JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 59, No. 1, Feb. 2024, pp. 221–248

doi:10.1017/S0022109022001521

Trader Competition in Fragmented Markets: Liquidity Supply Versus Picking-Off Risk

Alejandro Bernales
Universidad de Chile - Facultad de Economía y Negocios - Departamento de Administración
abernales@fen.uchile.cl

Nicolás Garrido
London School of Economics
N.garrido-sureda@lse.ac.uk

Satchit Sagade
Nasdaq and Leibniz Institute for Financial Research SAFE
satchit.sagade@nasdaq.com

Marcela Valenzuela
Pontificia Universidad Católica de Chile School of Management
mavalenb@uc.cl (corresponding author)

Christian Westheide
University of Vienna and Leibniz Institute for Financial Research SAFE
christian.westheide@univie.ac.at

Abstract

By employing a dynamic model with two limit order books, we show that fragmentation is associated with reduced competition among liquidity suppliers and lower picking-off risk of limit orders. Due to these countervailing channels, the impact of fragmentation on liquidity and welfare differs with asset volatility: When volatility is high (low), liquidity and aggregate welfare in a fragmented market are higher (lower) than in a single market. However, fragmentation always shifts welfare away from agents with exogenous trading motives and toward intermediaries. We empirically corroborate our model's predictions about liquidity. Our model reconciles the mixed results in the empirical literature.

I. Introduction

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 environment, where the same stock trades on multiple limit order books. In such an environment, traders compete with each other

For helpful comments and discussions, we thank Markus Baldauf, Evangelos Benos, Jonathan Brogaard, Sabrina Buti, Jean-Edouard Colliard, Thierry Foucault, Peter Gomber, Thomas Johann, Jan-Pieter Krahnen, Raman Kumar, Jens Lausen, Katya Malinova, Albert Menkveld, Josua Mollner, Robert Neal, Andreas Park, Seongkyu Gilbert Park, Christine Parlour, Giovanni Petrella, Talis Putnins, Uday Rajan, Yves Rannou, Ioanid Rosu, Patrik Sandas, Andriy Shkilko, Erik Theissen, Vincent van Kervel, Clara Vega, Shixuan Wang, Marius Zoican, 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, 2017 European Retail Investment Conference, 2017 Stern Microstructure Meeting, 2017 FIRS Conference, 2017 Spring International Conference of the French Finance

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

by dynamically choosing whether to supply or consume liquidity, and under what terms, across fragmented markets. These choices are determined by, among other things, investors' trading motives, market conditions, adverse selection, and security characteristics. Over the last 40 years, the literature examining the effects of fragmentation has provided mixed results and several relevant questions remain unanswered. For example, how does fragmentation shape agents' trading behavior? Is fragmentation good or bad for liquidity? What are the effects of fragmentation on welfare? In this article, we investigate these questions while allowing for the possibility that the answers may differ depending on market conditions.

We show that two intertwined channels drive traders' behavior, producing opposing effects on liquidity. First, as time priority in a fragmented market applies only within and not across limit order books, traders can circumvent time priority in one order book by submitting an identically priced limit order in the second order book. In a single-market setting, traders can jump ahead of the queue only by improving upon the existing price. This form of queue jumping leads to increased price competition among liquidity providers in a single-market setting as compared to a multi-market setup. We call this channel the *competition for time priority* channel.

At the same time, a fragmented market offers more protection against picking-off risk compared to a single market. This is because, if the asset's fundamental value moves against equally-priced limit orders in a fragmented market, the arbitrageur's aggressive orders do not necessarily arrive simultaneously in all the markets allowing market makers additional time to react and modify their unexecuted limit orders. Thus, fragmentation reduces the adverse selection component of the bid–ask spread due to a lower probability of limit orders being picked off. We call this channel the *adverse selection* channel.[1]

Our model builds on the single-market models of Goettler, Parlour, and Rajan ((2005), (2009)), which we extend to a multi-market setting. It is set up as a stochastic

Association, 2017 International Workshop in Financial Markets and Nonlinear Dynamics, 2017 FMA European Conference, BWL2017 Conference, 2017 AsianFA Conference (best paper award), 2017 Verein für Socialpolitik Jahrestagung, 2017 Northern Finance Association Conference, 2017 Annual Central Bank Conference on the Financial Microstructure of Financial Markets, 2017 Annual Meeting of the German Finance Association (DGF), 2018 SFS Cavalcade, and 2018 FMA Conference, and seminar participants at University of Birmingham, University of Mannheim, University of Manchester, University of Frankfurt, University of Chile, and 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. Sagade also gratefully acknowledges financial support from Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 329107530. Bernales acknowledges the support of Fondecyt Project No. 1231660. Valenzuela acknowledges the support of Fondecyt Project No. 1231630, and the support of Instituto Milenio ICM IS130002. We are grateful to Rodrigo Orellana and Italo Riarte for their research assistance. The support of the Economic and Social Research Council (ESRC) in funding the SRC is gratefully acknowledged (grant number ES/R009724/1). We are grateful to EUROFIDAI for access to the BEDOFIH database (for more information, see http://www.eurofidai.org). The views expressed in the article are solely the authors' and do not necessarily reflect those of Nasdaq.

[1]In our model, all agents can only submit one-share orders thereby restricting the arbitrageur's ability to simultaneously pick off limit orders across fragmented markets. However, the adverse selection channel is relevant even if this assumption is relaxed due to latencies associated with the geographical location of fragmented markets and the internal processing of orders within a market.

trading game in which a single asset, with a stochastically evolving fundamental value, can be traded in 2 limit order books by diverse (in their private values) and competing agents. Moreover, agents face a cost of delaying their execution, which represents an opportunity cost and incentivizes traders' competition.

Traders also have the possibility of reentering the market to revise existing limit orders, though they cannot do so instantaneously when the fundamental value changes. Thus, other traders can profit by picking off unexecuted unfavorably priced limit orders. This picking-off risk means that limit orders are exposed to the *adverse selection* channel. The adverse selection channel is relevant even in today's automated markets with high-frequency trading (HFT) firms. On the one hand, limit orders submitted by HFT firms are still exposed to picking-off risk, because there is no way to completely ensure that HFT firms will be able to modify their stale limit orders before the arrival of market orders from other HFT firms that would like to exploit such picking-off opportunities. On the other hand, HFT firms, looking for picking-off opportunities through market orders, still have delays in order transmissions and processing times in the exchanges (i.e., even HFT firms cannot simultaneously snipe all exchanges for picking-off opportunities), which can give time to other HFT firms to update their mispriced limit orders.

Goettler et al. ((2005), (2009)) parameterize their models based on Hollifield, Miller, Sandås, and Slive (2006), who in turn use order-level data from the early-1990s to compute trader arrival rates, fundamental value volatility, and private value distributions. We solve our model numerically by employing the same parameters used by Goettler et al. (2009) and compare a multi-market environment to a single-market setup under two scenarios involving different levels of asset volatility: A low (high) value of asset volatility corresponds to a low (high) level of picking-off risk. We also compare our model parameters with estimates obtained from order-level data for stocks traded on the London Stock Exchange in Jan. 2015, and confirm that the parameters from Goettler et al. (2009) are not too dissimilar to those in modern markets. This is because, while trading activity is much higher today than at the time of the data originally used to obtain those parameters, the parameters are expressed relative to each other, that is, while markets have undergone large changes when observed in clock time, the change is much smaller when the activity is measured in trading time.

Agents endogenously decide whether to provide or consume liquidity in a limit order book of their choice and in the presence of discrete prices and picking-off risk. Those with an intrinsic motive to trade balance the delay costs associated with limit orders and the immediacy costs of market orders. Those without any intrinsic trading motives generate their gains from liquidity provision or by picking off limit orders.

We provide evidence consistent with both channels of trader behavior. Consistent with the competition for time priority channel, we find that agents submit limit orders at the best quotes more frequently in the single-market setting. At the same time, consistent with the adverse selection channel, we find that the picking-off risk is lower in the multi-market setting. Furthermore, these two effects interact with each other: In the single-market setting, the increased aggressiveness of limit orders resulting from more competition for time priority increases the orders' exposure to the picking-off risk due to a higher probability of such orders becoming mispriced upon a change in the fundamental value of the asset.

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

The equilibrium impact on liquidity depends on which of the two channels dominates. The competition for time priority channel is more important when fundamental value volatility, or picking-off risk, is low. In this setting, liquidity, as measured by quoted and effective spreads, is higher in the single market as compared to the multi-market scenario. Conversely, when fundamental value volatility, or picking-off risk, is high, the adverse selection channel dominates. In this setting, liquidity is higher in the multi-market scenario than in the single market. The interplay between these two competing channels likely helps explain the mixed findings in the empirical literature regarding the relationship between liquidity and fragmentation.

Consistent with the liquidity effects, aggregate welfare is marginally higher in the consolidated (fragmented) market than in the fragmented (consolidated) market in the low (high) volatility setting. However, independent of the market conditions, intermediaries always extract higher welfare gains in fragmented markets. Conversely, agents with intrinsic trading motives are always better off in consolidated markets. The higher revenues earned by intermediaries in fragmented markets, without a commensurate increase in total welfare, strongly suggest that costly investments in intermediation capacities, such as the high-speed connections to venues and subscriptions to exchanges' real-time data feeds,[2] are socially wasteful. This raises the question of whether restricting fragmentation would lead to improvements in social welfare.[3]

We empirically test the model implications related to the liquidity effects of fragmentation using data from the second half of 2012 for German and French large-cap and mid-cap stocks. We employ panel regressions to determine how quoted and effective spreads depend on fragmentation – measured across the primary listing venue (Deutsche Börse or Euronext Paris) and the largest rival exchange, Chi-X – and within-stock variation in volatility. Consistent with our model, we find that, while there is an inverse relation between volatility and liquidity, an increase in fragmentation is associated with lower (higher) quoted and effective spreads on high (low) volatility days. As none of the venues implemented any major changes to their market structure during our sample period, these results lend empirical support to our model.

It is important to note that our focus is on providing a relatively simple but realistic model, in order to understand the impact of fragmentation on liquidity and welfare of two intertwined channels (i.e., the competition for time priority channel and the adverse selection channel), which produce opposing effects on market quality. We recognize that the impact of fragmentation on liquidity and welfare may be affected by other elements outside the scope of our modeling setup such as off-exchange venues, behavioral issues, and market regulations, amongst other factors. For example, our model does not capture agent behavior in a market

[2]Cespa and Foucault (2013) and Easley, O'Hara, and Yang (2016) further highlight the adverse effects associated with exchanges providing differential access to market data feeds.

