book: $\partial_x \rho_S^{(l)} \underset{x \to +\infty}{\longrightarrow} a > 0$, which corresponds to the latent liquidity parameter of Donier et al. (2015), and that $\rho_S^{(l)}$ does not diverge when $x \to -\infty$, which reflects the fact that the number of people willing to sell at a vanishing price is not infinite. We also impose that the revealed order book $\rho_S^{(r)}$ does not diverge for large prices $|x| \to +\infty$. This means that agents tend not to reveal their reservation price when it is far away from the indicative price. In practice, agents can send market orders or matchable limit orders far away from the indicative price in order to guarantee their participation in the auction volume. Market orders are not included in the densities $\rho_S^{(r)}$ and $\rho_B^{(r)}$ as latent liquidity models are defined in the reference frame of the indicative price.

Similar boundary conditions hold for the buy side.

In the original framework, $\nu_l = \nu_r$ is a constant rate, and Γ_r (resp. Γ_l) is conceived as a probability function of the relative price x for revealing a latent intention (resp. unrevealing a public intention), with $0 \leq \Gamma \leq 1$. In the context of equity auctions, the remaining time to the auction clearing plays an essential role Challet (2019). Thus, we allow for the quantities $\nu_r \Gamma_r$ and $\nu_l \Gamma_l$ to depend jointly on x and t. However, we will often posit that the variable separation is possible $(\nu\Gamma)(x,t) = \nu(t)\Gamma(x)$ for the sake of analytical tractability and interpretability.

In addition, the initial model comprises a reaction term that is formally written $\kappa R_{SB} = \kappa \rho_S^{(r)} \rho_B^{(r)}$, with $\kappa \to +\infty$ in order to make both densities interact and transactions happen; for sufficiently large κ, no overlap between the buy and sell densities is possible. In Walrasian auctions instead, no transaction takes place before the clearing, and buy and sell limit orders usually overlap. When the order book is partially opaque during the accumulation period, it is reasonable to believe that order densities interact solely through the knowledge of the indicative price. Therefore, when considered in the reference frame of the indicative price, buy and sell order densities should evolve independently, which leads us to set $\kappa = 0$.

Lastly, in open markets, the current price p_t is the point where vanishing supply meets vanishing demand; accordingly, Dall'Amico et al. (2019) define p_t as the point where $p \to \left(\rho_B^{(r)} - \rho_S^{(r)}\right)(p)$ changes sign. In auctions instead, the indicative price p_t is determined by equalizing supply and demand

$$S(p_t) + MO_{S,t} = D(p_t) + MO_{B,t}, \tag{4.3}$$

where $S(p) = \int_{p'<p} \rho_S^{(r)}(p')\mathrm{d}p'$, $D(p) = \int_{p'>p} \rho_B^{(r)}(p')\mathrm{d}p'$, and MO_B and MO_S is the volume of buy and sell auction market orders. Equation (4.3) makes it clear that only the difference of market order volume has an influence on the indicative price. As latent liquidity models are defined in the reference frame of the indicative price, market orders are not included in the densities $\rho_S^{(r)}$ and $\rho_B^{(r)}$.

When centering equations around $\log(p_t)$, we implicitly assume that the indicative price evolves independently of order densities. However, changes in the sell (resp. buy) order density for $x < 0$ (resp. $x > 0$) directly influence supply (resp. demand) subsequently changing the indicative price. We disregard this effect and consider that the dynamics of the indicative price are independent of average order densities around the indicative price. In the following, we focus on the sell-side equations (4.2), and denote ρ_S as ρ when there is no ambiguity.

4.2.2 A stationary solution

The simplest stationary solution of Eqs (4.2) is obtained when both $\nu_l\Gamma_l$ and $\nu_r\Gamma_r$ are time-independent and there is no diffusion in the revealed order book ($D_r = 0$), while orders may diffuse in the latent order book ($D_l \geq 0$). Eqs (4.2) reduce to

$$\begin{cases} 0 &= \nu_r\Gamma_r\rho^{(l)} - \nu_l\Gamma_l\rho^{(r)}, \\ 0 &= D_l\partial_{xx}\rho^{(l)} - \nu_r\Gamma_r\rho^{(l)} + \nu_l\Gamma_l\rho^{(r)}, \end{cases} \tag{4.4}$$

which further simplifies to

$$\begin{aligned} \rho^{(r)} &= \frac{\nu_r}{\nu_l} \cdot \frac{\Gamma_r}{\Gamma_l} \cdot \rho^{(l)} \\ D_l\partial_{xx}\rho^{(l)} &= 0. \end{aligned} \tag{4.5}$$

We solve Eq. (4.5) separately on $\mathbb{R}^+$ and $\mathbb{R}^-$ as the first derivative of $x \to \Gamma_{r/l}(x)$ might not be continuous at $x = 0$. Thus, whenever $D_l \neq 0$, the latent order book should be linear. Incorporating the boundary conditions $\partial_x\rho^{(l)} \underset{x\to+\infty}{\longrightarrow} a > 0$ and $\rho^{(l)} \underset{x\to-\infty}{\longrightarrow} b \geq 0$ yields

$$\rho^{(l)}(x) = \max(ax + b, b). \tag{4.6}$$

Subsequently, the stationary revealed order book is

$$\rho^{(r)}(x) = \frac{\nu_r}{\nu_l} \cdot \frac{\Gamma_r(x)}{\Gamma_l(x)} \cdot \max(ax + b, b). \tag{4.7}$$

As the revealed order book does not diverge for large prices, Eq. (4.7) imposes that Γ_r decays faster than Γ_l for $x \to -\infty$, and that $x\Gamma_r$ decays faster than Γ_l for $x \to +\infty$. For instance, this condition is satisfied by an exponentially decaying submission rate and a constant cancellation rate.

4.2.3 Dynamic solutions

Here, we derive non-stationary solutions of Eqs (4.2) in several cases. First, we suppose that diffusion is negligible in both latent and revealed order books ($D_r = D_l = 0$). This is a sound approximation as the calibrated orders of magnitude of D_l and D_r do not significantly influence the order book shape (see section 4.3.3).

Summing Eqs (4.2) yields

$$\partial_t(\rho^{(l)} + \rho^{(r)}) = 0, \tag{4.8}$$

which suggests defining the total density $\rho^{\Sigma} = \rho^{(l)}(x,t) + \rho^{(r)}(x,t)$. Thus, Eq. (4.8) implies that

$$\rho^{\Sigma}(x,t) = \rho^{\Sigma}(x,t=0). \tag{4.9}$$

Now, we posit the initial condition $\rho^{(r)}(x, t = 0) = 0$, which is a reasonable approximation because the revealed order book at the beginning of the auction is negligible in comparison with the final auction book. Additionally, we set $\rho^{(l)}(x, t = 0) = \max(ax + b, b)$ to satisfy the latent order book boundary conditions. Making these substitutions, we find that Eqs (4.2) can be decoupled into

$$\partial_t \rho^{(r)} + (\nu_r \Gamma_r + \nu_l \Gamma_l) \cdot \rho^{(r)} = \nu_r \Gamma_r \rho^{\Sigma}; \tag{4.10}$$

$$\partial_t \rho^{(l)} + (\nu_r \Gamma_r + \nu_l \Gamma_l) \cdot \rho^{(l)} = \nu_l \Gamma_l \rho^{\Sigma}, \tag{4.11}$$

where $\rho^{\Sigma} = \max(ax + b, b)$.

Time-independent rates

When $\nu_l\Gamma_l$ and $\nu_r\Gamma_r$ do not exhibit temporal dependencies, i.e., $(\nu\Gamma)(x,t) = \nu \cdot \Gamma(x)$, Eqs (4.11) yields the following expression for the revealed order book

$$\rho^{(r)}(x,t) = \rho_{\infty} \cdot \left[1 - e^{-(\nu_r\Gamma_r + \nu_l\Gamma_l)\cdot t}\right], \tag{4.12}$$

where $\rho_\infty = (\nu_r\Gamma_r \cdot \rho^\Sigma)/(\nu_r\Gamma_r + \nu_l\Gamma_l)$. The obtained solution converges to ρ_∞ in the long run, i.e., as $t \longrightarrow +\infty$. This solution replicates the dynamics of the revealed order book during auction phases when rates $\nu_l\Gamma_l$ and $\nu_r\Gamma_r$ can be considered time-independent. Moreover, it is valid for order book portions that are already in a stationary state. Finally when $(\nu_r\Gamma_r + \nu_l\Gamma_l) \cdot t \ll 1$, $\rho^{(r)}$ no longer depends on $\nu_l\Gamma_l$ anymore. This occurs for large values of $|x|$ where $\nu_r\Gamma_r + \nu_l\Gamma_l$ goes to zero while t remains bounded

$$\rho^{(r)}(x,t) \underset{|x|\gg 1}{\sim} \nu_r \cdot \Gamma_r(x) \cdot \rho^\Sigma(x) \cdot t. \tag{4.13}$$

Eq. (4.13) demonstrates that $\rho^{(r)}$ does not diverge when $|x| \to +\infty$ even when ρ_∞ does, e.g. when Γ_l decays faster than Γ_r. Note that this simple framework can not reproduce the accelerating auction dynamics as the clearing approaches (Challet, 2019).

Time-dependent rates

To replicate the accelerating order book activity as the auction time approaches, which usually results in a convex $\rho^{(r)}$ with respect to t as $t \longrightarrow T$, we need to introduce a pressure from the auction deadline T. The auction deadline is the final time that ensures trading with certainty. For liquid stocks listed in Euronext, $T =$ 17:35:00 for the closing auction, and 09:00:00 for the opening auction. The clearing randomly occurs in a thirty-second window starting at T. [1]

Drawing inspiration from Alfi et al. (2007, 2009), who argue that the probability of registering at a conference is inversely proportional to the remaining time to the registration deadline, we posit that the submission rate $\nu_r\Gamma_r$ (resp. the cancellation rate $\nu_l\Gamma_l$) is time-dependent for $t \geq t_r^{(0)}$ and constant for $t \leq t_r^{(0)}$ (resp. $t_l^{(0)}$), where $t_{r/l}^{(0)}$ is a cut-off time in $]0, T[$. More precisely, we set

$$\begin{aligned} (\nu_r\Gamma_r)(x,t) &= \frac{C_r}{\gamma_r + T - t} \cdot \Gamma_r(x), \qquad t \geq t_r^{(0)}; \\ (\nu_l\Gamma_l)(x,t) &= \frac{C_l}{\gamma_l + T - t} \cdot \Gamma_l(x), \qquad t \geq t_l^{(0)}, \end{aligned} \tag{4.14}$$

with $C_r, C_l > 0$, $\gamma_l, \gamma_r \geq 0$. A strictly positive γ_r (resp. γ_l) indicates that the perceived deadline for submitting (resp. canceling) limit orders around the indicative price is $T + \gamma_r$ (resp. $T + \gamma_l$).

Assuming that $\gamma_r = \gamma_l = \gamma$ and $t_r^{(0)} = t_l^{(0)} = t^{(0)}$, we substitute $\nu_r\Gamma_r$ and $\nu_l\Gamma_l$ of Eqs (4.14)

1. In the 28th of September 2015, Euronext introduced a random clearing window of thirty-second length for its equity auctions. The clearing randomly happens between 09:00:00 and 09:00:30 for the opening auction and 17:35:00 and 17:35:30 for the closing auction. This prevents fast agents from using low latency to take advantage of slower agents.

into the first equation of (4.11) and obtain for $t \geq t^{(0)}$

$$\partial_t \rho^{(r)} + \frac{C_r\Gamma_r + C_l\Gamma_l}{\gamma + T - t} \cdot \rho^{(r)} = \frac{C_r\Gamma_r}{\gamma + T - t} \cdot \rho^{\Sigma}, \tag{4.15}$$

whose solution is given by

$$\rho^{(r)}(x,t) = \rho_T - (\rho_T - \rho_0) \cdot \left(\frac{\gamma + T - t}{\gamma + T - t^{(0)}} \right)^{C_r\Gamma_r + C_l\Gamma_l}, \tag{4.16}$$

where $\rho_T = C_r\Gamma_r\rho^{\Sigma}/(C_r\Gamma_r + C_l\Gamma_l)$, and ρ_0 is obtained by substituting $t = t^{(0)}$ in Eq. (4.12).

The functional shape of Eq. (4.16) is convex w.r.t. t for all $0 \leq t \leq T + \gamma$ whenever $C_r\Gamma_r + C_l\Gamma_l < 1$. It converges to a finite solution ρ_T as $t \longrightarrow \gamma + T$. Assuming that Γ_l decays faster than Γ_r yields a divergent ρ_T for large $|x|$ and the solution of Eq. (4.16) becomes incompatible with the boundary conditions of the revealed order book. However, it should be noted that the time dependence of the reveal and unreveal rates in Eq. (4.14) is only valid in a limited region around the indicative price and that for large values of $|x|$ the time dependency vanishes.

Convex solutions of time around the auction deadline can be obtained even if $C_r, C_l > 1$ provided that $\gamma_l > \gamma_r$, i.e., when the perceived deadline of cancellations occurs later than that of submissions. The ordinary differential equation verified by the revealed order book reads

$$\partial_t \rho^{(r)} + \left(\frac{C_l}{\gamma_l + T - t} + \frac{C_r}{\gamma_r + T - t} \right) \cdot \rho^{(r)} = \frac{C_r}{\gamma_r + T - t}. \tag{4.17}$$

Likewise, when the rate of cancellations is constant $(\nu_l\Gamma_l)(x,t) = \nu_l > 0$, the resulting order book dynamic is convex as $t \longrightarrow T$, and the ordinary differential equation reads

$$\partial_t \rho^{(r)} + \left(\nu_l + \frac{C_r}{\gamma_r + T - t} \right) \cdot \rho^{(r)} = \frac{C_r}{\gamma_r + T - t}. \tag{4.18}$$

Eqs (4.17), (4.18) are challenging to solve analytically, and we present numerical solutions in Figure 4.2.1 by varying the parameters of interest.

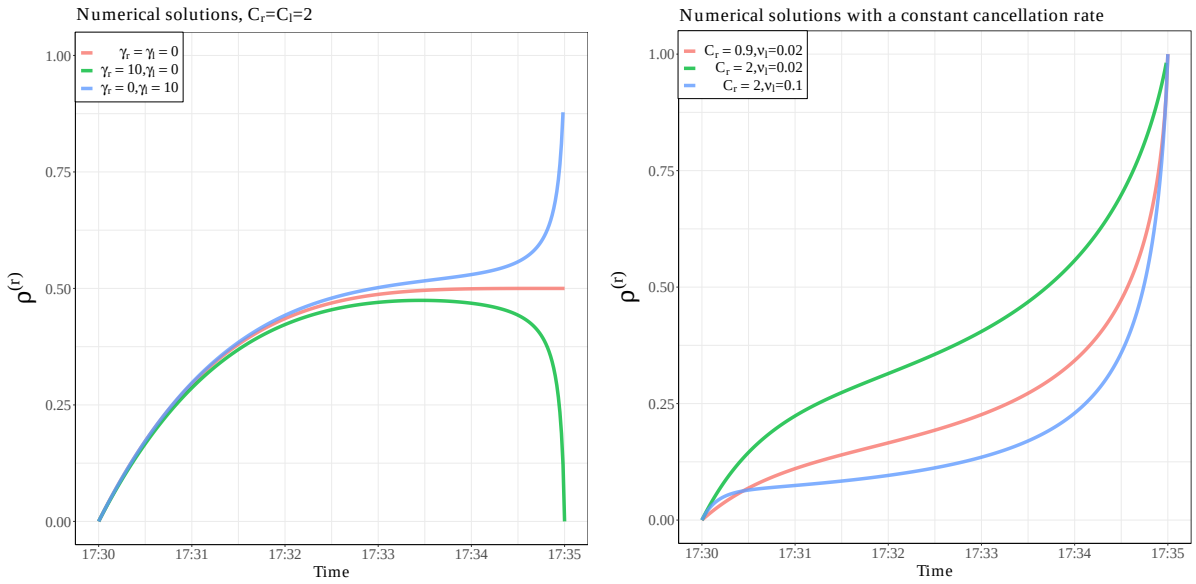

Figure 4.2.1 – Left panel: numerical solutions of Eq. (4.17) for $C_r = C_l = 2$. Right panel: numerical solutions of Eq. (4.18) with a constant cancellation rate ν_l.

Diffusion

The diffusion coefficient is the sum of two contributions: idiosyncratic price reassessments and reactions to changes in the indicative price. We explore the limiting case where latent price reassessments equal the revealed ones $D_r = D_l = D > 0$. Summing once again Eqs (4.2), we obtain

$$\partial_t \left(\rho^{(r)} + \rho^{(l)}\right) = D \cdot \partial_{xx} \left(\rho^{(r)} + \rho^{(l)}\right). \tag{4.19}$$

The solution $\rho^{\Sigma} = \rho^{(r)} + \rho^{(l)}$ of Eq. (4.19) is then

$$\begin{aligned} \rho^{\Sigma}(x,t) &= \frac{1}{\sqrt{4\pi D t}} \int_{\mathbb{R}} \rho^{\Sigma}(y, t=0) \cdot e^{-\frac{(x-y)^2}{4Dt}} \mathrm{d}y; \\ \rho^{\Sigma}(x, t=0) &= \max(ax+b, b). \end{aligned} \tag{4.20}$$

In this scenario, Eqs (4.2) can be decoupled, and order densities $\rho^{(r)}$ and $\rho^{(l)}$ satisfy

$$\begin{cases} \partial_t \rho^{(r)} &= D\partial_{xx}\rho^{(r)} - (\nu_r\Gamma_r + \nu_l\Gamma_l) \cdot \rho^{(r)} + \nu_r\Gamma_r\rho^{\Sigma}; \\ \partial_t \rho^{(l)} &= D\partial_{xx}\rho^{(l)} - (\nu_r\Gamma_r + \nu_l\Gamma_l) \cdot \rho^{(l)} + \nu_l\Gamma_l\rho^{\Sigma}. \end{cases} \tag{4.21}$$

If we further assume that $\nu_r\Gamma_r + \nu_l\Gamma_l$ is constant w.r.t. x and t, the revealed order density can be obtained in a closed-form formula Donier and Bouchaud (2016), and diffusion leads to non-trivial

interactions depending on the shape of the source terms. Outside this specific case ($D_r = D_l$ and $\nu_r\Gamma_r + \nu_l\Gamma_l$ constant), it is challenging to obtain closed-form solutions and we solve the general Eqs (4.2) numerically in section 4.3.3.

