Internet Appendix for “Trader Competition in Fragmented Markets: Liquidity Supply versus Picking-Off Risk ”

Alejandro Bernales Nicolás Garrido Satchit Sagade
Marcela Valenzuela Christian Westheide*

I. Model Solution

In this appendix, we describe in detail the agents' dynamic maximization problem in each state of the economy, the model equilibrium, the solution approach, and the implementation of the Pakes and McGuire (2001) algorithm.

A Agents' Dynamic Maximization

There is a set of states $s \in \{1, 2, \dots, S\}$ that describes the market conditions in the economy. These market conditions are observed by each agent before she makes any decision. The state s that an agent observes is described by the contemporaneous limit order books, L_1 and L_2, the agent's private value α and, in the case that the agent has previously submitted a limit order to any of the books, the status of that order in L_1 or L_2, i.e., its original submission price, its queue priority in the book, and its type (i.e., buy or sell). The fundamental value of the asset, v, is implicitly one of the variables that describe state s, since agents interpret limit order prices relative to the fundamental value. For convenience, we set the arrival time of an agent to zero in the following discussion.

Let $a \in \Theta(s)$ be the agent's potential trading decision, where $\Theta(s)$ is the set of all possible decisions that an agent can take in state s. Suppose that the optimal decision given state s is $\tilde{a} \in \Theta(s)$. Let $\eta(h|\tilde{a},\ s)$ be the probability density that the optimally submitted order is executed at time h. Let $\gamma(v|h)$ be the density function of v at time h, which is exogenous and characterized by the Poisson process of the fundamental value of the asset with rate λ_v. Thus, the expected value of the optimal order submission $\tilde{a} \in \Theta(s)$, if the order is executed prior to the agent's reentry time h_r, is

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

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

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

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

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

B Equilibrium

In equilibrium, each agent behaves optimally by maximizing her expected utility, based on the observed state that describes the market conditions (as in equation (2)). Thus, optimal decisions are state-dependent. They are also Markovian, because the state observed by an agent is a consequence of the previous states and the historical optimal decisions taken in the trading game. As there is competition between agents, the equilibrium is competitive (although there is no competition between venues). A competitive equilibrium means that competing agents will respond to an agent's local deviation in a way that leads to a reduction in the deviating agent's expected utility. We obtain a stationary and symmetric equilibrium, as in Doraszelski and Pakes (2007). In such an equilibrium, optimal decisions are time-independent, i.e., they stay the same when an agent faces the same state in the present or in the future.

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

As mentioned earlier, the state s is defined by the four-tuple ($L_{1,t}$, $L_{2,t}$, α, *status of previous limit order*), where all variables that describe the state are discrete.

Moreover, each agent's potential decision a is taken from $\Theta(s)$, which is the set of all possible decisions that can be taken in state s. This set of possible decisions is discrete and finite, given the features of the model. Consequently, the state space is countable and the decision space is finite; thus, the trading game has a Markov perfect equilibrium (see Rieder, 1979).

C Solution Approach

Given the large dimension of the state space, we use the Pakes and McGuire (2001) algorithm to compute a stationary and symmetric Markov perfect equilibrium. The intuition behind this algorithm is that the trading game by itself can be used as a tool through which agents learn how to behave in each state. Thus, we set the initial beliefs about the expected payoffs of potential decisions in each state. Agents take the trading decision that provides the highest expected payoff conditional on the state they observe. Subsequently, agents dynamically update their beliefs by playing the game and observing the realized payoffs of their trading decisions. In this sense, the algorithm is based on agents following a learning-by-doing mechanism. The Pakes and McGuire (2001) algorithm is able to deal with a large state space because it reaches the equilibrium only in the recurring states class.

The equilibrium is reached when there is nothing left to learn, i.e., when the beliefs about the expected payoffs have converged. We apply the same procedure as was used by Goettler, Parlour, and Rajan (2009) to determine whether the equilibrium has been reached. Once we reach the equilibrium, after making the agents play the game for at least 10 billion trading events, we fix the agents' beliefs and simulate a further 20 million events. All theoretical results presented in Section **??** are computed from the latter.