[3]For instance, due to national and international mergers between exchanges, individual market operators routinely operate several limit order books. In the United States, the three largest exchange operators – Intercontinental Exchange, Nasdaq OMX, and Cboe – operate a total of 12 lit equity exchanges as of Mar. 2022. Our results indicate that such within-operator fragmentation in the absence of, or under minimal, venue competition is harmful.

involving venues with no pretrade transparency such as dark pools or internalizing dealers who employ payment for order flow arrangements and pay retail brokers to route their customers' orders to such dealers instead of public exchanges. Instead, the objective of our study is to provide some light to understand the impact of fragmentation when competition among liquidity suppliers and adverse selection are simultaneously considered.

Our model allows for a potential explanation of the conflicting empirical results observed in the literature. For example, studies find that i) fragmentation increases liquidity (see, e.g., Boehmer and Boehmer (2003), Fink, Fink, and Weston (2006), Nguyen, Van Ness and Van Ness (2007), Foucault and Menkveld (2008), Hengelbrock and Theissen (2009), Chlistalla and Lutat (2011), Menkveld (2013), and He, Jarnecic and Liu (2015)); ii) fragmentation harms liquidity (see, e.g., Bessembinder and Kaufman (1997), Arnold, Hersch, Mulherin, and Netter (1999), Hendershott and Jones (2005), Bennett and Wei (2006), and Nielsson (2009)); and iii) fragmentation has mixed effects on liquidity (see, e.g., Boneva, Linton and Vogt (2016), Haslag and Ringgenberg (2023)). In our model, for high (low) volatility assets, the picking-off risk (competition for time priority) effect is dominant and, as a result, fragmentation is associated with smaller (higher) bid–ask spreads. Consistent with this intuition, O'Hara and Ye (2011) find that fragmentation is associated with significantly lower effective spreads only for small-cap stocks in the US.[4]

Our article is also related to the broader literature on market fragmentation in limit order markets.[5] Foucault and Menkveld (2008) model fee-based competition between 2 operators and predict that the entry of a second exchange will increase consolidated depth and that the increased use of smart order routers will increase liquidity in the entrant market. However, their model does not feature picking-off risk. Baldauf and Mollner (2021) also present a model for fee-based competition between trading venues. In their model, there is a single liquidity provider who quotes in all exchanges. They show that market fragmentation increases the picking-off risk. This is because they assume that the single liquidity provider maintains limit orders in all exchanges, and there are infinitely many snipers (arbitrageurs) looking for mispriced limit orders upon changes in market conditions. Thus, the single liquidity provider faces an increased risk of being picked off when there are more exchanges due to the exposure in each venue.

In our model, differently from Baldauf and Mollner (2021), we assume that liquidity is provided by diverse agents and that there are a finite number of potential snipers arriving randomly. Most importantly, as in reality, potential snipers cannot simultaneously check all exchanges to exploit picking-off opportunities due to delays in order transmissions and processing times (which is even observed by

[4]Small-cap stocks are associated with higher volatility than large-cap stocks (see, e.g., Table 1 in Brogaard, Hendershott, and Riordan (2014)).

[5]Early theories of fragmentation such as Mendelson (1987), Pagano (1989), and Chowdhry and Nanda (1991), while not explicitly modeling limit order markets, highlight the positive network externalities associated with consolidating trading in a single venue. However, such a consolidated market is no longer the equilibrium outcome in the absence of post-trade transparency (Madhavan (1995)) and in the presence of real-world frictions such as differences in markets' absorptive capacity and institutional mechanisms (Pagano (1989)), order-splitting behavior (Chowdhry and Nanda (1991)), and trader heterogeneity (Harris (1993)).

snipers with the high-frequency technology, as previously explained); thus, the probability of being picked off is lower for the different liquidity providers in each exchange, compared to what it would be in a single-market setting. Hence, the liquidity providers that have not been picked off after changes in market conditions, in expectation, get more time to update their limit orders in wrong positions. Therefore, differently to Baldauf and Mollner (2021), here, a multi-market environment offers more protection against picking-off risk compared to a single-market setup.

In other studies, Parlour and Seppi (2003) examine competition between a specialist market and a pure limit order book, Pagnotta and Philippon (2018) investigate the joint role of trading fees and speed of market access across competing venues, and Chao, Yao, and Ye (2019) focus on the role of tick sizes in the dispersion of fee schedules in fragmented markets. In contrast to these studies, our model simultaneously features competition for time priority and picking-off risk in fragmented markets. Consequently, it allows for more flexible agent behavior and better captures the dynamics of these two channels. Specifically, it allows for endogenous liquidity provision and consumption in the presence of real-world frictions (such as price discreteness) and in the absence of perfect competition between agents (as in Glosten (1998)). Using a dynamic equilibrium model for market fragmentation, we show that trader competition and picking-off risk, independently and through interaction with each other, can lead to heterogeneous effects of fragmentation on liquidity and investor welfare.

II. Multi-Market Model

Our aim is to study the effect of fragmentation on liquidity and welfare in the presence of competition for time priority and adverse selection (picking-off) risk. Our framework is based on the model of a single limit order market of Goettler et al. ((2005), (2009)), which we extend to a multi-market setting.[6] We set up a dynamic trading game in which agents make endogenous decisions to maximize their expected payoffs, taking into account their private reasons for trading the asset, market conditions, and the strategies employed by agents expected to arrive in the future.

A. Model Setting

We consider an economy in continuous time with a single financial asset that trades in two limit order books. The fundamental value of the asset, v_t, is stochastic, and its innovations follow a Poisson process with parameter λ_v. In the case of an innovation, the fundamental value increases or decreases by one tick, d, with equal probability. There is competition among agents. The economy is populated by agents who arrive sequentially following a Poisson process with intensity λ_a.

All agents observe both limit order books (i.e., prices and depths at each price) and the fundamental value of the asset v_t without delay. Agents can submit limit or market orders to either book. Moreover, agents can reenter the market to modify unexecuted limit orders. There is adverse selection. In other words, agents cannot

[6]We are indebted to Ron Goettler, Christine Parlour, and Uday Rajan for kindly providing the C codes for their models (Goettler et al. (2005), (2009)).

instantaneously modify their unexecuted limit orders after a change in market conditions, but instead, reenter the market following a Poisson process with parameter λ_r. Thus, agents submitting limit orders face picking-off risk.[7]

The limit order book at time t and in market m with $m \in \{1,2\}$, $L_{m,t}$, is characterized by a set of discrete prices denoted by $\{p_m^i\}_{i=-N}^{N}$, 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 we refer to as the tick size (i.e., $d = p_m^{i+1} - p_m^i$). The tick size is the same in both limit order books. Let $l_{m,t}^i$ be the queue of unexecuted limit orders in order book m at time t and price p_m^i. A positive (negative) $l_{m,t}^i$ denotes the number of buy (sell) limit orders, and represents the depth of the book $L_{m,t}$ at price p_m^i. In the book $L_{m,t}$ at time t, the best bid price is $B(L_{m,t}) = \sup\left\{p_m^i | l_{m,t}^i > 0\right\}$ and the best ask price is $A(L_{m,t}) = \inf\left\{p_m^i | l_{m,t}^i < 0\right\}$. $B(L_{m,t}) = -\infty$ or $A(L_{m,t}) = \infty$ if the order book $L_{m,t}$ is empty at time t on the buy side or on the sell side, respectively. Each limit order book independently respects price and time priority when executing the limit orders, that is, buy (sell) limit orders at higher (lower) prices have priority in the queue and limit orders submitted earlier at the same price are executed first.

Agents are risk-neutral, but heterogeneous in terms of their intrinsic economic motives for trading the asset. These motives are reflected in their private values. Each agent has a private value α, which is known to her. α is drawn from the vector $\Psi = \{\alpha_1, \alpha_2, \ldots, \alpha_g\}$ based on the cumulative distribution F_α, where g is a finite integer. Private values reflect the fact that agents want to trade for various reasons unrelated to the fundamental value of the asset (e.g., hedging needs, tax exposure, and/or wealth shocks). They are idiosyncratic and constant for each agent.

Furthermore, similarly to Goettler et al. (2009), agents face a cost of delaying, which represents an opportunity cost. This cost is denoted by $\rho \in [0,1]$, does not depend on the choice of the order book, and applies to agents' total payoff. It is the same for all agents and applies per unit of time, that is, all agents proportionally pay the same cost of delaying if they wait for the same amount of time until their order executes, while agents who wait for a longer amount of time until their order executes pay a proportionally higher cost of delaying.

Agent heterogeneity, delay costs, and the fundamental value of the asset determine agents' trading behavior. On the one hand, suppose agent i with a positive private value (i.e., $\alpha > 0$) arrives at time t_i. This agent is likely to be a buyer because she would like to have the asset to obtain the intrinsic benefit reflected by α. In this case, the agent's expected payoff is $(\alpha + v_{t'} - p)e^{-\rho(t'-t_i)}$, where p is the transaction price, t' is the time of the transaction, and $v_{t'}$ is the expected fundamental value of the asset at time t'. Moreover, if α is very high, her delay cost, denoted by $\left(e^{-\rho(t'-t_i)} - 1\right)\alpha$, is correspondingly high, and she may therefore prefer to buy the asset as soon as possible by using a market order. In this case, the agent will pay an immediacy cost denoted by $(v_{t'} - p)^{-\rho(t'-t_i)}$. The agent will accept this immediacy cost because she is mainly generating her profits from the large private value, α,

[7]As a robustness check, in untabulated results, we assume that agents observe the fundamental value of the asset, v_t, with a time lag Δ (i.e., agents observe $v_{t-\Delta}$). The outcomes of this robustness check are qualitatively similar to the results reported here. This is because such time lag simultaneously affects all agents, and thus only adds additional noise to the trading game.

rather than from the trading process per se. Accordingly, an agent with a high absolute private value will probably be a liquidity taker.[8]

On the other hand, suppose an agent j with a private value equal to 0 (i.e., $\alpha = 0$) arrives at time t_j. This agent needs to find a profitable opportunity purely in the trading process, by obtaining a good price relative to the fundamental value, because she does not obtain any intrinsic economic benefit from trading. Consequently, she may be patient and prefer to act as a liquidity provider, in turn earning the immediacy cost (bid–ask spread) paid by a liquidity taker. Alternatively, she may trade aggressively against a standing limit order that is mispriced relative to the fundamental value. Note that agents with $\alpha = 0$ are indifferent with respect to the side of the market they take because they can maximize their gains by either selling or buying the asset.

