Strategic Market Choice: Frequent Call Markets vs. Continuous Double Auctions for Fast and Slow Traders

Elaine Wah   Dylan R. Hurd   Michael P. Wellman

Computer Science & Engineering
University of Michigan
{ewah,drhurd,wellman}@umich.edu

ABSTRACT

Frequent call markets have been proposed as a market design solution to the latency arms race perpetuated by high-frequency traders in continuous markets, but the path to widespread adoption of such markets is unclear. If such trading mechanisms were available, would anyone want to use them? This is ultimately a question of market choice, thus we model it as a game of strategic market selection, where agents choose to participate in either a frequent call market or a continuous double auction. Our market environment is populated by fast and slow traders, who reenter to trade at different rates. We employ empirical game-theoretic methods to determine the market types and trading strategies selected in equilibrium. We also analyze best-response patterns to characterize the frequent call market's basin of attraction. Our findings show that in equilibrium, welfare of slow traders is generally higher in the call market. We also find strong evidence of a predator-prey relation between fast and slow traders: the fast traders chase agents into either market, and slow traders under pursuit seek the protection of the frequent call market.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics

Keywords

high-frequency trading; frequent batch auction; market mechanism

1. INTRODUCTION

High-frequency trading (HFT) is the practice of exploiting speed advantages in market access and execution to enhance profits in algorithmic trading in financial markets. HFT strategies typically operate at high volumes and low latencies (where fractions of a millisecond may be pivotal), and maintain positions with extremely short holding periods. HFT is estimated to account for over half of daily trading volume [7] and it has dominated financial news headlines in recent years. Incidents such as the "Flash Crash" of May 2010 and Knight Capital accidentally flooding the market with hundreds of orders in August 2012 [32] have piqued concerns, and controversies about computerized trading in today's financial markets have reached a critical point with the publication of *Flash Boys* by Michael Lewis [27].

Flash Boys tells the story of IEX, a trading venue designed specifically to protect institutional investor orders from anticipation by speed-advantaged algorithmic traders. Proponents of high-speed traders posit that HFT activity reduces trading costs for market participants. Others argue that these traders harm investors and that practices such as colocation, in which firms embed their HFT systems in the same data centers as an exchange's servers, contribute to a wasteful latency arms race [20].

Incremental speed yields advantage in trading due to the continuous nature of market mechanisms. Currently, most stock markets operate as *continuous double auctions* or CDAs [16]. In a CDA, orders are matched strictly on a first-come basis. This time priority rule induces a winner-take-all scenario, where the fastest trader can readily expropriate all gains from new information. The speed differential between high-frequency traders and slower, non-HF investors subjects the latter to adverse selection, in which the slower traders' resting orders are more likely to trade when information moves against them.

An alternative to continuous trading is a *frequent call market* or frequent batch auction, in which order matching is performed only at discrete, periodic intervals (e.g., on the order of tenths of a second). A discrete-time market facilitates more efficient trading by aggregating supply and demand and matching orders to trade at a uniform price [3, 18, 41]. In a frequent call market, there is no time-priority within each clearing interval. Each interval is a sealed-bid auction: Participants do not know what orders other traders have submitted, ergo orders in the frequent call market cannot be targeted specifically by incoming informed orders. Even if a fast trader knew somehow about a stale order sitting in the book, it could not exploit that completely because the prices are set via a competition among all traders able to submit orders within the clearing interval.

Allowing orders to accumulate over short time periods in a frequent call market has been advocated by several as a means to neutralize small speed advantages and mitigate detrimental effects of high-frequency trading [6, 36, 38, 45], and regulators are starting to take notice. For instance, New York Attorney General Eric Schneiderman endorsed frequent batch auctions in remarks during a March 2014 New York Law School panel on Insider Trading 2.0:

> *Currently, on our exchanges, securities are traded continuously, which...rewards high-frequency traders who*

continuously flood the market with orders, emphasizing speed over price.... If you had frequent batch auctions, there's no point in trying to get faster than whatever the interval is. It would discourage the risk taking that can cause flash crashes because, in the quest for greater and greater speed, there is, in and of itself, a threat to market stability [35].

U.S. Securities and Exchange Commission Chair Mary Jo White has also indicated receptiveness towards "flexible competitive solutions...[which] could include frequent batch auctions or other mechanisms designed to minimize speed advantages" [47].

Skeptics of frequent call markets raise various objections to their feasibility. Some doubt whether continuous- and discrete-time markets can coexist, and posit that it will be necessary to ensure that fragmented call markets clear in a synchronized manner [33]. Others question the ease of implementing these call markets [2]. Featherstone [14] argues that frequent batch auctions are an unattractive alternative to current continuous markets, contending that discrete-time markets will diminish trading and adversely affect price stability, while simultaneously creating the incentive to snipe within the clearing interval in the event of new information arriving before the clear. Similarly, Ross [34] surmises that the introduction of frequent batch auctions would engender a race to place the first order in the book for each call auction.

These and most other arguments we have encountered appear to be based on misconceptions or unfounded speculation. The call market does not need to give time priority for orders within the clearing interval, and so it is easy to avoid races to submit orders and to instead channel competition to the price dimension. Traders submitting the best price, whether fast or slow, will execute, and orders clear at a uniform price that no market participant knows in advance, making ties in price unlikely anyway (especially if prices are fine-grained). For the same reason, synchronization of multiple frequent call markets is unnecessary, given that the gain from a speed advantage is already reduced by the lack of visibility into the order book. In addition, implementation of frequent call markets is clearly feasible; many modern stock markets open and close trading each day with a call auction [28, 40].