4.3 Empirical observations and calibrations

In this section, we confront the model presented in section 4.2 to real auctions in Euronext Paris. Using high quality data from BEDOFIH, we process detailed tick-by-tick closing auction data for five active stocks (TotalEnergies, Sanofi, BNP Paribas, LVMH, Société Générale) in Euronext Paris between 2013 and 2017. First, we reconstruct order book snapshots at the auction time to calibrate the stationary solution of our model. Next, we leverage the level-3 tick-by-tick data to measure the submission, cancellation, and diffusion rates in the revealed order book. Lastly, we reconstruct 1-second successive snapshots to calibrate the full dynamics of the revealed order book during the auction.

4.3.1 Fitting order books at auction time

Using Eq. (4.7) we can fit the empirical order book at auction time $\rho^{(r)}(t = T)$. The empirical order book at auction time is obtained by averaging order book snapshots at the closing auction time across days. We choose the simplest functional forms for the stationary submission and cancellation rates, i.e., an exponentially decreasing revelation probability function $\Gamma_r(x) = e^{-|x|/x_r}$, $x_r > 0$ and a constant unrevelation probability function $\Gamma_l = 1$. Substituting these into Eq (4.7), we fit the empirical order density at auction time. The obtained fits suggest that accuracy is improved by allowing for more than one exponential term. Consequently, we use the following ansatz

$$\rho^{(r)}(x) = \frac{\nu_r}{\nu_l} \cdot \max(ax + b, b) \cdot \left[w \cdot e^{-|x|/x_r} + (1 - w) \cdot e^{-|x|/(k \cdot x_r)} \right], \tag{4.22}$$

where $0 \leq w \leq 1$, $k \geq 1$. This can be interpreted as having two types of agents with different price scales. In a similar framework, Benzaquen and Bouchaud (2018a) find that the typical price scale is proportional to the square root of the typical timescale. Thus, fast agents can be characterized as having a smaller price scale x_r and slow agents as having a larger price scale kx_r.

Note that the parameters a and b in Eq. (4.22) can only be determined up to a factor

Table 4.3.1 – Median value of ν_r/ν_l for $N = 1266$ closing auctions of TotalEnergies between 2013 and 2017.

ν_r/ν_l	Buy side	Sell side
last 1 second	1.17	1.2
last 10 seconds	1.13	1.15
last 30 seconds	1.03	1.05

ν_r/ν_l. Prior to delving into fits, we independently measure ν_r/ν_l as the ratio of the number of submissions to the number of cancellations, in the vicinity of $x = 0$, seconds before the clearing. We report in Table 4.3.1 the median values of this ratio in the last second, 10 seconds, and 30 seconds before 17:35:00 for each side. These results imply a median value around 1.

We now proceed to fitting buy and sell densities at auction time, as well as their breakdown by latency (HFT, MIX, and NON) for each of the five studied stocks. For instance, fitting the HFT order book yields an estimate of the HFT latent book, i.e., the latent book that contains trading intentions of HFT-flagged agents only. This breakdown will prove useful in section 4.3.2 where we measure the contribution by agent category to the global submission rate.

We run each minimization procedure from 18 different initializations in order to avoid local minima and use ordinary least squares to obtain a set of optimal parameters $(\hat{a}, \hat{b}, \widehat{x_r}, \hat{k}, \hat{w})$. We report the optimal estimates for each stock and side in Appendix 4.A, Table 4.A.1, and present fits for TotalEnergies in Fig. 4.3.1. Eq. (4.22) provides accurate fits for $x > 0$. However, it cannot reproduce the oscillations of liquidity at multiples of -0.5%. These fluctuations result from punctual order submissions in these locations that undergo diffusion with coefficient D_r seconds before the clearing.

4.3.2 The empirical dynamics of the auction book

Before providing measurements of the model parameters, we first present an overview of the dynamics of the empirical order density $(x, t) \to \rho^{(r)}(x, t)$. For that purpose, we perform 1-second order book snapshots during the closing auction averaged over days. We present in Fig. 4.3.2 the empirical functions $x \to \rho^{(r)}$ at round minutes and $t \to \rho^{(r)}$ at various prices. We observe that $\rho^{(r)}$ exhibits a skewed shape w.r.t. x, reaching its maximum for $x > 0$ (which corresponds to non-matched orders for both buy and sell sides by convention). Its temporal dependency is initially concave, then accelerates towards the clearing. The time acceleration is more pronounced around the maximum argument and vanishes for large values of $|x|$.

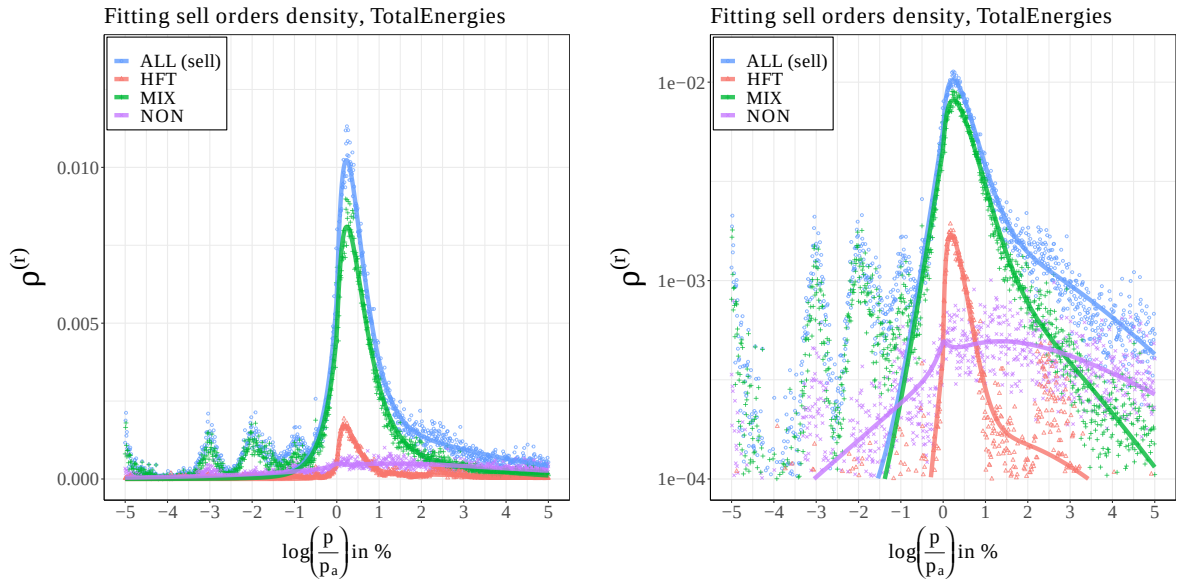

Figure 4.3.1 – Fits (solid lines) with the functional of Eq. (4.22) of the average orders' density $\rho^{(r)}$ just before the close clearing for TotalEnergies between 2013 and 2017. Left panel: ordinary scale. Right panel: Y-axis in log scale. The HFT flag denotes pure high-frequency traders, MIX denotes investment banks with high-frequency trading activities, and NON denotes traders without HFT activities.

Within the latent order book framework, the revealed order density $\rho^{(r)}$ evolves in time due to three mechanisms: submissions from latent order book $\rho^{(l)}$ with a rate $\nu_r\Gamma_r$, cancellations from the revealed book with a rate $\nu_l\Gamma_l$, and diffusion with a coefficient D_r. In this part, we provide empirical measurements of these rates.

The reconstructed tick-by-tick data from BEDOFIH allows us to track changes in quantity for each price level. We view pure price updates without a change in quantity as a diffusion mechanism. Thus, we do not categorize pure price updates as cancellations from the corresponding previous price limits nor as submissions to the new ones.

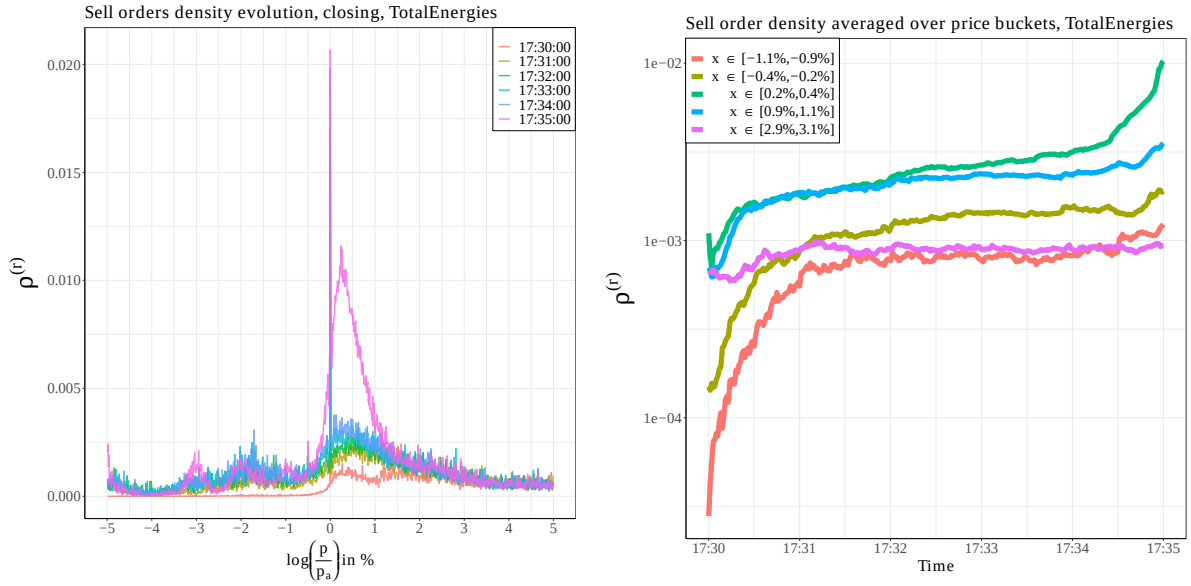

Figure 4.3.2 – Empirical time and price dependencies of the average sell order density $\rho_S^{(r)}$ throughout the closing auction. We averaged auction book data for TotalEnergies.

Measuring the cancellation rate: $\nu_l\Gamma_l$

The cancellation rate $(\nu_l\Gamma_l)(p, t-\delta t \to t)$ for price p between $t-\delta t$ and t is defined as

$$(\nu_l\Gamma_l)(p, t-\delta t \to t) = -\frac{1}{\delta t} \sum_{t-\delta t < t_i < t} \frac{\delta V(p,t_i)}{V_p(t_i)} \cdot \mathbb{1}_{\{\delta V(p,t_i)<0\}}, \tag{4.23}$$

where $\delta V(p,t_i)$ is the volume change at limit price p at time t_i excluding volumes that diffused to another price p', and $V_p(t_i)$ is the total volume at limit price p at time t_i^-.

For each aggregation period $[t-\delta t, t]$ and each asset, we average Eq. (4.23) over all the days in our dataset using equal log price intervals of length $\delta x = 2$ basis points and time step $\delta t = 2$ seconds. Figure 4.3.3 shows a non-trivial behavior of the cancellation rate w.r.t. x and t. At the beginning of the auction, a series of order cancellations yields a decreasing cancellation rate as a function of time. In addition, we have $\nu_l\Gamma_l \propto e^{-|x|}$. As the auction time approaches, $\nu_l\Gamma_l$ reaches a maximum at a value larger than the indicative price. We report in the left panel of Fig. 4.3.4 averaged values of $\nu_l\Gamma_l$ as a function of time over different price buckets in order to display the temporal dependency of cancellations: we confirm the cancellation decrease at the start of the auction, then the increase for prices around indicative price. We verify that $\nu_l\Gamma_l \propto 1/(\gamma_l + T - t)$ for $0 < x < 10$bps and $t > t_l^{(0)}$. To determine $t_l^{(0)}$, we use a change point detection criterion: we assume that when $t < t_l^{(0)}$ a constant model is a better fit than a fit to $t \to 1/(\gamma_l + T - t)$ and

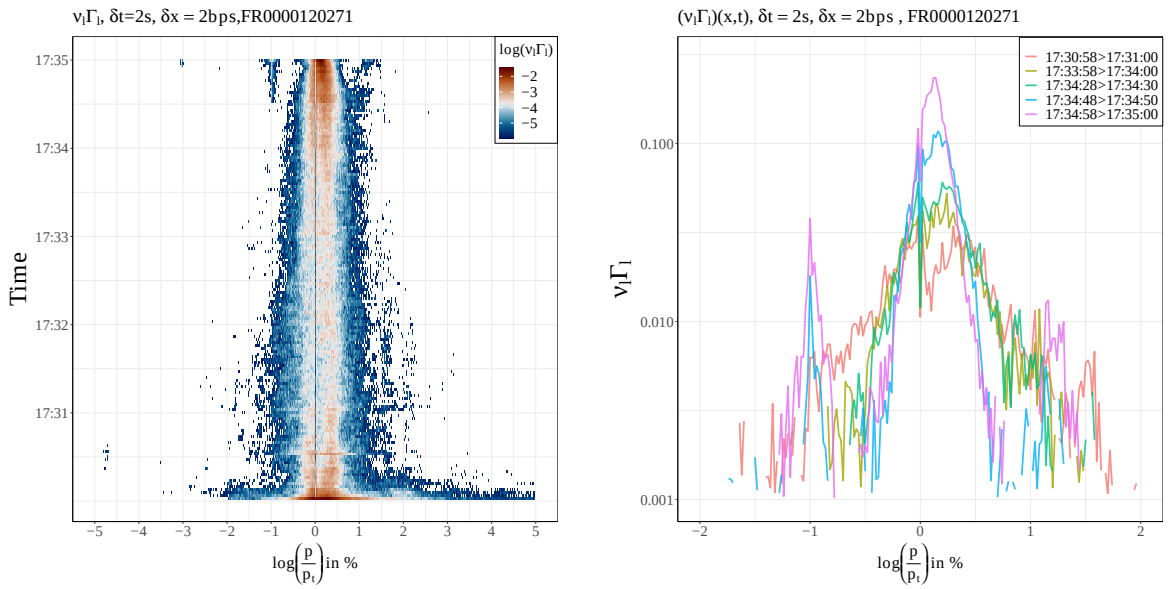

Figure 4.3.3 – Estimating the cancellation rate $\nu_l\Gamma_l$ from tick by tick data. Left panel: heatmap $(x,t) \to \nu_l\Gamma_l$. Right panel: $x \to \nu_l\Gamma_l$ at n-seconds before auction time T =17:35:00, $n \in \{0, 10, 30, 60, 240\}$.

vice versa when $t > t_l^{(0)}$. We exclude the first minute of the auction when cancellations decrease. The right panel of Fig. 4.3.4 displays such a time fit for TotalEnergies.

To examine the variation in the cancellation rates across market participants, we compute the cancellation rate separately for HFT-flagged traders and for non-HFTs. Although MIX includes the high-frequency activity of investment banks, we denote MIX and NON as non-HFTs for simplicity. Figure 4.3.5 reveals that the shape of cancellations for non-HFTs is markedly different from that of HFTs. The cancellation rate of non-HFTs is noisier and is only weakly dependent on the price most of the time, then peaking around the indicative price with exponential decay as the auction time approaches. On the contrary, HFT-flagged traders are the predominant contributors to the cancellation rate, significantly outweighing the cancellations initiated by non-HFTs. In reality, HFTs are highly active in auctions, accounting for more than 80% of all events on average, yet they contribute to only a minor portion of the closing volume, approximately 4% on average.

Lastly, notice a significant peak at $x = -1\%$ indicating a strategic behavior of non-HFTs that are canceling, on average, large matched volumes a few seconds prior to auction time.

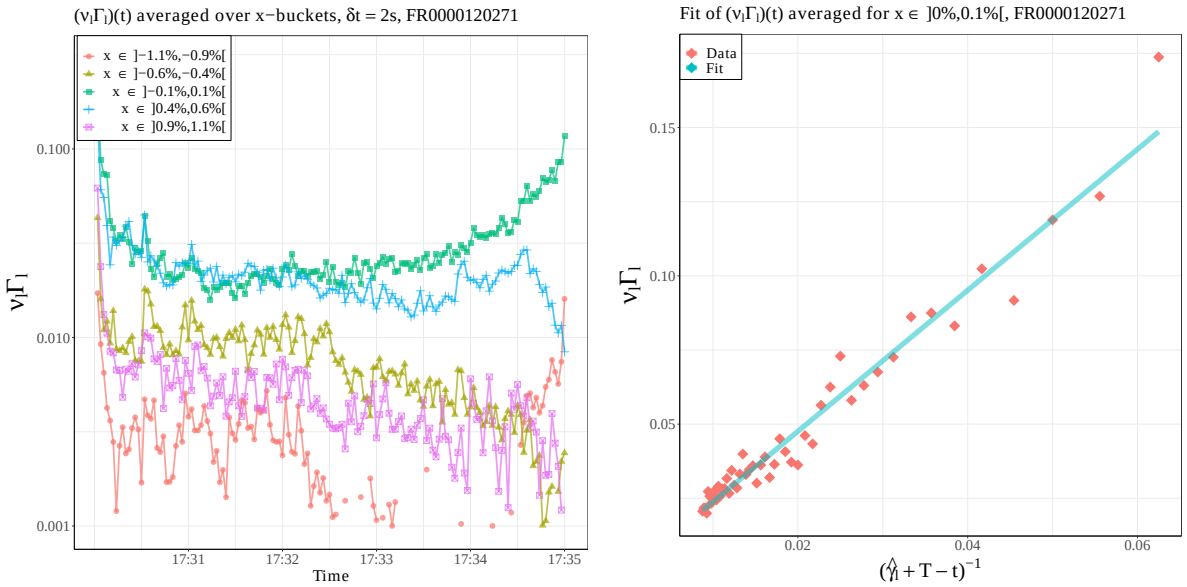

Figure 4.3.4 – Estimating $\nu_l\Gamma_l$. Left panel: $t \to \langle \nu_l\Gamma_l \rangle |_x$ averaged for different price buckets as a function of time. Right panel: average values of $\nu_l\Gamma_l$ for x in $]0, 0.1\%[$ (diamonds) as a function of $(\gamma_l + T - t)^{-1}$ for $t > t_l^{(0)}$. The solid line is a linear fit minimizing the L1 norm of errors. Fitted values: $t_l^{(0)} = 202$, $\hat{\gamma}_l = 16$, $\hat{C}_l = 2.38$.