Agents are exposed to the risk of being picked off, because limit orders can generate a negative payoff if they are in an unfavorable position relative to the fundamental value. For example, suppose an agent i with $\alpha = 0$ first arrives at time $t = 0$ and submits a buy limit order to set the best bid price, B, in market $m = 1$. Suppose further that, at time t^*, the fundamental value of the asset decreases to level v_{t^*}, such that $v_{t^*} < B$, and subsequently, another agent, denoted by j, with private value $\alpha = 0$, arrives. Since agent i cannot immediately modify her unexecuted limit order, agent j can submit a market sell order and pick off agent i's order, generating an instantaneous profit equal to $(B - v_{t^*})$. Agent i, on the other hand, has a negative realized payoff given by $(v_{t^*} - B)e^{-\rho t^*}$.

Each agent competes for the execution of her order of one share. This assumption captures the fact that even in a market environment with many HFT firms, they still have delays in order executions. For example, HFT firms cannot simultaneously exploit picking-off opportunities on all exchanges (due to delays in order transmissions and processing times in the exchanges), which can give time to other HFT firms to update their mispriced limit orders.[9]

Each agent takes four main trading decisions upon arrival: i) to submit an order to L_1 or L_2, ii) to submit a sell or a buy order, iii) to decide whether the order is a limit or a market order, and iv) to choose the submission price in the case of a limit order. As mentioned above, an agent can reenter the market and modify her unexecuted limit order. Hence, she has to take the following additional trading decisions after reentering: i) to keep her unexecuted limit order unchanged or to cancel it, ii) in the case of a cancelation, to submit a new order to L_1 or L_2, iii) to choose whether the new order will be a buy or a sell order, iv) to decide whether the order will be a limit or a market order, and v) to choose the submission price in the case of a limit order. The decision to leave the order unchanged has the advantage of maintaining its time priority in the respective queue. The disadvantage is the increased exposure to picking-off risk or nonexecution risk depending on the direction of the change in the fundamental value. Once a trader has submitted a limit order, she remains part of the trading game by revising her order until it is executed; however, the trader exits the market permanently after the execution of her order. Figure 1 describes the agent's trading decisions upon arrival and after reentering.

[8]Analogously, a similar example can be produced in the other direction in the case of an agent with a negative private value (i.e., $\alpha < 0$) having a preference to sell.

[9]We could include additional shares per agent in the trading decision. However, we assume one share per agent to make the model computationally tractable.

FIGURE 1

Agent's Trading Decisions

Figure 1 reflects the agent's trading decisions upon arrival and after reentering.

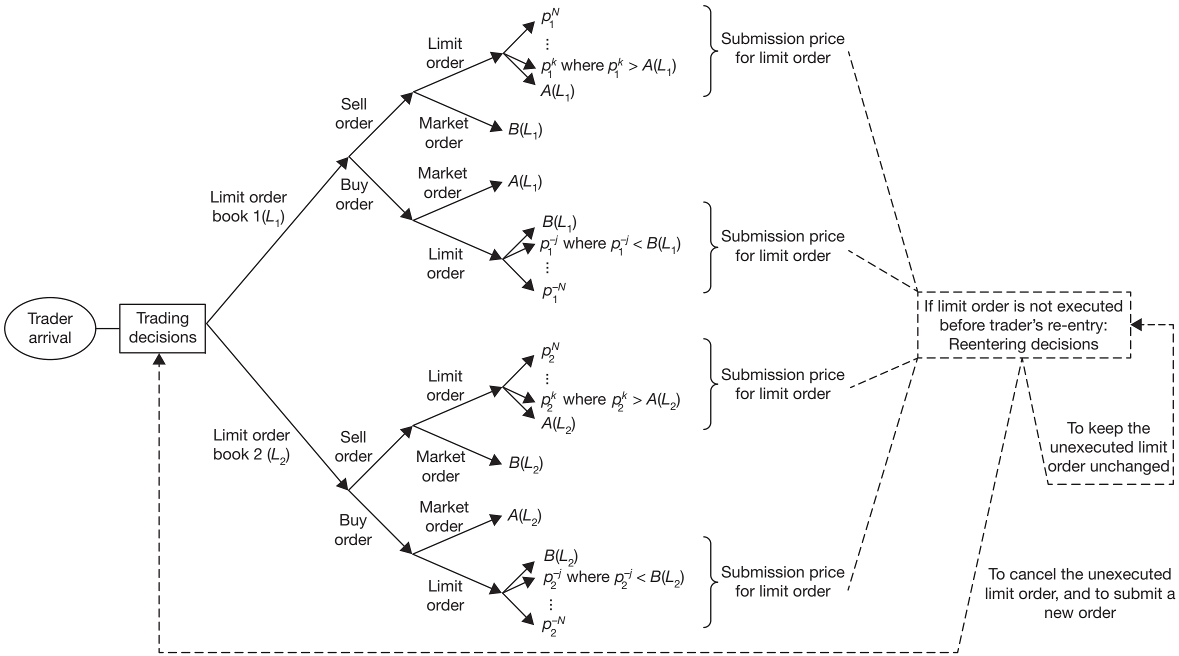

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

In the model, each trader behaves optimally by taking trading decisions that maximize her expected discounted payoff. Optimal decisions depend on the state of the economy. Each state of the economy is described by: market conditions (i.e., the status of both limit order books and the fundamental value of the asset); the trader's individual characteristics (i.e., the trader's private value); and, if the trader previously submitted a limit order, the status of her existing order, and otherwise the absence of such an order.

Since an analytic solution is not feasible, we numerically solve the model by using the Pakes and McGuire (2001) algorithm to compute a stationary and symmetric Markov perfect equilibrium. In Appendix I of the Supplementary Material, we describe in detail the agents' dynamic maximization problem in each state of the economy, the model's equilibrium, the solution approach, and the implementation of the Pakes and McGuire (2001) algorithm.

B. Model Parameterization

Goettler et al. (2009) rely on the empirical findings of Hollifield et al. (2006) to identify parameters that reasonably describe real market features. We mostly use the same parameters. Specifically, we set the intensity of the Poisson process followed by agents' arrival, λ, to 1. The intensity of the Poisson process followed by agents' reentry, λ_r, is set to 0.25; the intensity of the Poisson process followed by the innovations to the fundamental value, λ_v, is set to 0.125 (as in Goettler et al. (2009) and 0.625, to simulate scenarios of low and high volatility, respectively. We set the tick size, d, in both order books to 1, and the number of discrete prices available on each side of both order books to $N = 31$. The delay cost, 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 $F_\alpha = \{0.15, 0.35, 0.65, 0.85, 1.0\}$.

To alleviate concerns surrounding the suitability of our parameter choices, we employ the Hollifield et al. (2006) approach to compute the trader arrival rate, fundamental value volatility, and the distribution of private values using message-level data for the month of Jan. 2015 from the LSE for 2 FTSE-100 stocks. Appendix II of the Supplementary Material contains the implementation details of the Hollifield et al. (2006) approach, and the mean estimated parameters across the 21 days in our sample period. Evidently, these estimates are not substantively different from our model parameter choices.

Although the similarity of the parameters estimated from the 1990s and modern automated markets may appear surprising, it can be explained by the fact that they are expressed relative to the agents' arrival rate, which is set to 1. While to Goettler et al. (2009) a unit of time corresponds to 1 minute, it corresponds to a smaller period in modern markets. The parameters related to asset values are relative to the tick size, and the reduced magnitude of tick sizes in recent decades in combination with the reduced duration of a period in the model happens to make the volatility of the fundamental value correspond well to recent empirical data. Our calibration also shows that the private values, also expressed in ticks, appear to have reduced roughly proportionately to the reduction in tick size. Thus, even if financial markets have evolved over time, the relative relationships between trader arrival rates, fundamental value volatility, and the distribution of private values have remained quite stable.[10]

III. Theoretical Implications

A. Intuition Behind 2 Countervailing Channels

Before turning to the results, we provide an intuitive explanation for the two countervailing channels that help explain the effect of fragmentation on liquidity: The competition for time priority channel and the adverse selection channel.

On the one hand, to increase the probability of executing a limit order in the single-market setting, agents submit an order at a more aggressive price. However, in a fragmented-market setting, if there is a standing order at the best price on only one of the books, a trader can choose to submit an order at the same price in the second book and obtain a 50% probability of executing before the order submitted earlier. Hence, it is more likely to observe higher trading competition in a single-market than in a fragmented-market, as described in Foucault and Menkveld (2008).

On the other hand, upon a change in the asset's fundamental value, a given limit order cannot be modified instantaneously and hence can be picked off by incoming traders (i.e., there is an adverse selection). As these traders can potentially choose between stale orders present in both limit order books, the execution probability of individual limit orders in fragmented markets is lower allowing them, in expectation, more time to update their orders. Hence, compared to a single-market setting, a multi-market environment offers more protection against picking-off risk.

B. Trading Behavior

Our aim is to examine the effect of fragmentation on liquidity and welfare, specifically via the competition for time priority and adverse selection channels. To this end, we generate simulated data sets for the following two cases. First, we consider an environment with low volatility or, equivalently, low picking-off risk, by setting $\lambda_v = 0.125$. In this simulated environment, the importance of the competition for the time priority channel, compared to the adverse selection channel, is relatively high. Second, we consider an environment with high asset volatility by setting $\lambda_v = 0.625$, that is, 5 times the level used in the low-volatility case.[11] In this simulated environment with high picking-off risk, the adverse selection channel is

[10]This argument is similar to the use of business time instead of clock time in the so-called time deformation literature. Hasbrouck (1999) argues that business time transformations better capture price dynamics due to their ability to encompass variations in the rate of information arrival. For example, Clark (1973) proposes using volume time to better capture deviations of asset returns from normality. Jones, Kaul, and Lipson (1994) find a strong positive relationship between return volatility and number of transactions. More recently, Kyle and Obizhaeva (2016) propose that distributions of risk transfers and transaction costs are constant when time is measured in terms of the arrival rate of risk transfer bets. They refer to this as market microstructure invariance.

[11]We only report these two cases because, despite the use of the Pakes and McGuire (2001) algorithm, which reaches the equilibrium only in the recurring states class, the model solution is still computationally intensive. With modern hardware, the computation of the results for each specification takes around 30 days. Nevertheless, as a robustness check, we also run the model with an intermediate

TABLE 1
Trading Behavior

Panel A of Table 1 reports the distribution of buy limit and buy market orders executed by each agent type and the probability of submitting a limit order at the best bid price. Panel B reports the picking-off risk, defined as the proportion of buy limit orders executed above the fundamental value, for limit orders executed by each agent type. Panel C reports the average time to execution of limit orders, defined as the difference between the order execution time and the agent's market entry time, executed by each agent type. We generate simulated data sets for the following two cases: i) an environment with low picking-off risk, that is, low volatility, in which we set $\lambda_v = 0.125$ and ii) an environment with high picking-off risk, that is, high volatility, in which we set $\lambda_v = 0.625$. We report all statistics for a single and a fragmented market when volatility is low and when volatility is high. We omit standard errors for the differences because a large number of trader arrivals leads to a difference in means to the order of 10^{-3} being statistically significant.