Yet frequent call markets have hitherto not been widely adopted. This may be simply a matter of inertia; as markets have evolved from in-person to electronic, imposing an explicit time delay would take a deliberate intervention. Such time delays are intuitively retrograde to many, as they seem to compromise the general investor demand for trading immediacy [12].

This explains why existing continuous markets might not change their policies, but what about introducing new markets with the frequent call mechanism? We see no economic reason why a discrete-time market could not coexist alongside continuous market mechanisms [42], but admittedly the burden of demonstration may rest on those of us arguing for feasibility. To provide such a demonstration, we consider the question of market choice: given availability of both mechanisms, will traders elect to submit orders to a frequent call market over a continuous market, and if so, under what conditions?

These are the questions addressed by our study. We formulate the frequent call market vs. CDA scenario as a game of market choice in which fast and slow traders—who differ on the frequency with which they arrive to trade—specify, as part of their strategy, a selected market mechanism. This strategic market choice game is described in the following section. We discuss our market environment in Section 3, our methodology and game analysis in Section 4, and we survey additional related work in Section 5. Section 6 offers our conclusions.

2. STRATEGIC MARKET CHOICE

To determine whether a frequent call market operating alongside a continuous market can successfully attract investors, we propose a market choice game in which traders specify the preferred trading mechanism as part of their strategy.

The players in our game are traders, grouped in two *roles*: FAST and SLOW. These roles differ only in the frequency with which traders enter to submit an order. The game is *role-symmetric*, meaning that players in the same role share the same strategy set, and payoffs are completely determined by role membership and by the number playing each strategy in each role. That is, the specific strategy-to-player assignments within roles are irrelevant.

In our model, there is a single security available for trade on both a frequent call market and a continuous double auction. The fundamental value of this security is a mean-reverting stochastic process. A player's valuation for units of the security depends on this fundamental and the trader's private benefits of trading, with payoff defined as the total surplus accrued over the trading horizon.

Traders can elect to go to either the frequent call market or the continuous market. On each arrival, they observe the current fundamental value and submit a single-unit order to their selected market. Resting orders in the book are subject to adverse selection, since newly arriving traders have more current information about the fundamental, which they can exploit to pick off stale orders. Since SLOW traders arrive less frequently into the market than their FAST counterparts, SLOW-agent orders are on average based on older information, and thus are exposed to a greater degree of adverse selection.

We select a fixed, deterministic rate of clearing for the frequent call market in our market choice game. Though several have proposed randomizing the clearing interval to deter sniping [13, 22, 37], we have argued [42] that such randomization accomplishes no reduction in incentive for HFT speed advantages. A deterministic clear time offers the prospect of sniping within a small time window at the end of the clear interval, whereas a random clear time offers a small probabilistic prospect for advantage over the entire interval. In expectation, the value of this advantage is the same. Moreover, as the model in this paper does not include strategic timing, sniping is effectively ruled out by assumption.

In a market choice game with players who strategically decide among market mechanisms, there trivially exist equilibria in which all traders select any one given market, regardless of its merits. These equilibria arise because when all *other* agents are in that market, the remaining trader has no possibility to trade anywhere else. Thus the trader's only option for positive payoff is to join the focal market. To render these equilibria non-inevitable, we introduce to each market a set of *environment agents*, providing a base set of available trading partners. The environment traders follow designated strategies for their assigned market and are not considered players in our game model. As such, their behavior plays no part in game-theoretic analysis and their trading gains are ignored in surplus calculations.

We employ an empirical simulation-based approach to explore the strategy space. This facilitates identification of the market conditions under which traders may prefer one market mechanism over the other. From the empirical game induced over thousands of simulations of selected strategy profiles, we determine the market chosen in equilibrium, and we analyze the corresponding gains from trade. We characterize the frequent call market's basin of attraction through analysis of trader best responses that specify the frequent call market over the CDA.

Our findings show that in equilibrium, welfare of SLOW traders is generally higher in the frequent call market than in the continu-

ous double auction. We also find strong evidence of a predator-prey interaction between FAST and SLOW traders. The FAST traders follow their prey into either market, whereas the SLOW traders flee their pursuers, congregating in the frequent call market as long as it is sufficiently thick.

3. MARKET GAME SETUP

We construct a model of a single security traded simultaneously in a continuous double auction market (CDA) and a frequent call market (CALL). Prices are fine-grained but discrete, taking integer values. Time is likewise fine-grained and discrete, with finite horizon T. The environment is populated by multiple trading agents, representing investors, each associated with one of the two market mechanisms. Player agents choose between CDA and CALL, and environment agents are each assigned to one of these, in both cases for the duration of the trading period. Agents arrive according to a Poisson process, and on each arrival they submit a single-unit limit order to their associated market—replacing any prior outstanding order. Thus, at any given time, investors are restricted to a single order to buy or sell one unit.

Before placing orders, traders observe price quotes (BID and ASK) in their associated market. CDA price quotes reflect the best current outstanding orders, while the frequent call market quotes reflect the best outstanding orders immediately following the most recent market clear. Specifically, for the CDA, BID_t is the price of the highest buy offer at time t and ASK_t is the price of the lowest offer to sell. For the frequent call market, BID_c corresponds to the highest outstanding buy offer after the clear at the most recent clear time c, and ASK_c the lowest outstanding offer to sell. Other bids in either order book are not visible to traders.