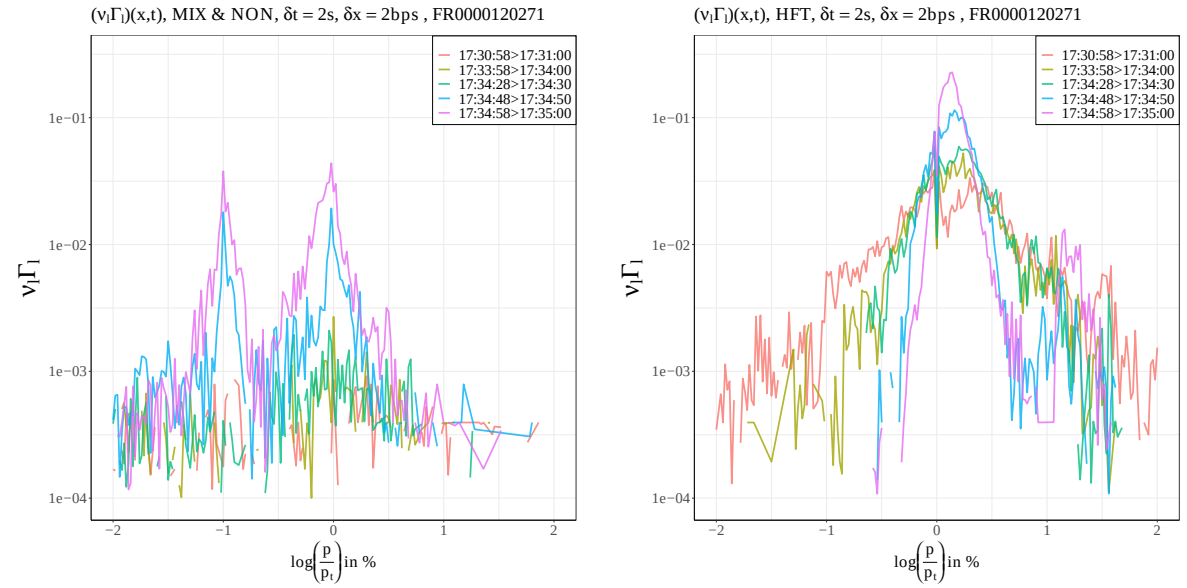

Figure 4.3.5 – Estimating the cancellation rate across market participants. Left panel: MIX & NON contribution. Right panel: HFT contribution.

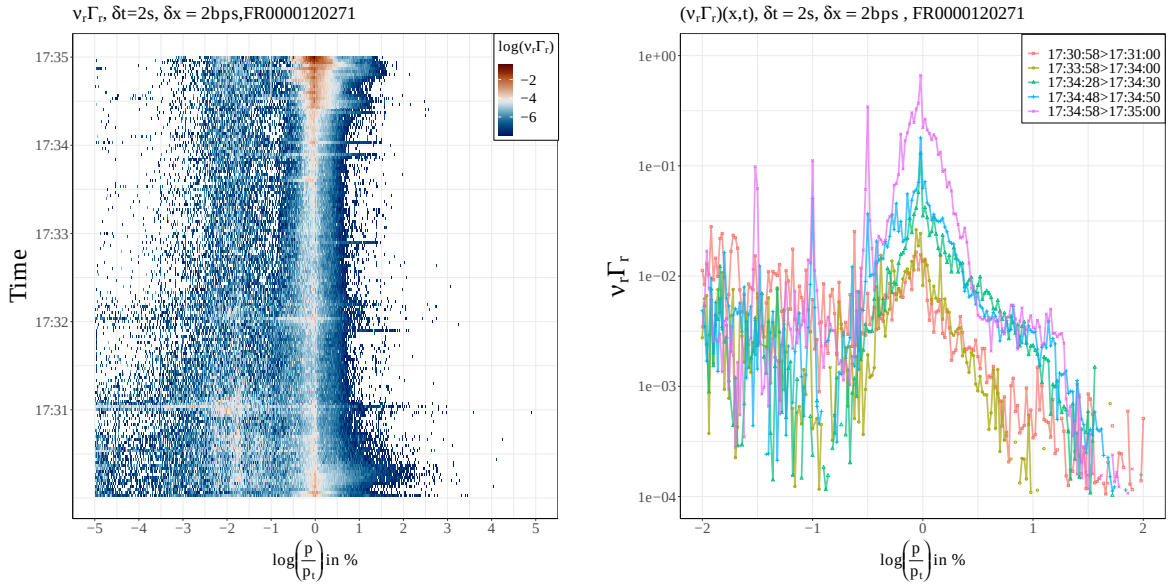

Figure 4.3.6 – Estimating the submission rate $\nu_r\Gamma_r$ from tick by tick data. Left panel: heatmap $(x,t) \to \nu_r\Gamma_r$. Right panel: $x \to \nu_r\Gamma_r$ at n-seconds before auction time T =17:35:00, $n \in \{0, 10, 30, 60, 240\}$.

Measuring the submission rate: $\nu_r\Gamma_r$

The product of the submission rate $\nu_r\Gamma_r$ by the latent order density $\rho^{(l)}$ for price p between $t - \delta t$ and t is defined as

$$(\nu_r\Gamma_r \cdot \rho^{(l)})(p, t - \delta t \to t) = \frac{1}{\delta t} \sum_{t-\delta t<t_i<t} \frac{\delta V(p,t_i)}{\delta x \cdot Q_a} \cdot \mathbb{1}_{\{\delta V(p,t_i)>0\}}, \tag{4.24}$$

where $\delta V(p,t_i)$ is the volume change in limit price p at time t_i excluding volumes that diffused from another price p', and Q_a is the final auction volume: we scale the submitted density $\delta V(p,t_i)/\delta x$ by the final auction volume Q_a to be able to average Eq. (4.24) over different days.

Likewise, we average Eq. (4.24) over days with equal log price intervals of $\delta x = 2$ basis points and time step $\delta t = 2$ seconds. Then, we infer $\nu_r\Gamma_r$ using the numerical estimate of the latent order book from section 4.3.1 where $\rho^{(l)}(x) = \max(\hat{a}, \hat{a}x + \hat{b})$. We implicitly assume that the latent order book remains stable throughout the accumulation period. This is the case when the latent book is significantly larger in comparison with the revealed order book at all times. Figure 4.3.6 shows that $\nu_r\Gamma_r$ displays a clear exponential decay for $x > 0$ with a constant price scale at most times, and a truncated exponential decay for $x < 0$.

We report in the left panel Fig. 4.3.7 averaged values of $\nu_r\Gamma_r$ as a function of time over

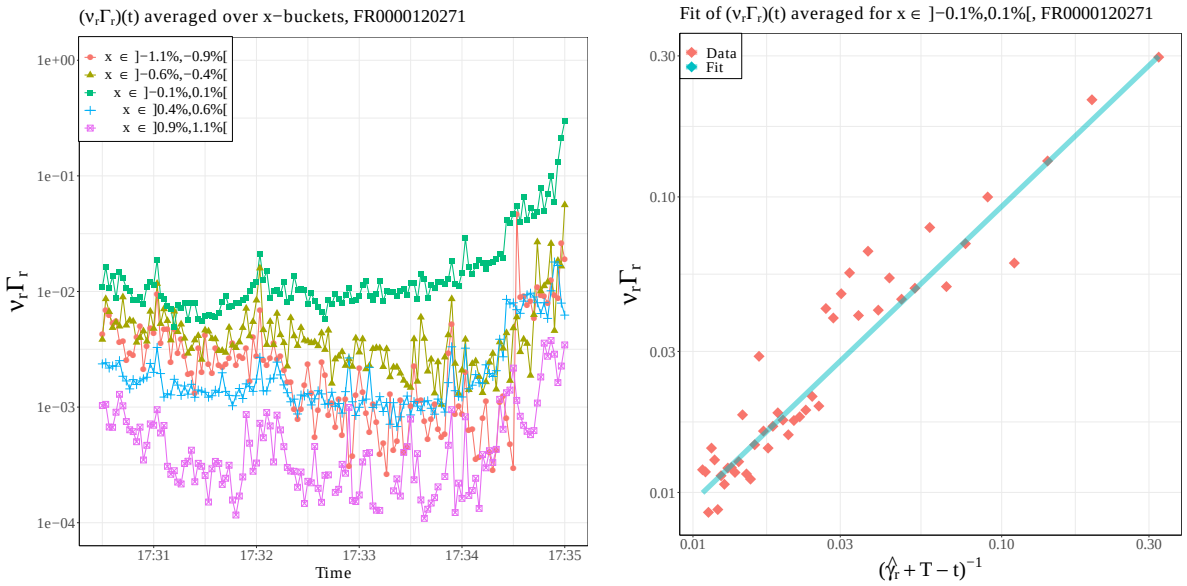

Figure 4.3.7 – Estimating $\nu_r\Gamma_r$. Left panel: $t \to \langle \nu_r\Gamma_r \rangle |_x$ averaged for different price buckets as a function of time. Right panel: average values of $\nu_r\Gamma_r$ for x in $]-0.1, 0.1\%[$ (diamonds) as a function of $(\gamma_r + T - t)^{-1}$ for $t > t_r^{(0)}$. The solid line is a linear fit minimizing the L1 norm of errors. Fitted values: $t_r^{(0)} = 210$, $\hat{\gamma}_r = 3.1$, $\hat{C}_r = 0.93$.

different price buckets to display the temporal dependency of submissions. We empirically verify that $\nu_r\Gamma_r$ grows proportionally to $1/(\gamma_r + T - t)$ for values of $t > t_r^{(0)}$, and determine $t_r^{(0)}$ using a change detection criterion. We report in the right panel of Fig. 4.3.7 such a time fit for TotalEnergies.

Considering HFTs and non-HFTs separately in Fig. 4.3.8, we find a similar shape of submissions for both with a larger price scale for non-HFTs. Note that we inferred each submission rate (HFTs and non-HFTs) from the respective latent book as computed section 4.3.1. Lastly, we notice large peaks at $x = 0$ and at multiples of -0.5% as the clearing approaches: these point to agents trying to pin the current indicative price when they send orders at $x = 0$, or to construct a barrier of matchable orders at a less favorable price when they send them at $x < 0$.

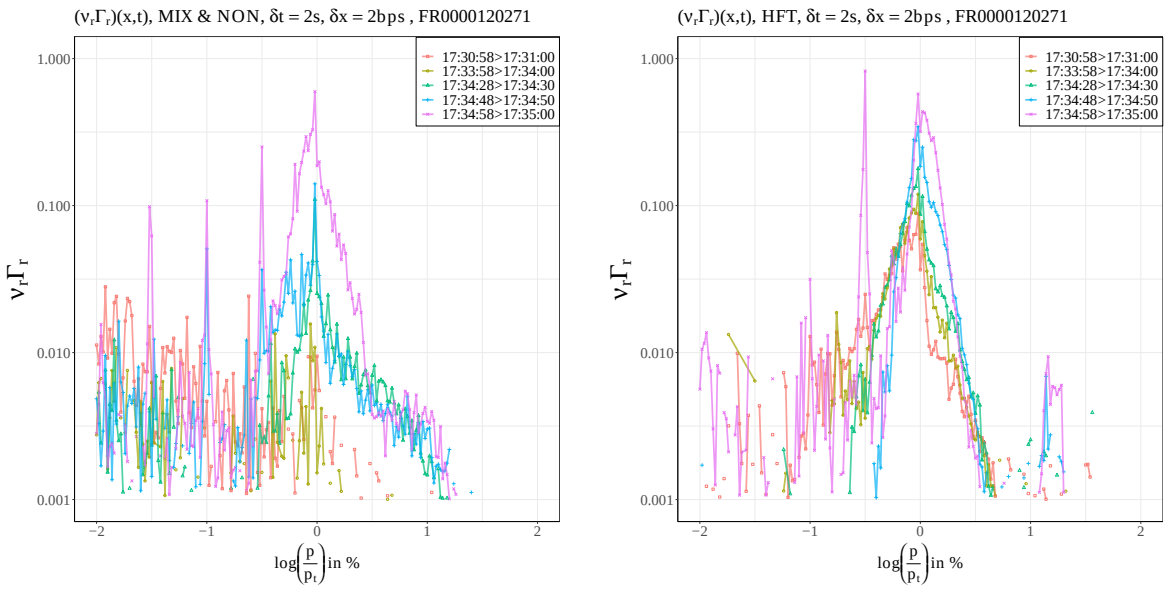

Figure 4.3.8 – Estimating the submission rate across market participants. Left panel: MIX & NON contribution. Right panel: HFT contribution.

Measuring diffusion in the revealed order book: D_r

When prices are diffusive, the diffusion coefficient (Donier and Bouchaud, 2016) $D(x)$ at a price $x \in \mathbb{R}$ is defined as

$$D(x) = \frac{1}{2}\left(\mathrm{Var}(\beta)\cdot\sigma^2 + \int_{\mathbb{R}}(x-y)^2\Gamma_D(x,y)\mathrm{d}y\right), \tag{4.25}$$

where σ is the volatility of the indicative price, $\mathrm{Var}(\beta)$ is a prefactor that encodes the heterogeneity of agent reactions to price movements: should agents neither over-react nor under-react to indicative price movements ($\beta = 1$ for all agents), the first term vanishes; Tóth et al. (2011) consider a unit prefactor $\mathrm{Var}(\beta) = 1$. The other term $\int(x-y)^2\Gamma_D(x,y)\mathrm{d}y$ corresponds to $\mathbb{E}_i[\Sigma_i^2]$ in Donier and Bouchaud (2016) and represents "the purely idiosyncratic noisy updates of agents": here $\Gamma_D(x,y)$ is the rate at which orders are updated from log-price x to log-price y. Historically, the derivation of the diffusion coefficient in latent models was introduced first in Tóth et al. (2011), then developed further in the appendix of Donier et al. (2015).

The definition of Eq. (4.25) corresponds to that of the revealed diffusion coefficient D_r in our framework, provided we only account for visible price reassessments. Note that the under-diffusive nature of the indicative price may result in a different first-term contribution to the diffusion coefficient.

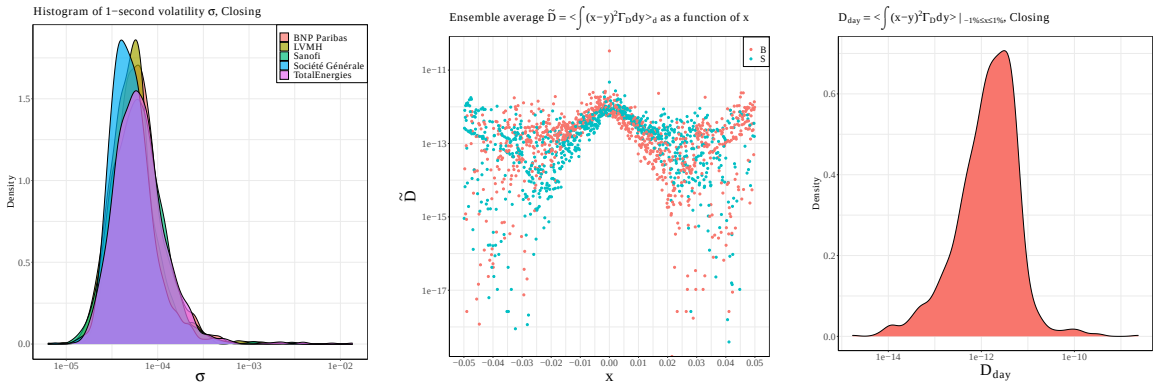

Figure 4.3.9 – Left panel: histogram of daily 1-second volatility during the closing auction. Middle panel: ensemble average of the price update contribution to diffusion as a function of x. Right panel: histogram of daily price update contribution (restricted to limit prices $-1\% \leq x \leq 1\%$).

First, let us consider the volatility per unit of time σ. The average 1-second volatility is of order 10^{-3} in the first minute of the accumulation period and of order 10^{-4} in the last four minutes (see section 4.4). To inspect daily measures, we plot in the left panel of Fig. 4.3.9 a histogram of the daily realized 1-second volatility $\sigma = \sqrt{(\sum_{t=1}^{N} \log(p_{t+1\mathrm{s}}/p_t)^2)/N}$ during an auction, with medians of order 10^{-4}. On an average day, the volatility contribution in the diffusion coefficient is of order 10^{-8}.

Let us now measure the contribution of pure price updates $\frac{1}{2}\int_{\mathbb{R}}(x-y)^2\Gamma_D(x,y)\mathrm{d}y$. We compute the update rate $\Gamma_D(x,y)$ as

$$(\Gamma_D)(x,y,t-\delta t \to t) = \frac{1}{\delta t}\sum_{t-\delta t<t_i<t}\frac{\delta V(p,p',t_i)}{V_p(t_i)}, \tag{4.26}$$

where $\delta V(p,p',t_i)$ is the amount of volume moving from limit price p to p' at time t_i, $x = \log(p/p_{t_i}^-)$, $y = \log(p'/p_{t_i}^+)$, and $V_p(t_i)$ is total volume in limit price p at time t_i^-. We plot in Fig. 4.3.9 the average contribution (across days) of the second term $\left\langle \frac{1}{2}\int_{\mathbb{R}}(x-y)^2\Gamma_D(x,y)\mathrm{d}y \right\rangle_d(x)$ as a function of x (middle panel), and a histogram of daily values (for $-1\% \leq x \leq 1\%$) of average price reassessments $\left\langle \frac{1}{2}\int_{\mathbb{R}}(x-y)^2\Gamma_D(x,y)\mathrm{d}y \right\rangle |_{-1\%\leq x\leq 1\%}$ (right panel). These measurements show that the second term has a mode of order 10^{-12}. Thus, the contribution of price reassessments in the revealed diffusion coefficient is negligible in comparison with that of the indicative price volatility provided unit prefactor, in line with the findings of Challet and Stinchcombe (2001).

It is worth noting that visible price reassessments might be undervalued, as agents may choose

to cancel a limit order entirely before resubmitting it at an updated price. This mechanism, which is difficult to observe in our data, should be accounted for as a form of revealed diffusion. Time priority is less important in auctions than in continuous trading; it is important only for limit orders whose price at the auction time equals the auction price.