		LOW_VOLATILITY: $\lambda_v = 0.125$ (Low Levels of Picking-off Risk)					HIGH_VOLATILITY: $\lambda_v = 0.625$ (High Levels of Picking-off Risk)				
			Private Value $\|\alpha\|$					Private Value $\|\alpha\|$			
	Order Type	Aggr. Prob.	0	4	8	Total	Aggr. Prob.	0	4	8	Total
Panel A. Order Submissions											
Single market	Limit	35.9%	77.1%	52.5%	19.5%	–	24.7%	21.4%	65.4%	58.2%	–
	Market	–	22.9%	47.5%	80.5%	–	–	78.6%	34.7%	41.9%	–
Frag market	Limit	28.5%	78.4%	50.7%	20.7%	–	22.8%	51.3%	57.5%	38.7%	–
	Market	–	21.6%	49.3%	79.3%	–	–	48.7%	42.5%	61.3%	–
Panel B. Picking-off Risk											
Single market	–	–	4.1%	26.8%	73.9%	21.8%	–	13.8%	53.0%	85.4%	59.3%
Frag market	–	–	3.0%	25.8%	72.1%	20.8%	–	10.9%	46.4%	75.7%	42.3%
Panel C. Time to Execution											
Single market	–	–	14.9	3.5	2.1	8.6	–	23.8	2.9	1.2	4.9
Frag market	–	–	11.5	3.7	1.9	7.1	–	12.1	3.7	1.7	5.8

relatively more important than the competition for the time priority channel. The low and high volatilities correspond to the low and high probabilities that the asset value changes, q_τ^{low} and q_τ^{high}, respectively. For both scenarios, we analyze two market setups: i) a consolidated market with one limit order book and ii) a fragmented market with two identical limit order books. We compute mean levels of the variables of interest for each of the four settings.[12]

Agents' order submission strategies determine the liquidity characteristics and the welfare of different agent types and that of the economy as a whole. Hence, we start by investigating the agents' trading behavior in consolidated and fragmented markets in the two volatility scenarios. Throughout this section (Section III.B), we report trader statistics for buy orders. The results for sell orders are identical and omitted for brevity.

Table 1 presents the results. In Panel A, we present the distribution of executed limit and market orders by each agent type, for low and high levels of volatility. When volatility is low, agents with private value $\alpha = 0$ execute 77.1% of their trades using limit orders in the single-market setting, whereas agents with private value $|\alpha| = 8$ execute 80.5% of their trades using market orders. Agents with private value $|\alpha| = 4$ use limit and market orders roughly equally. These frequencies remain

level of volatility by setting $\lambda_v = 0.375$. The results obtained from this robustness check are, as expected, in between those obtained for $\lambda_v = 0.125$ and $\lambda_v = 0.625$.

[12]We do not report standard errors because a large number of trader arrivals ensures that standard errors are sufficiently low and a difference in means even to the order of 10^{-2} is statistically significant.

largely unchanged when the market is fragmented. These results are consistent with the idea that agents with private value $\alpha = 0$ act as market makers and try to earn the bid–ask spread, whereas those with private reasons to trade tend to consume liquidity as their cost of waiting is relatively high.

When volatility is high, the behavior changes substantially. Speculators increase their frequency of market orders due to the increased exposure to picking-off risk, especially in the single market where they predominantly generate their trading gains by picking off mispriced limit orders. Conversely, agents with private value $|\alpha| = 8$ increase their frequency of limit orders, also to a larger extent in the single market than in the fragmented market. This is because their large private value allows them to bear some picking-off risk, especially when the costs of taking liquidity are relatively high due to a reduced supply of liquidity from other traders. The liquidity supplied by agents with private value $|\alpha| = 8$, in turn, is profitably picked off by the speculators. This is consistent with studies examining the behavior of high-frequency trading (HFT) firms across stocks with different volatility. For example, Brogaard et al. (2014) find that, in the US, the ratios of liquidity-consuming and -supplying trades by HFT firms for large-cap (or low volatility) stocks versus small-cap (or high volatility) stocks are 1.0 and 2.3, respectively. The differences in order choice between the low- and high-volatility scenarios are less pronounced in the multi-market setting, as it provides higher protection against picking-off risk, thereby making speculators more willing to submit limit orders than they are in the single-market setting. In conclusion, in a high volatility state, agents submit less aggressive orders in a consolidated market due to the higher adverse selection risk.

Panel A of Table 1 also reports the probability of submitting an aggressive limit order for each setting. This probability is computed as the proportion of buy limit orders that are submitted at the best bid price. For both volatility scenarios, fragmented markets decrease the level of traders' competition. In the low-volatility scenario, 35.9% of limit orders are aggressive in the consolidated market, whereas only 28.5% of limit orders are aggressive in the fragmented market. The difference persists even in the high-volatility scenario, where 24.7% (22.8%) of limit orders are aggressive in consolidated (fragmented) markets. These findings confirm that competition for time priority is less intense in fragmented than in consolidated markets.

In Panel B of Table 1, we report the level of picking-off risk for each of the two market settings and volatility scenarios for all agent types. Picking-off risk is measured as the proportion of buy limit orders executed above the fundamental value. Consistent with the intuitive explanation provided above, we find that the aggregate level of picking-off risk in the low-volatility setting is 21.8% in a single market and 20.8% in a fragmented market, whereas it increases to 59.3% in a single market and 42.3% in a fragmented market in the high-volatility setting. The picking-off risk is lower in the fragmented market for each agent type. For example, when asset volatility is low, agents with $\alpha = 0$ and $|\alpha| = 8$ have a picking-off risk of 4.1% (3.0%) and 73.9% (72.1%), respectively, in the consolidated (fragmented) market. The corresponding numbers in the high-volatility setting are 13.8% (10.9%) and 85.4% (75.7%), respectively. This shows that fragmented markets provide more protection against picking-off risk in both volatility settings.

Panels A and B of Table 1 also highlight the relative importance of the competition for time priority and picking-off risk channels in different volatility scenarios. The difference in the probability of submitting aggressive limit orders between the single- and multi-market settings is *less* pronounced (7.4% vs. 1.9%), and the difference in aggregate picking-off risk is *more* pronounced (1.0% vs. 17.0%) when the fundamental volatility is high. These differences suggest that the competition for time priority dominates the picking-off risk channel when asset volatility is low, whereas the effects associated with the picking-off risk channel dominate the effects associated with the competition for time priority channel when asset volatility is high.

Finally, in Panel C of Table 1, we report the time to execution of limit orders in aggregate and for each agent type. We define time to execution as the difference between the time an order executes and the market entry time of the agent submitting the order. The time to execution of agents with private value $\alpha = 0$ decreases from 14.9 units in a consolidated market to 11.5 units in a fragmented market when volatility is low because the need to reprice existing limit orders is reduced when queues are shorter and trader competition is lower. The difference is even more significant in the high-volatility scenario, where it decreases from 23.8 units in a consolidated market to 12.1 units in a fragmented market. For agents with private value $\alpha \neq 0$, there is little difference in the time to execution between the consolidated and fragmented market settings in both volatility scenarios. Specifically, for the agents with $|\alpha| = 4$, time to execution in the consolidated (fragmented) market is 3.5 (3.7) units in the low-volatility setting and 2.9 (3.7) units in the high-volatility setting. For the agents with $|\alpha| = 8$, time to execution in the consolidated (fragmented) market is 2.1 (1.9) units in the low-volatility setting and 1.2 (1.7) units in the high-volatility setting.

When we compare the two volatility scenarios, the time to execution for agents with $\alpha = 0$ ($\alpha \neq 0$) is lower in the low (high) volatility setting. This relates directly to the results in Panel A, discussed above. In the high-volatility scenario, intermediaries drastically reduce the frequency with which they provide liquidity. When they do submit limit orders, they quote very wide prices leading to an increase in the time to execution. The opposite holds for agents with $\alpha \neq 0$. They provide liquidity at attractive prices in the high-volatility setting leading to a lower time to execution.

C. Impact on Quoted and Effective Spreads

We compute time-weighted quoted bid–ask spreads based on *local* and *inside* quotes, where the former comprise the bid and ask prices in one of the two markets, whereas the latter comprises the highest bid and the lowest ask across the two limit order books. The two versions are obviously identical in a single-market setting. Differences in quoted spreads do not necessarily translate into commensurate differences in transaction costs for traders submitting market orders. Thus, we also compare the mean effective spread in the single and fragmented markets, which captures the actual transaction costs incurred by traders submitting market orders. We calculate the effective spread as follows:

(1) $$\text{EFFECTIVE\_SPREAD} = 2x_t(p_t - m_t)/m_t,$$

TABLE 2
Quoted and Effective Spreads

Table 2 reports the impact of fragmentation on spreads. We generate simulated data sets for the following two cases: i) an environment with low picking-off risk, that is, low volatility, in which we set $\lambda_V = 0.125$, and ii) an environment with high picking-off risk, that is, high volatility, in which we set $\lambda_V = 0.625$. We compute local and inside versions of spread measures. The former is based on the local quotes, that is, the bid and ask prices of a local market, whereas the latter is based on the inside quotes, that is, the highest bid and the lowest ask across the two limit order books. In the single-market setting, the local and inside quotes are identical. Quoted spread is the difference between the best bid and best ask prices. Effective spread is defined in equation (1). We report all statistics for a single and a fragmented market when volatility is low and when volatility is high. We omit standard errors for the differences because a large number of trader arrivals leads to a difference in means to the order of 10^{-3} being statistically significant.

	LOW_VOLATILITY: $\lambda_V = 0.125$ (Low Levels of Picking-off Risk)			HIGH_VOLATILITY: $\lambda_V = 0.625$ (High Levels of Picking-off Risk)		
	Single Market	Frag. Market	Difference	Single Market	Frag. Market	Difference
	1	2	2–1	1	2	2–1
Quoted spread: Local	1.50	2.56	1.06	5.13	4.93	−0.20
Quoted spread: Inside	1.50	1.86	0.36	5.13	3.66	−1.47
Effective spread: Local	1.34	1.89	0.55	3.58	3.53	−0.05
Effective Spread: Inside	1.34	1.57	0.23	3.58	2.89	−0.70

where x_t is +1 for a buyer-initiated order and −1 for a seller-initiated order, p_t is the transaction price, and m_t is the mid-quote immediately before the transaction. Table 2 reports the results.