We introduce an equal number of environment agents into both markets. Except for market selection, environment agents operate equivalently to traders (described below), but are not taken into account when calculating welfare. We denote the number of environment agents in each market by E.

3.1 Valuation Model

Each trader has an individual valuation for the security comprised of private and common components. We denote by r_t the common *fundamental value* for the security at time t. The fundamental time series is generated by a mean-reverting stochastic process:

$$r_t = \max \{0, \ \kappa \bar{r} + (1 - \kappa) r_{t-1} + u_t\}.$$

Parameter $\kappa \in [0, 1]$ specifies the degree to which the fundamental reverts back to the mean $\bar{r}$, and parameter $u_t \sim \mathcal{N}\left(0, \sigma_s^2\right)$ is a random shock at time t.

The private component for agent i is a vector Θ_i representing differences in private benefits of trading given the trader's net position, similar to the model of Goettler et al. [19]. These individual differences may arise due to portfolio considerations, hedging needs, or preferences regarding trading urgency. The vector is of size $2q_{\max}$, where $q_{\max} > 0$ is the maximum number of units the agent can be long or short at any time, with

$$\Theta_i = \left(\theta_i^{-q_{\max}+1}, \ldots, \theta_i^0, \theta_i^{+1}, \ldots, \theta_i^{q_{\max}}\right).$$

Element θ_i^q is the incremental private benefit obtained from selling one unit given current position q, where positive (negative) q indicates a long (short) position. Similarly, θ_i^{q+1} is the marginal private gain from buying an additional unit given current net position q.

We generate Θ_i from a set of $2q_{\max}$ values drawn independently from a Gaussian distribution. Let $\hat{\theta} \sim \mathcal{N}\left(0, \sigma_{PV}^2\right)$ denote one of

these drawn values. To ensure that the valuation reflects diminishing marginal utility, that is, $\theta^{q'} \geq \theta^q$ for all $q' \leq q$, we sort the $\hat{\theta}$ and set the θ^q to respective values in the sorted list.

Trader i's valuation v for the security at time t is based on its current position q_t and the value of the global fundamental at time T, the end of the trading horizon:

$$v_i(t) = r_T + \begin{cases} \theta_i^{q_t+1} & \text{if buying 1 unit} \\ \theta_i^{q_t} & \text{if selling 1 unit.} \end{cases}$$

For a single-quantity limit order transacting at time t and price p, a trader obtains surplus:

$$\begin{cases} v_i(t) - p & \text{for buy transactions, or} \\ p - v_i(t) & \text{for sell transactions.} \end{cases}$$

Since the price and fundamental terms cancel out in exchange, the total surplus achieved when agent B buys from agent S is $\theta_B^{q(B)+1} - \theta_S^{q(S)}$, where $q(i)$ denotes the pre-trade position of agent i.

3.2 Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs [10, 16, 46]. In this study, we consider trading strategies in the so-called *Zero Intelligence* (ZI) family [17].

The traders arrive at the market according to a Poisson process with rate λ. On each arrival, they are assigned to buy or sell (with equal probability), and accordingly submit an order to buy or sell a single unit. Agents may trade any number of times, as long as their net positions do not exceed $q_{\max}$ (either long or short). At the end of the simulation period, traders liquidate their accumulated inventory at r_T, the end-time fundamental.

A ZI trader assesses its valuation $v_i(t)$ at the time of market entry t, using an estimate $\hat{r}_t$ of the terminal fundamental r_T. The estimate is based on the current fundamental, r_t, adjusted to account for mean reversion:

$$\hat{r}_t = \left(1 - (1 - \kappa)^{T-t}\right) \bar{r} + (1 - \kappa)^{T-t} r_t. \quad (1)$$

The ZI agent then submits a bid shaded from this estimate by a random offset—the degree of surplus it demands from the trade. The amount of shading is drawn uniformly from range $[R_{\min}, R_{\max}]$. Specifically, a ZI trader i arriving at time t with current position q submits a limit order for a single unit of the security at price

$$p_i \sim \begin{cases} \mathcal{U}\left[\hat{r}_t + \theta_i^{q+1} - R_{\max}, \hat{r}_t + \theta_i^{q+1} - R_{\min}\right] & \text{if buying} \\ \mathcal{U}\left[\hat{r}_t + \theta_i^q + R_{\min}, \hat{r}_t + \theta_i^q + R_{\max}\right] & \text{if selling.} \end{cases}$$

We extend the baseline ZI strategy with a threshold parameter $\eta \in [0, 1]$, whereby if the agent could achieve a fraction η of its requested surplus at the current price quote, it would simply take that quote rather than posting a limit order to the book. Setting $\eta = 1$ is equivalent to the strategy without employing the threshold.

In our setting, traders repeatedly enter the market, thus we refer to the strategy as *ZI with Reentry* (ZIR). Upon each entry, the ZIR trader withdraws its previous order (if not transacted yet) before executing the strategy described above. Time between entries is distributed exponentially at rate λ. As described below, our game employs three rates, one each for FAST, SLOW, and environment traders.

The final strategy parameter indicates the selected market type: CALL or CDA. The market choice decision is made before trading commences at time 0, and once selected, the market for a given agent is fixed for the duration of the trading horizon.

$R_{\min}$	0	0	0	0	500	500	0
$R_{\max}$	125	250	500	1000	1000	1000	2500
η	1	1	1	1	0.4	1	1
Market type	Both	Both	Both	Both	CDA	CALL	Both

Table 1: ZIR strategy combinations included in empirical game-theoretic analysis. Market type indicates whether the strategy is available in the CDA, CALL, or both.