4.3.3 Solving the full dynamical equations

To replicate the full dynamics of the revealed auction book throughout the accumulation period, we return to the full model introduced in section 4.2 and solve Eqs (4.2) numerically. To this end, we take a constant cancellation rate $(\nu_l \Gamma_l)(x,t) = \nu_l > 0$, and model the submission rate $\nu_r \Gamma_r$ as follows: For positive prices $x \geq 0$, we take a weighted sum of two exponential terms, representing, respectively, the contribution of fast agents with price scale x_r and slow agents with price scale $k \cdot x_r$, $(k > 1)$ (see section 4.3.1). We allow for the fast agent contribution to increase proportionally to $1/(\gamma_r + T - t)$ when $t > t_r^{(0)}$. For negative prices, we choose a constant submission rate for prices $x < -x_0$, where $x_0 > 0$ is a threshold, then we employ one exponential term for prices $-x_0 \leq x < 0$ establishing continuity at $x = 0$ and $x = -x_0$. This yields

$$(\nu_r \Gamma_r)(x,t) = \begin{cases} \frac{w \cdot C_r}{\gamma_r + T - \max\{t, t_r^{(0)}\}} \cdot e^{-x/x_r} + \frac{(1-w)\cdot C_r}{\gamma_r + T - t_r^{(0)}} \cdot e^{-x/(k \cdot x_r)} & , x \geq 0; \\ A^* \cdot e^{x/x_r^*} & , -x_0 \leq x < 0; \\ m \cdot \frac{C_r}{\gamma_r} & , x < -x_0, \end{cases} \tag{4.27}$$

where $C_r, x_r, x_0, m > 0$, $0 < w < 1$, and $\gamma_r, t_r^{(0)}$ are computed in section 4.3.2. A^* and x_r^* are determined by the continuity conditions at $x = 0$ and $x = -x_0$. We complement Eqs (4.2) with the following initial conditions

$$\begin{aligned} \rho^{(r)}(x,0) &= 0; \\ \rho^{(l)}(x,0) &= \max(ax + b, b), \end{aligned} \tag{4.28}$$

where $a, b > 0$ are the stationary latent book parameters computed in section 4.3.1.

To calibrate Eqs (4.2) to order book data, we minimize the sum of squared errors at times 0,10,..., 290 seconds before the auction time $t \in J = \{10, 20, \ldots, 300\}$ and prices $x \in I = [-2\%, 2\%]$.

$$f(C_r, x_r, k, w, \nu_l, x_0, m, \ldots) = \sum_{x \in I} \sum_{t \in J} \left(\rho(x,t) - \widehat{\rho(x,t)} \right)^2. \tag{4.29}$$

Finally, to measure the influence of diffusion, we minimize (4.29) first taking $D_r = D_l = 0$

(Zero diffusion), then we allow for $D_r, D_l > 0$ as constant parameters to be calibrated (Constant diffusion). Lastly, as suggested by the temporal pattern of the indicative price volatility of section 4.4, we opt for a time-dependent D_r (Time diffusion)

$$D_r(t) = \begin{cases} D_0 & , \quad 0 < t \leq 1; \\ (D_T - D_0) \cdot \frac{t^{-1}-1}{T_s^{-1}-1} + D_0 & , \quad 1 < t < T_s; \\ D_T & , \quad t \geq T_s. \end{cases} \tag{4.30}$$

Eq. (4.30) states that the revealed diffusion coefficient is equal to D_0 at times $0 < t \leq 1$ then decreases as a power law with exponent -1 to a value $D_T < D_0$ until time $t = T_s$, where $T_s \leq T$ is a saturation time. Then for times $t \geq T_s$, D_r remains equal to D_T. We take $T_s = 180$s.

We present the obtained fits first in the plane $(\rho^{(r)}, x)$ at different times in Fig. 4.3.10, then in the plane $(\rho^{(r)}, t)$ at different prices in Fig. 4.3.11. Even under a constant cancellation rate and no diffusion, our model satisfactorily succeeds in replicating the dynamics of the empirical order density. This is largely attributed to the sophisticated submission rate of Eq. (4.27). The presence of two different price scales for $x > 0$ is of crucial importance. The time dependence of the first exponential term allows the order density around the auction price to increase as the clearing approaches. As the second exponential term is not time-dependent, the smooth price decay for larger x is not disrupted (Fig. 4.3.10). This translates into an expansion of the order density only around the indicative price (Fig 4.3.11). Thus, agents posting orders near the indicative price can be seen as more sensitive to the auction deadline compared with agents that act at larger prices.

We note that the constant diffusion fit ($D_r, D_l > 0$) permits the regularization of the obtained density around $x = 0$ and $x = -x_0$ compared with the zero diffusion fit ($D_r = D_l = 0$), where the discontinuity of the first derivative remains. Overall, the influence of diffusion is minimal within our framework, and adopting a time-varying revealed diffusion coefficient $D_r \propto 1/t$ yields slightly better fits around the indicative price. Lastly, locating the global minimum of the loss function f given by Eq. (4.29) poses a considerable challenge and our strategy has been to seek approximate optimal parameters. We present these in Table 4.3.2.

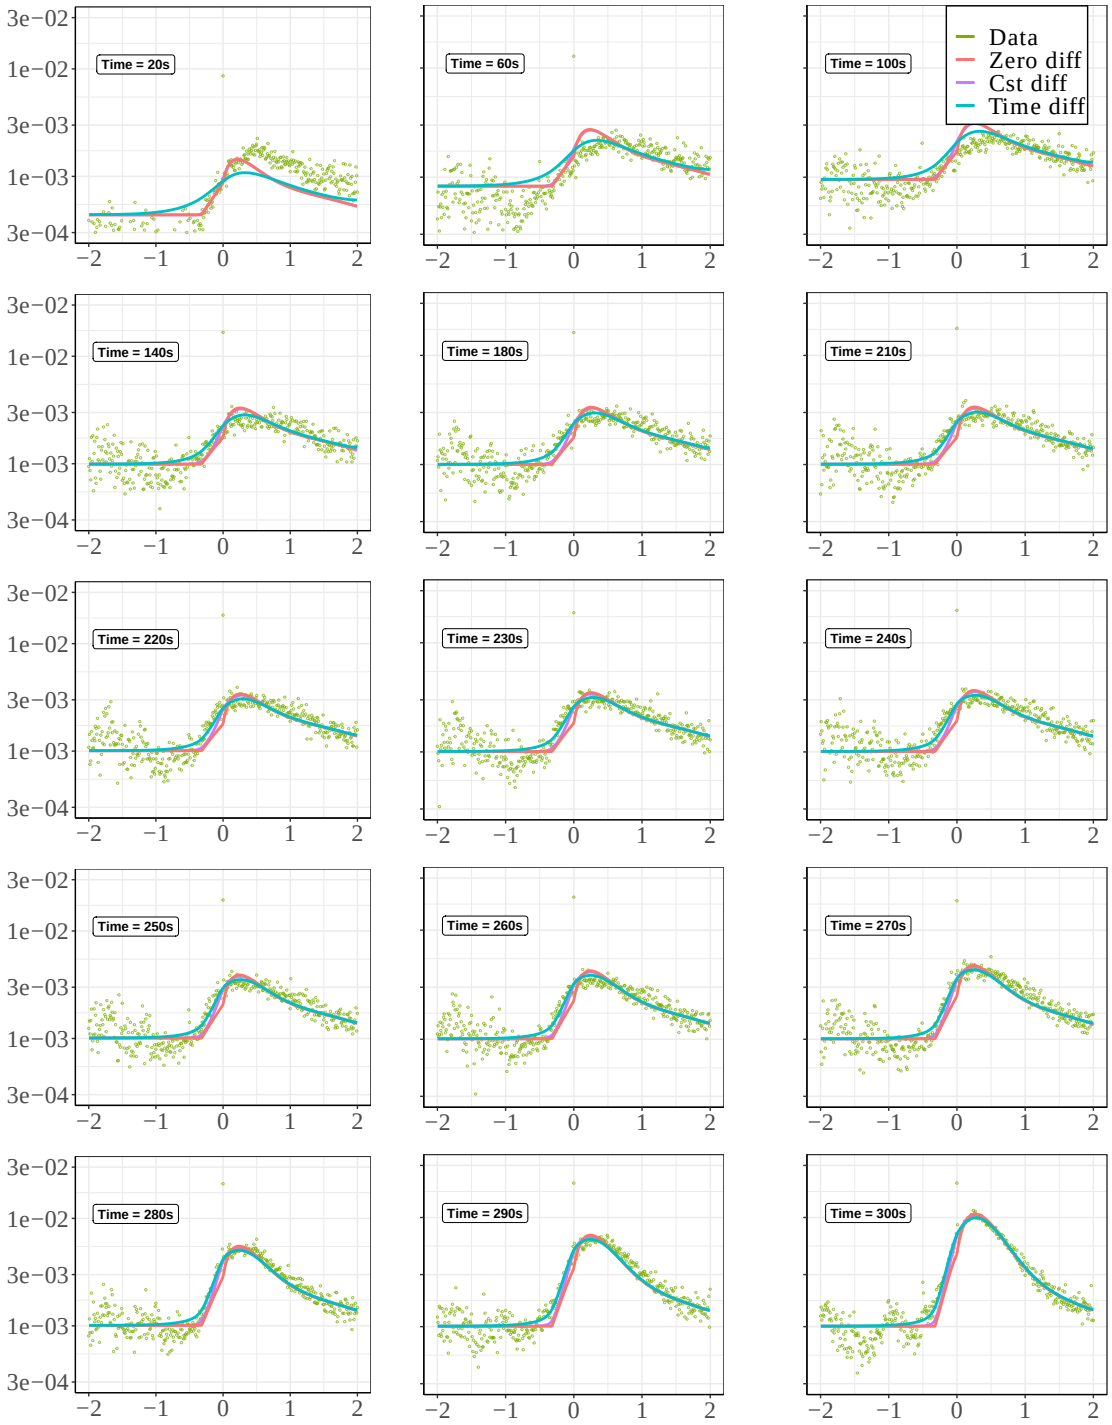

Figure 4.3.10 – Numerical fits $\rho_S^{(r)}$ (solid lines) of the average auction book (green dots) as a function of the centered log price $-2\% \leq x \leq 2\%$ at different instants during the accumulation period. The Y-axis is in log scale. The zero diffusion fit ($D_r = D_l = 0$) is in red lines, the constant diffusion fit ($D_r, D_l > 0$) is in purple lines, and the time diffusion fit ($D_r \propto 1/t$) is in blue lines.

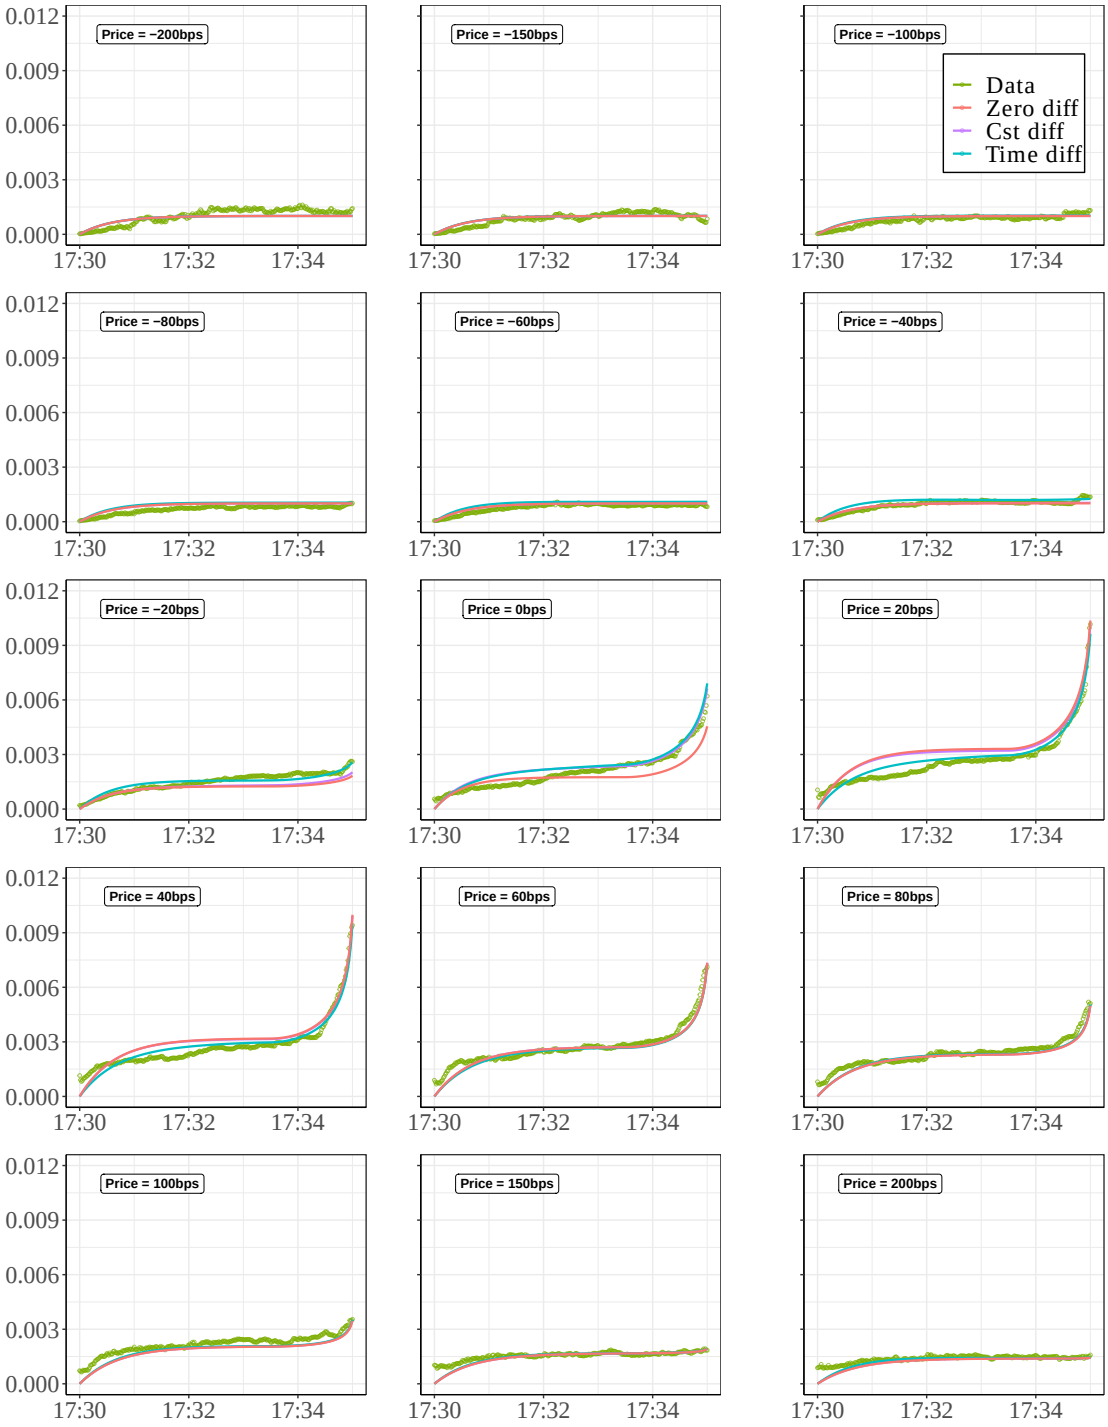

Figure 4.3.11 – Numerical fits $\rho_S^{(r)}$ (solid lines) of the average auction book (green dots) as a function of time $0 < t \leq T$ at different prices price $-2\% \leq x \leq 2\%$. The Y-axis is in ordinary scale. The zero diffusion fit ($D_r = D_l = 0$) is in red lines, the constant diffusion fit ($D_r, D_l > 0$) is in purple lines, and the time diffusion fit ($D_r \propto 1/t$) is in blue lines.

Table 4.3.2 – Fitted parameters from the minimization of Eq. (4.29). In the constant diffusion and the time diffusion cases, we fix the obtained parameters of the zero diffusion case.

	Zero diffusion	Constant diffusion	Time diffusion
$C_r \cdot 10^{-1}$	9.3	_	_
$x_r \cdot 10^{-3}$	2.3	_	_
$k \cdot 10^{0}$	4.9	_	_
$w \cdot 10^{-1}$	8.7	_	_
$\nu_l \cdot 10^{-2}$	2.3	_	_
$x_0 \cdot 10^{-3}$	3.2	_	_
$m \cdot 10^{-2}$	1.6	_	_
$D_l \cdot 10^{-9}$	0	2.4	4.8
$D_r \cdot 10^{-9}$	0	7	_
$D_0 \cdot 10^{-5}$	_	_	1.2
$D_T \cdot 10^{-8}$	_	_	2.2

4.3.4 Discussion

Even though the calibration procedure of section 4.3.3 yields fairly good results, it is important to note that some simplifications were made. The null initial condition for the revealed order book $\rho^r(x, t = 0)$ does not fully reflect reality, given that the order book is already populated by limit orders ($x > 0$) prior to the closing auction. A large fraction of these orders are canceled at the start of the accumulation period, which results in an initial decrease in the revealed order book. This downward trend at $t = 0$ is observable in Fig. 4.3.11 for prices $0 < x < 1\%$.

In addition, we have shown in section 4.3.2 that cancellations exhibit a U-shaped pattern over time. We checked, however, that the introduction of a time-decreasing cancellation rate $\nu_l \Gamma_l$ analogous to Eq. (4.30) does not significantly improve the quality of fits. Similarly, the time acceleration of cancellations around the indicative price does not enhance accuracy, given that the time acceleration of submissions can be adjusted to counterbalance it.

Note that plugging the empirical rates of submission and cancellation as measured in section 4.3.2 into our model does not yield an accurate shape of the empirical order book. While the shape of empirical submissions is similar to the proposed functional of Eq. (4.27), that of empirical cancellations seems to distort the order book at auction time. A likely reason for this discrepancy is the non-trivial behavior of high-frequency agents, which are responsible for intricate interactions and feedback loops that are not accounted for in our zero-intelligence framework.

A larger framework with fast and slow (or more) potentially interacting agents with markedly different submission, cancellation, and diffusion rates could be of interest but is beyond the scope of this paper. The total order density would then be a weighted sum of interacting individual order densities. Additionally, auctions are characterized by bursts of activity at specific round times typical of human behavior. These exogenous bursts of activity suggest segmenting the auction into distinct regimes (section 4.4). Lastly, our model does not capture the large peak of volumes at the indicative price: it is the result of strategic agents aiming to pin the auction price or simply sending orders at the current indicative price.

4.4 The anomalous scaling of the indicative price

When measuring the revealed diffusion coefficient D_r in section 4.3.2, the indicative price was assumed to be a diffusive process. However, it is known to be sub-diffusive (Challet, 2019). In this section, we examine the causes behind the anomalous diffusion of the indicative price.

We investigate the temporal pattern of the indicative price and find that it has non-stationary increments. Figure 4.4.1 depicts ensemble averages over days during the closing auction of 1-second absolute returns $\langle|\log(p_{t+1\text{s}}/p_t)|\rangle$ for five stocks on Euronext Paris. Here, we assume that the indicative price is a realization of the same stochastic process during each closing auction in order to proceed with ensemble averages for each stock. In particular, we observe volatility bursts and relaxations that suggest dividing the closing auction into different regimes: two 30-second regimes during the first minute, followed by a third regime between 17:31:00 and 17:32:00, a longer fourth regime between 17:32:00 and 17:34:00, and a final one-minute regime.