D Pakes and McGuire Algorithm

Here, we explain the implementation details associated with the Pakes and McGuire (2001) algorithm, along with some additional steps taken to reduce the state space of our model. In order to reduce the dimensionality of the state space, we center each limit order book at the contemporaneous fundamental value of the asset, i.e., by setting $p_m^0 = v_t$. Suppose, at time $t = 0$, the fundamental value is v_0, but after a period τ the fundamental value has experienced some innovations and is now v_τ, with $v_\tau - v_0 = qd$, where q is a positive or negative integer. In this case, we shift both books by q ticks to center them at the new level of the fundamental value v_τ. Thus, we move the queues of existing limit orders in both books to take the relative difference with respect to the new fundamental value into account. This implies that the prices of all orders are always relative to the current fundamental value of the asset and agents always make decisions in terms of prices relative to the fundamental value.

1 Updating Process to Reach the Equilibrium

For any state s of the economy, there is a set of possible actions, $\Theta(s)$, that a trader can take. Suppose that a given trader arrives for the first time or reenters the market at time t and observes the state s. In our model setup, the trader has beliefs about the expected payoff of each possible action that could be taken given the observed state s. Suppose that $U_t(\tilde{a}|s)$ is the expected payoff at time t that is associated with the action $\tilde{a} \in \Theta(s)$. Suppose that the trader decides at time t to take the optimal action $\tilde{a}^*$ that provides the maximum expected payoff out of all possible actions. As a first case, suppose that the optimal action $\tilde{a}^*$ is not a market order (e.g. a limit order, or a cancellation and resubmission). Later on, at time t_r, the same trader reenters the market, but the market conditions have changed. The trader observes a new state s_{t_r} in which she follows the optimal strategy $\tilde{a}^{**}$ that also gives a maximum payoff given the new market conditions. Consequently, the original decision $\tilde{a}^*$ induces a realized continuation of optimal actions and expected payoffs; and thus the updating process of beliefs can be written as:

$$U_{t_r}(\tilde{a}^*|s) = \frac{n_{\tilde{a}^*,s}}{n_{\tilde{a}^*,s}+1}U_t(\tilde{a}^*|s) + \frac{1}{n_{\tilde{a}^*,s}+1}e^{-\rho(t_r-t)}(U_{t_r}\tilde{a}^{**}|s_{t_r}),$$

where $n_{\tilde{a}^*,s}$ is a counter that increases by one when the action $\tilde{a}^*$ is taken in the state s.[1]

Alternatively, as a second case, suppose that the optimal decision $\tilde{a}^*$ is a market order (i.e. there is no future time t_r as in the previous case). Then, the updating process of the expected payoff of the optimal action $\tilde{a}^*$ in this scenario can be expressed as:

[1]The value of $n_{\tilde{a}^*,s}$ affects how quickly we reach the model equilibrium (a large value in $n_{\tilde{a}^*,s}$ is associated with a slow convergence). Therefore, we reset $n_{\tilde{a}^*,s}$ intermittently to improve the convergence speed.

$$U_{t_r}(\tilde{a}^*|s) = \frac{n_{\tilde{a}^*,s}}{n_{\tilde{a}^*,s}+1} U_t(\tilde{a}^*|s) + \frac{1}{n_{\tilde{a}^*,s}+1}(\alpha + v_t - \tilde{p})\tilde{x}.$$

Here $\tilde{p}$ is the submission price, α is the private value of the trader, v_t is the fundamental value of the asset, and $\tilde{x}$ is equal to one (minus one) when the trader submits a buy (sell) order.