4. EMPIRICAL GAME-THEORETIC ANALYSIS

We determine equilibria for our game of strategic market choice through *empirical game-theoretic analysis* (EGTA), a simulation-based process that allows us to perform strategy selection for traders by comparing the payoffs (i.e., surplus) of different combinations of trader and strategy assignments [44]. EGTA entails simulation of many strategy profiles, accumulating payoff observations, and inducing an empirical game model.

We apply EGTA in an iterative manner, interleaving exploration of the profile space with analysis of the empirical game model induced by average payoffs in simulation. Specifically, we generate equilibrium candidates by applying replicator dynamics to *complete subgames*, defined as sets of strategies (one set per role) for which we have simulated all profile combinations. This yields role-symmetric Nash equilibria (RSNE) of the subgames, which we can then test as candidate solutions for the overall game. If we can identify a strategy in the full strategy set that beneficially deviates from the candidate, we say the candidate is *refuted*. We say that a candidate profile is *confirmed* as an RSNE when all possible deviations have been evaluated, and none are beneficial.

We simulate additional profiles for a game until we have confirmed at least one RSNE, evaluated every pure-strategy symmetric profile (i.e., where the players in each role play a strategy with probability 1), and pursued with some degree of diligence every equilibrium candidate encountered. We confirm or refute each candidate by evaluating deviations to strategies outside their subgames. If a candidate is refuted, we construct a new subgame by adding the best response to its support, and proceed to explore the corresponding subgame. When this process reaches quiescence, we consider the search to have satisfied the diligence requirement.

In this study, we simulate the market model described in Section 3 using an extension of the financial market simulator we developed for our prior analysis of latency arbitrage [41]. We manage our experiments via the EGTAOnline infrastructure [8], and we run our simulations on the high-performance computing cluster at the University of Michigan.

We collect data for multiple combinations of the trader strategies: a minimum of 5,000 samples per profile evaluated, with 20,000 samples for most profiles and averaging at least 10,739 samples per profile in each environment. From these payoff estimates, we compute RSNE for each environment, and use these as a foundation for our analysis of the welfare effects in equilibrium for FAST versus SLOW traders, the attractiveness of the CALL over the CDA, and the loss in deviating from the equilibrium market type.

4.1 Game Reduction

As analysis of the full game is intractable due to the exponential growth in game size relative to the number of players, we employ *deviation-preserving reduction* (DPR) to construct a reduced game approximating the full role-symmetric game [48]. DPR preserves the payoffs from single-player, unilateral deviations, and maintains in the reduced game the same proportion of opponents playing each strategy as in the full game. In a deviation-preserving reduced game, each player views itself as controlling one full-game agent and views the other-agent profile in the reduced game as an aggregation of all other players in the full game.

In the market choice game, players are partitioned into roles $R = \{\text{FAST}, \text{SLOW}\}$, and players in either role can select among a set of strategies S. Consider the reduction of an $(N_{\text{FAST}}, N_{\text{SLOW}})$-player game to a $(k_{\text{FAST}}, k_{\text{SLOW}})$-player reduced game. Given the other FAST agents play strategies $(f_2, \ldots, f_{k_{\text{FAST}}})$ and the SLOW agents play strategies $(s_1, \ldots, s_{k_{\text{SLOW}}})$, the payoff for a FAST agent playing strategy $f_1 \in S$ in the reduced game is given by the payoff of playing f_1 in the full $(N_{\text{FAST}}, N_{\text{SLOW}})$-player game when the other $N_{\text{FAST}} - 1$ FAST traders are divided uniformly ($\frac{N_{\text{FAST}}-1}{k_{\text{FAST}}-1}$) among the strategies $f_2, \ldots, f_{k_{\text{FAST}}}$ and the other-role (i.e., SLOW) players are divided uniformly ($\frac{N_{\text{SLOW}}}{k_{\text{SLOW}}}$) among their strategies $s_1 \ldots s_{k_{\text{SLOW}}}$. The payoff for a single SLOW agent is analogous.

We deliberately select values for N_r and k_r, $r \in \{\text{FAST}, \text{SLOW}\}$, to ensure that the fractions above defining the game reduction come out as integers. That is, for each role r, we choose values such that N_r is evenly divisible by k_r and $N_r - 1$ is evenly divisible by $k_r - 1$. Specifically, our market choice game is comprised of 42 players, with $N_{\text{FAST}} = N_{\text{SLOW}} = 21$, which we approximate by a DPR game with $k_{\text{FAST}} = k_{\text{SLOW}} = 3$. We use simulation data from the full $(21, 21)$-player game to estimate the payoffs of the $(3, 3)$-player reduced game.

4.2 Environment Settings

We evaluate the performance of traders in four environments. Reentry rates are fixed across the environments, with FAST traders arriving in the market at rate $\lambda_F = 0.004$, and SLOW traders entering at rate $\lambda_S = 0.002$. In all settings, there is one CDA and one frequent call market, which clears every 100 time steps. Each simulation run lasts $T = 12000$ time steps. The mean-reverting global fundamental has a mean value $\bar{r} = 10^5$. The variance for the private value vector is $\sigma_{PV}^2 = 5 \times 10^6$. The fundamental shock variance is $\sigma_s^2 = 1 \times 10^6$.