As expected, there is a direct relationship between asset volatility and bid–ask spreads (as in Stoll (2000). Quoted bid–ask spreads are higher by a factor of approximately 3 (2) in a consolidated (fragmented) market combined with the high-volatility setting. This relationship also holds for effective spreads, though the differences between the consolidated and fragmented markets are smaller in magnitude.

We also observe that, in the low-volatility setting, local and inside quoted bid–ask spreads are higher in fragmented markets by 1.06 and 0.36 ticks, respectively. However, when asset volatility is high, local and inside quoted bid–ask spreads are lower by 0.2 and 1.47 ticks, respectively, in fragmented markets. The same relationship holds for effective spreads. In fragmented markets combined with low (high) picking-off risk, the local and inside effective spreads are 1.89 and 1.57 (3.53 and 2.89) ticks, respectively, whereas the corresponding numbers in consolidated markets are lower (higher) at 1.34 (3.58) ticks.

In our model, fragmentation reduces price competition because agents can gain priority by jumping the queue (with probability 1/2, as liquidity takers are indifferent between taking liquidity on either of the two order books) and entering orders in the second order book, thereby inducing an increase in the bid–ask spread. At the same time, fragmentation reduces the picking-off risk, since the probability that an order will be picked off is lower in the presence of a second limit order book, resulting in a lower bid–ask spread. As described in Section III.B, the level of asset volatility directly influences the relative importance of the competition for time priority and adverse selection channels. The former is more important when asset volatility is low, whereas the latter is more important when asset volatility is high. As a result, higher (lower) asset volatility is associated with lower (higher) bid–ask spreads in fragmented markets. In this sense, our model can potentially explain the mixed findings in the empirical literature regarding the relationship between liquidity and fragmentation. For example, O'Hara and Ye (2011) find that fragmentation is associated with lower effective spreads for small-cap (or high-volatility) stocks in the US. our model. In Section IV, we further test this model implication by examining the relationship between fragmentation and liquidity for different levels of volatility for European stocks.

D. Welfare Analysis

We next analyze the economic benefits per agent and for the whole market by examining the effect of fragmentation on welfare in the high- and low-volatility scenarios. Welfare is measured as the average realized payoff per agent. In addition, we decompose the realized payoffs into gains and losses associated with agents' private values and the trading process.

Suppose that an agent with private value α enters the market at time t. She submits an order (a limit order or a market order) to either of the books at price $\tilde{p}$, with order direction $\tilde{x}$ (to buy or sell). Suppose further that the agent does not modify the order and that it is finally executed at time t' when the fundamental value is $v_{t'}$. The agent's realized payoff is then given by

$$\Pi = e^{-\rho(t'-t)}(\alpha + v_{t'} - \tilde{p})\tilde{x}. \tag{2}$$

We can decompose the agents' payoffs and rewrite equation (2) as

$$\Pi = \text{GAINS\_FROM\_PRIVATE\_VALUE} + \text{WAITING\_COST} + \text{MONEY\_TRANSFER}, \tag{3}$$

where

$$\text{GAINS\_FROM\_PRIVATE\_VALUE} = \alpha\tilde{x}$$
$$\text{WAITING\_COST} = (e^{-\rho(t'-t)} - 1)\alpha\tilde{x}$$
$$\text{MONEY\_TRANSFER} = e^{-\rho(t'-t)}(v_{t'} - \tilde{p})\tilde{x}.$$

The first term in equation (3), GAINS_FROM_PRIVATE_VALUE, represents the gains obtained directly from the intrinsic reasons for trading the asset, $\alpha\tilde{x}$. The second term, WAITING_COST, reflects the cost associated with delaying the realization of the GAINS_FROM_PRIVATE_VALUE. Agents submitting limit orders do not trade immediately after arriving, but have to wait until the orders are executed. This is costly due to the delay cost ρ. The third term in equation (3), MONEY_TRANSFER, reflects the difference between the fundamental value $v_{t'}$ and the transaction price $\tilde{p}$, discounted back to the arrival time of the agent, which represents the welfare gain (or loss) associated with the trading process. In general, MONEY_TRANSFER is related to the immediacy cost incurred when an agent wants to realize her private value immediately. For example, an agent who submits a market order realizes her intrinsic private value without a delay, but she may have to pay a price for demanding immediacy, which would be reflected in a negative value of MONEY_TRANSFER. Alternatively, an agent submitting a limit order earns

TABLE 3
Decomposition of Welfare by Agent Type

Table 3 reports the welfare, defined as the average realized payoff. In addition, it presents the waiting cost and money transfer as defined in equation (3). These measures are reported for each agent type $|\alpha| = \{0,4,8\}$. We report all statistics for a single and a fragmented market under high ($\lambda_v = 0.125$) and low ($\lambda_v = 0.625$) levels of volatility. The last two columns report the average aggregate welfare per period and the deadweight loss. All measures are reported in ticks. We omit standard errors for the differences because a large number of trader arrivals leads to a difference in means to the order of 10^{-3} being statistically significant.

	Average Welfare per Trader				Waiting Cost per Trader				Money Transfer per Trader					
	Private Value $\|\alpha\|$				Private Value $\|\alpha\|$				Private Value $\|\alpha\|$					
	0	4	8	Total	0	4	8	Total	0	4	8	Total	Welfare Per Period	Deadweight Loss
Panel A. Low Levels of Picking-Off Risk														
Single market	0.543	3.510	7.265	3.745	0.000	−0.350	−0.162	−0.189	0.543	−0.140	−0.572	−0.065	3.745	0.255
Frag. market	0.626	3.479	7.202	3.740	0.000	−0.355	−0.172	−0.193	0.626	−0.166	−0.626	−0.066	3.740	0.260
Panel B. High Levels of Picking-Off Risk														
Single market	0.606	3.398	7.039	3.652	0.000	−0.367	−0.270	−0.228	0.606	−0.235	−0.691	−0.119	3.652	0.348
Frag. market	0.817	3.389	6.871	3.662	0.000	−0.417	−0.284	−0.252	0.817	−0.192	−0.845	−0.085	3.662	0.338

half of the bid–ask spread but may generate MONEY_TRANSFER losses if her order is picked off.

Table 3 presents the results. When volatility is high, aggregate welfare is higher in a fragmented market (3.662 vs. 3.652 ticks), while, when volatility is low, aggregate welfare is lower in a fragmented market (3.740 vs. 3.745 ticks). This result complements our previous finding: When the adverse selection effect dominates the competition for time priority effect, fragmented markets are beneficial in terms of liquidity and aggregate welfare is higher. Conversely, when the competition for time priority effect is more relevant, liquidity is reduced in a fragmented market and aggregate welfare is also lower. Another way to look at aggregate welfare is to consider the deadweight loss. Per period, on average there arrives one trader with a private value that, on average, is of absolute value 4. Thus, the deadweight loss, shown in the final column, is 4 minus the observed average welfare per period.

In addition, we find that the welfare shifts among the three categories of agents are substantial. Agents with no intrinsic motives for trading (i.e., $\alpha = 0$) generate larger payoffs in fragmented markets in the low-volatility setting (0.54 vs. 0.63 ticks). Their payoffs are even larger in fragmented markets combined with the high-volatility setting (0.61 vs. 0.82 ticks). The difference between the single and fragmented markets arises for three reasons. First, there is less limit order price competition in a fragmented market because time priority does not apply across order books; second, the risk of being picked off is reduced in a fragmented market; third, intermediaries' expected time to execution in a fragmented market is lower than in a consolidated market, which reduces the delay-cost component of their welfare.

Fragmentation has detrimental effects in terms of expected payoffs for agents with $|\alpha| = 4$ and $|\alpha| = 8$, for both levels of picking-off risk. In terms of the absolute waiting cost, they do not exhibit any significant difference between the single and fragmented settings. However, the negative welfare effects are generated due to

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

higher money transfer losses. The underlying reasons for these money transfer losses depend on the composition of limit versus market orders used by these agents in the two volatility scenarios. In the low-volatility setting, the higher fraction of market orders used by agents with $|\alpha| = 8$ combined with lower price competition leads to higher immediacy costs in fragmented markets. The same mechanism also works in the high-volatility setting for $|\alpha| = 8$ traders, but the higher bid–ask spreads due to higher picking-off risk lead to substantially higher money transfer losses. In the single-market setting combined with high asset volatility, these same agents predominantly use limit orders and earn the bid–ask spread, but are adversely selected by $|\alpha| = 0$ traders, leading to money transfer losses. Nevertheless, they are still better off than they would be in the multi-market setting with high asset volatility.[13]

In conclusion, fragmentation induces a shift in trading gains away from agents with intrinsic motives to trade and toward intermediaries with little effect on aggregate welfare. This indicates that costly investments in intermediation capacities such as colocation charges, subscriptions to data feeds, and investments in high-speed networks are socially wasteful. These investments are outside the scope of our model but in reality, are endogenously determined by investors. In a computationally simpler alternative, we double the population of agents with $\alpha = 0$ in a multi-market setting combined with low asset volatility to mimic the increased participation by intermediaries resulting from their high profitability in fragmented markets. These results further confirm that intermediaries' entry reduces the welfare of agents with private reasons to trade. The trading gains of intermediaries as a group are higher, even though, individually, they extract lower welfare.[14,15]

E. Alternative Parameterizations

In this section, we consider three variations of our model to further elaborate on which features of the parameterization are crucial for determining our main results. Table 4 reports the results for spreads, which will be our main focus.[16]

First, we solve the model when the reentry rate for agents with zero private values is tripled. These agents resemble HFT firms so this analysis allows us to examine the effect of the zero-private-value traders being faster than others. Panel A of Table 4 shows the results. Unsurprisingly, this change in the reentry rate of zero alpha traders leads to a change in equilibrium spreads. When market volatility is low (high), the single market spreads are marginally lower (higher) in this setting

[13]Note that, in Table 3, the total money transfers do not add up to 0, as the expected payoff in a single transaction of the limit order and corresponding market order is discounted back to different times. This is due to the asynchronous arrivals of the agents who submit these two orders. However, the instantaneous money transfer, before it is discounted back, is equal to 0.

[14]The welfare decomposition results under this parameterization are reported in Table C4 in the Supplementary Material.

[15]In Appendix II of the Supplementary Material, where we estimate the model parameters based on data from 2015 for two FTSE-100 stocks trading on the LSE, consistent with this alternative parameter choice, the estimated population of agents with $\alpha = 0$ is higher compared to our baseline parameterization.

[16]Results for the other variables of interest are available in Appendix III of the Supplementary Material.