As a third case, suppose that the optimal decision $\tilde{a}^*$ is a limit order; however, later on at time t_r this limit order is executed because another trader submits a market order. The updating process for the first trader with the optimal action $\tilde{a}^*$ is reflected in the following equation:

$$U_{t_r}(\tilde{a}^*|s) = \frac{n_{\tilde{a}^*,s}}{n_{\tilde{a}^*,s}+1} U_t(\tilde{a}^*|s) + \frac{1}{n_{\tilde{a}^*,s}+1} e^{-\rho(t_r - t)}(\alpha + v_{t_r} - \tilde{p})\tilde{x},$$

where α is the private value of the first trader. Similarly, for the second trader, who submits the market order that executes the limit order of the first trader, the updating process can be expressed as:

$$U_{t_r}(\tilde{a}'|s_{t_r}) = \frac{n_{\tilde{a}',s_{t_r}}}{n_{\tilde{a}',s_{t_r}}+1} U_t(\tilde{a}'|s_{t_r}) + \frac{1}{n_{\tilde{a}',s_{t_r}}+1}(\alpha' + v_{t_r} - \tilde{p})(-\tilde{x}),$$

where α' and $\tilde{a}'$ are the private value and the optimal decision of the second trader, respectively. In this case, $\tilde{a}'$ is a market order which is chosen at time t_r by the second trader when the state s_{t_r} is found. In this last case, it is important to observe that any market order implies the execution of a previously submitted limit order. Thus, in the presence of market orders the updating process in beliefs always involves two traders: the trader who submits the market order, and the trader who submitted the limit order which is executed by the market order.[2]

2 Convergence Criteria

We use the same convergence criteria as Goettler et al. (2009); thus, further details can be found in their study. We check for convergence after running the trading game for at least 10 billion trading events. Subsequently, we check the evolution of agents' beliefs for convergence after every 500 million simulations. Let us assume that the first group of 500 million simulations after we start checking for convergence finishes at time t_1 and the second group of 500 million simulations finishes at time t_2. Let $U_{t_1}(\tilde{a}|s)$ and $U_{t_2}(\tilde{a}|s)$ be the expected payoffs that are associated with the action $\tilde{a}$ when the state s is present at times t_1 and t_2, respectively. In addition, suppose that $k_{\tilde{a},s}^{t_1,t_1}$ is the number of times that the action $\tilde{a}$ was taken between t_1 and t_2 when traders face s. We evaluate the change in the expected value of the expression $|U_{t_2}(\tilde{a}|s) - U_{t_1}(\tilde{a}|s)|$ for all pairs $(\tilde{a}, s)$ weighted by $k_{\tilde{a},s}^{t_1,t_1}$ after every 500 million simulations. Once this weighted absolute difference is smaller than 0.01 (suggesting that the model has converged), we apply two further convergence criteria in line with Pakes and McGuire (2001) and Goettler et al. (2009).

[2] The initial beliefs about the expected payoffs $U_0(\tilde{a}|s)$ of the possible actions $\tilde{a} \in \Theta(s)$ that a trader can take given that she faces state s are set as follows. Suppose one of the possible actions for a trader with private value α in the state s is to submit a limit sell order at price p when the fundamental value is v. We set the initial expected payoff of this action as $p - v - \alpha$ discounted by ρ until the expected time at which a new fast trader will arrive in the market. This value is only a first approximation since we assume that v is constant, which is not true in the model, and there is a chance that the next trader may submit another limit order instead of a market order that executes the limit order of the previous trader. In the case of a market sell order, the expected payoff is simply $p - v - \alpha$ without any discount. Similar values are obtained for buy orders.

