

A Tale of One Exchange and Two Order Books: Effects of Fragmentation in the Absence of Competition

Alejandro Bernales* Italo Riarte[†] Satchit Sagade[‡] Marcela Valenzuela[§]
Christian Westheide[¶]

First version: August 2016

This version: May 2017

Abstract

Exchanges nowadays routinely operate multiple limit order markets for the same security that are almost identically structured. We study the effects of such fragmentation on market performance using a dynamic model of fragmented markets where agents trade strategically across two identically-organized limit order books. We show that fragmented markets, in equilibrium, offer higher welfare to intermediaries at the expense of investors with intrinsic trading motives, and lower liquidity than consolidated markets. Consistent with our theory, we document improvements in liquidity and lower profits for liquidity providers when Euronext, in 2009, consolidated its order flow for stocks traded across multiple, country-specific, and identically-organized limit order books onto a single order book. Our results suggest that competition in market design, not fragmentation, drives previously documented improvements in market quality when new trading venues emerge; in the absence of such competition, market fragmentation is harmful.

Keywords: Fragmentation, Competition, Liquidity, Price Efficiency

JEL Classification: G10, G12

*University of Chile (DII), abernales@dii.uchile.cl

[†]University of Chile (DII)

[‡]Department of Finance and Research Center SAFE, Goethe University Frankfurt, sagade@safe.uni-frankfurt.de

[§]University of Chile (DII), mvalenzuela@dii.uchile.cl

[¶]Finance Area, University of Mannheim, and Research Center SAFE, Goethe University Frankfurt, westheide@uni-mannheim.de

For helpful comments and discussions we thank Jonathan Brogaard, Peter Gomber, Jan-Pieter Krahen, Katya Malinova, Albert Menkveld, Andreas Park, Talis Putnins, Ioanid Rosu, Erik Theissen, and Vincent van Kervel, conference participants at the 2014 Market Microstructure: Confronting Many Viewpoints Conference, 2016 SAFE Market Microstructure Workshop, 2016 CMStatistics Conference, 2016 India Finance

Conference, 2017 Securities Markets: Trends, Risks and Policies Conference, and seminar participants at University of Birmingham, University of Mannheim, University of Manchester, University of Frankfurt, University of Chile, Pontifical Catholic University of Chile. Sagade and Westheide gratefully acknowledge research support from the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE. Valenzuela acknowledges the support of Fondecyt Project No. 11140541 and Instituto Milenio ICM IS130002.

When you split these liquidity pools [...] what happens is that overall volumes tend to go up because the market starts to arbitrage and tries to put the market back together, the value of data goes up. And the whole thing for us turns out to be very good business [...] we don't think it's in the best interest of the market [...]

– Jeffrey Sprecher, Chairman and CEO, Intercontinental Exchange during the Q1 2017 Earnings Call dated 03 May 2017

1. Introduction

Increased fragmentation of trading activity has been one of the most significant changes experienced by equity markets in recent years. Equity markets in the United States, the European Union, and elsewhere have evolved from national/regional stock exchanges being the dominant liquidity pools to a fragmented multi-market environment where a stock now trades on multiple exchanges. These markets have simultaneously also experienced a process of consolidation as a result of national and international mergers of exchanges such that only a small number of operators, each running several exchanges, now compete with one another. For example, in the United States, the three large exchange operators – Intercontinental Exchange, Nasdaq OMX, and BATS – currently operate a total of ten lit equity exchanges. While it is possible that exchange operators allow a certain degree of competition between the different exchanges they own, it appears implausible that such competition would be similar to that between exchanges run by different operators. In most cases, the individual exchanges operated by a single operator employ almost identical rules and use the same technology such that differences between exchanges are minimal. This raises the question as to the effects of fragmentation when competition between venues is absent or minimal.

In this paper, we examine the effects of fragmentation on market performance through a dynamic equilibrium model which characterizes such a multi-market environment. Our

model is set up as a stochastic trading game in which a single asset can be traded in two identically-organized limit order markets. Agents, who are heterogeneous in terms of their intrinsic economic reasons to trade the asset, enter the market following a Poisson process, and make endogenous trading decisions depending on market conditions (e.g. where to submit an order, the type of order, and the limit price). Agents can reenter the market to revise or cancel previously submitted limit orders. They make optimal decisions depending on the state of both limit order books, the stochastically evolving fundamental value of the asset, their private values, and costs of delaying order execution. Limit orders in both order books are independently executed based on price and time priority. By comparing a multi-market environment to a consolidated market setup, we analyze the effects of fragmentation across multiple venues when these venues do not actively compete with each other.

Our model builds on those developed by [Goettler et al. \(2005, 2009\)](#) to characterize a single limit order market. They present a dynamic model in which investors make asynchronous trading decisions based on the prevailing market conditions. We extend their model to describe a fragmented limit order market setting. This is a non-trivial task as the diversity of trading options and trading rules in this setting significantly increases the decision-state space. Furthermore, in contrast to [Goettler et al. \(2005, 2009\)](#), we do not rely on model simplifications to reduce this large state space.

We focus on liquidity, price efficiency, and welfare. In the model, agents endogenously decide whether they provide or consume liquidity. Agents who have an intrinsic motive to trade balance the delay costs associated with submitting limit orders and immediacy costs associated with submitting market orders when determining their optimal strategy. Agents with large absolute private values are more likely to submit market orders because of the proportionally higher expected delay costs. Agents with no intrinsic trading motives generate their profits solely from the trading process. Consequently, they are more patient and hence act as intermediaries by either submitting new limit orders, or sniping mispriced limit orders as in [Budish et al. \(2015\)](#).

In a fragmented environment, agents who provide liquidity submit less aggressive limit orders than in a consolidated market because they can submit an order to one market in order to avoid the time priority of standing limit orders in the second market. This reduction in competition among liquidity providers in a fragmented market translates into higher immediacy costs for liquidity demanding agents.

A comparison of welfare observed in the two different market setups shows that aggregate welfare does not differ markedly between a consolidated and fragmented market. However, the distribution of welfare between the different agent types changes, primarily due to lower price competition in fragmented markets. Agents without any intrinsic trading motive are better off in a fragmented market; their expected payoffs are significantly higher as they obtain better terms of trade. Conversely, fragmented markets are welfare-reducing for agents with exogenous trading motives due to higher costs of obtaining immediacy.

Agents' order submission strategies in fragmented versus consolidated markets have a direct impact on liquidity and price discovery. We find that quoted spread and top-of-book depth are higher in the multi-market environment. We also observe that actual trading costs, proxied using effective spreads, and liquidity providers trading gains, proxied using realized spreads, are lower in a single market setup. At the same time, microstructure noise, defined as the absolute difference between quote midpoint and the fundamental value of the asset, is also higher when markets are fragmented. The above results hold irrespective of whether we measure liquidity and microstructure noise using local or inside quotes. These results also assume exogenous market entry and constant agent populations in both scenarios.

If we were to endogenize market entry of different agent types by allowing them to make entry decisions based on the trade-off between expected trading profits and participation costs, the higher profits in fragmented markets earned by agents without any intrinsic motive should lead to their increased participation. In a computationally simpler alternative, we re-parameterize the model by doubling the number of such agents in the fragmented market and compare its outcomes to those observed under the original parameterization. We find

that quoted bid-ask spreads – albeit lower than in the earlier discussed fragmented market – remain higher than in the single market. Quoted depth in this setup is also highest across the three scenarios. Effective and realized spreads remain higher than in the single market. Conversely, price efficiency improves in this setup because the presence of a higher number of intermediaries leads to prices reacting faster to the arrival of public information. Finally, we observe an incremental shift in welfare towards agents without intrinsic trading motives when their arrival rate is doubled, while aggregate welfare does not change significantly.

We empirically test the model predictions by examining a unique event in which Euronext, starting 14 January 2009, implemented a single order book per asset for their Paris, Amsterdam, and Brussels markets. Euronext previously operated multiple independent order books for stocks cross-listed on these markets. The event led to a decrease in fragmentation for the affected stocks. Existing empirical studies, such as [Foucault and Menkveld \(2008\)](#), [Hengelbrock and Theissen \(2009\)](#) and [Chlistalla and Lutat \(2011\)](#), examining the effects of new exchange operators entering a market can be viewed as joint tests of fragmentation and competition. This is because the entry of a new market, in addition to increasing fragmentation, also materially alters the competitive environment. The new operator typically attempts to differentiate its platform along critical features such as trading speed, transaction fees, or the ability to execute large blocks. In contrast, the multiple order books operated by Euronext had exactly identical trading protocols before the implementation of a single order book.

The empirical analysis broadly confirms the theoretical results. We find quoted spreads in the consolidated market to be lower by 30% than local spreads in an individual order book before the event. Quoted depth (both local and at the inside quotes) is also higher after consolidation but the results are statistically insignificant. This is consistent with the empirical level of intermediation in fragmented markets being in between the two theoretically modeled scenarios. Consistent with our theoretical results, effective spreads, both measured using local and inside quotes, are smaller after consolidation. Higher competition in the single order book reduces the potential for rent extraction by liquidity providers, resulting

in 35% lower realized spreads after consolidation. Price impact, the other component of the effective spread, in the absence of private information measures the extent of trading at stale prices, and remains unchanged when compared to the price impact based on inside quote midpoints in the fragmented market. Price efficiency, measured using autocorrelations and variance ratios, also improves after consolidation, although the improvements are weakly significant at best.

While we are unable to empirically compute welfare effects, we find that the introduction of a single order book leads to a weakly significant increase in trading volume. This is despite the elimination of arbitrage trades between the multiple Euronext markets, which are responsible for up to 7.8% of the trading volume before the introduction of a single order book. This is likely due to reduced transaction costs allowing more participation by investors with intrinsic trading motives and is consistent with our theoretical results.

Our results contribute to the literature on equity market fragmentation.¹ Early theories on fragmentation such as [Mendelson \(1987\)](#), [Pagano \(1989\)](#), [Chowdhry and Nanda \(1991\)](#) highlight the positive network externalities generated by consolidating trading on a single venue. [Harris \(1993\)](#) argues that fragmentation can emerge as a consequence of real-world frictions and heterogenous trading motives. Even in some of the above models, a consolidated market is no longer the equilibrium outcome when the fragmented markets differ in their absorptive capacity and institutional mechanisms ([Pagano, 1989](#)), and when traders are allowed to split their orders over time ([Chowdhry and Nanda, 1991](#)). [Madhavan \(1995\)](#) argues that markets fragment only if there is a lack of trade disclosure. Fragmentation in his model benefits dealers and large traders, and increases volatility and price inefficiency. In possibly the most relevant study to today's competitive landscape of equity markets, [Foucault and Menkveld \(2008\)](#) model competition between two limit order books and predict that the entry of a second market increases consolidated depth, and that increased use of smart order routers leads to an increase in liquidity in the entrant market.

¹ See [Gomber et al. \(2016\)](#) for a detailed survey of this literature.