We test for the presence of anomalous scaling of the indicative price in each regime by computing average Hurst exponents. Assuming that the indicative price is a self-similar process, we can estimate an average Hurst exponent H as

$$\left\langle \left\langle (\log(p_{t+\tau}/p_t))^2 \right\rangle_t \right\rangle_d \sim \tau^{2H}, \tag{4.31}$$

where $\langle\cdot\rangle_t$ is the average over one realization of the indicative price series, and $\langle\cdot\rangle_d$ is the ensemble average over days. The underlying process undergoes normal diffusion when $H = 1/2$ and anomalous diffusion when $H \neq 1/2$ (super-diffusion when $H > 1/2$, and sub-diffusion when $H < 1/2$).

Chen et al. (2017) show that H can be decomposed as $H = J + L + M - 1$, where each of the exponents J, L, M is associated with the failure of one condition of the central limit

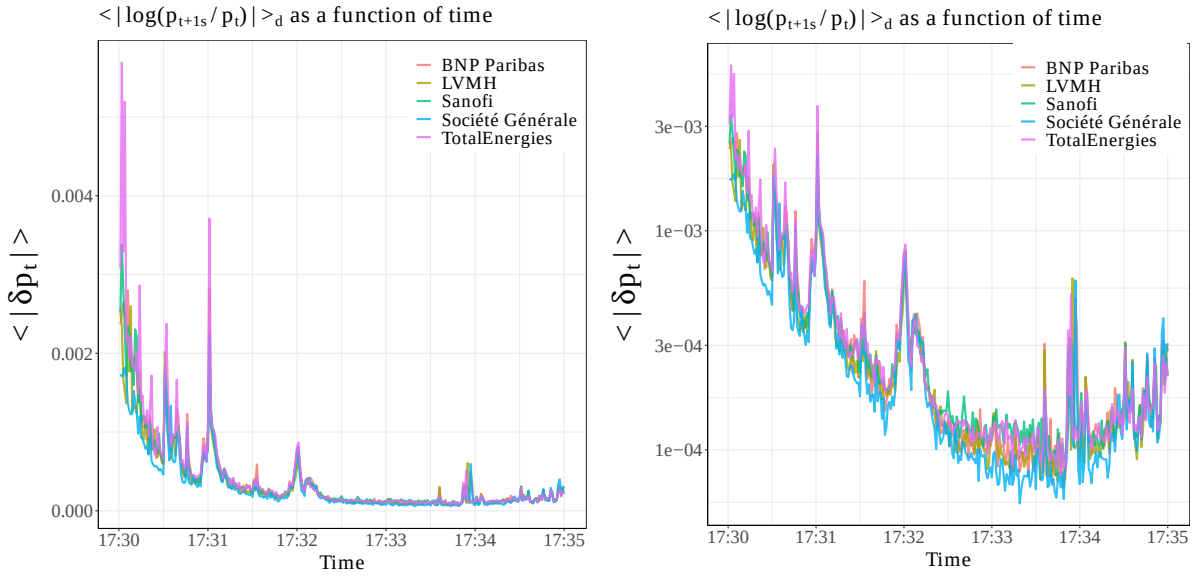

Figure 4.4.1 – Ensemble average of the absolute value of 1-second log price increments $\mathbb{E}[|\log(p_{t+1\text{s}}/p_t)|]$ in the closing auction of five stocks traded on Euronext Paris between 2013 and 2017: BNP Paribas, LVMH, Sanofi, Société Générale, and TotalEnergies. Left panel: ordinary scale. Right panel: Y-axis in log scale.

theorem: (i) the presence of correlations/long-term memory, (ii) infinite variance, and (iii) the non-stationarity of increments.

1. Joseph exponent J quantifies the long-term memory of increments. It is defined as the scaling of the ensemble average of a rescaled range statistic

$$\langle R_t/S_t \rangle_d \sim t^J, \tag{4.32}$$

where the considered range statistic is expressed as $R_t = \max_{1\leq s\leq t}[X_s - s/tX_t] - \min_{1\leq s\leq t}[X_s - s/tX_t]$, the deviation being $S_t^2 = Z_t/t - (X_t/t)^2$, and $X_t = \log(p_t/p_0)$ is the log indicative price centered around the origin of the considered regime. If $J > 1/2$ increments are positively correlated, $J < 1/2$ corresponds to negatively correlated increments, and $J = 1/2$ refers to the absence of correlations; Chen et al. (2017) argue that the Joseph exponent J is the appropriate measure to test the efficient market hypothesis and not the Hurst exponent H.

2. Noah exponent L quantifies whether increments have finite variance ($L = 1/2$), or infinite variance ($L > 1/2$). Assuming that increments have power law tails with exponent γ, i.e.

$\mathbb{P}_>(|x|) \underset{|x|\to+\infty}{\sim} |x|^{-\gamma}$, then

$$L = \max(\frac{1}{2}, \frac{1}{\gamma}). \tag{4.33}$$

3. Moses exponent M quantifies whether increments are stationary ($M = 1/2$) or non stationary ($M \neq 1/2$). It is defined as the scaling of the ensemble average of centered absolute increments (see Fig. 4.4.1)

$$\langle |\log(p_{t+1}/p_t)| \rangle_d \sim t^{M-\frac{1}{2}}. \tag{4.34}$$

We estimate H, J, L, and M for each regime separately. For M and L, we use robust methods. Namely, M is estimated using the median of the sum of absolute increments

$$m\left[\sum_{k=0}^{t-1} |\log(p_{k+1}/p_k)|\right] \sim t^{M+\frac{1}{2}}, \tag{4.35}$$

and L is estimated using the median of the sum of square increments

$$m\left[\sum_{k=0}^{t-1} \log(p_{k+1}/p_k)^2\right] \sim t^{2L+2M-1}, \tag{4.36}$$

where each sum starts at the beginning of the considered regime. For an in-depth discussion and proofs regarding the estimation of each exponent, see Chen et al. (2017).

We report in Table 4.4.1 numerical estimates of the scaling exponents in each regime for five different stocks. We draw the following remarks and conclusions from our results:

— the indicative price is a sub-diffusive process during the closing auction as $H < 1/2$, except in the final minute, where it switches to a diffusive (or an over-diffusive) behavior $H \gtrsim 1/2$;

— the increments of the indicative price do not exhibit long-term memory during the accumulation period as $J \approx 1/2$ for most stocks and regimes. Note that as the auction end approaches, the rescaled range statistic R_t/S_t flattens due to the stabilization of the indicative price. As a result, J is not defined in the last regime. We thus conclude that the indicative price is efficient in the sense of Chen et al. (2017): even if the apparent Hurst exponent is not $1/2$, $J = 1/2$ precludes price predictions. Finally, Besson and Fernandez (2021) report an overreaction of the indicative price based on the following definition of the indicative jump on close $J(t)$

$$J(t) = \frac{p_t - p_{\text{ref}}}{p_{\text{auction}} - p_{\text{ref}}}. \tag{4.37}$$

However, this incorporates the first jump from p_{ref} (the last price of the continuous trading

Table 4.4.1 – Average Hurst, Joseph, Levy, and Moses exponents for the studied stocks in each regime.

Regime	Stock	H	J	L	M	H-J-L-M+1
	BNP Paribas	_	0.47***	_	_	_
	LVMH	_	0.56***	_	_	_
1	Sanofi	_	0.48***	_	_	_
	Société Générale	_	0.41***	_	_	_
	TotalEnergies	_	0.4***	_	_	_
	BNP Paribas	0.35***	0.54***	0.72***	0.09***	0
	LVMH	0.43***	0.55***	0.71***	0.13***	0.04
2	Sanofi	0.38***	0.49	0.73***	0.13***	0.03
	Société Générale	0.39***	0.49	0.78***	0.11***	0.01
	TotalEnergies	0.38**	0.52.	0.69***	0.18***	-0.01
	BNP Paribas	0.37***	0.63***	0.72***	0.03***	-0.01
	LVMH	0.34***	0.54***	0.73***	0.06***	0.01
3	Sanofi	0.37***	0.52***	0.68***	0.1***	0.07
	Société Générale	0.32***	0.53***	0.76***	-0.01***	0.04
	TotalEnergies	0.34***	0.53***	0.71***	0.05***	0.04
	BNP Paribas	0.35***	0.55***	0.73***	0.12***	-0.04
	LVMH	0.36***	0.48***	0.69***	0.14***	0.05
4	Sanofi	0.34***	0.5	0.61***	0.26***	-0.03
	Société Générale	0.37***	0.55***	0.74***	0.05***	0.03
	TotalEnergies	0.36***	0.51**	0.68***	0.15***	0.01
	BNP Paribas	0.58***	_	_	_	_
	LVMH	0.55***	_	_	_	_
5	Sanofi	0.49	_	_	_	_
	Société Générale	0.65***	_	_	_	_
	TotalEnergies	0.51*	_	_	_	_

1. Values of H in the first regime are omitted as the mean square displacement of the indicative price is not a power law of the lag in this regime and is rather noisy. We suspect the activation of VFA,VFC (Valid For Auction, Valid For Closing) to be the cause of this large noise.
2. Values of J in the last regime are omitted as the ensemble average of the rescaled range statistic is not a power law of time. Instead, $E[R_t/S_t]$ flattens with time: this results from the stabilization of the indicative price during the last moments of the accumulation period.
3. The symbols ***,**, and * indicate significance at the 0.1%, 1%, and 5% level, respectively for testing the null hypothesis $\{S = 0.5\}$, where $S \in \{H, J, L, M\}$.

phase) to p_0 (the first indicative price of the accumulation period). When the reference price is the first indicative price of the accumulation period p_0, $J(t)$ becomes $\tilde{J}(t)$

$$\tilde{J}(t) = \frac{p_t - p_0}{p_{\text{auction}} - p_0}, \tag{4.38}$$

which does not display any overreaction. In addition, we find that the $p_{\text{auction}} - p_{\text{ref}}$ does not have systematically the same sign as $p_{\text{auction}} - p_0$ (only ≈50% of the time). Thus, the first jump from p_{ref} to p_0 is not predictive of the direction of the auction price.

— the increments of the indicative price are highly non-stationary as M significantly differs from 1/2. Considering that $M \approx 0$ implies that the indicative price volatility is a decreasing time function in each regime $\sigma_t \propto t^{-1/2}$, and equivalently that the revealed diffusion coefficient is time-dependent $D_r \propto 1/t$. This appears to be the major cause of the indicative price anomalous scaling;

— the increments of the indicative price exhibit infinite variance as $L \approx 0.7$.

4.5 Conclusion

Zero intelligence models are surprisingly able to reproduce non-trivial stylized facts in financial markets (Farmer et al., 2005). Here, we showed that by adapting the zero intelligence latent/revealed liquidity framework of Dall'Amico et al. (2019) to auctions, we are able to replicate complex price-time dynamics of the average order book throughout the accumulation period. Within our framework, the skewed shape of the order book emerges from the product of the linear latent book by the exponentially decreasing submission rate. The time acceleration around the indicative price arises from inversely proportional rates to the remaining time to the deadline, analogously to typical human behavior when facing a deadline Alfi et al. (2009). These results were confirmed by the estimation of the submission, cancellation, and diffusion rates. These represent a new piece of evidence advocating for the relevance of the latent order book of Tóth et al. (2011).

Although successful at reproducing many of the complex patterns observed during auctions, our model can only describe average order books where large daily fluctuations are neglected. Additionally, price-changing events (sell limit orders above the indicative price and buy limit orders below the indicative price that are breaching zero impacts) were supposed not to have an influence on the indicative price when they directly impact the supply/demand equilibrium similarly to market orders. Finite-size effects such as the large peak of volumes at the indicative

price cannot be captured by continuous models and result from a possible strategic behavior. Finally, the heterogeneous nature of the agents involved and the reasons why they take part in auctions should be of interest: these range from manually trading agents for idiosyncratic reasons to more sophisticated trading algorithms minimizing impact and/or maximizing profits. A general model accounting for the overall auction ecology, price discreteness, volume fluctuations, and the strategic behavior of agents is needed to explain daily deviations, e.g., during index rebalancing and derivatives expiry days, or after the release of a significant piece of news.

Acknowledgments

We thank Michele Vodret for fruitful discussions and acknowledge the use of the EUROFIDAI BEDOFIH's database acquired through "Equipex PLADIFES ANR-21-ESRE-0036 (France 2030)".

Appendix

4.A Calibration of stationary order densities

Table 4.A.1 – Estimate values of the optimal parameters $(\nu_r/\nu_l \cdot a, \nu_r/\nu_l \cdot b, x_r, k, w)$ fitting the orders' density using the stationary setting. Fitting range: $x = 0 \pm 5\%$.

Side	Agent type	Stock	$\nu_r/\nu_l \cdot a$	$\nu_r/\nu_l \cdot b \cdot 10^{-2}$	x_r	k	w
Buy	ALL	BNP Paribas	5.72	0.56	0.37%	5.6	0.984
		LVMH	9.23	0.62	0.3%	5.4	0.978
		Sanofi	7.89	0.58	0.29%	5.2	0.983
		Société Générale	6.32	0.63	0.38%	6.1	0.974
		TotalEnergies	6.08	0.50	0.33%	6.8	0.989
	MIX	BNP Paribas	5.24	0.47	0.36%	3.9	0.980
		LVMH	7.47	0.53	0.31%	5.4	0.986
		Sanofi	5.89	0.47	0.31%	4.2	0.986
		Société Générale	5.94	0.50	0.37%	6.8	0.988
		TotalEnergies	4.68	0.40	0.34%	5.0	0.990
	HFT	BNP Paribas	0.87	0.04	0.04%	7.8	0.227
		LVMH	1.65	0.03	0.28%	5.6	0.977
		Sanofi	2.45	0.05	0.2%	11.5	0.995
		Société Générale	0.21	0.05	0.33%	3.3	0.707
		TotalEnergies	1.91	0.04	0.21%	10.5	0.993
	NON	BNP Paribas	0.01	0.05	0.44%	6.8	0.341
		LVMH	0.07	0.07	0.2%	11.2	0.591
		Sanofi	0.08	0.07	0.14%	11.6	0.631
		Société Générale	0.08	0.08	0.2%	10.4	0.317
		TotalEnergies	0.01	0.03	0.35%	7.6	_
Sell	ALL	BNP Paribas	6.39	0.51	0.36%	7.4	0.980
		LVMH	8.96	0.67	0.3%	4.7	0.971
		Sanofi	7.83	0.54	0.29%	5.4	0.979
		Société Générale	6.83	0.60	0.37%	6.2	0.977
		TotalEnergies	6.77	0.58	0.3%	5.1	0.969
	MIX	BNP Paribas	5.38	0.44	0.35%	3.7	0.962
		LVMH	7.78	0.55	0.3%	5.3	0.987
		Sanofi	5.92	0.43	0.31%	4.8	0.984
		Société Générale	6.23	0.48	0.35%	4.1	0.976
		TotalEnergies	5.12	0.46	0.32%	3.8	0.971
	HFT	BNP Paribas	1.00	0.02	0.29%	12.1	0.987
		LVMH	1.77	0.02	0.28%	7.2	0.983
		Sanofi	2.30	0.04	0.21%	11.2	0.995
		Société Générale	1.82	0.09	0.02%	21.2	0.585
		TotalEnergies	2.05	0.04	0.2%	7.7	0.987
	NON	BNP Paribas	0.06	0.06	0.18%	16.5	0.452
		LVMH	0.08	0.06	0.27%	8.6	0.516
		Sanofi	0.03	0.04	0.21%	9.6	0.106
		Société Générale	0.10	0.07	0.44%	4.8	0.291
		TotalEnergies	0.05	0.05	0.2%	11.7	0.258

Chapter 5

Discussion and outlook

5.1 Price impact in equity auctions

Using the shape of the order book, we computed price impact at auction time as a function of the order size for a given day and a given auction. Although closely related to the concept of virtual impact, price impact at auction time is a faithful measure of the incurred cost, assuming, for instance, that a low-latency agent manages to act last in the accumulation period. In fact, the market has little time to react to later submissions. Additionally, the shape of the 1-lag response function of the indicative price shows that price impact is mostly mechanical and linear as a function of the order size submission (or cancellation) during the accumulation period.

5.1.1 Impact-related cost for early submissions

Whereas the 1-lag response function describes how order submissions impact the indicative price, it does not tell the impact-related cost, e.g., for an early submission during the accumulation period. A matched submission at the clearing transacts at the auction price. Thus, the impact-related cost is the difference in the auction price between a world in which this submission occurred and a parallel world in which this submission did not occur; this is the reaction impact (Bouchaud et al., 2018a). The observed impact is the difference between the auction price and the indicative price right before an order submission. The observed impact is a reasonable proxy for the reaction impact, and equality holds when prices are unpredictable. The impact-related cost I^{cost} measured on the auction price reads similarly to that of Challet and Gourianov (2018)

$$I^{\text{cost}}(\omega) = \langle \varepsilon_t \cdot (p_a - p_t) | \omega \rangle \,. \tag{5.1}$$

Whereas the instantaneous impact is linear on average, we find no dependence of the I^{cost} on the order size. Actually, the impact on the auction price has a strong dependence on the time of the submission with early submissions having a larger auction impact, on average, than later submissions. Figure 5.1.1 reports the average mechanical impact $R^M = \left\langle \varepsilon_t \cdot (p_t^+ - p_t) \right\rangle$, the lagged indicative price response $R^1 = \langle \varepsilon_t \cdot (p_{t+1} - p_t) \rangle$, and the average impact on the

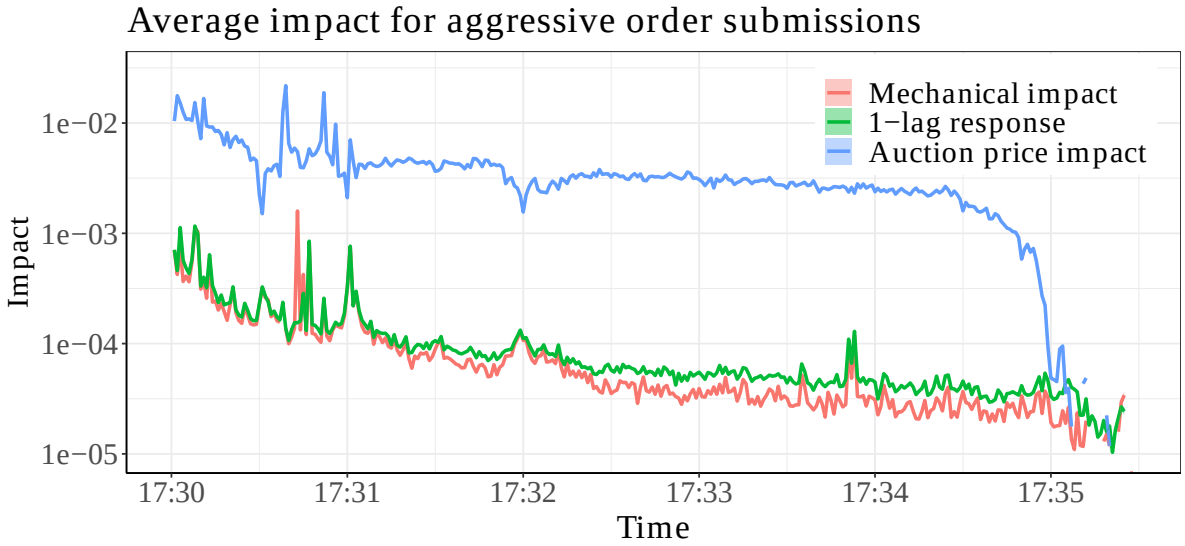

Figure 5.1.1 – Average impact as a function of the time in the accumulation period independently of the order size. The impact on the auction price (blue), the mechanical impact (red), and the 1-lag response function (green) of aggressive order submissions. Closing auction data for TotalEnergies between 2013 and 2017.

auction price $I^{\text{cost}} = \langle \varepsilon_t \cdot (p_a - p_t) \rangle$ as a function of the time in the accumulation period and independently of the order size. Note that the average is taken over grouped submissions in 1-second time bins. As one may expect, when the auction time approaches, the impact on the auction price for aggressive order submissions decreases to match their instantaneous impact as well as their 1-lag impact function. This confirms the relevance of acting last during the auction when looking to minimize the price impact of large one-shot submissions.