The strategy of environment agents is fixed; they play a ZIR strategy with range $[0, 1000]$ and $\eta = 1$. The environment agents enter their respective markets with rate $\lambda_E = 0.005$. The environments differ in the value of the mean-reversion parameter (κ) and the number of environment agents $E \in \{8, 14, 42\}$. The configurations are as follows:

Environment I $E = 8, \kappa = 0.05$

Environment II $E = 8, \kappa = 0.01$

Environment III $E = 14, \kappa = 0.01$

Environment IV $E = 42, \kappa = 0.01$

The empirical games for these environments include 12 strategies (Table 1) for traders, 6 in each market.

4.3 Social Optimum

We assess efficiency by comparison of market outcomes with the social optimum. We define this optimum for a population of
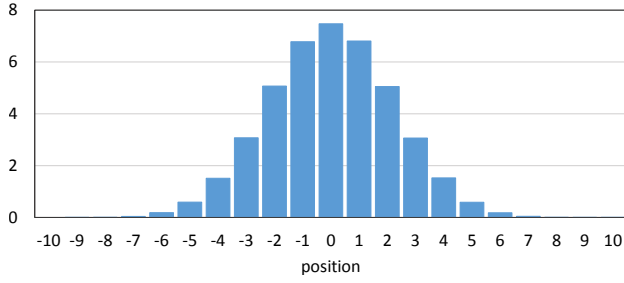
Figure 1: Histogram of the net position (i.e., number of units traded) of traders in the socially optimal allocation. The distribution is compiled from 20,000 samples.

42 traders, based on the distribution of the private component of agents' valuations, with parameters $q_{max} = 10$ and $\sigma^2_{PV} = 5 \times 10^6$. To calculate an optimal allocation for a particular array of draws from this distribution, we simply find the competitive equilibrium using the call market clearing function. Each trader submits its valuation vector as a demand curve, with q_{max} sell orders at prices $\bar{r} + \theta^s_i$, $s \in \{-q_{max} + 1, \ldots, 0\}$ and q_{max} buy orders at prices $\bar{r} + \theta^b_i$, $b \in \{+1, \ldots, q_{max}\}$, with each order for a single unit of the security. Over 20,000 samples, we find a mean social welfare of 27887. As they are not considered players in the market choice game, we do not include environment agents in the determination of the socially optimal allocation. Figure 1 shows the histogram of trades per player in the social optimum.

4.4 Basin of Attraction

The main results of this study are shown in the heat maps of Figure 2, which illustrate the trader population conditions under which the CALL serves as an attractor. We characterize the frequent call market's basin of attraction by categorizing the market type a trader selects when the other traders' strategies are fixed; in other words, we identify and classify the trading mechanism selected in the trader's best response.

Given a trader, we fix the other-agent profile (i.e., the set of strategy counts for the 20 players in the same role and for the 21 players in the other role), and we identify the market type selected in the trader's best response. Since we selectively sample full-game profiles for the $(3, 3)$-player DPR approximation, we can bucket all other-agent profiles into 12 unique categories by role, based on the population of traders in each market type.

For example, if we examine a SLOW trader in the CALL, the 20 other SLOW traders may all be in the same market (CDA or CALL) or they may be equally split between the two markets (10 in the CALL and 10 in the CDA). No other cross-market divisions of same-role players are possible because we selectively collect profiles to reduce via DPR to a $(3, 3)$-player game. The 21 traders in the other role (FAST) may all be in the same market (CDA or CALL), or they may be split between the two markets, with 7 agents in one market and 14 in the other.

For each of the 12 categories of trader population distributions across markets, we count the number of other-agent profiles for which the given player's best response specifies the CALL over the CDA. We report the corresponding percentages in two best-response heat maps, one per role, for each market choice game.

To ensure full coverage of all population categories, we construct a complete subgame[1] for each environment. Our results for these

subgames are illustrated in the heat maps of Figure 2; these characterize the frequent call market's basin of attraction from the perspective of a single trader in each role (SLOW on the left, FAST on the right). For example, the top left entry in a SLOW trader heat map reports the percentage of all other-agent profiles comprised of 20 SLOW traders in the CDA and 21 FAST traders also in the CDA in which a SLOW trader's best response specifies the CALL.

The higher mean reversion in environment I implies that slower traders are less likely to be picked off by speed-advantaged traders, and therefore we find that the SLOW traders display no strong preference to switch to the CALL unless the majority of other traders (regardless of speed) are in the frequent call market as well. When the degree of mean reversion is reduced, the SLOW agents face greater risk of being picked off by FAST agents with newer and better information. Therefore, environments II through IV are much more salient in answering questions about strategic market choice under adverse selection, and we focus the rest of the following discussion in this section on those corresponding subgames.

We see from the environment II–IV heat maps that there is safety in numbers for a single SLOW trader deciding between the CALL and CDA: if 20 of the SLOW traders are in a given market, the best response is more often than not to pick the same market as everyone else, whether that is the CDA or CALL. When the SLOW agent population is equally divided between the two markets, however, we observe a gradual mass exodus of SLOW traders from the frequent call market as more FAST traders enter the CALL. The percentage of SLOW-trader best responses selecting the CALL decreases monotonically from around 90% to below 40% as FAST traders leave the CDA for the frequent call market. Despite the protection afforded to them in the CALL, the SLOW traders would rather take their chances in the CDA than remain in the same market as the FAST traders. However, if the CALL is sufficiently thick (as in environment IV), the SLOW traders prefer the sanctuary of the frequent call market, regardless of where the FAST traders are.