The empirical study closest to our paper is [Amihud et al. \(2003\)](#) who study the reduction in fragmentation on the Tel Aviv Stock Exchange resulting from the exercise of deep in-the-money share warrants and find an increase in stock price and improvement in liquidity. However, their results cannot be extended to modern equity markets because: (i) the stocks and warrants traded periodically in single or multiple batch auctions as opposed to continuously in limit order markets; (ii) the warrant and the underlying stock cannot be considered as perfectly fungible assets such that investors are indifferent between holding the two.

[Hengelbrock and Theissen \(2009\)](#) and [Chlistalla and Lutat \(2011\)](#) analyze the market entry of Turquoise and Chi-X, respectively, in the European markets and find positive effects on liquidity in the main market. [Boehmer and Boehmer \(2003\)](#) and [Nguyen et al. \(2007\)](#) examine the impact of NYSE's entry in the ETF market and also find improvements in different measures of liquidity. [Riordan et al. \(2010\)](#) find that new entrants contribute to the majority of quote-based price discovery for the FTSE100 stocks in the UK. [Kohler and von Wyss \(2012\)](#) and [Hellström et al. \(2013\)](#) find that fragmentation in the Swedish market increases liquidity, for all but large stocks, and price efficiency for all stocks. [O'Hara and Ye \(2011\)](#) analyze overall fragmentation in the US equity markets and find that it is not harmful to market quality. [Degryse et al. \(2015\)](#) and [Gresse \(2017\)](#) differentiate between lit and dark fragmentation and find that the former improves liquidity, but disagree on the effects of the latter.

We contribute to this literature by analyzing the impact of fragmentation across multiple, identically-organized limit order books on market performance. We consider a dynamic model of multiple limit order markets that incorporates several real-world features and allows for more flexible agent behavior as compared to previous models (see for example [Mendelson, 1987](#); [Pagano, 1989](#); [Chowdhry and Nanda, 1991](#); [Biais, 1993](#); [Parlour and Seppi, 2003](#)). We provide evidence that fragmentation has detrimental effects on market quality and welfare, benefiting intermediaries at the expense of agents who trade for intrinsic motives.

The remainder of the paper is structured as follows. Section 2 describes the theoretical model central to our analyses. In Section 3, we analyze the theoretical implications of consolidated versus fragmented markets on welfare and market quality. In Section 4 we present the empirical results from the event study. Finally, we conclude in Section 5.

2. Multi-Market Model

2.1 Model Setting

Consider an economy in continuous-time with a single financial asset that is traded on two independent financial markets. The economy is populated by risk-neutral agents trading the asset. Agents arrive sequentially following a Poisson process with intensity λ , and they can use either of the two financial markets to trade the asset. Agents do not cooperate, and they make trading decisions based on a maximization of expected payoffs. Hence, trading activity in the two financial markets reflects a sequential non-cooperative game, where agents make asynchronous decisions by taking into account private reasons to trade the asset, market conditions and the potential strategies employed by other agents arriving in the future.

The two financial markets in the economy, denoted by $m \in \{1, 2\}$, are organized as limit order markets. Agents can submit limit orders and market orders. A limit order is a commitment made by an agent to trade the asset at a price p in the future, where the value of p is decided by the agent at order submission time. A market order is an order to buy or sell immediately at the best available price, where this price is provided by a previously submitted limit order. Hence, a buy (sell) market order submitted by an agent is always matched with a sell (buy) limit order previously submitted by another agent. Agents submitting limit orders are liquidity providers, whereas agents submitting market orders are liquidity consumers.

As in limit order markets found in the real world, the order books are described by a discrete set of prices at which orders can be submitted. The limit order book at time t and in

market m , $L_{m,t}$, is characterized by the set of prices denoted by $\{p_m^i\}_{i=-N_m}^{N_m}$, where $p_m^i < p_m^{i+1}$ and N is a finite number. Let d be the distance between any two consecutive prices, which will be referred to as tick size (i.e. $d = p_m^{i+1} - p_m^i$). The tick size is assumed to be equal for both limit order books. In both limit order books, there is a queue of unexecuted buy or sell limit orders associated with each price. Let $l_{m,t}^i$ be the queue in the limit order market m at time t associated with price p_m^i . A positive (negative) number in $l_{m,t}^i$ denotes the number of buy (sell) unexecuted limit orders, and it represents the depth of the book $L_{m,t}$ at price p_m^i . Thus, in the book $L_{m,t}$ at time t , the best bid price is $B(L_{m,t}) = \sup\{p_m^i | l_{m,t}^i > 0\}$ and the best ask price is $A(L_{m,t}) = \inf\{p_m^i | l_{m,t}^i < 0\}$. If the order book $L_{m,t}$ is empty at time t on the buy side or on the sell side, $B(L_{m,t}) = -\infty$ or $A(L_{m,t}) = \infty$, respectively. All agents observe both limit order books (i.e. prices and depths at each price) before making any trading decision.

In each market, the limit order book respects price and time priority for the execution of limit orders. In the book $L_{m,t}$, limit orders submitted earlier at the same price p_m^i are executed first, and buy (sell) limit orders at higher (lower) prices have priority in the queue, even if other orders with less competitive prices are submitted earlier. Time and price priority apply independently for each limit order book.² The limit order price determines whether an order is a market order: an order to buy (sell) at a price equal to or above (below) the best ask (bid) price is a market order and is executed immediately at the best ask (bid) price.

Agents can monitor both limit order books. However, due to limited cognition, they cannot immediately modify their unexecuted limit orders after a change in market conditions. In that sense, decisions regarding limit order submissions are sticky. Traders re-enter the market to modify unexecuted limit orders according to a Poisson processes with parameter λ_r , which is the same for both markets and is independent of the arrival process.

Agents are heterogeneous in terms of their intrinsic economic motives to trade the asset.

² The existence of an order protection rule ensuring price priority across order books does not affect the outcomes of the model.

These motives are reflected in their private values. Each agent has a private value, α , which is known by the agent. α is drawn from the discrete vector $\Psi = \{\alpha_1, \alpha_2, \dots, \alpha_g\}$ using a discrete distribution, F_α , where g is a finite integer. Private values reflect the fact that agents would like to trade for various reasons unrelated to the fundamental value of the asset (e.g. hedging needs, tax exposures and/or wealth shocks). They are idiosyncratic and constant for each agent.

Agents face a cost when they cannot immediately trade the asset, which is called a delaying cost. The delaying cost is reflected by a discount rate ρ applied to the agent's payoff (with $0 < \rho < 1$). The value ρ is constant and has the same value whether orders are executed in $L_{1,t}$ or $L_{2,t}$. This delaying cost does not represent the time value of the money. Instead, it reflects opportunity costs and the cost of monitoring the market until an order is executed.

The fundamental value of the asset, v_t , is stochastic and known by agents; its innovations follow an independent Poisson process with parameter λ_v . In case of an innovation, the fundamental value increases or decreases by d , both with an equal probability of 0.5, where d is the tick size of the limit order books.

The heterogeneity of agents (in terms of private values), the delaying costs and the fundamental value of the asset all play an important role in agents' trading behavior. On the one hand, suppose agent x with a positive private value (i.e. $\alpha > 0$) arrives at time t_x . This agent has to be a buyer because she would like to have the asset to obtain the intrinsic benefit given by α . In this case, the agent's expected payoff of trading one share is: $(\alpha + v_{t'} - p)e^{-\rho(t' - t_x)}$, where p is the transaction price, t' is the expected time of the transaction, and $v_{t'}$ is the expected fundamental value of the asset at time t' . Moreover, if the value of α is very high, the agent may also prefer to buy the asset *as soon as possible* in order to avoid a high delaying cost (i.e. the agent has a discount on the level of α given by $(e^{-\rho(t' - t_x)} - 1)\alpha$). She may even prefer to buy the asset immediately using a market order. Consequently, an agent with a high positive private value will probably be a liquidity consumer. However,

there is no free lunch for the liquidity consumer. The agent will probably have to pay an immediacy cost that is given by $(v_{t'} - p)^{-\rho(t' - t_x)}$, since it is likely that $v_{t'} - p < 0$. The agent will accept this immediacy cost because she is mainly generating her profits from the large private value, α , rather than from the transaction *per se*.³

On the other hand, suppose an agent y with a private value equal to zero (i.e. $\alpha = 0$) arrives at time t_y . This agent needs to find a profitable opportunity purely in the transaction process because she does not obtain any intrinsic economic benefits from trading. Consequently, she is willing to wait until she obtains a good price relative to the fundamental value. Thus, this agent will probably act as a liquidity provider and receive the immediacy cost paid by the liquidity consumer. It is important to note that agents with $\alpha = 0$ are indifferent with respect to taking either side of the market because they can maximize their benefits by either selling or buying (i.e. by respectively maximizing $(p - v_{t''})e^{-\rho(t'' - t_y)}$ or $(v_{t''} - p)e^{-\rho(t'' - t_y)}$, where t'' is the expected time of the transaction).

Liquidity providers are also affected by the so-called picking-off risk because limit orders can also generate a negative payoff if they are in an unfavorable position relative to the fundamental value. A limit buy (sell) order executed above (below) the fundamental value of the asset generates a negative economic benefit in the transaction. For example, suppose that the agent I with $\alpha = 0$ first arrives at time $t = 0$. Additionally, suppose that this agent has a standing limit buy order at the best bid price, B in market $m = 1$. Suppose that the current time is t^* and v_{t^*} is the current fundamental value of the asset, such that $v_{t^*} > B$. In this case, the agent can make a positive profit if the order is executed immediately at time t^* ; this potential profit is given by $(v_{t^*} - B)e^{-\rho t^*}$. Now suppose at time t^{**} , the fundamental value of the asset decreases to level $v_{t^{**}}$, which is below B (i.e. $v_{t^{**}} < B$) and simultaneously agent II with private value $\alpha = 0$ arrives in the market. Since agent I cannot immediately modify her unexecuted limit order, agent II can submit a market sell order, and pick off the limit

³ A similar example can be explained in the other direction in case of an agent with a negative private value (i.e. $\alpha < 0$) having a preference to sell.

buy order submitted by agent *I*. Agent *II* is thus able to generate an instantaneous profit equal to $(B - v_{t^{**}})$ whereas agent *I* has a negative realized payoff given by $(v_{t^{**}} - B)e^{-\rho t^{**}}$.⁴ Consequently, limit buy orders generally have prices below v_t while limit sell orders have prices above v_t . If that were not the case, a newly arriving agent could pick off limit buy (sell) orders above (below) v_t . This also implies that limit orders in unfavorable positions should disappear quickly from both limit order books.

We center each limit order book at the contemporaneous fundamental value of the asset, i.e. by setting $p_m^0 = v_t$. Suppose at time $t = 0$ the fundamental value is v_0 , but after a period τ the fundamental value experiences some innovations and its new value is v_τ , with $v_\tau - v_0 = qd$, where q is a positive or negative integer. In this case, we shift both books by q ticks to center them at the new level of the fundamental value v_τ . Thus, we move the queues of existing limit orders in both books to take the relative difference with respect to the new fundamental value into account. This implies that prices of all orders are always relative to the current fundamental value of the asset. This transformation allows us to greatly reduce the dimensionality of the state-space because agents always make decisions in terms of relative prices regarding the fundamental value of the asset.⁵

Each agent can trade one share and has to make three main trading decisions upon arrival: i) to submit an order either to $L_{1,t}$ or $L_{2,t}$; ii) to submit either a buy or a sell order; and iii) to choose the limit price, which implies the decision to submit either a market or a

⁴ A similar example, but in opposite direction, can be explained for the cost of being picked off with a limit sell order below the fundamental value of the asset.