After reaching a small weighted absolute difference in the change in the expected values as described in the previous paragraph, we fix the agents' beliefs concerning the expected payoffs, $U^*(\cdot)$, and simulate the trading game for another 500 million events. Then, we calculate the realized payoffs of all order submissions after they have been executed. Let us denote these realized payoffs as $\tilde{J}(\cdot)$. $\tilde{J}(\cdot)$ is a direct measure of the realized benefits of trading. First, we require that the correlation between beliefs $U^*(\cdot)$ and realized outcomes $\tilde{J}(\cdot)$ is higher than 0.99. Second, we require that the mean absolute error in beliefs, i.e. the difference between $U^*(\cdot)$ and $\tilde{J}(\cdot)$ weighted by the number of times that a specific action has been selected in a given state within the last 500 million simulated events, is less than 0.01 (i.e. in a similar way to the previous paragraph when we evaluated the change in the expected value between $U_{t_2}(\tilde{a}|s)$ and $U_{t_1}(\tilde{a}|s)$ weighted by $k_{\tilde{a},s}^{t_1,t_1}$). If any convergence criterion is not reached, we continue simulating the trading game and updating the beliefs until all convergence criteria are satisfied.

II. Parameter Benchmarking

In order to ascertain the suitability of our model parameter choices, we employ the Hollifield, Miller, Sandås, and Slive (2006) approach to empirically estimate the trader arrival rate, asset volatility, and private value in limit order markets for two FTSE-100 stocks listed on the London Stock Exchange (LSE): Vodafone (VOD) and Hargreaves Lansdown (HRGV). The two stocks belong to the largest and smallest market value quartile of the FTSE-100 index, respectively. We use message-level data from the LSE's rebuild order book service for the month of January 2015. The data contain a record of the submission, cancellation, and execution of each visible order submitted to the LSE. Each message is time-stamped to the millisecond.

Hollifield et al. (2006) model traders' order submission decisions in limit order markets as a function of their private valuations, expected execution probabilities and picking-off risk, and market conditions. They employ a two-step process to compute the model parameters. In the first step, they estimate the execution probabilities and picking-off risk of different orders as a function of variables capturing the state of the limit order book, past order submission activity, order characteristics, and market conditions. In the second step, they use the execution probabilities and picking-off risk to estimate the trader arrival rates, private value distributions, and execution costs using maximum likelihood estimation.

Hollifield et al. (2006) estimate the gains from trading using this approach with data from the Vancouver Stock Exchange – primarily a venture capital exchange – for the period between May 1990 and November 1993. Due to differences in the specific markets, data, and time periods analyzed, we deviate from their approach in the following ways:

1. We identify all unexecuted immediate-or-cancel and fill-or-kill limit orders as those that have the same order submission and order cancellation timestamp and exclude them from our analysis.

2. We exclude multi-day orders from our analysis.

3. For computational reasons, we estimate the parameters independently for each

day of our sample. Consequently, we exclude the exogenous variables used by Hollifield et al. (2006), as these are only updated on a daily basis.

4. When estimating the conditional distributions of time to execution and time to cancellation, we treat all limit orders that survive longer than two hours as censored observations. The corresponding time period in Hollifield et al. (2006) is two days.

5. Similarly, we compute the probability of execution within two hours instead of two days for each limit order type.

6. We set execution costs to zero.

7. We exclude all opening, closing, intraday, and volatility auctions from our analysis and focus only on the continuous trading session.

Apart from these deviations, we exactly follow the Hollifield et al. (2006) two-step estimation approach. We do not compute individual and market-wide trading gains, as our focus is to identify the optimized model parameters. Specifically, we compute the trader arrival rate λ, fundamental volatility λ_v, and private value distribution F_α.

We start by calibrating the speed of time clock update. In our model, this value is linked with the trader arrival rate for the asset. The fundamental value volatility and traders' reentry rate are further determined relative to this trader arrival rate. Goettler et al. (2009), based on data from Hollifield et al. (2006), use a time clock update frequency of approximately one minute and link this rate to the daily number of trades. We rely on the same intuition, but instead use the daily number of trades to compute the *new* trader arrival rate in seconds. In our model – similar to that of Goettler, Parlour, and Rajan (2005); Goettler et al. (2009) – all traders remain part of the trading game until their orders are executed. Hence the number of trades acts as a proxy for the number of *new* trader arrivals. We compute the average time clock speed, $ClockSpeed_t$, on day t as:

$$ClockSpeed_t = \frac{8.5 \times 3600}{2 \times No.\ Of\ Trades_t} \tag{3}$$

where the numerator corresponds to the number of seconds in a trading day. We multiply the number of trades in the denominator by 2 as each trade involves an aggressive order trading against a passive order sitting in the limit order book.