TABLE 4

Quoted and Effective Spreads: Alternative Parametrizations

Table 4 reports the bid–ask spread for changes to our baseline parametrizations. We consider 4 sets of alternative parametrizations: In Panel A, we triple the reentry rate for agents with private value $\alpha = 0$. In Panel B, we relax the agents' capacity constraint by allowing an agent to be followed by a second identical agent with a probability of 5%. Finally, in Panel C, we keep 10% of agents as captives in one of the books (book 1). Similar to Table 2, we generate simulated data sets for the following two cases: i) an environment with low picking-off risk, that is, low volatility, in which we set $\lambda_v = 0.125$, and ii) an environment with high picking-off risk, that is, high volatility, in which we set $\lambda_v = 0.625$. We compute local and inside versions of spread measures. The former is based on the local quotes, that is, the bid and ask prices of a local market, whereas the latter is based on the inside quotes, that is, the highest bid and the lowest ask across the two limit order books. In the single-market setting, the local and inside quotes are identical. Quoted spread is the difference between the best bid and best ask prices. Effective spread is defined in equation (1). We report all statistics for a single and a fragmented market when volatility is low and when volatility is high. We omit standard errors for the differences because a large number of trader arrivals leads to a difference in means to the order of 10^{-3} being statistically significant.

	LOW_VOLATILITY: $\lambda_v = 0.125$ (Low Levels of Picking-off Risk)			HIGH_VOLATILITY: $\lambda_v = 0.625$ (High Levels of Picking-off Risk)		
	Single Market	Frag. Market	Difference	Single Market	Frag. Market	Difference
	1	2	2 − 1	1	2	2 − 1
Panel A. Change in the Reentry Rate						
Quoted spread: Local	1.46	3.83	2.37	5.30	4.57	−0.73
Quoted spread: Inside	1.46	2.82	1.36	5.30	3.43	−1.87
Effective spread: Local	1.31	2.68	1.37	3.48	3.12	−0.35
Effective spread: Inside	1.31	2.15	0.83	3.48	2.53	−0.95
Panel B. Capacity Constraints						
Quoted spread: Local	1.53	2.81	1.28	5.11	4.86	−0.26
Quoted spread: Inside	1.53	2.04	0.51	5.11	3.61	−1.50
Effective spread: Local	1.36	2.06	0.70	3.60	3.49	−0.11
Effective spread: Inside	1.36	1.70	0.33	3.60	2.86	−0.74
Panel C. Clientele Effect						
Quoted spread: Local1	1.50	2.43	0.78	5.13	4.73	−0.40
Quoted spread: Local2	1.50	3.33	2.60	5.13	4.98	−0.33
Quoted spread: Inside	1.50	2.05	1.62	5.13	3.65	−1.48
Effective spread: Local1	1.34	1.91	0.57	3.58	3.42	−0.16
Effective spread: Local2	1.34	2.32	0.98	3.58	3.50	−0.08
Effective spread: Inside	1.34	1.69	0.35	3.58	2.85	−0.73

compared to the base-case scenario. The results are the opposite in fragmented markets. On the one hand, ceteris paribus, the faster reentry rate allows zero private value traders to reduce their picking-off risk. At the same time, it also allows them to snipe existing limit orders on the other side of the limit order book. The former effect dominates in the low-volatility single market setting and the high-volatility fragmented market setting. The latter effect dominates in the low-volatility fragmented market setting and the high-volatility single market setting. However, we still observe that in the low-volatility setting spreads in the single market are smaller than those in the fragmented market. The converse is true in the high-volatility setting.

Second, in our baseline specification, agents can submit an order for one share. We examine whether our findings change if we relax this constraint of agents' capacity. However, allowing flexible amounts of orders would make the state space increase exponentially, making the simulation computationally infeasible. To overcome this issue, we instead change our baseline setup as follows: When an agent enters the market, we allow another identical agent to enter the market at the same time. We further require the two agents to also reenter the market together. This allows us to mimic the effect of agents submitting an order for two shares instead of one. This relaxation of capacity constraints not only captures the effect of agents' submitting larger orders to a single order book, but also that of them splitting a large order across the two order books. We allow 5% of the agent population to have this ability to submit two-share orders. Panel B of Table 4 shows the results. In addition to the limit order at the front of the queue, the second limit order in the order book is exposed to picking-off risk. Additionally, in fragmented markets, equally priced limit orders' exposure to picking-off risk across the two order books also increases. These additional exposures result in higher bid–ask spreads. However, our main result remains qualitatively unchanged: single (fragmented) market spreads are smaller than the fragmented (single) market spreads in the low (high) volatility setting.

Finally, we examine clientele effects, that is, a setting where some agents are only able to trade on one exchange. While in the US and several other markets, all market participants can access all venues, this modification makes the specification resemble European equity markets, where some participants access only the primary exchange. Nevertheless, even in the US, clientele effects can also be observed when there is a payment for order flow (PFOF), in which a broker receives a compensation for routing trades to a particular market maker. In our modified specification, we treat 10% of the agents as captives in Book 1. Noncaptive agents can trade freely in both books. Panel C of Table 4 shows the results. Due to a loss of symmetry across the two order books, we report local spreads separately for Book 1 and Book 2. Unsurprisingly, Book 1 turns out to be more liquid than Book 2. Other results remain qualitatively unchanged.

IV. Empirical Application

In this section, we empirically examine the main implications of our theoretical findings. The testable predictions suggest that fragmentation has beneficial effects on the quoted and effective bid–ask spread when volatility is high, whereas the effect is the opposite when volatility is low. More generally, we predict that the increase in spread that is associated with an increase in volatility is smaller in a fragmented than in a consolidated market.

Our data comprise the constituents of the French and German large and mid-cap indices (CAC40, Next20, DAX, MDAX). After removing the smallest MDAX stocks to get a sample that is balanced across the two countries and eliminating stocks with incomplete data, we obtain a sample of 111 stocks. We employ high-frequency data from EUROFIDAI BEDOFIH for Euronext Paris, Deutsche Börse, and Chi-X. Euronext Paris and Deutsche Börse are the main listing exchanges for French and German firms, respectively. Chi-X established itself as the largest competitor to the incumbent exchanges subsequent to the implementation of the Markets in Financial Instruments Directive (MiFID) in 2007, which eliminated all rules prohibiting trading outside the national markets in the European Union.[17] Our data set contains all lit order messages, allowing us, also, to accurately sign and match trades to the prevailing aggregate order book.[18] Furthermore, compared to

[17]Chi-X is operated by Cboe today.

[18]We do not use data from other smaller lit trading venues or off-exchange trading, for which equally precise data are unavailable. However, the trading activity contained in our data comprises the large majority of lit trading for the stocks contained in our sample.

some other commonly used high-frequency databases, our data have the advantage of containing accurate exchange time-stamps, which is important when consolidating limit order books to obtain the best prices across multiple venues. We use data from the continuous trading session from 9:00 to 17:30 CET for each day. In other words, our data begin immediately after the opening auction clears, end immediately before the closing auction starts, and also exclude the period of the mid-day call auction for German stocks. Our sample period comprises the second half of 2012. To the best of our knowledge, none of the three exchanges implemented any changes to their market structures during this period.[19] Hence, our analysis can be interpreted as what happens when exchange competition remains unchanged, consistent with the fact that our model does not feature strategic competitive behavior by the individual trading venues.[20]

A. Empirical Approach

We compute the quoted bid–ask spread for each stock i on day t ($\text{Quoted}_{i,t}$) as the difference between the best ask and bid prices divided by their midpoint, time-weighted over day t. We compute this measure both for the consolidated order book, that is, using the best prices across both limit order books, and locally, that is, considering only the best bid and ask prices in one order book. We also compute the effective spread ($\text{EFFECTIVE}_{i,t}$) as twice the difference between the trade price and the prevailing midpoint times the trade indicator variable (1 for buyer-initiated, -1 for seller-initiated trades) divided by the midpoint, and express it in basis points. We then aggregate the observations by value-weighting across all trades executed during day t. Effective spreads are computed with respect to both the inside midpoint and the local midpoint.

To measure the level of fragmentation ($\text{FRAGMENTATION}_{i,t}$), we follow previous literature (e.g., Degryse, de Jong, and Kervel (2015), Haslag and Ringgenberg (2023)) and first compute the Herfindahl–Hirschman Index of traded volumes for every stock i and day t ($\text{HHI}_{i,t}$) across the two exchanges. As HHI measures the concentration of trading, we define $\text{FRAGMENTATION}_{i,t}$ as one minus $\text{HHI}_{i,t}$. Specifically, we compute the measure as follows:

$$\text{FRAGMENTATION}_{i,t} = 1 - \sum_{k} s^2_{i,k,t}, \tag{4}$$

where $s^2_{i,k,t}$ is the square of the share of the total trades of stock i on day t that take place in trading venue k.

[19] In Oct. 2012, a commercial microwave link between Frankfurt and London became operational. Our results remain unchanged after controlling for this fact. Sagade, Scharnowski, Theissen, and Westheide (2019) analyze the market quality implications of this event.

[20] The exchanges differ in their trading fees. Trading fees on European stock exchanges are defined as a proportion of a transaction's euro value. Chi-X offers a maker-taker model where liquidity takers pay 0.30 bps whereas liquidity makers are rebated 0.15 bps. The standard fees on the other exchanges do not differ between liquidity makers and takers, though there are special market-making programs that provide lower fees and rebates, such that the trading fees for the average executed orders are unclear. Importantly, there are no important changes to the fees during our sample period, and all stocks are equally affected by trading fees. Thus, we conduct our analyses based on gross-of-fee measures.

In our model, rather than the continuous measure of the level of fragmentation, we used two extreme cases: i) a single market, which is equivalent to a market environment that is not fragmented, that is, $\text{FRAGMENTATION}_{i,t} = 0$; and ii) a fully fragmented market with two identical limit order markets, where agents trade equally in both venues, and thus $\text{FRAGMENTATION}_{i,t} = 0.5\,(\text{i.e.}\, 1 - (0.5^2 + 0.5^2))$.[21]

We also estimate volatility as the standard deviation of 1-second mid-quote returns, denoted by $\sigma_{i,t}$, for each stock i and day t. To eliminate the effects of any other stock characteristics, we consider the variation of volatility over time within each stock rather than across stocks. We use the continuous fragmentation measure defined above and, in order to facilitate the interpretation of interaction effects, create five dummy variables corresponding to five volatility quintiles:

$$(5)\quad \text{VOLATILITY}(j-\text{th quintile})_{i,t} = \begin{cases} 1 & \text{if } \sigma_{i,t} \text{ is between the } 20(j-1)\% \text{ and } 20j\% \text{ percentiles} \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we include the natural logarithm of the stock's trading volume measured in euros across the two exchanges ($\text{VOLUME}_{i,t}$) as an additional explanatory variable, which proxies for variations in the arrival rate of traders.[22]

We estimate panel regressions of the two liquidity measures (LIQ), that is, the quoted bid–ask spread and the effective spread, on the level of fragmentation, the volatility quintile dummies, the interaction between fragmentation and the volatility dummies, and the trading volume, as follows:

$$(6)\quad \begin{aligned} \text{LIQ}_{i,t} = {} & \alpha + \beta\,\text{FRAGMENTATION}_{i,t} + \sum_{j=1, j\neq 3}^{j=5} \gamma_j \text{VOLATILITY}\,(j-\text{th quintile})_{i,t} \\ & + \sum_{j=1, j\neq 3}^{j=5} \delta_j \text{FRAGMENTATION}_{i,t} \times \text{VOLATILITY}\,(j-\text{th quintile})_{i,t} \\ & + \zeta \log(\text{VOLUME}_{i,t}) + v_i + \eta_t + \varepsilon_{i,t+1}, \end{aligned}$$

where v_i and η_t denote the stock and day-fixed effects, respectively. To avoid multicollinearity, we exclude the dummy variable for the middle quintile.