⁵ It is important to note that under this normalization, we can still observe limit orders being picked-off. For example, suppose that the current time is t and the fundamental value is v_t ; hence $p_m^0 = v_t$. Suppose, that the current bid price is $B(L_{m,t}) = p_m^{-1}$ and the ask price is $A(L_{m,t}) = p_m^2$. Subsequently, at time t_{po} , if the fundamental value decreases by twice the amount of the tick size (i.e. $q = -2$), after re-centering the book, the bid and ask prices are $B(L_{m,t_{po}}) = p_m^1$ and $A(L_{m,t_{po}}) = p_m^4$, respectively. Thus, a newly arriving agent can submit a market order against the limit order at the bid price to generate a profit. Subsequently, the limit order at p_m^1 will disappear, and the new bid price will be below the price at the center of the book (i.e. $B(L_{m,t_{po}+\Delta t}) = p_m^0$, where Δt is the time until the limit buy order above the fundamental value is picked-off).

limit order, depending on whether the price is inside or outside the quotes.^{6,7} As mentioned above, an agent can re-enter the market and modify her unexecuted limit order. Hence, she has to make the following additional trading decisions after re-entering: i) to keep her unexecuted limit order unchanged or to cancel it; ii) in case of a cancellation, to submit a new order to $L_{1,t}$ or $L_{2,t}$; iii) to choose whether the new order will be a buy or a sell order; and iv) to choose the price of the new order.

The decision to leave the order unchanged has the advantage of maintaining the it's time priority in the respective queue. The negative side of leaving an order in any of the books unchanged is the potential costs agents can incur when the fundamental value of the asset moves in directions that affect the expected payoff. For example, in the case of a reduction in v_t , a limit buy order could be priced too high. This possibility represents an implicit cost of being picked off. Conversely, when the asset value increases, a buy limit order has the risk of waiting for a long period before being executed.

Therefore, agents have to take the possibility of re-entry into account when they make their initial decision after arriving in the economy. Once an agent submits a limit order, she remains part of the trading game until her order is executed; she exits the market forever after trading the asset.

2.2 Agents' Dynamic Maximization Problem and Equilibrium

There is a set of states $s \in \{1, 2, \dots, S\}$ that describes the market conditions in the economy. These market conditions are observed by each agent before making any decision. The state s that an agent observes is described by the contemporaneous limit order books, L_1 and L_2 ; the agent's private value α ; and in the case that the agent previously submitted a limit order

⁶ We can include additional shares per agent in the trading decision. However, similarly to [Goettler et al. \(2009\)](#), we assume one share per trader to make the model computationally tractable.

⁷ A potential decision to wait outside any of the markets (without submitting an order) is not optimal because there are no transaction fees, submission fees or cancellation fees. An agent can always submit a limit order far away from the fundamental value such that it is unlikely to be executed, but if executed, the potential economic benefit is high.

to any of the books, the status of that order in L_1 or L_2 , i.e. its original submission price, its queue priority in the book, and its type (i.e. buy or sell). The fundamental value of the asset, v , is implicitly part of the variables that describe the state s , since agents interpret limit prices relative to the fundamental value. For convenience, we set the arrival time of an agent to zero in the following discussion.

Let $a \in \Theta(s)$ be the agent's potential trading decision, where $\Theta(s)$ is the set of all possible decisions that an agent can take in state s . Suppose that the optimal decision given state s is $\tilde{a} \in \Theta(s)$. Let $\eta(h|\tilde{a}, s)$ be the probability that an optimally submitted order is executed at time h . The probability $\eta(\cdot)$ depends on future states and potential optimal decisions taken by other agents up to time h . The probability $\eta(0|\tilde{a}, s)$ is equal to one if the agent submits a market order, while $\eta(h|\tilde{a}, s)$ converges to zero as the agent submits a limit order further away from the fundamental value. Let $\gamma(v|h)$ be the density function of v at time h , which is exogenous and characterized by the Poisson process of the fundamental value of the asset at rate λ_v . Thus, the expected value of the optimal order submission $\tilde{a} \in \Theta(s)$, if the order is executed prior to the agent's re-entry time h_r , is:

$$\pi(h_r, \tilde{a}, s) = \int_0^{h_r} \int_{-\infty}^{\infty} e^{-\rho h} ((\alpha + v_h - \tilde{p})\tilde{x}) \cdot \gamma(v_h|h) \cdot \eta(h|\tilde{a}, s) dv_h dh \quad (1)$$

where \tilde{p} and \tilde{x} are components of the optimal decision \tilde{a} , in which \tilde{p} is the submission price and \tilde{x} is the order direction indicator (i.e. $\tilde{x} = 1$ if the agent buys and $\tilde{x} = -1$ if the agent sells). The expression $(\alpha + v_h - \tilde{p})\tilde{x}$ is the instantaneous payoff, which is discounted back to the trader's arrival time at rate ρ .

Let $\psi(s_{h_r}|h_r, \tilde{a}, s)$ be the probability that state s_{h_r} is observed by the agent at her re-entry time h_r , given her decision \tilde{a} taken in the previous state s . The probability $\psi(\cdot)$ depends on the states and potential optimal decisions taken by other agents up to time h_r . In addition, let $R(h_r)$ be the cumulative probability distribution of the agent's re-entry time, which is exogenous and described by the Poisson process governing agents' re-entry with rate λ_r . Thus, the Bellman equation that describes the agent's problem of maximizing her total

expected value, $V(s)$, after arriving in state s is given by:

$$V(s) = \max_{\tilde{a} \in \Theta(s)} \int_0^\infty \left[\pi(h_r, \tilde{a}, s) + e^{-\rho h_r} \int_{s_{h_r} \in S} V(s_{h_r}) \cdot \psi(s_{h_r} | h_r, \tilde{a}, s) ds_{h_r} \right] dR(h_r) \quad (2)$$

where S is the set of possible states. The first term is defined in Equation (1), and the second term describes the subsequent payoffs in the case of re-entries.

The intuition for the equilibrium is that each agent behaves optimally by maximizing her expected utility, based on the observed state that describes market conditions (as in Equation (2)). In this sense, optimal decisions are state dependent. They are also Markovian, because the state observed by an agent is a consequence of the previous states and the historical optimal decisions taken in the trading game. We obtain a stationary and symmetric equilibrium, as in [Doraszelski and Pakes \(2007\)](#). In such an equilibrium, optimal decisions are time independent, i.e., they are the same when an agent faces the same state in the present or in the future.

The trading game is also Bayesian in the sense that an agent knows her intrinsic private value to trade (α), but she does not know the private values of other agents that are part of the game. Hence, our solution concept is a Markov perfect Bayesian Equilibrium (see [Maskin and Tirole, 2001](#)). In the trading game, there is a state transition process where the probability of arriving in state s_{h_r} from state s is given by $\psi(s_{h_r} | \tilde{a}, s, h_r)$.⁸ Thus, two conditions must hold in the equilibrium: agents solve equation (2) in each state s , and the market clears.

As mentioned earlier, the state s is defined by the four-tuple $(L_{1,t}, L_{2,t}, \alpha, \textit{status of previous limit order})$, where all variables that describe the state are discrete. Moreover, each agent's potential decision a is taken from $\Theta(s)$, which is the set of all possible decisions that can be taken in state s . This set of possible decisions is discrete and finite given the features of the model. Consequently, the state space is countable and the decision space is finite; thus

⁸ It is important to note that $\psi(s_{h_r} | \tilde{a}, s, h_r) = \psi(s_{h_r} | s)$, since optimal decisions are state dependent and Markovian, and we focus on a stationary and symmetric equilibrium.

the trading game has a Markov perfect equilibrium (see [Rieder, 1979](#)). Despite the fact that the model does not lend itself to a closed-form solution, we check whether the equilibrium is computationally unique by using different initial values.

2.3 Solution approach and model parametrization

Given the large dimension of the state space, we use the [Pakes and McGuire \(2001\)](#) algorithm to compute a stationary and symmetric Markov-perfect equilibrium. The intuition behind the [Pakes and McGuire \(2001\)](#) algorithm is that the trading game by itself can be used, at the beginning, as a learning tool in which agents learn how to behave in each state. At the beginning, we set the initial beliefs about the expected payoffs of potential decisions in each state. Agents take the trading decision that provides the highest expected payoff conditional on the state they observe. Subsequently, agents dynamically update their beliefs by playing the game and observing the realized payoffs of their trading decisions. Thus, the algorithm is based on agents following a learning-by-doing mechanism.

The equilibrium is reached when there is nothing left to learn, i.e., when beliefs about expected payoffs have converged. We apply the same procedure used by [Goettler et al. \(2009\)](#) to determine whether the equilibrium is reached. The [Pakes and McGuire \(2001\)](#) algorithm is able to deal with a large state space because it reaches the equilibrium only on the recurring states class. Once we reach the equilibrium after making agents play in the game for at least 10 billion trading events, we fix the agents' beliefs and simulate a further 600 million events. Therefore, all theoretical results presented in this paper are calculated from the last 600 million simulated events, after the equilibrium has already been reached.

The multi-market model involves a higher level of complexity than a single market setup. First, the state space increases enormously in a multi-market environment, because all combinations of variable values across the two order books have to be considered. Second, in contrast to [Goettler et al. \(2005, 2009\)](#), we do not use model simplifications to reduce the large state space generated by our multimarket model. [Goettler et al. \(2005\)](#) assume that

cancellations are exogenous, and [Goettler et al. \(2009\)](#) reduce the dimension of the state space by using information aggregation (in the spirit of [Krusell and Smith, 1998](#) and [Ifrach and Weintraub, 2016](#)). [Goettler et al. \(2009\)](#) also describe the limit order book by only considering the bid and ask prices, the depth at the top of the book, and the cumulative buy and sell depths in the book. We avoid such model simplifications as they may induce the kernel of state variables to be non-Markovian. We instead solve the model by only employing the [Pakes and McGuire \(2001\)](#) algorithm.⁹ While parameterizing our model, we use the same market characteristics for both limit order markets. In addition, since our model is an extension of the dynamic model of a single market presented in [Goettler et al. \(2009\)](#), we use the same parameters as in their study.

We set the intensity of the Poisson process followed by the agents’ arrivals to one. A unit of time in our model is equal to the average time between new trader arrivals. The intensity of the Poisson process followed by the agents’ re-entry is set to 0.25; the intensity of the Poisson process followed by the innovations of the fundamental value is set to 0.125. We set the tick size in both order books to one, and the number of discrete prices available on each side of the order book on both markets to $N_1 = N_2 = 31$. The delaying cost reflected by the rate ρ is set to 0.05. The private value α is drawn from the discrete vector $\Psi = \{-8, -4, 0, 4, 8\}$ using the cumulative probability distribution $F_\alpha = \{0.15, 0.35, 0.65, 0.85, 1.0\}$.¹⁰

While market entry is exogenous in our model, we posit that, if entry were exogenous, higher profits generated by any agent type in fragmented markets would likely increase their participation. In a computationally simpler alternative, we create an additional parameter configuration by keeping the arrival rates of agents with non-zero private value unchanged

⁹ The implementation of the [Pakes and McGuire \(2001\)](#) algorithm, applied to our multi-market model, requires between 600GB and 800GB of RAM, depending on the parameters used. We relied on a high performance computing facility with latest generation processors and 1TB of RAM, which ran over 5-6 weeks to obtain the equilibrium.