On the other hand, FAST traders clearly stand to gain from the informationally disadvantaged orders submitted by their slower counterparts. Therefore, they exhibit a strong preference for the market selected by the majority of SLOW traders, and they readily follow the SLOW traders to either market. We observe that their preference for the CALL increases strictly monotonically, from 0% to nearly 100%, as the number of SLOW traders in the CALL increases.
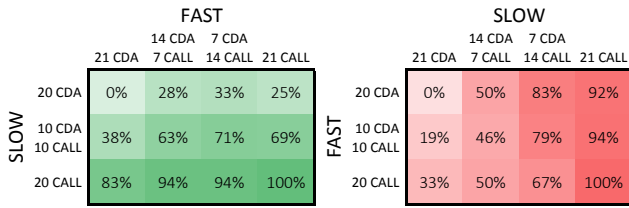
These results reveal the dynamics of the predator-prey interaction between the FAST and SLOW traders. The SLOW traders face less risk as part of a large group, but once they are split up between the two markets, those in the CDA tend to flee to the CALL to get away from the FAST traders, while the FAST traders relentlessly pursue the SLOW traders, regardless of market.

We also analyze the collected profiles in the full games, shown in Figure 3. The heat maps for the SLOW traders are similar to those in the complete subgames, but the results for FAST traders are markedly different. This is due to the bias in sampling full-game profiles for our game-theoretic analysis. As sampling all 681,264 profiles in the full game (given two roles, with 12 strategies each) is intractable, our coverage of the profile space is primarily determined by the more promising subgames identified during EGTA.

4.5 Equilibrium Analysis

Our equilibrium results are shown in Table 2. For each RSNE, we compute surplus for traders in each role by sampling 10,000

---

[1]In each subgame we include the strategies played with the highest probabilities across all RSNE found in that environment. The full

details of all equilibria in the four games, as well as the strategy sets of the complete subgames used to characterize the frequent call market's basin of attraction, are available in an online appendix (http://hdl.handle.net/2027.42/111897).

Figure 2

(a) Environment I

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	28%	33%	25%
10 CDA 10 CALL	38%	63%	71%	69%
20 CALL	83%	94%	94%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	50%	83%	92%
10 CDA 10 CALL	19%	46%	79%	94%
20 CALL	33%	50%	67%	100%

(b) Environment II

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	17%	0%	6%	0%
10 CDA 10 CALL	100%	83%	67%	31%
20 CALL	100%	100%	100%	92%

SLOW (FAST trader perspective, red)

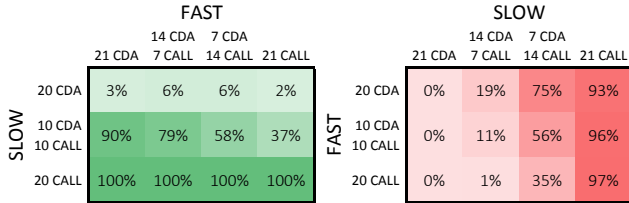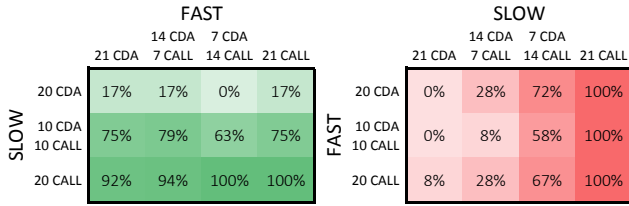| FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	33%	83%	100%
10 CDA 10 CALL	0%	13%	67%	100%
20 CALL	0%	0%	39%	100%

(c) Environment III

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	3%	6%	6%	2%
10 CDA 10 CALL	90%	79%	58%	37%
20 CALL	100%	100%	100%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	19%	75%	93%
10 CDA 10 CALL	0%	11%	56%	96%
20 CALL	0%	1%	35%	97%

(d) Environment IV

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	17%	17%	0%	17%
10 CDA 10 CALL	75%	79%	63%	75%
20 CALL	92%	94%	100%	100%

SLOW (FAST trader perspective, red)

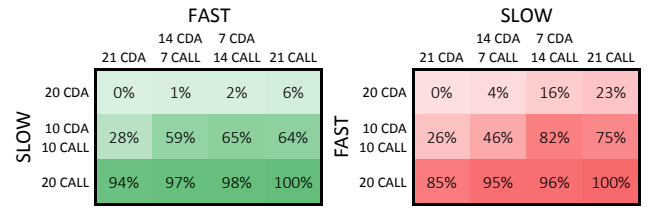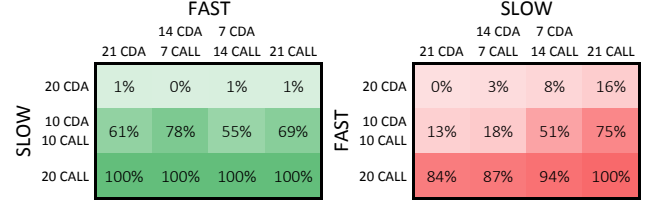| FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	28%	72%	100%
10 CDA 10 CALL	0%	8%	58%	100%
20 CALL	8%	28%	67%	100%

Figure 2: Basin of attraction for CALL, as characterized by best-response heat maps of complete subgames for each environment. The subgame for environment III has a 6×6 strategy space, with three strategies in each market, for each role; the subgames for the other environments have strategy spaces of size 4×4. The matrices on the left (in green) are from the perspective of a single SLOW trader; the matrices on the right (in red) are from that of a FAST trader. The rows in each matrix specify the distribution of same-role agents across the two markets, and the columns specify the cross-market distribution of other-role agents. Each entry in the heat map matrix gives the percentage of all other-agent profiles in which a single agent's best response specifies CALL. Heat map colors follow a scale where light corresponds to 0% and dark to 100%.