The intensity of the Poisson process governing the change in the fundamental value λ_v corresponds to the expected number of time units after which the fundamental value v changes by one tick. This translates to the return volatility on day t, defined as the standard deviation of open-to-close returns, denoted by:

$$Volatility_t = \sqrt{\frac{8.5 \times 3600}{ClockSpeed_t \times k}} \times Rel.\ Tick\ Size_t \tag{4}$$

where k denotes the frequency with which the fundamental value changes. In our simulations, $k = 8$ ($k = 1.6$) in the low (high) volatility setting, i.e., the fundamental value changes, on average, after every 8 (1.6) trader arrivals. We compare this with the standard deviation of the midpoint prices on day t.

In our model, the distribution of private values F_α is discrete. However, Hollifield et al. (2006) parameterize the agents' private values as a mixture of two normal distributions with standard deviation denoted by σ_1 and σ_2, and their corresponding weights denoted by ρ and $(1-\rho)$, respectively. Specifically, the private value distribution F_α is parameterized as:

$$F_\alpha = \rho \left(\frac{\alpha}{y_t \sigma_1} \right) + (1-\rho) \left(\frac{\alpha}{y_t \sigma_2} \right) \tag{5}$$

where y_t is the common value of the stock on day t, proxied by the opening price. We assign the probability mass for all private values in the interval $[-\infty, -6]$, $[-6, -2]$, $[-2, 2]$, $[2, 6]$, and $[6, \infty]$, respectively to the five discrete values of $\alpha \in \{-8, -4, 0, 4, 8\}$ in our model.

Table B1 contains the parameter values. Consistent with VOD and HRGV belonging to the largest and smallest market value quartiles of the FTSE-100 index, respectively, the former has a lower time between trader arrivals and lower volatility than the latter. For VOD (HRGV), the average time between new trader arrivals λ_t is 1.3 (7.7) seconds. Based on these estimates, we obtain average low (high) estimates of the standard deviation of open to close returns of 1.2% (2.7%) for VOD and 1.4% (3.2%) for HRGV.[3] These estimates are comparable to realized intraday volatility of 1.4% for VOD and 1.8% for HRGV and suggest that our arrival rate and fundamental value volatility scale reasonably well under the assumption of a unit of time in our model being equivalent to a few seconds of calendar time. This is consistent with an increase in trading speed due to widespread electronification, investments in

[3]These estimates are based on the actual value of one tick, which is approximately 2bps for VOD and 6bps for HRGV.

technology by exchanges and market participants, and the use of algorithmic trading in global equity markets. Finally, the fractions of traders with absolute private value $|\alpha| = 8$, $|\alpha| = 4$, and $|\alpha| = 0$ are respectively equal to 24%, 42%, and 34% for VOD, and 14%, 38%, and 48% for HRGV. On the one hand, in our parameter set, traders with $|\alpha| = 8$ ($|\alpha| = 0$) have a higher (lower) weight than the corresponding fractions for VOD and HRGV. On the other hand, our chosen value for traders with $|\alpha| = 4$ is close to the corresponding values for VOD and HRGV. In conclusion, our chosen parameters appear largely comparable to those of modern electronic equity exchanges such as the LSE.

Table B1. Calibrated Model Parameters

This table presents the model parameters calibrated using message-level data for two FTSE-100 stocks from the London Stock Exchange. We report mean values of all parameters for the 21 trading days of January 2015. Trade Count is the number of transactions during the daily continuous trading session. $ClockSpeed_t$ is the number of seconds between two trader arrivals. Relative Tick Size is the ratio of one tick (in GBP) and the daily opening price. Low (High) Volatility λ_t^L (λ_t^H) is the average daily volatility assuming a change in the fundamental value, on average, after every 8 (1.6) trader arrivals. Empirical Volatility is the daily volatility computed based on one-minute returns. $F_{|\alpha|=8}$, $F_{|\alpha|=4}$, $F_{\alpha=0}$ is the fraction of traders with $|\alpha| = 8$, $|\alpha| = 4$, $\alpha = 0$, respectively. We report the mean values for all parameters across the 21 trading days in January 2015.