¹⁰ As a robustness check, we multiply the following original [Goettler et al. \(2009\)](#) parameters by 0.8 and 1.2: the delaying cost, ρ ; the agents’ arrival intensity λ ; the innovation arrival intensity of the fundamental value, λ_v ; and the re-entering intensity λ_r . The results obtained are qualitatively similar to the results presented here.

and doubling the arrival rate of agents with private value equal to zero. In other words, we set the intensity of agent arrival to 1.3 and draw the different agent types from the cumulative distribution $F_\alpha = \{0.15/1.3, 0.35/1.3, 0.95/1.3, 1.15/1.3, 1.0\}$. In addition to the above rationale, this alternative configuration allows to proxy for a second empirical fact observed in real-world markets. It is often the case that liquidity providers are active in multiple limit order books. [van Kervel \(2015\)](#) describes a model of order cancellations in fragmented markets where high-frequency liquidity providers duplicate their orders across multiple order books to improve execution probabilities while simultaneously managing adverse selection risk. A comparison of the different market outcomes across the three (two fragmented and one consolidated) scenarios allows us to highlight potential effects, if any, associated with increased intermediation in fragmented markets.

3. Theoretical Implications

We are interested in examining the theoretical implications of the effects of market fragmentation on trading behavior, welfare, and market quality. To do so, we generate a dataset of trades and order book updates by simulating 10 million events for the following three specifications: (i) a consolidate market with one limit order book; (ii) a fragmented market with two limit order books; and (iii) a fragmented market with two limit order books and twice as many agents with no intrinsic value as the first two specifications. We compute mean levels of the measures of interest under all three market settings.

3.1 Trading Behavior

The order submission strategy determines the price formation of an asset and the liquidity of the market, and as a consequence, it has a direct effect on the welfare of individuals and society. Hence, it is important to analyze how the introduction of a second limit order book affects the trading behavior of agents. We study the trading patterns of agents in single and

fragmented markets. For the latter, we provide results for two scenarios: when the arrival rate of agents without exogenous reasons to trade is the same as in a single market and when the rate is twice as large. Table 1 presents the results.^{11,12}

We find that agents submit more aggressive limit orders in a single market compared to a fragmented market. In a single market setting, about 36% of the orders are placed at the best ask price, whereas this is the case for only about 28% of orders in the fragmented market. If the arrival rate of traders without exogenous reasons to trade is doubled, almost 33% of limit order are submitted at the best ask price, probably because of the higher degree of competition among limit order traders in this setup. More aggressive limit orders in a single market compared to a multiple market setting with same arrival rates lead to a higher picking-off risk, i.e., the share of executed limit orders that are picked off, inducing agents to cancel their orders more often. We find that the picking-off risk is indeed lower in a fragmented market. The results in Table 1 indicate that the picking-off risk declines from 21.80% in a single market to 20.82% in multiple markets. When the arrival rate of intermediaries is doubled, the picking-off risk is even lower. Untabulated results reveal that the picking-off risk, in this setting, is higher for each agent type, which is consistent with a higher competition between speculators. However, this measure decreases on average compared to a single setting, because of the higher share of agents of type $\alpha = 0$, who have the lowest picking-off risk. A higher picking-off risk induces agents to cancel their limit order more often, increasing the execution time from her arrival time until the execution of her limit order. Consistent with this intuition, the average number of limit order cancellations per trader is 1.2 in a single order book as compared to 1.01 when there is a second book. We also corroborate that limit orders execute faster in a multi-market setting. The average execution time is 8.61 in a single market, whereas in a fragmented market the time is reduced

¹¹As the model is symmetric we focus on the sell side of the market. The results for the buy side of the market are analogous.

¹²We do not report standard errors because the large number of trader arrivals implies that the standard errors on the sample means are sufficiently low such that a difference in means of an order of 10^{-2} is significantly different from zero.

to 7.15 units of time in a fragmented market with the same distribution of agents.

If we double the arrival rate of market-makers, the number of cancelations is 1.58 and the execution time is 13.10, which is also consistent with higher competition of limit orders inducing agents to cancel more often, increasing, in turn, their execution time. The much longer time until execution can be explained by the fact that an overwhelming share of limit order traders in this setup are intermediaries, who are patient traders.

Table 2 shows the proportions of limit orders and market orders submitted by each trader type. We report the distribution of limit orders and market orders for a given trader type. As expected, we find that agents with intrinsic motives to trade (i.e., $|\alpha| \neq 0$) act as liquidity demanders, whereas agents with no intrinsic motives to trade (i.e., $|\alpha| = 0$) act as liquidity suppliers. Almost all of the agents without intrinsic motive to trade (i.e., $|\alpha| = 0$) act as speculators submitting limit orders. Only about 5% of them submit market orders to take advantage of mispriced limit orders. Conversely, about 72% agents with private value $|\alpha| = 8$ submit market orders.

The behavior of agents with private value $|\alpha| = 4$ is in between those of the other types. The choice between limit and market orders does not markedly differ between the single and multi-market setups with the same trader populations. However, differences in order choice between the trader types are more pronounced when we doubled the arrival rates of zero private value agents, as traders with non-zero α use limit orders much less frequently.

Our findings are consistent with the study of [Goettler et al. \(2009\)](#) who examine the trading behavior in a single market setting. They also find that agents with $|\alpha| = 0$ supply liquidity to the market, agents with extreme valuation ($|\alpha| = 8$) are more likely to demand liquidity, and the behavior of agents with $|\alpha| = 4$ is in between that of the more extreme types.

Although our findings reveal that fragmentation does not change the main strategies adopted by traders, it is interesting to notice that, assuming an unchanged population of

traders, agents with private value $|\alpha| = 8$ submit a higher proportion of limit orders when there are two limit order markets. We will show later that market fragmentation leads to wider spreads. As market orders are more expensive in such a setting, some agents with exogenous reasons to trade prefer to submit more limit orders when there are two limit order books. However, when we increase the arrival rate of market makers, the latter appear to crowd out the limit order submissions of other types of traders.

3.2 Market Quality

In this subsection, we compare consolidated and fragmented markets in terms of the major determinants of market quality, i.e., liquidity and price efficiency.

We begin by estimating the effect of market fragmentation on various measures of quoted and traded liquidity. We calculate liquidity measures employing either *local* or *inside quotes*. Local quotes comprise the bid and ask prices of one of the markets whereas inside quotes are combine the highest bid and the lowest ask across the two limit order books.

We measure daily quoted liquidity by time-weighted quoted spreads and time-weighted top-of-book depth. We also report the total number of limit orders waiting to be executed on the sell side of the market. Panel A of Table 3 provides the results. Our theoretical findings indicate that fragmentation by and large impairs liquidity. This is illustrated by wider spreads and lower depth when there are two limit order markets. In particular, both local and inside quoted spreads decrease about 1.04 and 0.34 ticks, respectively, when the market moves from a fragmented to a single market and the arrival rates of all trader types are the same as in the single market. Spreads are also reduced in the single market compared to when the arrival rate of zero private value agents in the fragmented market is twice as large, although the effect is smaller.

Naturally, because of order flow fragmentation between the two markets, fragmented markets also show a decrease in the top-of-book depth. Local top-of-book depth is reduced

by more than 30% in a fragmented market. Inside depth is also lower as compared to the single market. The results change if we double the participation of agents of type $\alpha = 0$: the increased number of liquidity providers leads to a substantial increase in inside depth, and local depth is also slightly higher than in the single market scenario. Thus, our results with respect to quoted liquidity show that spreads are unambiguously smaller in a single market whereas the results for depth are ambiguous.

Improvements in quoted liquidity do not necessarily translate into actual transaction cost savings for traders submitting market orders. Thus, we next compare differences in traded liquidity in single and fragmented markets. We measure traded liquidity by the trade-weighted effective spreads, which capture the actual transaction costs incurred by traders submitting marketable orders. The effective spread is calculated as follows:

$$\text{effective spread} = x_t(p_t - m_t)/m_t, \quad (3)$$

where x_t is +1 for a buyer-initiated order, p_t is the traded price, and m_t is the mid-quote. We further decompose effective spread into realized spread and price impact (adverse selection). The former is calculated as follows:

$$\text{realized spread} = 2x_t(p_t - m_{t+k})/m_t, \quad (4)$$

where k is the number of seconds in the future. As the results are qualitatively similar, we only report the findings for 30 seconds. Finally, price impact is effective spread minus realized spread. The price impact captures the level of information in a trade, whereas the realized spread measures liquidity providers' compensation after accounting for adverse selection losses associated with informed orders. As our model does not contain private information, the price impact measure captures picking-off risk associated with stale limit orders when new (public) information arrives in the market. Just like quoted liquidity, we compute local and inside variants of all three measures using the inside quote midpoints across the two books and local quote midpoints in the order book where a transaction is

executed.

Panel B of Table 3 reports the results. Transaction costs in terms of effective spread and realized spread are higher when there are two limit order books. When the arrival rate of zero private value agents remains the same, effective spreads decrease from 1.80 ticks and 1.45 ticks based on local and inside quotes, respectively, in fragmented markets, to 1.31 ticks in the single market. The differences are even larger if we double the arrival rate of agents of type $\alpha = 0$.

Realized inside and local spreads are higher in the fragmented market by approximately 0.15 ticks with the same population of agents, and higher by about one half of a tick if we double the participation of intermediaries.

Price impact measured relative to local quotes is lower in the single market whereas it is similar when measured relative to inside quotes. This is because, in fragmented markets, a newly arriving trader is more likely to trade in an order book containing a stale quote, leading to a higher local price impact. The inside price impact is smaller because the inside quote midpoint already reflects part of the information. Price impacts are lower if we double the arrival rate of agents of type $\alpha = 0$, likely because the increased arrival rates leads to the exploitation of even mispricing of even small magnitudes. The local price impact in this scenario is still slightly larger than that in the single market, though the inside price impact is substantially smaller.

Finally, we analyze the degree of inefficiency in prices when the market consists of one order book as opposed to multiple ones. If an asset is traded on multiple markets, the degree of price dislocations may be exacerbated *ceteris paribus*, making prices on each book less efficient than they would be if all demand and supply were to meet on a single order book. In the context of our model, the effect of these frictions is measured as the deviation of the quote midpoint from the fundamental value v_t . In Panel C, we present the mean absolute difference between the quote midpoint and the fundamental value. This value changes from 0.67 ticks

in a fragmented market to 0.46 ticks in a single market. The corresponding differences based on inside quotes are in the same direction although the magnitudes are lower. However, in fragmented markets with doubled arrival rate of zero private value agents, microstructure noise is lower than in the single market, suggesting that prices are more efficient in this case. This result is expected as in the absence of private information, a higher number of traders with no intrinsic reasons to trade results in a faster adjustment of quotes when public information arrives. Thus, the degree of mispricing depends on the number of intermediaries in the market and, because their number in real fragmented markets is likely larger but not twice as large as it is in consolidated markets, our model makes no strong predictions about differences in price efficiency between such markets.