Figure 3

(a) Environment I

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	1%	2%	6%
10 CDA 10 CALL	28%	59%	65%	64%
20 CALL	94%	97%	98%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	4%	16%	23%
10 CDA 10 CALL	26%	46%	82%	75%
20 CALL	85%	95%	96%	100%

(b) Environment II

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	1%	0%	1%	1%
10 CDA 10 CALL	61%	78%	55%	69%
20 CALL	100%	100%	100%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	3%	8%	16%
10 CDA 10 CALL	13%	18%	51%	75%
20 CALL	84%	87%	94%	100%

(c) Environment III

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	2%	4%	2%	0%
10 CDA 10 CALL	78%	84%	67%	68%
20 CALL	100%	100%	100%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	6%	25%	26%
10 CDA 10 CALL	9%	10%	48%	82%
20 CALL	75%	69%	84%	99%

(d) Environment IV

FAST (SLOW trader perspective, green)

SLOW	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	1%	1%	1%	3%
10 CDA 10 CALL	72%	74%	70%	59%
20 CALL	100%	100%	100%	100%

SLOW (FAST trader perspective, red)

FAST	21 CDA	14 CDA / 7 CALL	7 CDA / 14 CALL	21 CALL
20 CDA	0%	3%	8%	18%
10 CDA 10 CALL	21%	31%	52%	90%
20 CALL	80%	96%	96%	100%

Figure 3: Basin of attraction for CALL, as characterized by best-response heat maps of sampled full-game profiles for each environment. Data presented is as for Figure 2.

full-game profiles based on the equilibrium mixture probabilities, with one simulation run per sampled profile. We successfully find at least one and up to six RSNE in each environment; each equilibrium has one to three strategies played with positive probability for a given role. There is at least one all-CALL RSNE in each environment; all but one environment has at least one all-CDA equilibrium.
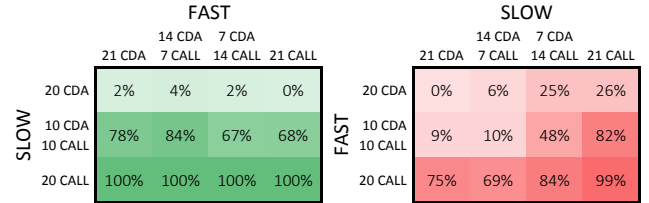
We find empirical support for the general welfare benefits of the CALL market, but primarily for SLOW traders: the mean total SLOW-agent surplus accrued over the all-CALL equilibria in a given environment is uniformly higher than that over the all-CDA equilibria in the 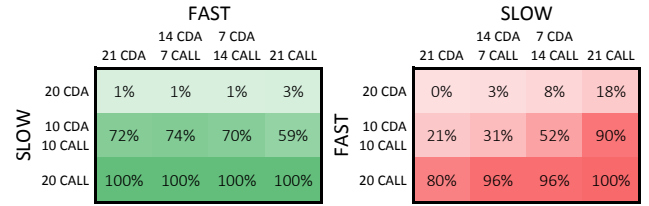same environment. Environments I and II have the same environment-agent population, but the lower mean reversion in the latter makes SLOW traders more susceptible to adverse selection. This is reflected in the significant reduction in total SLOW-agent surplus in environment II versus environment I. FAST traders accrue approximately the same level of surplus in both environments. We also observe that although the total welfare in environment I is close to the social optimum described in Section 4.3, increased adverse selection reduces overall surplus, and it is in this setting that the CALL provides significant welfare improvement over the CDA.

Within the same environment and with reduced mean reversion, FAST traders also generally shade their bids less in the frequent call market versus the CDA, as can be evidenced by reduced R_{mid} values. This effect does not hold for the SLOW traders, who shade approximately the same regardless of market type. The reduction in FAST-trader bid shading is indicative of the shift from a competition on speed in the CDA to a competition on price in the CALL.

We find at least one all-CDA RSNE in environments I through III. This is due to the low number of environment agents in these games. When there are only 4 environment agents in each market, as in environments I and II, the CALL market is not thick enough— sufficient volume is required for the call auction to deliver on its

Env	Total surplus	FAST				SLOW			
		market	surplus	R_{mid}	η	market	surplus	R_{mid}	η
I	27288	CALL	14469	129	1	CALL	12819	230	1
I	26697	CALL	14384	210	1	CALL	12314	486	1
I	27261	CDA	14598	250	1	CDA	12662	198	1
I	26785	CDA	14136	435	0.769	CDA	12649	250	1
I	25321	CDA	13502	418	0.943	CDA	11819	750	0.4
I	26133	CDA	13969	559	0.630	CDA	12165	500	1
II	21050	CALL	14697	703	1	CALL	6353	1250	1
II	21242	CDA	15355	710	0.448	CDA	5887	1250	1
III	19992	CALL	13790	644	1	CALL	6202	1250	1
III	20441	CALL	13909	500	1	CALL	6532	1111	1
III	19734	CDA	14483	750	0.4	CDA	5251	1250	1
IV	18067	CALL	12856	970	1	CALL	5211	1250	1

Table 2: **Role-symmetric equilibria for the four strategic market choice games, one per environment, calculated from the $(3, 3)$-player DPR approximation. Each row of the table describes one equilibrium found, including, for each role in the RSNE, the selected market mechanism (CALL or CDA) and the average values for total surplus of players in the role and for two strategy parameters: R_{mid} (the midpoint of ZIR range $[R_{\min}, R_{\max}]$) and threshold η. Values presented are averages over strategies in the profile, weighted by mixture probabilities. There is at least one all-CALL RSNE in each environment and one all-CDA RSNE in environments I to III, but we did not find any all-CDA equilibria in environment IV.**

promise of welfare improvement. But in environment IV, where $E = 42$, there is ample volume and order activity in the CALL market for traders to strongly prefer it over the CDA, hence we find no all-CDA RSNE in this environment.