	VOD	HRGV
Trade Count	12,156	2,149
$ClockSpeed_t$ in seconds	1.3	7.7
Relative Tick Size	2bps	6bps
Low Volatility ($Volatility_t^L$)	1.2%	1.4%
High Volatility ($Volatility_t^H$)	2.7%	3.2%
Empirical Volatility	1.4%	1.8%
$F_{\|\alpha\|=8}$	24%	14%
$F_{\|\alpha\|=4}$	42%	38%
$F_{\alpha=0}$	34%	48%

III. Alternative Model Parameterizations

In this appendix, we report the results of the following four alternative parameterizations of our model:

1. Triple reentry rate for $\alpha = 0$ traders: This setting captures the impact of having faster intermediaries (HFTs) compared to the speed of natural liquidity trades ($|\alpha| > 0$). Table C1 contains the results.

2. Capacity constraints: This setting investigates the effect of relaxing the constraint of agents trading one share only. Table C2 reports the results.

3. Clientele effect: In this setting, we run our simulations by requiring 10% of the agents to only trade in one of the two order books in the fragmented setting. Table C3 reports the results.

4. Double $\alpha = 0$ agents: In this setting we double the population of intermediaries. Table C4 contains the results.

Under all alternative specifications, intermediaries extract a larger welfare compared to our main specificatio. This is because these agents are not exposed to delay costs, can trade on either side of the market, and can exploit small price changes. In the setting involving the faster reentry of such agents, they additionally benefit from their inherent speed advantage over other agents. In the specification featuring two-share traders, they are the primary beneficiaries of the high picking-off risk in the market. In the specification featuring captive traders, the order routing friction allows them to charge a wider spread in the less liquid market while trading aggressively in the more liquid market.

Table C1. Trading Behavior and Welfare: tripled reentry rate

This table presents the trading behavior and welfare results when the reentry rate is tripled. Specifically, Panel A 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 ask price. Panel B reports the picking-off risk, defined as the proportion of buy limit orders executed below 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. Panel D reports the welfare, defined as the average realized payoff, the waiting cost, and money transfer as defined in equation (3). These measures are reported for each agent type. 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. 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)				
	Order	Aggr.	Private value $\|\alpha\|$				Aggr.	Private value $\|\alpha\|$			
	Type	Prob.	0	4	8	Total	Prob.	0	4	8	Total
Panel A: Order Submissions											
Single Market	Limit	34.3%	78.9%	53.8%	15.9%	–	22.2%	12.8%	69.3%	61.5%	–
	Market	–	21.1%	46.2%	84.1%	–	–	87.2%	30.7%	38.5%	–
Frag Market	Limit	22.14%	72.5%	50.9%	26.3%	–	22.0%	59.1%	56.4%	32.4%	–
	Market	–	27.5%	49.1%	73.7%	–	–	40.9%	43.6%	67.6%	–
Panel B: Picking-Off Risk											
Single Market	–	–	2.9%	23.6%	72.9%	18.5%	–	9.9%	54.1%	82.8%	61.3%
Frag Market	–	–	1.0%	29.6%	70.0%	23.5%	–	6.5%	42.4%	71.8%	35.4%
Panel C: Time to Execution											
Single Market	–	–	16.0	3.5	2.0	9.3	–	19.1	2.8	1.1	3.4
Frag Market	–	–	13.0	3.6	1.5	7.3	–	12.2	3.8	1.6	6.4

Table C1. Trading Behavior and Welfare: Change in the reentry rates (cont.)