3.3 Welfare Analysis

In order to analyze the potential economic benefits per agent and for the whole market, we examine the effect on welfare of both single and fragmented markets. Welfare is measured as the average realized payoff per agent. In addition, we decompose the realized payoffs of investors to analyze the gains and losses from the trading process.

Suppose that an agent with a private value α and delaying discount rate of ρ arrives on the market at time t . She submits an order (i.e., a limit order or a market order) to any of the books at price \tilde{p} with order direction \tilde{x} (i.e., to buy or to sell). Suppose that the agent does not modify the order, and it is finally executed at time t' when the fundamental value is $v_{t'}$ (in the case of a market order $t = t'$). Then the realized payoff of the agents from the order execution is given by:

$$\Pi = e^{-\rho(t'-t)} (\alpha + v_{t'} - \tilde{p}) \tilde{x}. \quad (5)$$

We can decompose the agents' payoffs and rewrite (5) as:

$$\begin{aligned} \Pi &= \textit{Gains from private value} + \textit{Waiting cost} + \textit{Money Transfer}, \textit{ where} \\ \textit{Gains from private value} &= \alpha\tilde{x} \\ \textit{Waiting cost} &= (e^{-\rho(t'-t)} - 1)\alpha\tilde{x} \\ \textit{Money Transfer} &= e^{-\rho(t'-t)}(v_{t'} - \tilde{p})\tilde{x} \end{aligned} \tag{6}$$

The first term in (6), *gains from private value*, represents the gains obtained directly from the exogenous reasons to trade for each agent, $\alpha\tilde{x}$. Agents initially submitting a limit order do not trade immediately after arriving on the market. and, thus have to wait until they obtain their private values. This waiting process is costly due to the delaying cost ρ . The second term in (6), *waiting cost*, reflects the cost paid by agents in terms of delaying the *gains from private value*.

The realized payoff in (5) results from a transaction in which one agent buys the asset and another agent sells it at a price that may differ from the fundamental value. The third term in (6), *money transfer*, reflects the difference between the fundamental value $v_{t'}$ and the transaction price \tilde{p} , and thus the money gained (or lost) in the transaction. It is discounted depending on the arrival time of the trader. In general, the *money transfer* is associated to the immediacy cost incurred when an agent wants to immediately realize her private value. For example, an agent who submits a market order realizes her intrinsic private value without delay. Thus, this trader does not have any *waiting cost*, but she may have to pay a cost for demanding immediacy, which would be reflected in a negative *money transfer*.

Table 4 presents the results. In the first set of columns, we present the results of (5), i.e., the average payoff for each trader type in each market scenario. We find a similar global welfare in the three setups. While the aggregate welfare effects are altogether negligible, the shifts among categories of agents are substantial. Agents with non intrinsic motives to trade (i.e., $|\alpha| = 0$) take more advantages from fragmented markets and, as a consequence, have

higher profits. When we double the arrival rate of agents without intrinsic motive to trade, the welfare for each such agent decreases, but their aggregate welfare is larger than in the other two scenarios. In a single market scenario, agents place more aggressive orders to jump the queue and thus to raise their probability of execution. This fact generates competition on price due to the consolidation of all order flow in a single trading venue. Contrarily, in the presence of multiple markets, as there is no time priority across order books, traders can, with a positive probability, jump ahead of the standing limit orders in one market by submitting an order to the other market. The lack of time priority reduces competition on price such that agents with $|\alpha| = 0$ obtain better terms of trade.

As aggregate welfare effects in the model are not quantitatively meaningful, any interpretation regarding the overall desirability of fragmentation needs to go beyond the model. Market participation in the model is exogenous. In real markets, one would expect that traders endogenously decide about their market entry based on the trade-off between expected trading profits and participation costs. Thus, the higher profit earned by liquidity providers in fragmented markets should lead to an increased participation of this group of traders. If participation in markets is costly - a realistic assumption considering the investments made by intermediaries in modern equity markets - these additional traders incur costs that do not increase aggregate welfare, i.e., the privately optimal decisions are not socially optimal. This suggests there are welfare losses resulting from market fragmentation.

Next we analyze the second and third components of total payoff described in (6). In the next set of columns, we report the waiting cost and money transfer per trader.¹³ Agents with intrinsic motives to trade (i.e., $|\alpha| \neq 0$) exhibit a reduction in absolute waiting costs in the fragmented market, even more so if market maker participation is doubled. However, they obtain worse terms of trade, which is reflected in high money transfer costs. For example,

¹³Note that in Table 4, the total money transfer do not add up to zero because they are discounted back to time t and $t' - t$ is different for the trader who submits the market order and the trader who submits the corresponding limit order due to traders' asynchronous arrivals. However, the instantaneous money transfer not discounted back does add up to zero.

agents with private value $|\alpha| = 8$ experience smaller money transfer losses in consolidated markets as compared to fragmented markets (-0.572 ticks versus -0.626 or -0.835 ticks). This is because lower waiting costs do not compensate for the losses associated with money transfer. Finally, agents with $|\alpha| = 0$ obtain higher gains from trading in fragmented markets primarily through higher money transfer gains.

In conclusion, the welfare of agents with non intrinsic motives to trade is increased under the presence of a second limit order book and this gain is to the detriment of traders with non-zero private values, likely because price competition in fragmented markets is less severe. Hence, they pay the cost associated with higher profit for agents with $|\alpha| = 0$ in fragmented markets.

4. Empirical Application

In this section, we test the empirical predictions generated by our model in Section 3.2. We also indirectly address the predictions about welfare described in Section 3.3.¹⁴ To this end, we conduct an event study based on Euronext’s decision to implement a single order book per asset for their Paris, Amsterdam, and Brussels markets. Pagano and Padilla (2005) and Nielsson (2009) analyze the effects of integrating trading on Euronext for stocks listed in these three markets.

4.1 Euronext’s Institutional Background

We begin by describing Euronext’s institutional arrangements leading up to the introduction of the Single Order Book. Euronext was formed in 2000 following a merger of the Paris, Amsterdam and Brussels stock exchanges. In 2002, the Lisbon Stock Exchange became the fourth exchange to merge with Euronext.¹⁵ Stock listings on Euronext pertain to a

¹⁴ We cannot empirically test the predictions from Section 3.1 pertaining to trading behavior due to data limitations.

¹⁵ In 2007, Euronext merged with the NYSE to form NYSE Euronext, which was taken over by Intercontinental Exchange in 2012. In 2014, Euronext was spun off through an IPO.

listing on one or more national markets.¹⁶ Until 13 January 2009, each national listing corresponded to the operation of one limit order book. For example, a stock listed on the Paris market would be traded on the limit order book of Euronext Paris. Firms cross-listed in multiple Euronext markets traded in parallel on multiple Euronext order books, besides other competing markets. On 16 August 2007, the exchange announced its intention to eliminate this arrangement for their Paris, Amsterdam and Brussels markets by unifying all trading in these markets on to a single order book, the so-called “Market of Reference” (MoR). Cross-listed firms had to choose one MoR that continued operating after the implementation of a single order book. This new arrangement was implemented on 14 January 2009.

The existence of multiple order books led to fragmentation of order flow routed to Euronext. As the rules and trading protocols governing the individual order books were identical, the introduction of a single order book decreased fragmentation for the stocks without any corresponding change in the competitive environment. Pagano and Padilla (2005) describe the steps taken by Euronext to standardize its trading protocols and technological platform as the source of the efficiency gains generated through the merger. This is particularly relevant as it allows us to test the isolated effects of fragmentation. Euronext, in its press release announcing the event, made clear that the trading environment remained unchanged: “The Single Order Book will have no impact on the NSC system as the market rules and order book management will remain unchanged [...] In practice, from a trading perspective, Single Order Book implementation simply means the end of order book trading on marketplaces other than the market of reference.”¹⁷ Moreover, as it was based on a business decision by Euronext, all multi-listed stocks received the same treatment such that there was no selection bias. Finally, the announcement was made more than one year before the event date in order to allow market participants to adapt and test their trading systems. This eliminates potential concerns about the event date confounding with other

¹⁶ With the implementation of the Markets in Financial Instruments Directive (MiFID), all rules prohibiting trading outside the national markets were repealed such that investors can now trade these stocks in any regulated market.

¹⁷ See Euronext press release dated 14 January 2009.

market events around the same time.¹⁸ Thus, this empirical analysis of a transition from a multi-market environment to a single market setup can be viewed as a natural experiment, allowing us to compare the outcomes with those obtained from our theoretical model.

4.2 Sample Selection

A total of 45 instruments, cross-listed on at least two of the three Euronext markets, are affected by the event. However, we reduce the sample of treated stocks used in our study for several reasons. First, we remove stocks whose primary listing is not on Euronext. These include stocks whose main trading activity takes place in other European markets or in the United States. Second, we eliminate exchange-traded mutual funds because we do not expect their trading activity to be comparable to that of listed firms. Finally, we require that there not only exist multiple Euronext order books before the event, but also that the total share of trading activity on the less active order books is at least equal to 1% of the respective stock's total Euronext trading volume. This reduces the list of instruments to ten. We further exclude one additional stock due to data errors, reducing our final sample to nine stocks.¹⁹ The number of stocks is small due to the unique nature of the event we study. Nonetheless, our sample consists of the whole population of stocks affected by the event, except a subset of stocks which are excluded through objective criteria.

We construct a matched control group of stocks based on stock price and market capitalization obtained from Compustat Global using the distance metric employed by [Huang and Stoll \(1996\)](#), and subsequently, in many other market microstructure studies. Specifically, for each stock in our treatment group, we identify the stock that is its closest match in terms of these two criteria as on the last trading day of 2008 (30 December 2008). The population of stocks from which the control group is constructed comprises all stocks with a primary

¹⁸Although the original date of implementation was postponed, this was due to technical reasons as opposed to concerns about market conditions. The final implementation date was announced more than 60 days in advance.

¹⁹One stock in our sample was listed in all three Euronext markets. However, we exclude the least active limit order book as it had a market share of 0.3%.

listing on Euronext not affected by the event.

Davies and Kim (2009) simulate the matching performance of a control group constructed using multiple criteria such as price volatility, trading volume and industry classification, in tests of differences for variables typically used in the microstructure literature, and conclude that one-to-one sampling without replacement based on stock price and market capitalization provides the best results. They also show that results obtained by matching based on the distance metric employed by Huang and Stoll (1996) are similar to those obtained when using the Mahalanobis distance measure.

Using a control group allows us to identify the effects of reduced fragmentation, implicitly controlling for market-wide changes in variables such as liquidity and volatility. It also allows us to control for two additional market-wide changes implemented by Euronext close in time to the introduction of a single order book. First, a harmonized settlement platform known as the Euroclear Settlement for Euronext-zone Securities for all French, Dutch and Belgian stocks was implemented on 19 January 2009. Second, the Universal Trading Platform, having “superior functionality, faster speed and much greater capacity”, was introduced on 16 February 2009.²⁰ These were market-wide events that affected both the control and treatment stocks. Consequently, we can attribute any difference in trading activity and market quality between the two groups exclusively to market consolidation resulting from the introduction of a Single Order Book.