5.1.2 Impact-related cost for metaorders

An interesting research direction would be to assess the impact of metaorders, if any, during auctions. This requires agent-labeled data such as proprietary trading firm or brokerage data. Executing a metaorder by slicing it into two or several child orders during the closing is less straightforward than during the continuous trading phase. Due to the distinct temporal pattern of the closing, an agent can choose, for instance, to submit child orders during activity bursts at round minutes to go undetected. The large liquidity build-up towards the clearing allows for agents to complete their metaorder with large one-shot submissions. If similar one-shot submissions were sent during the trading day, they would have caused abnormal price jumps.

Note that even when a large one-shot submission results in an abnormal indicative price move at the start of the accumulation period, agents can mitigate this effect, and the indicative price mean reverts to reasonable values before the clearing.

Another interesting direction is to provide response elements to the following "optimal split problem". Suppose an agent is willing to liquidate a large volume during a given day, including at the closing auction, and is looking to minimize the total impact costs. Fund managers are often confronted with a similar problem, looking to execute large orders at the closing price without substantially impacting the price formation in the closing auction. A way around this is to divide the metaorder into two fractions. The first fraction would be traded during the continuous trading phase using, for instance, a target close algorithm (Guéant, 2016). The remaining fraction would be traded in the closing auction. In appendix A, we derive the optimal fraction $x^* \in [0, 1]$ to be traded in the continuous trading phase as a function of the intraday volatility, the auction liquidity, and the market impact prefactors during both trading phases, while minimizing the total cost of trading.

5.2 Modeling auction dynamics

5.2.1 Upgrading the latent/revealed liquidity framework

When adapted to equity auctions, the latent/revealed liquidity framework is successful in reproducing the average dynamics of the limit order book throughout the accumulation period. In Section 4.3.3, the submission rate of Eq. (4.27) implicitly implements a multi-time scale approach; it is the sum of two exponential terms. The first term with price scale x_r represents the contribution of fast agents, and the second term with a larger price scale $k\, x_r$, $k > 1$ represents the contribution of slow agents. We allowed for the contribution of fast agents' to be pressured by the auction deadline.

As pointed out in Section 4.3.4, the model can be upgraded by adopting a global multi-time/price-scale approach not only for the submission rate (Benzaquen and Bouchaud, 2018a). For instance, one can model the order book evolution for each agent category, say fast and slow agents or more. These order books may influence each other, and the resulting order book is their sum.

Additionally, the behavior of high-frequency agents likely requires accounting for additional memory terms. In fact, the actions of HFTs during the accumulation period mainly involve updating the volume of their existing orders either by increasing or decreasing the number of

shares in reaction to the order flow. Within our model, these quantity updates are viewed as (independent) order submissions or cancellations, which mechanically inflates the cancellation rate. To account properly for these events, we need to take into consideration their dependence on existing orders as well as the order flow via a memory kernel resulting in complex integro-differential equations.

The indicative price dynamics, that are implicit in our model, can be further investigated. The indicative price is determined by equalizing supply and demand and is driven by market orders and aggressive limit orders. A general framework where limit prices are not centered around the current indicative price may be of interest. Thus, modeling market orders as well as their influence on the price through, e.g., linear instantaneous impact or propagators is required.

5.2.2 Questioning the closing auction liquidity

At the closing auction time, most of the non-matchable orders expire at the end of the trading day. These non-matchable orders, usually a few ticks away from the auction price, represent a substantial fraction of the total order book volume at auction time. As liquidity begets liquidity, it is interesting to question the presence of these large orders bound to expire at the end of the auction. Are they used as a resistance barrier to adverse price moves? Or do they represent revealed intentions (reservation price) at the auction time, waiting for a discount in the last moments of the accumulation period?

Similarly, one might wonder why high-frequency market makers display a large activity during auctions (more than 55% of total events) when their share in the closing volume is around 1%. A likely explanation is that they continuously minimize their impact by profiting from zero-impact volumes. Another likely explanation is to influence the price as they surprisingly account for 25% of price-changing events.

On the one hand, high-frequency agents are less likely to be interested in auction price pinning as the average HFT book shape does not display a peak around the indicative price. On the other hand, non-high frequency agents clearly engage in an apparent "price pinning war" by sending large limit orders at the indicative price as the auction end approaches. In Fig. 5.1.1), the auction impact significantly decreases in the last five seconds of the accumulation period and converges towards its instantaneous counterpart. This significant decrease can partly be attributed to the stabilization of the indicative price towards the clearing which itself results from this price-pinning competition.

5.2.3 Towards more realistic models

Continuous price models can not reproduce some important finite-size effects as limit orders are considered to have an infinitesimal size. For instance, if an aggressive buy (or sell) limit order has a large size, it impacts the current indicative price all the way up (or down) to its limit price. Shifting to a discrete price model will likely result in different yet rich aggregate behavior.

An interesting direction would be to design a game in which an agent plays against zero intelligence agents, rational agents, or a mix of both. The idea would be to find the optimal strategy that maximizes the agent's profit, minimizes the impact costs, or fixes the auction price at a given level. Should the agent split his submission throughout the accumulation period? Or should he send it in one shot? When is the optimal time to send his order(s)?

A parallel direction would be to investigate the effect of exogenous information on the price formation process during auctions. Kyle's model (Kyle, 1985) describes the auction price formation in the presence of a market maker, an insider, and a noise trader. In contrast with private information, public information can be accessed by almost all agents in the market. However, the news interpretation may differ from agent to agent, as well as their reactions, which translates into heterogeneous investment decisions. Thus, the release of a piece of news may change the way agents act during auctions vs when there is no relevant news.

Appendices

Appendix A

Trading costs reduction: continuous trading and close split

Based on an unpublished work
by Charles-Albert Lehalle, Mohammed Salek, Damien Challet, and Ioane Muni Toke

A.1 Introduction

Open markets can absorb large meta-orders during the day with impact-related costs scaling as the square root of the metaorder size. In parallel, the closing auction allows the matching of large submissions, particularly towards the auction clearing, with impact-related costs scaling linearly with the order size. The cost of trading during the closing auction is typically smaller than in the continuous trading phase for reasonable order sizes. When the daily volatility is large, the exchanged volume during the continuous trading phase tends to be smaller, and that of the closing auction tends to be larger. Indeed, an increase of 1% in the intraday volatility yields, on average, a decrease of 3% in the market share of the closing volume.

Thus, it is unclear whether a meta-order is better off traded during the continuous trading phase, the closing auction, or a mix of both. This decision may depend on the cost of trading during both phases, intraday volatility, and auction liquidity, among other considerations. In this Chapter, we assume that a large (meta)order of size V is split with fraction $x \in [0, 1]$ that trades at the continuous trading phase and the remaining fraction $1 - x$ trades at the closing auction. We look for the optimal x^* that minimizes the total trading costs. The following framework was originally designed to check if price manipulation can accidentally occur in closing auctions (Lehalle, 2022). Indeed, there exists a degenerate case where an agent minimizes his impact costs and simultaneously results in the maximum price impact.

The price impact of a meta-order of size v during the continuous trading phase is given by

$$\delta I_c(v) = Y \, \sigma \sqrt{\frac{v}{V_d}}, \tag{A.1}$$

where σ is the intraday volatility, V_d is the daily exchanged volume, and Y is a constant of order unity. Henceforth, we take $Y = 1$.

Additionally, the price impact of an order submission of size v late in the closing auction is given by

$$\delta I_a(v) = \tilde{S} \cdot \frac{v}{Q_a}, \tag{A.2}$$

where $\tilde{S}$ is the impact slope (see Section 3.C), and Q_a is the auction volume. Even though a late one-shot submission in the accumulation period may result in zero impact part, the price impact of Eq. (A.2) represents the average instantaneous impact which is linear and converges to the impact on the auction price as the clearing approaches.

Assuming that the fraction x of the meta-order is traded before the start of the closing auction, one should account for the relaxation of the impact by including a factor $1 - r \propto t^{-1/2}$, where t is the measured time from the last continuous trade of the meta-order until the start of the accumulation period. $r = 0$ corresponds to no relaxation and the closing auction starts right after the last continuous trade; $r = 1$ corresponds to a case where the last continuous trade occurs long before the start of the accumulation period and the peak impact totally relaxes before the start of the closing accumulation period. Therefore, the meta-order is responsible for a total price impact

$$\begin{aligned} \delta I(V) &= (1 - r) \cdot \delta I_c(x \cdot V) + \delta I_a((1 - x) \cdot V) \\ &= (1 - r) \cdot \sigma \cdot \sqrt{\frac{x \cdot V}{V_d}} + \tilde{S} \cdot \frac{(1 - x) \cdot V}{Q_a}. \end{aligned} \tag{A.3}$$

We introduce $\phi = V/V_d$, the participation rate of the meta-order in the total exchanged volume, and $\psi = Q_a/V_d$, the ratio of the auction volume to the total exchanged volume. Thus

$$\delta I(x) = (1 - r) \cdot \sigma \cdot \sqrt{\phi} \cdot \sqrt{x} + \tilde{S} \cdot \frac{\phi}{\psi} \cdot (1 - x). \tag{A.4}$$

Figure A.1.1 provides a general overview of the auction liquidity ψ, which is the fraction of the closing volume in the daily exchanged volume vs the intraday volatility σ for the TotalEnergies stock between 2013 and 2017.

If the impact relaxation is weak $r \to 0$, a minimal price impact is reached when $x = 0$, i.e., when the meta-order fully trades at the closing auction. This is because the auction impact is assumed to be much lower than the continuous trading impact ($\delta I_a \leq \delta I_c$). If $r \to 1$, price impact is zero when the metaorder trades fully at the continuous trading phase $\delta I(x = 1, r = 1) = 0$: the market totally forgets about the metaorder.

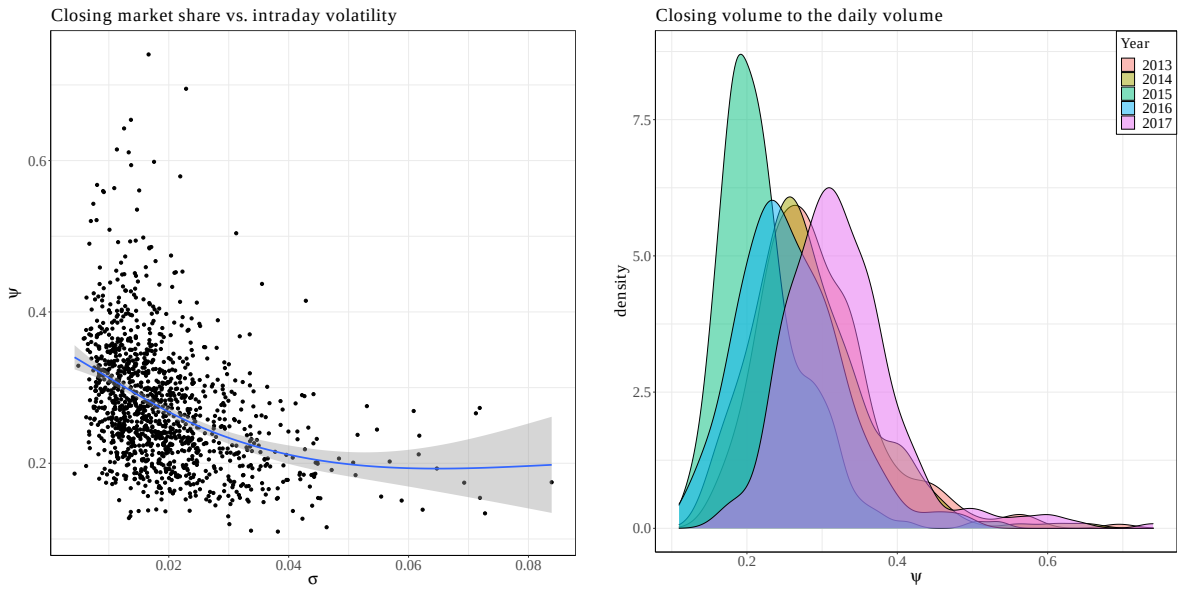

Figure A.1.1 – Left panel: $\psi = Q_a/V_d$ the fraction of the closing volume in the daily exchanged volume as a function of the Parkinson volatility estimator $\sigma = \log(S_{\text{high}}/S_{\text{low}})$. Right panel: smoothed histogram of ψ broken down by year.

We can compute an equivalent relaxation r_e such that price impact is the same whether the meta-order is fully traded during the continuous trading phase or fully traded during the closing auction

$$r_e = 1 - \frac{\tilde{S}}{\sigma} \cdot \frac{\sqrt{\phi}}{\psi} \approx 1 - \frac{\sqrt{\phi}}{\psi}. \tag{A.5}$$

Typical values $\phi = 10^{-3}, \psi = 25\%$ yield $r_e = 87\%$, meaning that 87% of the peak impact in the continuous phase should relax in order for the price impact to be equal in both phases in which the meta-order is traded.

A.2 Maximal impact

From now on, we suppose $r < r_e$, i.e., trading during the auction yields the least price impact. Since $x \to \delta I(x)$ is a concave function, we can compute a critical fraction x_c such that the total price impact is maximal

$$x_c = \min\left\{1; \frac{\psi^2\sigma^2(1-r)^2}{4\tilde{S}^2\phi}\right\}, \quad r < r_e. \tag{A.6}$$

If r is sufficiently close to 1 (strong relaxation), we can have $x_c < 1$. Figure A.2.1 shows the value of x_c when $r \in [0, 1]$ in the case of usual values of ϕ, ψ (top panel) and in the case of a large participation rate ϕ (bottom panel). As the relaxation factor grows, the impact of the continuous trading phase is lost, and trading fully at the auction ($x = 0$) yields a larger impact than trading fully during the continuous phase (see $r = 0.8$ for instance).

A.3 Minimal impact-related cost

The total cost of trading is given by the product of respective price impacts and traded volumes.

$$\delta C(V) = x \cdot V \cdot \delta I_c(x \cdot V) + (1 - x) \cdot V \cdot \delta I_a((1 - x) \cdot V). \tag{A.7}$$

Note that the cost of trading during the continuous trading phase does not involve any relaxation factor. Therefore, the cost per unit volume (dividing the total cost by V) is

$$\delta C(x) = \sigma \cdot \sqrt{\phi} \cdot x^{\frac{3}{2}} + \tilde{S} \cdot \frac{\phi}{\psi} \cdot (1 - x)^2, \tag{A.8}$$

and the optimal fraction x^* that minimizes the total cost of trading is

$$x^* = \min \left\{ 1; 1 + \frac{9\psi^2\sigma^2 - 3\psi\sigma\sqrt{64\tilde{S}^2\phi + 9\psi^2\sigma^2}}{32\tilde{S}^2\phi} \right\}. \tag{A.9}$$

Figure A.3.1 shows the optimal fraction x^* as a function of ψ (top panel) and of ϕ (bottom panel). Intuition matches the obtained results: the larger the auction liquidity ψ, the smaller is x^*, and the more should be traded in the auction phase. Similarly, the larger the participation rate ϕ, the larger is x^*, and the larger is the traded fraction in the continuous trading phase. However, notice that even for large participation rates, e.g. $\phi = 1\%$, and a usual value of the auction liquidity $\psi = 25\%$, $x^* \approx 19\%$: this means that the vast majority of the volume is executed in the auction phase. If the agent chooses however to execute the entire volume in the continuous trading phase, he will incur 13bps additional cost for being far from optimality.

Finally, even when the closing has low liquidity $\psi = 5\%$, for a participation rate of 1%, roughly more than 25% of the metaorder should be executed during the auction. If, however, the entire metaorder trades in the continuous phase, 5 bps of additional cost are incurred compared to the optimal split case.