According to our theoretical results, we expect higher fragmentation to lead to a wider bid–ask spread when volatility is low, whereas when volatility is high, the spread is expected to be reduced. Therefore, this theoretical implication translates into testing that the interaction coefficient δ_1 is significantly higher than δ_5.

In an alternative specification, to make the results more easily comparable with those predicted by the model, each day, we sort stocks into two halves based on whether their fragmentation is lower or higher than the median and define the corresponding dummy variables, LOW_FRAGMENTATION and HIGH_FRAGMENTATION. We also create two dummy variables LOW_VOLATILITY

[21] Of course, it would be preferable to use an empirical setting in which fragmentation also only takes 2 extreme values. Unfortunately, such instances known to us are problematic because of small sample sizes, a change in competitive behavior of market operators coinciding with the change in fragmentation, other parallel changes to trading venues – such as dark pools – not captured in our model, or regulatory issues preventing some market participants from being active in all venues.

[22] The exclusion of trading volume from the regressions does not qualitatively affect the results.

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

and HIGH_VOLATILITY that indicate whether volatility is lower or higher than the median for a stock on each day.

We then estimate panel regressions where we regress the two liquidity measures (LIQ), that is, the quoted bid–ask spread and the effective spread, on the combinations of low and high fragmentation and volatility, respectively, and the trading volume, as follows:

$$\text{(7)} \quad \begin{aligned} \text{LIQ}_{i,t} = {} & \beta_1 \mathbf{1}_{\text{HIGH\_FRAGMENTATION} \times \text{LOW\_VOLATILITY}_{i,t}} \\ & + \beta_2 \mathbf{1}_{\text{LOW\_FRAGMENTATION} \times \text{HIGH\_VOLATILITY}_{i,t}} \\ & + \beta_3 \mathbf{1}_{\text{HIGH\_FRAGMENTATION} \times \text{HIGH\_VOLATILITY}_{i,t}} \\ & + \gamma \log(\text{volume}_{i,t}) + \nu_i + \eta_t + \varepsilon_{i,t+1}, \end{aligned}$$

where ν_i and η_t denote the stock and day-fixed effects, respectively. We define the combination of low volatility and low fragmentation to be the base category such that the other β_k coefficients are relative to this base category. We cluster standard errors by both stock and day.

We estimate the models both for local liquidity measures on the primary listing exchange (Euronext Paris or Deutsche Börse), which is empirically the larger of the two exchanges under consideration for each stock, and for inside liquidity measures, that is, considering the highest bid and lowest ask prices across the primary listing exchange and Chi-X.

B. Empirical Results

Table 5 shows the results of the regressions using volatility quintiles, and the continuous measure of fragmentation. Columns 1 and 2 show that, unconditionally, fragmentation does not significantly affect either the local bid–ask spread or the effective spread. We find that volatility increases the quoted and effective spreads, consistent with the theoretical prediction that the increased exposure to adverse selection in high-volatility states leads to wider spreads. In particular, the quoted bid–ask spreads monotonically increase with the volatility quintiles and the difference between the coefficients for the most and least volatile quintiles suggests that the bid–ask spread is higher by about 3.5 basis points for a stock in its most volatile quintile of trading days compared to its least volatile quintile of trading days. The impact on the effective spreads is similar, with a difference of about 4.2 basis points between the most and least volatile trading days. These effects are economically large compared to average quoted bid–ask spreads of 7.3 basis points and average effective bid–ask spreads of 6.2 basis points in our sample (untabulated).

Most importantly, we find that the effect of fragmentation on the spread varies with the level of asset volatility. Specifically, the coefficients of the interaction terms between fragmentation and the volatility dummies are positive (negative) for low (high) levels of volatility. The tests for the equality of the coefficients for the first and fifth quintiles, shown at the bottom of the table, show that the differences are statistically significant. Note that the results do not suggest that more fragmented stocks become more liquid as volatility rises, as the variation in fragmentation is not that large, but instead show that fragmented markets compare more favorably to consolidated markets when volatility is high than when it is low.

TABLE 5

Panel Regressions of Local Liquidity on Fragmentation and Volatility

Table 5 presents the results of panel regressions of relative quoted and effective spreads on the listing exchange for the sample of French and German stocks (columns 1 and 2) or across the order books (columns 3 and 4). FRAGMENTATION is 1 minus the Herfindahl–Hirschman Index of trading volumes on the listing exchange and Chi-X. VOLATILITY quintiles are computed within stock. log(VOLUME) is the natural logarithm of the euro trading volume. All regressions include stock- and day-fixed effects. We double-cluster standard errors by stock and day. *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. *t*-statistics are in parentheses.

	Local		Inside	
	Quoted	Effective	Quoted	Effective
	1	2	3	4
FRAGMENTATION	1.722 (1.05)	0.489 (0.36)	−1.983** (−2.13)	−4.363*** (−4.20)
VOLATILITY_1ST_QUINTILE	−1.516*** (−2.69)	−1.658*** (−3.01)	−1.560*** (−3.14)	−2.095*** (−3.57)
VOLATILITY_2ND_QUINTILE	−0.666** (−2.52)	−0.142 (−0.45)	−0.918** (−2.45)	−1.475*** (−2.65)
VOLATILITY_4TH_QUINTILE	0.627 (1.05)	1.183** (2.40)	0.145 (0.84)	−0.149 (−0.48)
VOLATILITY_5TH_QUINTILE	1.967** (2.33)	2.577*** (3.44)	1.207** (2.56)	1.342** (2.54)
VOLATILITY_1ST_QUINTILE × FRAGMENTATION	2.423* (1.84)	2.757** (2.20)	2.755** (2.42)	4.021*** (3.02)
VOLATILITY_2ND_QUINTILE × FRAGMENTATION	1.048* (1.80)	−0.153 (−0.22)	1.769** (2.02)	3.068** (2.40)
VOLATILITY_4TH_QUINTILE × FRAGMENTATION	−1.108 (−0.84)	−2.565** (−2.32)	−0.057 (−0.13)	0.514 (0.67)
VOLATILITY_5TH_QUINTILE × FRAGMENTATION	−3.701* (−1.96)	−5.211*** (−3.05)	−2.080* (−1.96)	−2.551** (−2.04)
log(VOLUME)	−1.453*** (−8.15)	−0.219 (−1.16)	−1.254*** (−8.32)	−0.231 (−1.39)
Constant	31.441*** (11.09)	9.737*** (2.88)	28.616*** (11.26)	11.051*** (3.71)
VOLA_5 × FRAG − VOLA_1 × FRAG	−5.466***	−7.838***	−4.836***	−6.571***
p-value	0.009	<0.001	0.010	0.001
No. of obs.	14,097	14,097	14,097	14,097

Columns 3 and 4 report the analogous results for the inside liquidity measures. Fragmentation here is unconditionally associated with smaller quoted and effective spreads. The results for the effect of volatility on inside liquidity are similar to those observed for local liquidity. Quoted (effective) spreads are about 2.8 (3.4) basis points larger when stocks are most versus least volatile, which is large compared to the average inside quoted (effective) spreads of 6.4 (3.7) basis points in our sample. The results for the interaction effect between volatility and fragmentation are similar to those for the local liquidity measures, confirming the predictions of our model. In each model, higher fragmentation, as compared to lower fragmentation, is associated with lower spreads on high-volatility days, whereas it is associated with higher spreads on low-volatility days, and the difference in these results between the high- and low-volatility days is statistically significant.

We next turn to the results for the regression model introduced in equation (7). The first two columns of Table 6 contain the results for the local liquidity measures and the next two those for the inside liquidity measures. Below the regression results, the table also shows the results of F-tests for the differences between coefficients.

TABLE 6

Panel Regressions of Liquidity on Fragmentation and Volatility

Table 6 presents the results of panel regressions of relative quoted and effective spreads on the listing exchange for the sample of French and German stocks (columns 1 and 2) or across the order books (columns 3 and 4). High and low volatility indicates volatility above and below the median for the respective stock. High and low fragmentation indicates fragmentation above and below the median of the cross-section of stocks on each day. log(VOLUME) is the natural logarithm of the euro trading volume. All regressions include stock- and day-fixed effects. We double-cluster standard errors by stock and day. *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively. *t*-statistics are in parentheses.

	Local		Inside	
	Quoted	Effective	Quoted	Effective
	1	2	3	4
LOW_VOLATILITY × LOW_FRAGMENTATION	0.000	0.000	0.000	0.000
LOW_VOLATILITY × HIGH_FRAGMENTATION	0.163**	0.072	−0.066	−0.112**
	(2.22)	(1.13)	(−1.10)	(−2.06)
HIGH_VOLATILITY × LOW_FRAGMENTATION	0.601***	0.583***	0.476***	0.494***
	(4.64)	(5.68)	(5.09)	(5.80)
HIGH_VOLATILITY × HIGH_FRAGMENTATION	0.547***	0.419***	0.243***	0.144**
	(5.10)	(4.88)	(3.00)	(2.19)
log(VOLUME)	−1.321***	−0.030	−1.105***	−0.038
	(−7.52)	(−0.15)	(−7.51)	(−0.21)
LOW_VOLATILITY × HIGH_FRAGMENTATION− LOW_VOLATILITY × LOW_FRAGMENTATION	0.163**	0.072	−0.066	−0.112**
p-value	0.029	0.263	0.272	0.042
HIGH_VOLATILITY × HIGH_FRAGMENTATION− HIGH_VOLATILITY × LOW_FRAGMENTATION	−0.054	−0.164**	−0.233***	−0.349***
p-value	0.504	0.021	0.001	<0.001
(HIGH_VOLATILITY × HIGH_FRAGMENTATION − HIGH_VOLATILITY × LOW_FRAGMENTATION)− (LOW_VOLATILITY × HIGH_FRAGMENTATION − LOW_VOLATILITY × LOW_FRAGMENTATION)	−0.217*	−0.237**	−0.167*	−0.238***
p-value	0.068	0.013	0.067	0.004
No. of obs.	14,097	14,097	14,097	14,097

For the local liquidity measures, our empirical results reveal that, if volatility is low, the quoted spreads are statistically significantly higher when fragmentation is high, which is consistent with our theoretical prediction. For the effective spreads, while the coefficient has the same sign, it is not statistically significant. Conversely, as predicted by our model, when volatility is high, spreads are lower in fragmented markets: quoted spreads become indistinguishably different, and effective spreads are significantly smaller when fragmentation is high. Consistent with the previous results, our last test shows that the difference in liquidity between the cases of high and low fragmentation is reduced when volatility is high, that is, liquidity in fragmented markets is comparably better than in less fragmented markets when volatility is high.