	Average Welfare per Trader				Waiting Cost per Trader				Money Transfer per Trader				Aggr. Welfare per Period	Deadweight loss
	Abs.	Private Value		Total	Abs.	Private Value		Total	Abs.	Private Value		Total		
	0	4	8		0	4	8		0	4	8			
Panel D: Welfare														
Low levels of picking-off risk														
Single Market	0.562	3.498	7.247	3.741	0.000	-0.350	-0.128	-0.179	0.562	-0.151	-0.625	-0.079	3.741	0.259
Frag Market	0.752	3.410	6.979	3.684	0.000	-0.368	-0.178	-0.201	0.752	-0.222	-0.843	-0.117	3.684	0.316
High levels of picking-off risk														
Single Market	0.683	3.375	7.032	3.663	0.000	-0.364	-0.264	-0.225	0.683	-0.261	-0.704	-0.110	3.663	0.337
Frag Market	0.865	3.346	6.814	3.642	0.000	-0.421	-0.235	-0.239	0.865	-0.232	-0.951	-0.118	3.642	0.358

Table C2. Trading Behavior and Welfare: Capacity constraints

This table presents the trading behavior and welfare results with the contraint that agents can submit only one share relaxed. We allow an agent to be followed by a second identical agent with probability 5%. Panel A 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 ask price. Panel B reports the picking-off risk, defined as the proportion of buy limit orders executed below 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. Panel D reports the welfare, defined as the average realized payoff, the waiting cost, and money transfer as defined in equation (3). These measures are reported for each agent type. 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. 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)				
	Order	Aggr.	Private value $\|\alpha\|$				Aggr.	Private value $\|\alpha\|$			
	Type	Prob.	0	4	8	Total	Prob.	0	4	8	Total
Panel A: Order Submissions											
Single Market	Limit	36.4%	74.9%	51.8%	22.7%	–	25.0%	18.9%	66.1%	59.6%	–
	Market	–	25.1%	48.2%	77.3%	–	–	81.1%	33.9%	40.4%	–
Frag Market	Limit	28.5%	76.6%	49.2%	24.5%	–	22.8%	50.5%	57.5%	39.4%	–
	Market	–	23.4%	50.8%	75.5%	–	–	49.5%	42.5%	60.6%	–
Panel B: Picking-Off Risk											
Single Market	–	–	3.9%	29.9%	77.5%	24.7%	–	14.0%	55.1%	86.6%	61.7%
Frag Market	–	–	2.6%	29.0%	74.4%	23.5%	–	10.8%	47.9%	76.8%	43.5%
Panel C: Time to Execution											
Single Market	–	–	13.8	3.8	2.2	8.1	–	23.2	3.0	1.2	4.7
Frag Market	–	–	10.3	4.0	2.1	6.6	–	11.8	4.0	1.8	5.8

Table C2. Trading Behavior and Welfare: Capacity constraints (cont.)

	Average Welfare per Trader				Waiting Cost per Trader				Money Transfer per Trader				Aggr. Welfare per Period	Deadweight loss
	Abs. Private Value 0	4	8	Total	Abs. Private Value 0	4	8	Total	Abs. Private Value 0	4	8	Total		
Panel D: Welfare														
Low levels of picking-off risk														
Single Market	0.590	3.481	7.224	3.738	0.000	-0.368	-0.196	-0.206	0.590	-0.151	-0.580	-0.058	3.738	0.262
Frag Market	0.698	3.439	7.124	3.722	0.000	-0.379	-0.216	-0.216	0.698	-0.182	-0.660	-0.061	3.722	0.278
High levels of picking-off risk														
Single Market	0.640	3.372	7.003	3.642	0.000	-0.387	-0.292	-0.243	0.640	-0.240	-0.705	-0.115	3.642	0.358
Frag Market	0.841	3.366	6.842	3.652	0.000	-0.436	-0.307	-0.267	0.841	-0.196	-0.851	-0.081	3.652	0.348