For the purpose of our analysis, we define all days from the beginning of December 2008 to 13 January 2009 as the pre-event period and all days from 26 January 2009 to the end of February 2009 as the post-event period. We exclude all trading days from the event date until the end of the subsequent calendar week in order to eliminate any effect associated with the transition.

²⁰See press release titled “NYSE Euronext’s European Equities Trading Successfully Migrates to the Universal Trading Platform” dated 17 February 2009.

4.2.1 Data Description and Summary Statistics

We use high-frequency data from Thomson Reuters Tick History between December 2008 and February 2009 for the purpose of our analysis. This data contains trades and order book updates time-stamped with a millisecond resolution. We apply two filters to the data. First, we eliminate all order book updates where the best bid or ask prices are zero or the bid-ask spread is negative. Next, we exclude trades that are executed during the opening and closing auctions as well as trades within the first and last minute of the continuous trading session.

Table 5 describes the characteristics of stocks in the treatment and [Huang and Stoll \(1996\)](#) control group. The average market capitalization across stocks in the treatment (control) group of €4.4 (€4.8) billion and the average stock price of €18.4 (€18.2) are suggestive of high matching quality based on these two variables. The share of the more active venue as a percentage of total Euronext volume across all the days before the event ranges from 54% to 98% across the nine stocks in the treatment group. The simple (volume-weighted) average across all stocks is 78% (62%). This implies that almost 40% of the total Euronext volume was executed on the less active market. The market share of the sole listing Euronext venue for the stocks in the control group is, by construction, 100%. Trading activity when measured in terms of number of trades also provides a similar picture.

4.2.2 Estimation Methodology

In order to test the main implications of our model, we compute several variables capturing the trading activity, liquidity and price efficiency of the stock, as described in Section 3.2. Similar to our numerical results, we calculate both local and inside measures. Unsophisticated investors who choose to trade only on a single order book are likely to select the more active and liquid one. Hence, we compute the local measures for the market having higher trading volume during the pre-event period. The inside measures use the highest bid and the lowest ask across the two limit order books. We estimate a panel difference-in-differences

regression with stock and day fixed effects and standard errors double clustered by stock and day. We estimate the regression for levels and natural logarithms of the variables of interest in order to account for the wide dispersion in the levels of these variables across the stocks in our sample.

4.3 Empirical Results

4.3.1 Quoted Liquidity

We begin by analyzing the effect on quoted spread and top-of-book depth. Table 6 presents the results. Consistent with our theoretical findings, we observe an overall improvement in quoted spreads on Euronext after the introduction of a single order book. The more active of the two Euronext markets experiences a significant reduction in local spreads of 81bps or approximately 30%. The effect on inside spreads depends on the test specification and is statistically insignificant. The absence of a significant improvement in the inside spread can be explained by the fact that in real markets, different from our model, market participants do not always route their orders optimally, i.e. to the market offering the highest bid or lowest ask.²¹ Thus, while in the model inside spreads correctly reflects the gains, before adverse selection, that liquidity providers expect to earn, a non-zero probability of traders routing their orders to the market not offering the best price means expected gains earned by liquidity providers are in reality larger. This effect of suboptimal order routing vanishes after consolidation, *ceteris paribus* leading to an increase in inside quoted spreads. Conversely, an increase in price competition among liquidity providers, as predicted by the theory, leads to a decrease in quoted spreads after consolidation. These effects empirically cancel out such that the coefficients for inside spreads appear insignificant.

²¹In European markets, best execution requirements allow brokers to consider other criteria besides price when making order routing decisions. In contrast to the US, European markets also do not have an order protection rule that requires exchanges to re-route orders to venues offering a superior price. Even in the US, communication latencies between geographically dispersed exchanges and exceptions to the order protection rule result in liquidity takers obtaining sub-optimal prices.

We observe a positive though statistically insignificant effect of order flow consolidation on local and inside top-of-book depth. These results lie in between those observed in the simulations with different participation rates of market makers. In other words, they are consistent with an amount of market-making in fragmented markets that is larger than, but less than twice as large as that in a consolidated market. The results for local and inside depth do not markedly differ, which is in contrast to the theoretical predictions where inside depth in the fragmented market is relatively higher. Differences between the tick sizes on Euronext as compared to those in the theory may drive this result. The empirical tick size, relative to the price fluctuation, is substantially smaller than that in the simulations,²² such that instances with the same best prices offered on the two order books are infrequent, leading to a relatively smaller inside depth than in a large-tick market. Ye (2017) illustrates the negative relationship between flickering quotes and tick size in a single market setup. Extending this argument to fragmented markets, prices across multiple markets will be synchronized less frequently when the tick size is small.

4.3.2 Traded Liquidity

Table 7 presents the results for effective spreads and their decomposition. Effective inside (local) spreads decrease by an economically large 14.5% (37.5%) after the introduction of a single order book, though only the results for local spreads are unambiguously statistically significant. Realized inside spreads significantly decrease by 31.3%, 42.1%, and 45.9%, at the 10-, 30-, and 60-second horizon, respectively. The corresponding decrease in realized local spreads is larger in magnitude and also significant. Local price impacts decrease across all specifications, even though the statistical significance varies. All the above mentioned results on traded liquidity are consistent with our theoretical predictions.

The empirically observed change in inside price impacts differs depending on the empirical

²²Tick sizes on Euronext during our sample period are smaller than in most international markets. A simulation applying parameters that would closely match Euronext tick sizes is infeasible because the large number of possible prices would lead to a corresponding increase of the state space.

specification and is never significant, whereas our theory tells us the effect should be near-zero or positive. The previously-mentioned fact that liquidity takers empirically sometimes do not trade on the market offering the best price may explain why inside price impacts are not larger in the fragmented market. An order trading against a standing limit order at a price inferior to the lowest ask or highest bid does not mechanically generate an inside price impact even if it executes against the entire limit order, leading to a relatively smaller inside price impact compared to the theoretical predictions. Additionally, the empirical results, in contrast to the model, also capture the effects associated with private information possessed by traders. Lower transaction costs potentially allow traders with small amounts of private information to profitably participate, leading to a decrease in average price impact. The latter channel may cancel out the positive effect of consolidation on price impact predicted by our theory.

4.3.3 Price Efficiency

In our numerical simulation, we examine the price efficiency by measuring the extent to which the mid quote deviates from the fundamental value v_t . Empirically, as we cannot observe the fundamental value, we measure price efficiency using return autocorrelations and variance ratios. Return autocorrelations are measured at 30 second and 5 minute intervals. Variance ratios capture the deviation between long-term and short-term return variance and are calculated as one minus the ratio of long-term and short-term return variance, each scaled by the respective time periods. We calculate variance ratios between 30 second and 5 minute returns variances. As in [Boehmer and Kelley \(2009\)](#), we measure the impact of consolidation on absolute values of both measures because we are interested in departures from a random walk in either direction. The closer these measures are to zero, the more closely does the price path resemble a random walk. Similar to the liquidity measures, we calculate price efficiency based on local and inside quotes.

Table 8 presents the results.²³ The variance ratio becomes closer to one after the implementation of the single order book, though the change is statistically insignificant. The results for autocorrelations also point to improved price efficiency although only results for the 5-minute autocorrelation are significant at the 10% level. These results generally provide evidence for unchanged or higher price efficiency in consolidated markets. When compared to the theoretical results, this appears consistent with a fragmented market containing more but less than twice as much intermediation as a consolidated market.

4.3.4 Trading Volume and Arbitrage

The existence of multiple order books empirically allows market participants to earn arbitrage profits by exploiting occasions of crossed markets, i.e. situations where the bid price on one order book is higher than the ask price on the other. These situations would otherwise be immediately resolved by the adjustment of limit order prices. In other words, such trades do not contribute to an increase in price efficiency, but only result in losses for limit order traders who consequently impose higher trading costs on liquidity seekers. This arbitrage-driven rent extraction may lead to welfare losses if otherwise beneficial trades are crowded out (Foucault et al., 2017; Budish et al., 2015).

We measure trades associated with such “toxic” arbitrage and the resulting costs to market-makers in the empirical data as follows. We start by identifying instances of a crossed order book. Such a situation can arise as a result of new order(s) submitted to either or both order book(s). Next, we identify whether these instances are resolved through a trade, quote-update, or both. This approach is similar to Foucault et al. (2017) who define the resolution through trades as toxic arbitrage if the following two conditions are fulfilled: (i) prices offered in different markets allow aggressive traders to earn a profit by trading against the bid on one market and ask on the other; (ii) they are able to do so because of liquidity providers’ slow reaction to new information, rather than them offering attractive prices to manage their

²³The empirical results based on returns measured at other frequencies are qualitative similar to those reported here and are available upon request.

inventories. Fragmentation is an obvious precondition for such arbitrage trades to occur. For each stock-day, we calculate the number of unique crossed instances, the fraction of a day when inside spreads are on average negative, and the total trading volume contributing to the resolution of a crossed market. Panel A of Table 9 reports the mean values for each stock across all days in the pre-event period. The frequency of unique instances of a crossed market for an average day ranges from 0.3 to 622 across all stocks, with an average value of 124, which corresponds to one instance every four minutes. An average stock has a negative inside spread for 6.4% of the continuous trading session. Finally, 7.8% of the total trading volume on Euronext for an average stock can be attributed to the resolution of instances where the two markets are crossed. Approximately 50% of this, or almost 4% of the total Euronext trading volume, is associated with toxic arbitrage as defined in [Foucault et al. \(2017\)](#).

Since, by construction, arbitrage trades between multiple Euronext order books are eliminated after the introduction of a Single Order Book, trading volume should, everything else equal, be reduced. Panel B of Table 9 shows that the actual change in trading volume is in fact weakly positive. This suggests that the volume transacted by investors with intrinsic motives to trade increases in the consolidated order book. This welfare gain is consistent with our theory. The amount of volume traded by agents with intrinsic reasons to trade is constant in the model because private values are assumed to be sufficiently large such that they never refrain from trading. The increase in trading by such agents suggested by our empirical results indicates that some traders who were earlier crowded out in a less liquid fragmented market, now participate, leading to an overall welfare gain after consolidation.

5. Conclusion

We examine the effects of market fragmentation when competition between markets is non-existent or at best minimal. Such fragmentation is routinely observed after exchange mergers,

when a single exchange operator continues operating multiple order books to trade the same asset post merger. In an attempt to extract synergies from the merger, the operator typically eliminates structural and technological differences across the merging markets resulting in operator-level order flow fragmenting across (nearly) identical limit order books.

Our model allows us to examine the effects on several aspects of market performance such as liquidity, price efficiency, agents' payoffs and overall welfare. As limit order priority is not enforced across markets, fragmentation leads to reduced competition between intermediaries. This results in the deterioration of liquidity in fragmented markets as compared to the consolidated market benchmark. While overall welfare remains largely unchanged under both market setups, the distribution of welfare across the heterogeneous agent types in the model is markedly different. Agents with intrinsic trading motives extract lower payoffs in fragmented markets whereas agents acting as intermediaries are better off in fragmented markets.