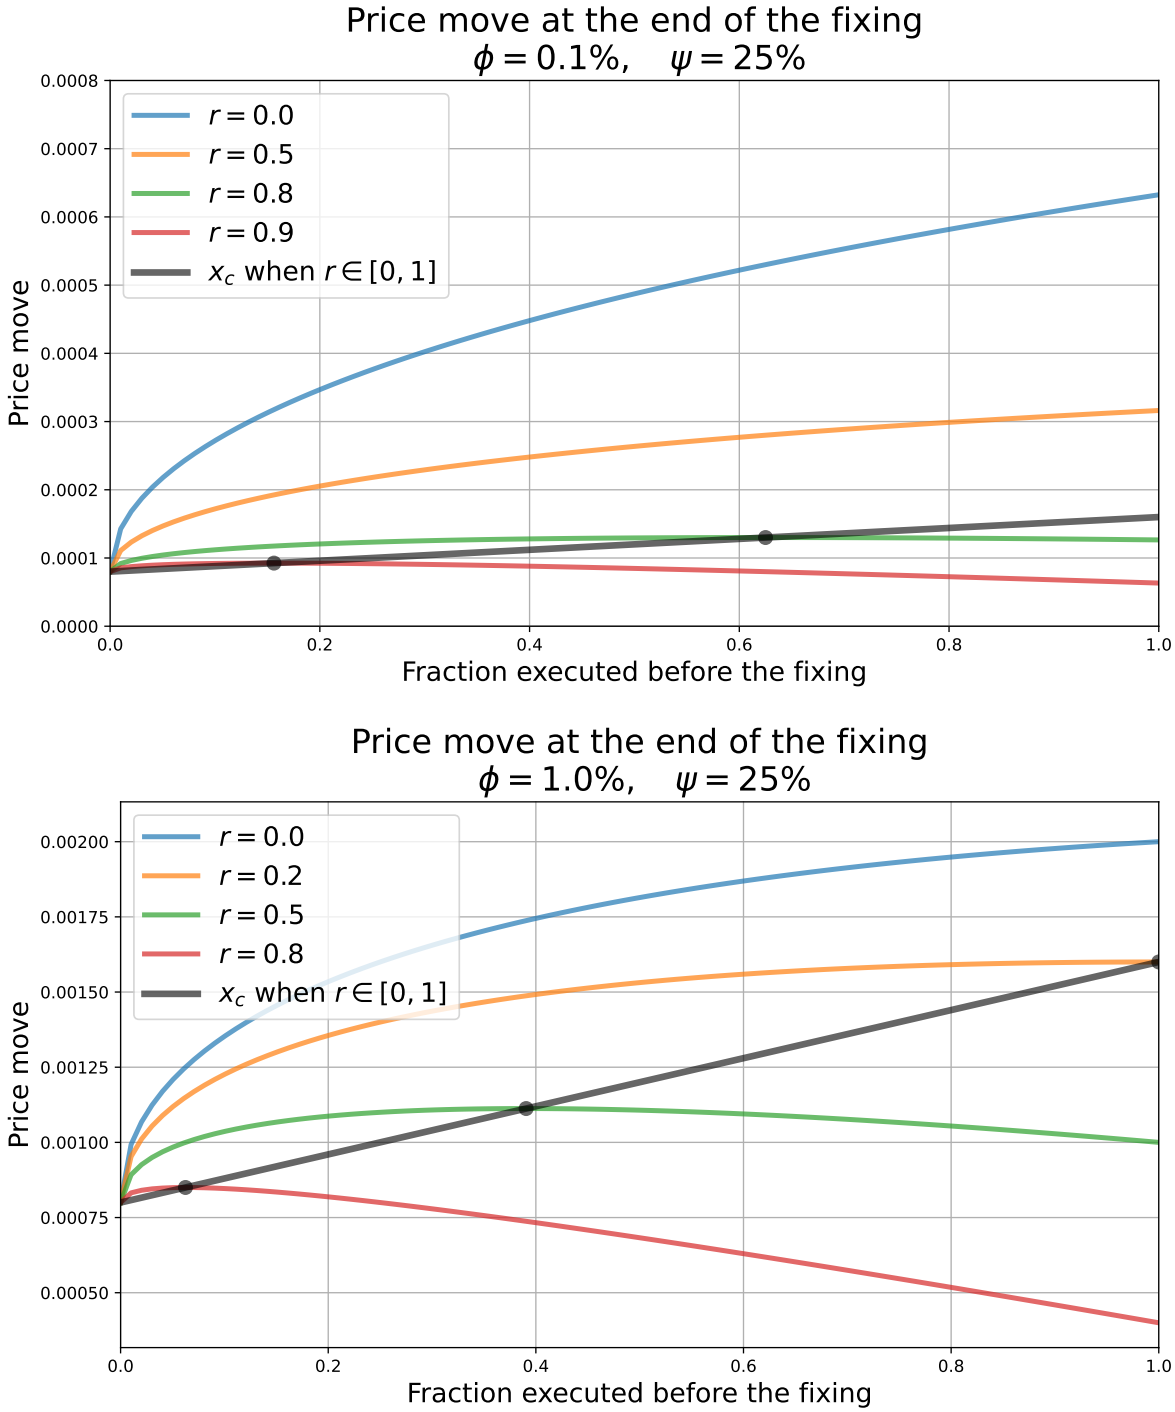

Figure A.2.1 – Price impact of the meta-order at the end of the closing auction as a function of the executed fraction x during the continuous trading phase. Top panel: usual values $\phi = 0.1\%, \psi = 25\%$. Bottom panel: large participation rate $\phi = 1\%$. The critical fraction x_c yielding a maximum impact for a given relaxation factor r is drawn in black.

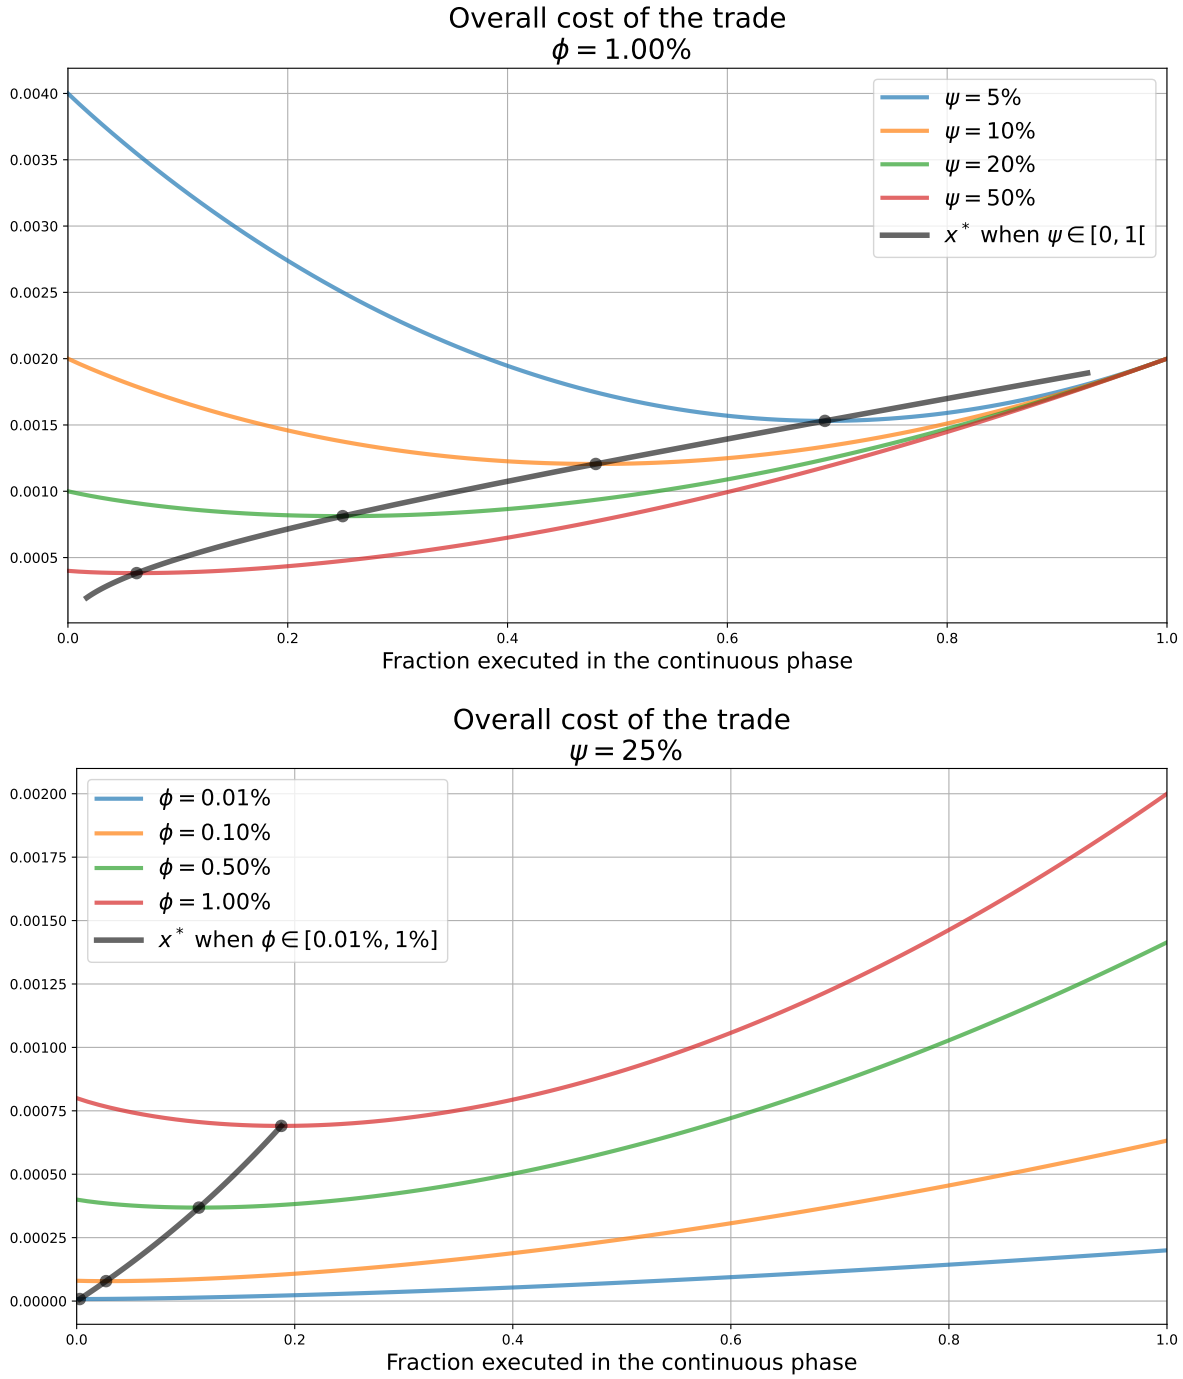

Figure A.3.1 – The total cost of trading (per unit volume) as a function of the executed fraction x during the continuous trading phase. Top panel: large participation rate $\phi = 1\%$ and different values of the auction liquidity ψ. Bottom panel: $\psi = 25\%$ and different values of the participation rate ϕ. The optimal fraction x^* to be executed in the continuous trading phase is drawn in black.

A.4 Conclusion

Observe that ϕ and σ play symmetric roles in Eqs (A.6) and (A.9), and keeping the product $\psi\sigma$ constant does not alter the critical values of x. A degenerate case occurs when an agent minimizes his total cost of trade and accidentally results in the highest possible price impact. In this case $x^* = x_c$ and the corresponding critical relaxation factor r_c is

$$r_c = 1 - \frac{\sqrt{64\tilde{S}\phi + 18\psi^2\sigma^2 - 6\psi\sigma\sqrt{64\tilde{S}^2\phi + 9\psi^2\sigma^2}}}{4\psi\sigma}. \tag{A.10}$$

Using the typical values from above, $r_c \approx 65\%$.

To sum up, in order to optimally split a large order between the continuous trading phase and the closing auction, one should gauge price impact during both phases by accurately measuring the impact prefactors, i.e., the intraday volatility σ, and the (average) auction impact slope $\tilde{S}$. Even if the auction liquidity is extremely large or extremely low, thereby favoring a full execution during the auction or the continuous phase, respectively, one may incur additional costs for not optimally splitting the order size as seen above. Finally, predicting the ratio of the auction volume to the total exchanged volume ψ, as well as its own participation ratio ϕ is not an easy task. Predictive models can approach the desired quantities, e.g., by controlling for seasonality and the effect of special days such as index rebalancing or derivative expiry days.

Bibliography

Frédéric Abergel, Marouane Anane, Anirban Chakraborti, Aymen Jedidi, and Ioane Muni Toke. Limit order books. Cambridge University Press, 2016.

Frédéric Abergel, Côme Huré, and Huyên Pham. Algorithmic trading in a microstructural limit order book model. In Commodities, pages 691–730. Chapman and Hall/CRC, 2022.

Eduardo Abi Jaber and Eyal Neuman. Optimal liquidation with signals: the general propagator case. Available at SSRN 4264823, 2022.

Michael Aitken, Carole Comerton-Forde, and Alex Frino. Closing call auctions and liquidity. Accounting & Finance, 45(4):501–518, 2005.

Christos Alexakis, Vasileios Pappas, and Emmanouil Skarmeas. Market abuse under different close price determination mechanisms: A European case. International Review of Financial Analysis, 74:101707, 2021.

Valentina Alfi, Giorgio Parisi, and Luciano Pietronero. Conference registration: how people react to a deadline. Nature Physics, 3(11):746–746, 2007.

Valentina Alfi, Andrea Gabrielli, and Luciano Pietronero. How people react to a deadline: time distribution of conference registrations and fee payments. Open Physics, 7(3):483–489, 2009.

Aurélien Alfonsi, Antje Fruth, and Alexander Schied. Optimal execution strategies in limit order books with general shape functions. Quantitative Finance, 10(2):143–157, 2010.

Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. Journal of Risk, 3:5–40, 2001.

Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. Direct estimation of equity market impact. Risk, 18(7):58–62, 2005.

AMF. Study of the behaviour of high-frequency traders on Euronext Paris. 2017. Available at https://www.amf-france.org/en/news-publications/publications/reports-research-and-analysis/study-behaviour-high-frequency-traders-euronext-paris, last accessed March 2024.

Matteo Aquilina, Eric Budish, and Peter O'neill. Quantifying the high-frequency trading "arms race". The Quarterly Journal of Economics, 137(1):493–564, 2022.

Fatemeh Aramian and Carole Comerton-Forde. Closing mechanisms in European equities. Available at SSRN 4533545, 2023.

Marco Avellaneda and Michael D. Lipkin. A market-induced mechanism for stock pinning. Quantitative Finance, 3(6):417, 2003.

Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. Quantitative Finance, 8(3):217–224, 2008.

Marco Avellaneda, Gennady Kasyan, and Michael D. Lipkin. Mathematical models for stock pinning near option expiration dates. Communications on Pure and Applied Mathematics, 65 (7):949–974, 2012.

Louis Bachelier. Théorie de la spéculation. In Annales scientifiques de l'École normale supérieure, volume 17, pages 21–86, 1900.

Per Bak, Maya Paczuski, and Martin Shubik. Price variations in a stock market with many agents. Physica A: Statistical Mechanics and its Applications, 246(3-4):430–453, 1997.

Norges Bank. The role of closing auctions in well-functioning markets. 2020. Available at https://www.nbim.no/en/publications/asset-manager-perspectives/2020/the-role-of-closing-auctions-in-well-functioning-markets/, last accessed February 2024.

Mario Bellia. Essays on empirical market microstructure and high frequency data. PhD thesis, Università Ca'Foscari Venezia, 2018.

Mario Bellia, Loriana Pelizzon, Marti G Subrahmanyam, Jun Uno, and Darya Yuferova. Coming early to the party. 2017. SAFE Working Paper. Available at SSRN 3038699.

Michael Benzaquen and Jean-Philippe Bouchaud. Market impact with multi-timescale liquidity. Quantitative Finance, 18(11):1781–1790, 2018a.

Michael Benzaquen and Jean-Philippe Bouchaud. A fractional reaction–diffusion description of supply and demand. The European Physical Journal B, 91:1–7, 2018b.

Paul Besson and Raphaël Fernandez. Better trading at the close thanks to market impact models. 2021. Euronext quantitative research report.

Paul Besson and Matthieu Lasnier. Cumulative market impact of consecutive orders over one and two days: how long does the market remember past trades? Quantitative Finance, 22(1): 1–21, 2022.

Paul Besson, Matthieu Lasnier, and Antoine Falck. The benefits of European periodic auctions beyond MiFID dark trading caps. Journal of Investing, 28(6):91–108, 2019.

Bruno Biais, Pierre Hillion, and Chester Spatt. An empirical analysis of the limit order book and the order flow in the Paris Bourse. The Journal of Finance, 50(5):1655–1689, 1995.

Bruno Biais, Pierre Hillion, and Chester Spatt. Price discovery and learning during the preopening period in the Paris Bourse. Journal of Political Economy, 107(6):1218–1248, 1999.

Fischer Black. Noise. The Journal of Finance, 41(3):528–543, 1986.

Blackrock. A global perspective on market-on-close activity. 2020. Available at https://www.blackrock.com/corporate/literature/whitepaper/viewpoint-a-global-perspective-on-market-on-close-activity-july-2020.pdf, last accessed February 2024.

Jean-Philippe Bouchaud. Price impact. In Encyclopedia of Quantitative Finance, pages 57–160. John Wiley & Sons, 2010.

Jean-Philippe Bouchaud. The inelastic market hypothesis: a microstructural interpretation. Quantitative Finance, 22(10):1785–1795, 2022.

Jean-Philippe Bouchaud and Damien Challet. Why have asset price properties changed so little in 200 years. In Econophysics and Sociophysics: Recent Progress and Future Directions, pages 3–17. Springer, 2017.

Jean-Philippe Bouchaud, Marc Mézard, and Marc Potters. Statistical properties of stock order books: empirical results and models. Quantitative Finance, 2(4):251, 2002.

Jean-Philippe Bouchaud, Yuval Gefen, Marc Potters, and Matthieu Wyart. Fluctuations and response in financial markets: the subtle nature of random price changes. Quantitative Finance, 4(2):176, 2003.

Jean-Philippe Bouchaud, Julien Kockelkoren, and Marc Potters. Random walks, liquidity molasses and critical response in financial markets. Quantitative Finance, 6(02):115–123, 2006.

Jean-Philippe Bouchaud, J Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In Handbook of financial markets: dynamics and evolution, pages 57–160. Elsevier, 2009.

Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier, and Martin Gould. Trades, quotes and prices: financial markets under the microscope. Cambridge University Press, 2018a.

Jean-Philippe Bouchaud, Stefano Ciliberti, Yves Lempérière, Adam Majewski, Philip Seager, and Kevin Sin Ronia. Black was right: Price is within a factor 2 of value. Risk.net, 2018b.

Selma Boussetta, Laurence Daures-Lescourret, and Sophie Moinas. The role of pre-opening mechanisms in fragmented markets. In Paris December 2017 Finance Meeting EUROFIDAI-AFFI, 2017.

Frédéric Bucci, Michael Benzaquen, Fabrizio Lillo, and Jean-Philippe Bouchaud. Crossover from linear to square-root market impact. Physical Review Letters, 122(10):108302, 2019a.

Frédéric Bucci, Iacopo Mastromatteo, Michael Benzaquen, and Jean-Philippe Bouchaud. Impact is not just volatility. Quantitative Finance, 19(11):1763–1766, 2019b.

Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. The Quarterly Journal of Economics, 130(4): 1547–1621, 2015.

Eric Budish, Peter Cramton, Albert S Kyle, Jeongmin Lee, and David Malec. Flow trading. Technical report, National Bureau of Economic Research, 2023.

Francesco Capponi and Rama Cont. Trade duration, volatility, and market impact. Available at SSRN 3351736, 2019.

Mark Cartwright. Trade in medieval europe. World History Encyclopedia, 2019. Available at https://www.worldhistory.org/article/1301/trade-in-medieval-{e}urope/, last accessed January 2024.