Table C3. Trading Behavior and Welfare: Clientele effect

This table presents the trading behavior and welfare results when we keep 10% of agents as captives in one of the books (book 1). Panel A 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 ask price. Panel B reports the picking-off risk, defined as the proportion of buy limit orders executed below 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. Panel D reports the welfare, defined as the average realized payoff, the waiting cost, and money transfer as defined in equation (3). These measures are reported for each agent type. 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. 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)				
	Order	Aggr.	Private value $\|\alpha\|$				Aggr.	Private value $\|\alpha\|$			
	Type	Prob.	0	4	8	Total	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%	–
Frag Market with clientele	Limit	34.4%	77.6%	49.5%	23.0%	–	26.5%	50.1%	57.6%	39.8%	–
	Market	–	22.4%	50.5%	77.0%	–	–	49.9%	42.4%	60.2%	–
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%	48.0%	76.8%	43.7%
Frag Market with clientele	–	–	2.5%	27.4%	72.2%	22.0%	–	10.9%	49.9%	78.6%	46.2%
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
Frag Market with clientele	–	–	10.4	3.7	1.9	6.6	–	11.8	3.8	1.7	5.7

Table C3. Trading Behavior and Welfare: Clientele effect (cont.)

	Average Welfare per Trader				Waiting Cost per Trader				Money Transfer per Trader				Aggr. Welfare per Period	Deadweight loss
	Abs. 0	Private Value 4	8	Total	Abs. 0	Private Value 4	8	Total	Abs. 0	Private Value 4	8	Total		
Panel D: Welfare														
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
Frag Market captive agents	0.686	3.456	7.139	3.730	0.000	-0.363	-0.191	-0.203	0.686	-0.181	-0.671	-0.068	3.730	0.270
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
Frag Market captive agents	0.837	3.378	6.851	3.660	0.000	-0.424	-0.296	-0.258	0.837	-0.197	-0.854	-0.084	3.660	0.340

Table C4. Decomposition of Welfare by Trader Type

This table reports the welfare defined as the average realized payoff. In addition, this table presents waiting cost and money transfer as defined in equation (3). These measures are reported for each agent type. We report all statistics for a single and a fragmented market under the low-volatility parameterization. In addition to reporting the welfare statistics in the single and fragmented market settings with $F_\alpha = \{0.15, 0.35, 0.65, 0.85, 1.0\}$ (baseline parameterization), we also report the welfare statistics for a fragmented market with the population of agents with $\alpha = 0$ doubled, i.e. $F_\alpha = \{0.15/1.3, 0.35/1.3, 0.95/1.3, 1.15/1.3, 1.0\}$. The last two columns report the average aggregate welfare per period and the deadweight loss. All measures are reported in ticks.

	Average Welfare per Trader				Waiting Cost per Trader				Money Transfer per Trader				Aggr. Welfare per Period	Deadweight loss
	Abs. 0	Private 4	Value 8	Total	Abs. 0	Private 4	Value 8	Total	Abs. 0	Private 4	Value 8	Total		
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
Frag Market $2 \times \alpha = 0$	0.485	3.312	7.137	2.890	0.000	-0.127	-0.029	-0.046	0.485	-0.561	-0.835	-0.142	3.757	0.243

References

Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in io. *Handbook of Industrial Organization 3*, 1887–1966.

Goettler, R. L., C. A. Parlour, and U. Rajan (2005). Equilibrium in a Dynamic Limit Order Market. *Journal of Finance 60*(5), 2149–2192.

Goettler, R. L., C. A. Parlour, and U. Rajan (2009). Informed traders and limit order markets. *Journal of Financial Economics 93*(1), 67–87.

Hollifield, B., R. A. Miller, P. Sandås, and J. Slive (2006). Estimating the gains from trade in limit-order markets. *The Journal of Finance 61*(6), 2753–2804.

Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory 100*(2), 191–219.

Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse'of dimensionality. *Econometrica 69*(5), 1261–1281.

Rieder, U. (1979). Equilibrium plans for non-zero-sum markov games. *Game Theory and Related Topics*, 91–101.