These higher intermediation gains should, under conditions of endogenous entry, lead to more intermediaries entering the market. We mimic these conditions by doubling the population of intermediaries in the model while keeping all other market parameters constant. We observe that under these conditions the allocation of trading gains between intermediaries and non-intermediaries shift further in favour of the former while still not altering overall market welfare materially. These results point to fragmentation-induced investment in intermediation capacities, such as high-speed connections required to access the trading systems and real-time data feeds from multiple venues, being socially wasteful.

We empirically test the model implications by investigating the effects of Euronext's decision to introduce a single order book for their Paris, Amsterdam, and Brussels markets. As opposed to existing empirical research on this question which necessarily investigates the joint impact of changes in fragmentation and competition (say, when a new trading center venue the market), this event allows us examine the effects associated with the consolidation of multiple non-competing order books. The empirical analysis broadly confirms the

theoretical predictions related to the effects on liquidity, price efficiency, and market makers' profits. Additionally, we also obtain evidence that trading volume after consolidation does not decrease even though the amount of arbitrage trading in the market mechanically reduces after the event. This suggests that, while the (substantial) revenues generated by modern exchanges' from the sale of market data may decrease after consolidation, improvements in market quality need not come at the expense of reduced trading fees for the exchange operators.

Overall our results suggest that the positive externalities associated with consolidating order flow in a single location (or fewer locations) still exist and are substantial. This is true even in modern electronic limit order markets where the activities of high-frequency traders serve to integrate fragmented order books. The adverse effects of fragmentation are significantly larger for unsophisticated investors who do not possess the technological ability to route their trades to the most advantageous trading center. For such investors consolidation of order flow, at least between non-competing markets, likely results in transaction cost reductions.

Our results also have important policy implications. Regulators may be able to improve the welfare of investors who trade for intrinsic motives by: (i) preventing individual market operators from keeping an artificially high(er) level of order flow fragmentation in the absence of commensurate benefits; and (ii) limiting excessive investment in intermediation capacities necessary to link multiple order books which come at a cost to end investors.

References

- Amihud, Y., B. Lauterbach, and H. Mendelson (2003). The value of trading consolidation: evidence from the exercise of warrants. *Journal of Financial and Quantitative Analysis* 38(04), 829–846.
- Biais, B. (1993). Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets. *Journal of Finance* 48(1), 157–185.
- Boehmer, B. and E. Boehmer (2003). Trading your neighbor’s ETFs: Competition or fragmentation? *Journal of Banking and Finance* 27(9), 1667–1703.
- Boehmer, E. and E. K. Kelley (2009). Institutional Investors and the Informational Efficiency of Prices. *Review of Financial Studies* 22(9), 3563–3594.
- Budish, E., P. Cramton, and J. Shim (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics* 130(4), 1547–1621.
- Chlistalla, M. and M. Lutat (2011). Competition in securities markets: the impact on liquidity. *Financial Markets and Portfolio Management* 25(2), 149–172.
- Chowdhry, B. and V. Nanda (1991). Multimarket trading and market liquidity. *Review of Financial Studies* 4(3), 483–511.
- Davies, R. J. and S. S. Kim (2009). Using matched samples to test for differences in trade execution costs. *Journal of Financial Markets* 12(2), 173–202.
- Degryse, H., F. de Jong, and V. v. Kervel (2015). The Impact of Dark Trading and Visible Fragmentation on Market Quality. *Review of Finance* 19(4), 1587–1622.
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in io. *Handbook of industrial organization* 3, 1887–1966.
- Foucault, T., R. Kozhan, and W. W. Tham (2017). Toxic Arbitrage. *Review of Financial Studies* 30(4), 1053–1094.
- Foucault, T. and A. J. Menkveld (2008). Competition for Order Flow and Smart Order Routing Systems. *Journal of Finance* 63(1), 119–158.
- Goettler, R. L., C. A. Parlour, and U. Rajan (2005). Equilibrium in a Dynamic Limit Order Market. *Journal of Finance* 60(5), 2149–2192.
- Goettler, R. L., C. A. Parlour, and U. Rajan (2009). Informed traders and limit order markets. *Journal of Financial Economics* 93(1), 67–87.

- Gomber, P., S. Sagade, E. Theissen, M. C. Weber, and C. Westheide (2016). Competition Between Equity Markets: A Review Of The Consolidation Versus Fragmentation Debate. *Journal of economic surveys*. forthcoming.
- Gresse, C. (2017). Effects of Lit and Dark Market Fragmentation on Liquidity. *Journal of Financial Markets*. forthcoming.
- Harris, L. E. (1993). Consolidation, Fragmentation, Segmentation and Regulation. *Financial Markets, Institutions & Instruments* 2(5), 1–28.
- Hellström, J., Y. Liu, and T. Sjögren (2013). Time-varying return predictability and equity market fragmentation.
- Hengelbrock, J. and E. Theissen (2009). Fourteen at One Blow: The Market Entry of Turquoise. Available at SSRN 1570646.
- Huang, R. D. and H. R. Stoll (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics* 41(3), 313–357.
- Ifrach, B. and G. Y. Weintraub (2016). A framework for dynamic oligopoly in concentrated industries.
- Kohler, A. and R. von Wyss (2012). Fragmentation in European Equity Markets and Market Quality – Evidence from the Analysis of Trade-Throughs. Technical Report 1210.
- Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy* 106(5), 867–896.
- Madhavan, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies* 8(3), 579–603.
- Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory* 100(2), 191–219.
- Mendelson, H. (1987). Consolidation, Fragmentation, and Market Performance. *Journal of Financial and Quantitative Analysis* 22(02), 189–207.
- Nguyen, V., B. F. Van Ness, and R. A. Van Ness (2007). Short- and Long-Term Effects of Multi-market Trading. *Financial Review* 42(3), 349–372.
- Nielsson, U. (2009). Stock exchange merger and liquidity: The case of Euronext. *Journal of Financial Markets* 12(2), 229–267.
- O’Hara, M. and M. Ye (2011). Is Market Fragmentation Harming Market Quality? *Journal of Financial Economics* 100(3), 459–474.

- Pagano, M. (1989). Trading Volume and Asset Liquidity. *Quarterly Journal of Economics* 104(2), 255–274.
- Pagano, M. and J. A. Padilla (2005). Gains from Stock Exchange Integration: The Euronext Evidence. *Working Paper*.
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the ‘curse’ of dimensionality. *Econometrica* 69(5), 1261–1281.
- Parlour, C. A. and D. J. Seppi (2003). Liquidity-Based Competition for Order Flow. *Review of Financial Studies* 16(2), 301–343.
- Rieder, U. (1979). Equilibrium plans for non-zero-sum markov games. *Game theory and related topics*, 91–101.
- Riordan, R., A. Storkenmaier, and M. Wagener (2010). Fragmentation, competition and market quality: A post-mifid analysis. *Working Paper*.
- van Kervel, V. (2015). Competition for Order Flow with Fast and Slow Traders. *Review of Financial Studies* 28(7), 2094–2127.
- Ye, M. (2017). Who Provides Liquidity, and When: An Analysis of Price vs Speed Competition on Liquidity and Welfare.

Table 1. Impact on Trading Behavior

This table shows measures of trader behavior, such as, the percentage of limit orders executed among all limit orders submitted, the probability of being picked-off after submitting a limit order, the number of limit orders submitted per trader, the number of limit order cancellations per trader, the average time between the instant in which a trader arrives and his execution (in time units of our model), the time between the instant in which a trader arrives and the execution of his limit order (in time units of our model) and the probability of submitting a limit sell order at the ask price (i.e., an aggressive limit sell order). All the measures are calculated for a consolidated market and two versions of fragmented market: one with the same distribution of zero private value agents as in a single market and the other one with double arrival rate of zero private value agents. Since the model is symmetric on both sides of the book it is not necessary to also report the probability of submitting a limit buy order at the bid price. The probability of being picked-off is calculated with executed limit orders: we take the number of limit sell (buy) orders that are executed when their execution price is below (above) the fundamental value of the asset, which is divided by all the limit orders executed in the market. All trading behavior measures are determined as mean of 20 million market new entries in equilibrium. Standard errors for all trader behavior measures are small enough since we use a large number of simulated events. The Markov equilibrium is obtained independently for each scenario.

	Single Market	Fragmented Markets	Fragmented Markets Double $\alpha = 0$
Prob. of submitting a limit sell order at the ask price	35.87%	28.45%	32.62%
Execution time of a limit order	8.61	7.15	13.10
Prob. of being picked-off for a limit order	21.80%	20.82%	10.92%
Number of limit order cancelations per trader	1.20	1.01	1.58

Table 2. Impact on Trading Behavior by Agent's Type

This table shows the distribution of limit orders and market orders separated by private value α . The results are reported for a consolidated market and two versions of fragmented market: one with the same distribution of zero private value agents as in a single market and the other one with double arrival rate of zero private value agents. The first three columns show the proportion of limit orders and market orders for a given agents' type α . The next set of columns present how the orders are distributed through the different private values $|\alpha| = 0, 4, 8$. LO denotes limit orders, whereas MO market orders. All trading behavior measures are determined as mean of 20 million market new entries in equilibrium. Standard errors for all trader behavior measures are small enough since we use a large number of simulated events. The Markov equilibrium is obtained independently for each scenario.

$ \alpha $		0	4	8
Single Market	LO	94.6%	68.6%	27.7%
	MO	5.4%	31.4%	72.3%
	Total	100%	100%	100%
Fragmented Market	LO	93.9%	67.8%	29.2%
	MO	6.1%	32.2%	70.8%
	Total	100%	100%	100%
Fragmented Market (Double $\alpha = 0$)	LO	97.7%	34.8%	7.4%
	MO	2.3%	65.2%	92.6%
	Total	100%	100%	100%

Table 3. Impact on Liquidity

Panel A shows the quoted spread and depth for a market containing a single order book and two order books considering arrival rates of zero private value agents being the same and double as in a single market. We present both local and inside liquidity measures. The former refers to measures employing local quotes, i.e., the bid and the ask prices of a local market, whereas the latter refers to liquidity measures using inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. Panel B describes the difference in traded liquidity in consolidated and fragmented markets. We report the level of effective spread, and its decomposition into realized spreads and price impact based on 30 second future quote midpoints. We calculate effective spread as defined in (3) and realized spread as defined in (4). The price impact is then given by the difference. Finally, Panel C presents the difference in price (quote midpoint) efficiency in consolidated and fragmented markets. We report the mean and the standard deviation of the microstructure noise which is defined as the absolute difference between quote midpoint and fundamental value v_t .