Mark Casson and John S Lee. The origin and development of markets: A business history perspective. Business History Review, 85(1):9–37, 2011.

CBOE. CBOE Europe equities guidance note for periodic auctions book. 2020. Available at https://www.cboe.com/europe/equities/trading/periodic_auctions_book/resources, last accessed March 2024.

Anirban Chakraborti, Ioane Muni Toke, Marco Patriarca, and Frédéric Abergel. Econophysics review: I. Empirical facts. Quantitative Finance, 11(7):991–1012, 2011a.

Anirban Chakraborti, Ioane Muni Toke, Marco Patriarca, and Frédéric Abergel. Econophysics review: II. Agent-based models. Quantitative Finance, 11(7):1013–1041, 2011b.

Damien Challet. Strategic behaviour and indicative price diffusion in Paris Stock exchange auctions. In New Perspectives and Challenges in Econophysics and Sociophysics, pages 3–12. Springer, 2019.

Damien Challet and Nikita Gourianov. Dynamical regularities of US equities opening and closing auctions. Market Microstructure and Liquidity, 4(01n02):1950001, 2018.

Damien Challet and Robin Stinchcombe. Analyzing and modeling 1+1d markets. Physica A: Statistical Mechanics and its Applications, 300(1-2):285–299, 2001.

Donald R Chambers, Keith H Black, and Nelson J Lacey. Alternative investments: A primer for investment professionals. CFA Institute Research Foundation, 2018.

Lijian Chen, Kevin E Bassler, Joseph L McCauley, and Gemunu H Gunaratne. Anomalous scaling of stochastic processes and the moses effect. Physical Review E, 95(4):042141, 2017.

Carole Comerton-Forde and Barbara Rindi. Trading @ the close. Available at SSRN 3903757, 2022.

Rama Cont, Sasha Stoikov, and Rishi Talreja. A stochastic model for order book dynamics. Operations Research, 58(3):549–563, 2010.

Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. Journal of Financial Econometrics, 12(1):47–88, 2014.

David M Cutler, James M Poterba, and Lawrence H Summers. What moves stock prices? The Journal of Portfolio Management, 15(3):4–12, 1989.

João da Gama Batista, Jean-Philippe Bouchaud, and Damien Challet. Sudden trust collapse in networked societies. The European Physical Journal B, 88:1–11, 2015.

Joao da Gama Batista, Domenico Massaro, Jean-Philippe Bouchaud, Damien Challet, and Cars Hommes. Do investors trade too much? A laboratory experiment. Journal of Economic Behavior & Organization, 140:18–34, 2017.

Lorenzo Dall'Amico, Antoine Fosset, Jean-Philippe Bouchaud, and Michael Benzaquen. How does latent liquidity get revealed in the limit order book? Journal of Statistical Mechanics: Theory and Experiment, 2019(1):013404, 2019.

Joffrey Derchu, Philippe Guillot, Thibaut Mastrolia, and Mathieu Rosenbaum. AHEAD: Ad-Hoc Electronic Auction Design. arXiv preprint arXiv:2010.02827, 2020.

Joffrey Derchu, Dimitrios Kavvathas, Thibaut Mastrolia, and Mathieu Rosenbaum. Equilibria and incentives for illiquid auction markets. arXiv preprint arXiv:2307.15805, 2023.

Mike Derksen, Bas Kleijn, and Robin De Vilder. Clearing price distributions in call auctions. Quantitative Finance, 20(9):1475–1493, 2020.

Mike Derksen, Bas Kleijn, and Robin De Vilder. Heavy tailed distributions in closing auctions. Physica A: Statistical Mechanics and its Applications, 593:126959, 2022a.

Mike Derksen, Bas Kleijn, and Robin De Vilder. MiFID II and European Closing Auctions. Available at SSRN 4074752, 2022b.

Jonathan Donier and Julius Bonart. A million metaorder analysis of market impact on the bitcoin. Market Microstructure and Liquidity, 1(02):1550008, 2015.

Jonathan Donier and Jean-Philippe Bouchaud. Why do markets crash? Bitcoin data offers unprecedented insights. PloS one, 10(10):e0139356, 2015.

Jonathan Donier and Jean-Philippe Bouchaud. From Walras' auctioneer to continuous time double auctions: A general dynamic theory of supply and demand. Journal of Statistical Mechanics: Theory and Experiment, 2016(12):123406, 2016.

Jonathan Donier, Julius Bonart, Iacopo Mastromatteo, and Jean-Philippe Bouchaud. A fully consistent, minimal model for non-linear market impact. Quantitative Finance, 15(7):1109–1121, 2015.

Mouhamad Drame. Limit Order Book (LOB) shape modeling in presence of heterogeneously informed market participants. Market Microstructure and Liquidity, 5(01n04):2050007, 2019.

Zoltan Eisler, Jean-Philippe Bouchaud, and Julien Kockelkoren. The price impact of order book events: market orders, limit orders and cancellations. Quantitative Finance, 12(9):1395–1419, 2012.

Sofiene El Aoud and Frédéric Abergel. A stochastic control approach to option market making. Market Microstructure and Liquidity, 1(01):1550006, 2015.

EUROFIDAI. European High Frequency database - data description guide, 2020. Accessible at https://www.eurofidai.org/sites/default/files/inline-files/High%20Frequency%20data%20descriptions%20guide.pdf, last accessed March 2024.

Euronext. Euronext frequently asked questions. 2019. Available at https://www.euronext.com/sites/default/files/2019-09/52118_Euronext-FAQ-2019_v07_0.pdf, last accessed March 2024.

Euronext. Euronext rule book, Book 1: Harmonised Rules. 2023a. Available at https://www.euronext.com/en/regulation/euronext-regulated-markets, last accessed February 2024.

Euronext. Trading manual for the OPTIQ trading platform. 2023b. Available at https://www.euronext.com/sites/default/files/2022-03/Notice%204-01%20Trading%20Manual%20with%20universal%20reservation_2022_03_14.pdf, last accessed February 2024.

Eugene F Fama. Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2):383–417, 1970.

J Doyne Farmer, Laszlo Gillemot, Fabrizio Lillo, Szabolcs Mike, and Anindya Sen. What really causes large price changes? Quantitative Finance, 4(4):383–397, 2004.

J Doyne Farmer, Paolo Patelli, and Ilija I Zovko. The predictive power of zero intelligence in financial markets. Proceedings of the National Academy of Sciences, 102(6):2254–2259, 2005.

J Doyne Farmer, Austin Gerig, Fabrizio Lillo, and Szabolcs Mike. Market efficiency and the long-memory of supply and demand: Is price impact variable and permanent or fixed and temporary? Quantitative Finance, 6(02):107–112, 2006.

Antoine Fosset. Endogenous liquidity crises in financial markets. PhD thesis, Institut Polytechnique de Paris, 2020.

Karl Frauendorfer and Louis Müller. Liquidity-related price sensitivities of closing auctions in equity markets. 2020. Working paper accessible at https://www.alexandria.unisg.ch/handle/20.500.14171/111583, last accessed March 2023.

Christoph Frei and Joshua Mitra. Optimal closing benchmarks. Finance Research Letters, 40: 101674, 2021.

John Kenneth Galbraith. A short history of financial euphoria. Penguin, 1994.

Xuefeng Gao and SJ Deng. Hydrodynamic limit of order-book dynamics. Probability in the Engineering and Informational Sciences, 32(1):96–125, 2018.

Matthieu Garcin. Estimation of time-dependent hurst exponents with variational smoothing and application to forecasting foreign exchange rates. Physica A: Statistical Mechanics and its Applications, 483:462–479, 2017.

Gao-Feng Gu, Fei Ren, Xiao-Hui Ni, Wei Chen, and Wei-Xing Zhou. Empirical regularities of opening call auction in Chinese stock market. Physica A: Statistical Mechanics and its Applications, 389(2):278–286, 2010.

Olivier Guéant. The Financial Mathematics of Market Liquidity: From optimal execution to market making, volume 33. CRC Press, 2016.

Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Dealing with the inventory risk: a solution to the market making problem. Mathematics and Financial Economics, 7:477–507, 2013.

Bernard Guerrien and Ozgur Gun. L'étrange silence du Nobel Prize Committee sur la «théorie des marchés efficients». Revue de la régulation. Capitalisme, institutions, pouvoirs, (14), 2013.

Larry Harris. Trading and exchanges: Market microstructure for practitioners. OUP USA, 2003.

Joel Hasbrouck. Measuring the information content of stock trades. The Journal of Finance, 46 (1):179–207, 1991.

Natascha Hey, Jean-Philippe Bouchaud, Iacopo Mastromatteo, Johannes Muhle-Karbe, and Kevin Webster. The cost of misspecifying price impact. arXiv preprint arXiv:2306.00599, 2023.

Pierre Hillion and Matti Suominen. The manipulation of closing prices. Journal of Financial Markets, 7(4):351–375, 2004.

Narasimhan Jegadeesh and Yanbin Wu. Closing auctions: Nasdaq versus NYSE. Journal of Financial Economics, 143(3):1120–1139, 2022.

Eugene Kandel, Barbara Rindi, and Luisella Bosetti. The effect of a closing call auction on market quality and trading strategies. Journal of Financial Intermediation, 21(1):23–49, 2012.

Vijay Krishna. Auction theory. Academic press, 2009.

Albert S Kyle. Continuous auctions and insider trading. Econometrica: Journal of the Econometric Society, pages 1315–1335, 1985.

Aimé Lachapelle, Jean-Michel Lasry, Charles-Albert Lehalle, and Pierre-Louis Lions. Efficiency of the price formation process in presence of high frequency participants: a mean field game analysis. Mathematics and Financial Economics, 10:223–262, 2016.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. Japanese Journal of Mathematics, 2(1):229–260, 2007.

Alexandre Laumonier. 6/5. Zones Sensibles Editions, 2014.

Charles-Albert Lehalle. Financial markets (and more) for quants, 2022. Available at https://bit.ly/lehalle-calendar-2022, last accessed March 2024.

Charles-Albert Lehalle and Sophie Laruelle. Market microstructure in practice. World Scientific, 2018.

Ismael Lemhadri. Price impact in a latent order book. Market Microstructure and Liquidity, 5 (01n04):2050004, 2019.

Jiayi Li, Sumei Luo, and Guangyou Zhou. Call auction, continuous trading and closing price formation. Quantitative Finance, 21(6):1037–1065, 2021.

Fabrizio Lillo and J Doyne Farmer. The long memory of the efficient market. Studies in nonlinear dynamics & econometrics, 8(3), 2004.

Fabrizio Lillo and J Doyne Farmer. The key role of liquidity fluctuations in determining large price changes. Fluctuation and Noise Letters, 5(02):L209–L216, 2005.

Fabrizio Lillo, J Doyne Farmer, and Rosario N. Mantegna. Master curve for price-impact function. Nature, 421(6919):129–130, 2003.

Fabrizio Lillo, Szabolcs Mike, and J Doyne Farmer. Theory for long memory in supply and demand. Physical Review E, 71(6):066122, 2005.

London Stock Exchange. Guide to the trading system. 2018. Available at https://docs.londonstockexchange.com/sites/default/files/documents/mit201-guidetotradingservicesv146andtradecho.pdf, last accessed March 2024.

Ananth Madhavan. Trading mechanisms in securities markets. The Journal of Finance, 47(2): 607–641, 1992.

Rosario N Mantegna and H Eugene Stanley. Introduction to econophysics: correlations and complexity in finance. Cambridge University Press, 1999.

Riccardo Marcaccioli, Jean-Philippe Bouchaud, and Michael Benzaquen. Exogenous and endogenous price jumps belong to different dynamical classes. Journal of Statistical Mechanics: Theory and Experiment, 2022(2):023403, 2022.

Sergei Maslov. Simple model of a limit order-driven market. Physica A: Statistical Mechanics and its Applications, 278(3-4):571–578, 2000.

Iacopo Mastromatteo, Bence Toth, and Jean-Philippe Bouchaud. Agent-based models for latent liquidity and concave price impact. Physical Review E, 89(4):042805, 2014.

R Preston McAfee and John McMillan. Auctions and bidding. Journal of Economic Literature, 25(2):699–738, 1987.

Haim Mendelson. Market behavior in a clearing house. Econometrica: Journal of the Econometric Society, pages 1505–1524, 1982.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. Journal of Economic Theory, 26(1):17–27, 1982.

Paul Robert Milgrom. Putting auction theory to work. Cambridge University Press, 2004.

Esteban Moro, Javier Vicente, Luis G Moyano, Austin Gerig, J Doyne Farmer, Gabriella Vaglica, Fabrizio Lillo, and Rosario N Mantegna. Market impact and trading profile of hidden orders in stock markets. Physical Review E, 80(6):066102, 2009.

Ioane Muni Toke. The order book as a queueing system: average depth and influence of the size of limit orders. Quantitative Finance, 15(5):795–808, 2015a.

Ioane Muni Toke. Exact and asymptotic solutions of the call auction problem. Market Microstructure and Liquidity, 1(01):1550001, 2015b.

NobelPrize.org. The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel, 2020. Available at https://www.nobelprize.org/prizes/economic-sciences/2020/prize-announcement/, last accessed February 2024.

Anna A Obizhaeva and Jiang Wang. Optimal trading strategy and supply/demand dynamics. Journal of Financial Markets, 16(1):1–32, 2013.

Michael S. Pagano and Robert A. Schwartz. A closing call's impact on market quality at Euronext Paris. Journal of Financial Economics, 68(3):439–484, 2003.

Seongkyu Gilbert Park, Wing Suen, and Kam-Ming Wan. Call auction mechanism and closing price manipulation: evidence from the Hong Kong stock exchange. Available at SSRN 3482351, 2020.

Jusselin Paul, Mastrolia Thibaut, and Rosenbaum Mathieu. Optimal auction duration: A price formation viewpoint. Operations Research, 69(6):1734–1745, 2021.

Vasiliki Plerou, Parameswaran Gopikrishnan, Xavier Gabaix, and H Eugene Stanley. Quantifying stock-price response to demand fluctuations. Physical Review E, 66(2):027104, 2002.

Marc Potters and Jean-Philippe Bouchaud. More statistical properties of order books and price impact. Physica A: Statistical Mechanics and its Applications, 324(1-2):133–140, 2003.

Franck Raillon. The growing importance of the closing auction in share trading volumes. Journal of Securities Operations & Custody, 12(2):135–152, 2020.

Juan C Reboredo. The switch from continuous to call auction trading in response to a large intraday price movement. Applied Economics, 44(8):945–967, 2012.

Emilio Said, Ahmed Bel Hadj Ayed, Alexandre Husson, and Frédéric Abergel. Market impact: A systematic study of limit orders. Market Microstructure and Liquidity, 3(03n04):1850008, 2017.

Mohammed Salek, Damien Challet, and Ioane Muni Toke. Price impact in equity auctions: zero, then linear. arXiv preprint arXiv:2301.05677, 2023. To appear in Market Microstructure and Liquidity.

Mohammed Salek, Damien Challet, and Ioane Muni Toke. Equity auction dynamics: latent liquidity models with activity acceleration. arXiv preprint arXiv:2401.06724, 2024. To appear in Quantitative Finance.

Yuki Sato and Kiyoshi Kanazawa. Inferring microscopic financial information from the long memory in market-order flow: A quantitative test of the Lillo-Mike-Farmer model. Physical Review Letters, 131(19):197401, 2023.

Robert J Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? The American Economic Review, 71(3):421–436, 1981.

Germain Sicard. Les moulins de Toulouse au Moyen Age: aux origines des sociétés anonymes. Number vol. 6 in Affaires et gens d'affaires. A. Colin, 1953.

Eric Smith, J Doyne Farmer, László Gillemot, and Supriya Krishnamurthy. Statistical theory of the continuous double auction. Quantitative Finance, 3(6):481, 2003.

Karline Soetaert and Filip Meysman. Solving partial differential equations, using R package ReacTran. 2009.

Karline Soetaert, Thomas Petzoldt, and R Woodrow Setzer. Solving differential equations in R: package deSolve. Journal of Statistical Software, 33:1–25, 2010.

Serge Svizzero and Clement Tisdell. Barter and the origin of money and some insights from the ancient palatial economies of Mesopotamia and Egypt. hal preprint hal-02274856, 2019.

Richard H Thaler. Anomalies: The winner's curse. Journal of Economic Perspectives, 2(1): 191–202, 1988.

Erik Theissen and Christian Westheide. Call of duty: Designated market maker participation in call auctions. Journal of Financial Markets, 49:100530, 2020.

Bence Tóth, Yves Lemperiere, Cyril Deremble, Joachim De Lataillade, Julien Kockelkoren, and Jean-Philippe Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. Physical Review X, 1(2):021006, 2011.

Bence Toth, Imon Palit, Fabrizio Lillo, and J Doyne Farmer. Why is equity order flow so persistent? Journal of Economic Dynamics and Control, 51:218–239, 2015.

Bence Tóth, Zoltán Eisler, and J-P Bouchaud. The square-root impace law also holds for option markets. Wilmott, 2016(85):70–73, 2016.

Jean-Claude Trichet. Reflections on the nature of monetary policy non-standard measures and finance theory, 2010. URL https://www.ecb.europa.eu/press/key/date/2010/html/sp101118.en.html. Speech by the President of the ECB, at the 20th Frankfurt European Banking Congress.

Michele Vodret, Iacopo Mastromatteo, Bence Tóth, and Michael Benzaquen. A stationary Kyle setup: microfounding propagator models. Journal of Statistical Mechanics: Theory and Experiment, 2021(3):033410, 2021.

Michele Vodret, Iacopo Mastromatteo, Bence Tóth, and Michael Benzaquen. Do fundamentals shape the price response? A critical assessment of linear impact models. Quantitative Finance, 22(12):2139–2150, 2022.

Dmitry Voyakin. The great silk road. UNESCO. Available at https://en.unesco.org/silkroad/knowledge-bank/great-silk-road, last accessed January 2024.

Léon Walras. Éléments d'économie politique pure: ou, Théorie de la richesse sociale. F. Rouge, 1900.

Ernst Juerg Weber. A short history of derivative security markets. In Vinzenz Bronzin's option pricing models: Exposition and appraisal, pages 431–466. Springer, 2009.

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

Kevin T Webster. Handbook of Price Impact Modeling. CRC Press, 2023.

Xetra. Market model for the trading venue Xetra. 2021. Available at https://www.xetra.com/resource/blob/2762072/d8e4d932939c6826e4f8391b568e754c/data/T7_Release_10.0_-_Market_Model-_Xetra.pdf, last accessed March 2024.

Elia Zarinelli, Michele Treccani, J Doyne Farmer, and Fabrizio Lillo. Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. Market Microstructure and Liquidity, 1(02):1550004, 2015.