The results for the inside liquidity measures are similar to those for the local ones, with the exception that fragmented markets perform comparatively better when volatility is low. The latter may be explained by the fact that, even if exchanges do not actively change their degree of competition during our sample period, higher fragmentation increases the competitiveness of the smaller exchange and thus the overall degree of competition, which reduces trading costs (Rust and Hall (2003), O'Hara and Ye (2011), and Haslag and Ringgenberg (2023)).

These findings provide empirical support for our main theoretical implications concerning the impact of fragmentation on liquidity. The effect of fragmentation is heterogeneous and greatly depends on the level of picking-off risk resulting from the volatility of the stock's value. The effect of high picking-off risk can exceed that of price competition, and agents respond by submitting less aggressive orders. As fragmented markets provide greater protection against adverse selection, agents are able to submit more aggressive orders in a fragmented market than in a single market, leading to a reduction in the bid–ask spread. On the other hand, if we only considered the effect of traders' competition, fragmentation would lead to an increase in the bid–ask spread since more fragmented markets lead to reduced price competition. Hence, when the level of picking-off risk is low, the competition for time priority effect becomes stronger, leading to an increase in the bid–ask spread in fragmented markets.

V. Conclusion

We model a fragmented market for an asset that trades in two limit order books populated by heterogeneous agents who endogenously choose to supply or consume liquidity, and compare the results with a single-market setting. Two channels – competition for time priority and adverse selection – drive the results. As time priority is not enforced across markets, fragmentation reduces price competition between intermediaries. At the same time, fragmentation provides better protection from adverse selection risk for limit orders, as incoming arbitrageurs who can only trade one unit may trade against mispriced orders in either of the two order books.

The former effect dominates when asset volatility is low, whereas the latter effect dominates when asset volatility is high. Hence, market fragmentation can be beneficial or harmful for market liquidity depending on the level of asset volatility. These heterogeneous liquidity effects associated with market fragmentation also affect total welfare in our model: A consolidated (fragmented) market is associated with higher total welfare when asset volatility is low (high). However, in both volatility settings, the distribution of welfare across the heterogeneous agent types differs markedly across consolidated and fragmented markets. Agents with intrinsic trading motives extract lower payoffs in fragmented markets whereas agents acting as intermediaries are better off in fragmented markets.

Our central theoretical prediction concerning the relationship between market fragmentation and liquidity can potentially reconcile the differences in the empirical literature. We also provide empirical support for this prediction by investigating the relationship between liquidity and fragmentation for French and German stocks in the second half of 2012.

Overall, our results suggest that the positive effects of consolidating order flow in a single (or fewer) location(s) still exist 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 borne by investors who trade for intrinsic trading motives.

Supplementary Material

To view supplementary material for this article, please visit http://doi.org/10.1017/S0022109022001521.

https://doi.org/10.1017/S0022109022001521 Published online by Cambridge University Press

References

Arnold, T.; P. Hersch; J. H. Mulherin; and J. Netter. "Merging Markets." *Journal of Finance*, 54 (1999), 1083–1107.

Baldauf, M., and J. Mollner. "Trading in Fragmented Markets." *Journal of Financial and Quantitative Analysis*, 56 (2021), 93–121.

Bennett, P., and L. Wei. "Market Structure, Fragmentation, and Market Quality." *Journal of Financial Markets*, 9 (2006), 49–78.

Bessembinder, H., and H. M. Kaufman. "A Cross-Exchange Comparison of Execution Costs and Information Flow for NYSE-Listed Stocks." *Journal of Financial Economics*, 46 (1997), 293–319.

Boehmer, B., and E. Boehmer. "Trading Your Neighbor's ETFs: Competition or Fragmentation?" *Journal of Banking and Finance*, 27 (2003), 1667–1703.

Boneva, L.; O. Linton; and M. Vogt. "The Effect of Fragmentation in Trading on Market Quality in the UK Equity Market." *Journal of Applied Econometrics*, 31 (2016), 192–213.

Brogaard, J.; T. Hendershott; and R. Riordan. "High-Frequency Trading and Price Discovery." *Review of Financial Studies*, 27 (2014), 2267–2306.

Cespa, G., and T. Foucault. "Sale of Price Information by Exchanges: Does It Promote Price Discovery?" *Management Science*, 60 (2013), 148–165.

Chao, Y.; C. Yao; and M. Ye. "Why Discrete Price Fragments U.S. Stock Exchanges and Disperses Their Fee Structures." *Review of Financial Studies*, 32 (2019), 1068–1101.

Chlistalla, M., and M. Lutat. "Competition in Securities Markets: The Impact on Liquidity." *Financial Markets and Portfolio Management*, 25 (2011), 149–172.

Chowdhry, B., and V. Nanda. "Multimarket Trading and Market Liquidity." *Review of Financial Studies*, 4 (1991), 483–511.

Clark, P. K. "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices." *Econometrica*, 41 (1973), 135–155.

Degryse, H.; F. de Jong; and V. Kervel. "The Impact of Dark Trading and Visible Fragmentation on Market Quality." *Review of Finance*, 19 (2015), 1587–1622.

Easley, D.; M. O'Hara; and L. Yang. "Differential Access to Price Information in Financial Markets." *Journal of Financial and Quantitative Analysis*, 51 (2016), 1071–1110.

Fink, J.; K. E. Fink; and J. P. Weston. "Competition on the Nasdaq and the Growth of Electronic Communication Networks." *Journal of Banking & Finance*, 30 (2006), 2537–2559.

Foucault, T., and A. J. Menkveld. "Competition for Order Flow and Smart Order Routing Systems." *Journal of Finance*, 63 (2008), 119–158.

Glosten, L. R. "Competition, Design of Exchanges and Welfare." Working Paper, Columbia University (1998).

Goettler, R. L.; C. A. Parlour; and U. Rajan. "Equilibrium in a Dynamic Limit Order Market." *Journal of Finance*, 60 (2005), 2149–2192.

Goettler, R. L.; C. A. Parlour; and U. Rajan. "Informed Traders and Limit Order Markets." *Journal of Financial Economics*, 93 (2009), 67–87.

Harris, L. E. "Consolidation, Fragmentation, Segmentation and Regulation." *Financial Markets, Institutions & Instruments*, 2 (1993), 1–28.

Hasbrouck, J. "Trading Fast and Slow: Security Market Events in Real Time." NYU Working Paper No. FIN-99-012 (1999).

Haslag, P. H., and M. Ringgenberg. "The Demise of the NYSE and NASDAQ: Market Quality in the Age of Market Fragmentation." *Journal of Financial and Quantitative Analysis*, forthcoming (2023).

He, P. W.; E. Jarnecic; and Y. Liu. "The Determinants of Alternative Trading Venue Market Share: Global Evidence from the Introduction of Chi-X." *Journal of Financial Markets*, 22 (2015), 27–49.

Hendershott, T., and C. M. Jones. "Island Goes Dark: Transparency, Fragmentation, and Regulation." *Review of Financial Studies*, 18 (2005), 743–793.

Hengelbrock, J., and E. Theissen. "Fourteen at One Blow: The Market Entry of Turquoise." Working Paper, University of Mannheim (2009).

Hollifield, B.; R. A. Miller; P. Sandås; and J. Slive. "Estimating the Gains from Trade in Limit-Order Markets." *Journal of Finance*, 61 (2006), 2753–2804.

Jones, C. M.; G. Kaul; and M. L. Lipson. "Transactions, Volume, and Volatility." *Review of Financial Studies*, 7 (1994), 631–651.

Kyle, A. S., and A. A. Obizhaeva. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, 84 (2016), 1345–1404.

Madhavan, A. "Consolidation, Fragmentation, and the Disclosure of Trading Information." *Review of Financial Studies*, 8 (1995), 579–603.

Mendelson, H. "Consolidation, Fragmentation, and Market Performance." *Journal of Financial and Quantitative Analysis*, 22 (1987), 189–207.

Menkveld, A. J. “High Frequency Trading and the New Market Makers.” *Journal of Financial Markets*, 16 (2013), 712–740.

Nguyen, V.; B. F. Van Ness; and R. A. Van Ness. “Short- and Long-Term Effects of Multimarket Trading.” *Financial Review*, 42 (2007), 349–372.

Nielsson, U. “Stock Exchange Merger and Liquidity: The Case of Euronext.” *Journal of Financial Markets*, 12 (2009), 229–267.

O’Hara, M., and M. Ye. “Is Market Fragmentation Harming Market Quality?” *Journal of Financial Economics*, 100 (2011), 459–474.

Pagano, M. “Trading Volume and Asset Liquidity.” *Quarterly Journal of Economics*, 104 (1989), 255–274.

Pagnotta, E. S., and T. Philippon. “Competing on Speed.” *Econometrica*, 86 (2018), 1067–1115.

Pakes, A., and P. McGuire. “Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the ‘Curse’ of Dimensionality.” *Econometrica*, 69 (2001), 1261–1281.

Parlour, C. A., and D. J. Seppi. “Liquidity-Based Competition for Order Flow.” *Review of Financial Studies*, 16 (2003), 301–343.

Rust, J., and G. Hall. “Middlemen Versus Market Makers: A Theory of Competitive Exchange.” *Journal of Political Economy*, 111 (2003), 353–403.

Sagade, S.; S. Scharnowski; E. Theissen; and C. Westheide. “A Tale of Two Cities - Inter-Market Latency, Market Integration, and Market Quality.” Working Paper, University of Mannheim (2019).

Stoll, H. R. “Friction.” *Journal of Finance*, 55 (2000), 1479–1514.