Panel A: Quoted Liquidity

	Single Market	Fragmented Market	Fragmented Market Double ($\alpha = 0$)
Quoted Spread: Local	1.565	2.601	2.240
Quoted Spread: Inside	1.565	1.904	1.860
Quoted Depth: Local	1.584	1.082	1.692
Quoted Depth: Inside	1.584	1.445	2.751

Panel B: Traded Liquidity

	Single Market	Fragmented Market	Fragmented Market Double ($\alpha = 0$)
Effective Spread: Local	1.312	1.799	1.862
Effective Spread: Inside	1.312	1.452	1.613
Realized Spread 30: Local	0.865	1.013	1.372
Realized Spread 30: Inside	0.865	1.011	1.371
Price Impact 30: Local	0.441	0.789	0.487
Price Impact 30: Inside	0.441	0.442	0.242

Panel C: Price Efficiency

	Single Market	Fragmented Market	Fragmented Market Double ($\alpha = 0$)
Microstructure Noise Local: Mean $ v_t - p_t $	0.464	0.670	0.369
Microstructure Noise Inside: Mean $ v_t - p_t $	0.464	0.570	0.350

Table 4. Decomposition of Welfare by Trader Type

We report the average welfare defined in 5, waiting cost and money transfer defined in 6, all of them per trader differentiated by private value. The first row reports the results for a single market, i.e., a market organized as a single limit order market, whereas the second row for a fragmented market, i.e., a market organized as two limit order markets. Finally, the third row reports the results when the arrival rate of zero private value agents is double. The three measures are reported in ticks and calculated as the mean over 20 million new arrivals in equilibrium. Standard errors for all measures are small enough due to the large number of simulated events and are hence omitted. The Markov equilibrium is obtained independently for each scenario.

	Total		Average welfare per trader		Waiting cost per trader		Money transfer per trader						
	Welfare per Period	0	Private Value	8	Private Value	4	Private Value	0	Private Value	4	8	Total	
Single Market	3.742	0.543	3.510	7.265	3.745	0.000	-0.350	-0.162	-0.189	0.543	-0.140	-0.572	-0.065
Fragmented Market	3.740	0.626	3.479	7.202	3.740	0.000	-0.355	-0.172	-0.193	0.626	-0.166	-0.626	-0.066
Fragmented Market (Double $\alpha = 0$)	3.757	0.485	3.312	7.137	2.890	0.000	-0.127	-0.029	-0.049	0.485	-0.561	-0.835	-0.142

Table 5. Stock Characteristics

This table reports the characteristics of the treatment stocks and the corresponding control stocks generated based on [Huang and Stoll \(1996\)](#). Market Capitalization is the product of shares outstanding and Stock Price as on 31 December 2008, Trading Volume and Number of Trades is the average daily trading volume and number of trades for each stock between 1 December 2008 and 13 January 2009. We also report the market share of the two Euronext order books. Large (Small) order book is the order book with higher (lower) trading volume.

Panel A: Treatment Stocks

	Market Cap	Stock	Trading Volume			Number of Trades		
	€ million	Price (€)	€ '000	% Large	% Small	Count	% Large	% Small
DEXI	3,355	2.9	8,514	54.3%	45.7%	2,843.1	52.6%	47.4%
FOR	2,187	0.9	25,239	71.5%	28.5%	6,613.4	71.7%	28.3%
ISPA	24,985	17.2	180,346	55.6%	44.4%	18,231.8	52.7%	47.3%
UNBP	8,598	104.9	37,687	87.7%	12.3%	4,268.1	86.4%	13.6%
GLPG	80	3.8	183	77.1%	22.9%	67.9	27.9%	72.1%
ONCOB	87	6.6	23	90.7%	9.3%	3.5	73.5%	26.5%
RCUS	193	6.2	119	98.0%	2.0%	68.3	98.5%	1.5%
VRKP	105	20	43	94.4%	5.6%	19.5	94.3%	5.7%
THEB	59	3.5	11	68.4%	31.6%	7.5	77.1%	22.9%
MEAN	4,405	18.4	28,018	77.5%	22.5%	3,569.2	70.5%	29.5%

Panel B: Control Stocks ([Huang and Stoll, 1996](#))

	Market Cap	Stock	Trading Volume			Number of Trades		
	€ million	Price (€)	€ '000	% MoR	% Alternate	Count	% MoR	% Alternate
STM	3,355	4.6	15,157	100.0%	0.0%	2,635.1	100.0%	0.0%
CNAT	3,635	1.3	5,333	100.0%	0.0%	2,265.3	100.0%	0.0%
ABI	25,439	15.9	75,884	100.0%	0.0%	7,245.6	100.0%	0.0%
HRMS	10,652	101	12,193	100.0%	0.0%	1,546.2	100.0%	0.0%
OMT	84	4	15	100.0%	0.0%	9.2	100.0%	0.0%
TAM	81	6.8	29	100.0%	0.0%	23.5	100.0%	0.0%
AMG	184	6.9	2,989	100.0%	0.0%	903.5	100.0%	0.0%
SMTPC	117	20	20	100.0%	0.0%	9.4	100.0%	0.0%
DEVG	63	3.5	156	100.0%	0.0%	111.0	100.0%	0.0%
MEAN	4,846	18.2	12,420	100.0%	0.0%	1,638.8	100.0%	0.0%

Table 6. Empirical Findings: Impact on Quoted Liquidity

This table presents the results on the impact of the introduction of a single order book on quoted liquidity. We calculate quoted spread and depth in single and fragmented markets. We present both local and inside liquidity measures. The former refers to measures employing local quotes, i.e., the bid and the ask prices of a local market, whereas the latter refers to liquidity measures using inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. We estimate a difference-in-difference regression for quoted spread and quoted depth, in level and logarithm, and report the coefficient of the variable interacting the event dummy (which equals one for all days on or after 26 January 2009 and zero otherwise) with the treatment dummy (which equals one for all treatment stocks and zero for all control stocks). We employ stock and day fixed effects and double cluster standard errors by stock and day. In order to calculate local liquidity we choose one of the two order books in the simulated data and the venue with the larger trading volume in the pre-event period in the empirical analysis. Inside liquidity, in fragmented markets, is measured based on the best quotes (highest bid and lowest ask) across the two order books, and in consolidated markets, it is equal to the local liquidity. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Quoted Spread: Local	-0.542	0.266	-0.808**	-0.322**
Quoted Spread: Inside	-0.056	0.266	-0.323	0.090
Quoted Depth: Local	-878	-5,435	4,589	0.017
Quoted Depth: Inside	-485	-5,435	4,979	0.005

Table 7. Empirical Findings: Impact on Traded Liquidity

This table describes the difference in traded liquidity in consolidated and fragmented markets. We report the impact of the introduction of single order book on traded liquidity. We estimate a difference-in-difference regression for effective spreads, realized spreads, and price impact, in level and logarithm, and report the coefficient of the variable interacting the event dummy (which equals one for all days on or after 26 January 2009 and zero otherwise) with the treatment dummy (which equals one for all treatment stocks and zero for all control stocks). We employ stock and day fixed effects and double cluster standard errors by stock and day. In both panels, we compute local and inside traded liquidity. We measure local liquidity based on quotes on the order books where a transaction is executed. Inside liquidity, in fragmented markets, is measured based on the inside quotes across the two order books, and in consolidated markets, it is equal to the local liquidity. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Effective Spread: Local	-1.228	0.151	-1.394*	-0.375**
Effective Spread: Inside	-0.764	0.151	-0.926*	-0.145
Realized Spread 10: Local	-1.115	0.079	-1.211*	-0.517***
Realized Spread 30: Local	-1.110	0.064	-1.190*	-0.643***
Realized Spread 60: Local	-1.109	0.031	-1.153	-0.725***
Realized Spread 10: Inside	-0.803	0.079	-0.894	-0.376***
Realized Spread 30: Inside	-0.817	0.064	-0.893	-0.547***
Realized Spread 60: Inside	-0.796	0.031	-0.836	-0.615***
Price Impact 10: Local	-0.112	0.072	-0.183	-0.164
Price Impact 30: Local	-0.118	0.086	-0.204*	-0.133
Price Impact 60: Local	-0.120	0.121	-0.241**	-0.214
Price Impact 10: Inside	0.040	0.072	-0.032	0.123
Price Impact 30: Inside	0.053	0.086	-0.033	0.161
Price Impact 60: Inside	0.032	0.121	-0.090	0.024

Table 8. Empirical Findings: Impact on Price Efficiency

This table describes the difference in price (quote midpoint) efficiency in consolidated and fragmented markets. We report the impact of the introduction of single order book on price efficiency. We estimate a difference-in-difference regression for absolute values of return autocorrelation measured at 30-second and 5-minute intervals and the variance ratio based on 30-second and 5-minute returns, in level and logarithm, and report the coefficient of the variable interacting the event dummy (which equals one for all days on or after 26 January 2009 and zero otherwise) with the treatment dummy (which equals one for all treatment stocks and zero for all control stocks). We employ stock and day fixed effects and double cluster standard errors by stock and day. In order to calculate local price efficiency measures we choose one of the two order books in the simulated data and the venue with the larger trading volume in the pre-event period in the empirical analysis. Inside price efficiency, in fragmented markets, is measured based on the inside quotes across the two order books, and in consolidated markets, it is equal to the local price efficiency. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Autocorrelation 30: Inside	-0.007	0.007	-0.015	0.201
Autocorrelation 30: Local	-0.014	0.007	-0.021*	-0.233
Autocorrelation 300: Inside	-0.013	-0.002	-0.011	-1.444*
Autocorrelation 300: Local	-0.022	-0.002	-0.020	-1.545*
Variance Ratio 30/300: Inside	-0.045	-0.017	-0.028	-0.224
Variance Ratio 30/300: Local	-0.052	-0.017	-0.035	-0.206

Table 9. Impact on Trading Volume and Cross Market Arbitrage Analysis

Panel A reports the impact of the introduction of single order book on total Euronext trading volume. We estimate a difference-in-difference regression for the trading volume, in level and logarithm, and report the coefficient of the variable interacting the event dummy (which equals one for all days on or after 26 January 2009 and zero otherwise) with the treatment dummy (which equals one for all treatment stocks and zero for all control stocks). We employ stock and day fixed effects and double cluster standard errors by stock and day. Panel B summarizes the arbitrage opportunities arising on the two Euronext order books during the pre-event period i.e., between 1 December 2008 and 13 January 2009, and their resolution. Section 4.3.4 describes how we identify each arbitrage opportunity. Unique Instances are the average daily frequency of arbitrage opportunities on the two order books between 08:01 and 16:29, Negative Spread Time is the total amount of time during a trading session when the markets are crossed, Magnitude of Negative Spread is the frequency with which the negative bid-ask spread is equal to one tick, two ticks, three ticks, four ticks, and five or more ticks, and Trading Volume is the average daily volume which can be attributed towards resolution of the arbitrage opportunities. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Arbitrage Analysis

Stock	Unique Instances	Negative Spread Time	Trading Volume
DEXI	138.1	12.8%	1,145,164
FOR	183.8	11.5%	2,361,274
ISPA	622.0	6.0%	14,185,211
UNBP	166.1	3.0%	1,968,394
GLPG	2.4	4.3%	7,558
ONCOB	0.3	1.6%	990
RCUS	0.8	9.5%	9,020
VRKP	1.2	6.7%	4,351
THEB	0.6	2.6%	1,553
MEAN	123.9	6.4%	2,187,057

Panel B: Trading Volume

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Total Volume	2,628	-948	3,638*	0.080