Notably, we only find RSNE in which both FAST and SLOW agents select the same market. We can definitively rule out two-market equilibria—in which all agents in one role choose the CALL and all those in the other role choose the CDA—by exploiting an independence property of market choice games. In these games, the payoff for any given strategy depends only on the strategies of traders in the same market. We identify CALL-CDA equilibrium candidates by exploring four subgames for each environment. In each of these subgames, we limit the 21 traders in one role to a single strategy in the first market, while permitting traders in the other role to select any of the six strategies specifying the other market. In essence, we limit these subgames to one market and one role. We compute the equilibria in each of these subgames and form equilibrium candidates of the target form; we can then confirm or refute these candidates within the full strategy space.

For example, we explore a subgame with SLOW traders playing CALL strategies from Table 1 and FAST traders playing some strategy s_{CDA} in the CDA. We also explore a subgame with FAST agents playing CDA strategies and SLOW agents playing a fixed strategy s_{CALL} in the CALL market. Analysis of the first (second) subgame gives the equilibria for FAST (SLOW) traders in the CDA assuming no SLOW (FAST) traders are present. We can then form a CALL-CDA equilibrium candidate from any equilibrium in the first subgame (which specifies the strategies for FAST traders in the CDA) and any equilibrium in the second subgame (which specifies the strategies for SLOW agents in the CALL).

We refute all such candidate equilibria in all four environments, hence there are no RSNE in which all FAST traders are in the CDA and all SLOW traders are in the CALL. Such equilibria might be expected given that FAST traders benefit from picking off stale orders in the CDA. Ultimately, we find no such RSNE because the FAST traders benefit from being in the same market as the SLOW traders. The SLOW traders face greater risk of adverse selection in the CDA, so they select the frequent call market, followed close behind by the FAST traders.

4.6 Regret Analysis

We also evaluate the degree to which a trader is attracted to the CALL versus the CDA. To that end, we compute NE regret [23], which captures the loss of utility for a player who deviates from a Nash equilibrium to a specified strategy. The NE regret of a given strategy s is defined as the utility to the player in equilibrium less the payoff it accrues when it deviates to s. Accordingly, the NE regret of any equilibrium strategy is zero.

To compute the NE regret of deviating to the other market, we use the sampled surplus values in Table 2 to determine the per-agent surplus for a trader in a given role, and we subtract from that the payoff of the best-performing strategy in the other market. Again, we can exploit the independence of the two markets in our model, this time to determine the best other-market strategy. For example, for an all-CALL RSNE, we can measure the payoff to an agent that deviates to strategy s_{CDA} in the CDA via the payoff in any profile in which a single trader plays s_{CDA} and the other traders are in the CALL. We average the payoffs accumulated across all such profiles to determine the maximum-payoff other-market strategy for each RSNE, and we use these to compute the minimum NE regret for deviating to the non-RSNE market.

Our results are shown in Figure 4. SLOW traders generally have lower regret if deviating to the CALL from an all-CDA RSNE than if deviating to a CDA from an all-CALL RSNE. This is indicative of the greater loss they face if they leave the CALL market, as they are at high risk of being picked off by the faster traders in the CDA. The FAST traders, on the other hand, stand to lose more if they deviate from an all-CDA RSNE to the CALL, versus deviating to the CDA from an all-CALL RSNE, because their payoffs are based on exploiting their speed advantage over the SLOW traders. In short, SLOW traders would much rather stay in the CALL market, while FAST traders exhibit a stronger preference for the continuous market. We observe that FAST traders have universally greater regret than the SLOW traders; this is because they already accrue the lion's share of overall welfare, hence they have greater profits to lose. The negative regrets in our results are indicative of the limitations of the DPR approximations we use in deriving equilibria.

Also notable is that the best strategy when deviating to the CDA from an all-CALL RSNE is always the one strategy in which the
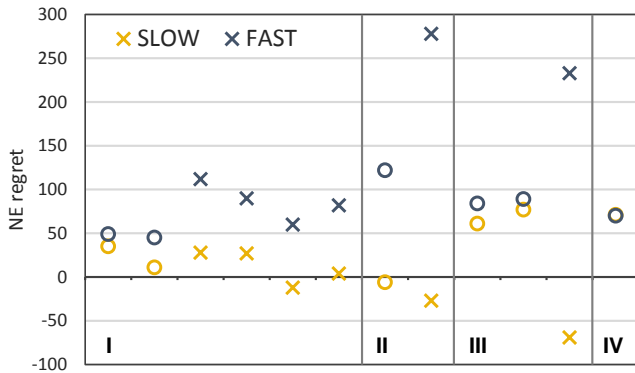
Figure 4: NE regret of equilibria in the four environments (I–IV), computed for each RSNE as the per-agent surplus in a role, less the maximum payoff possible if a player in that role deviates to the other market. Os indicate the NE regret in deviating to the CDA from an all-CALL RSNE; Xs indicate the NE regret in deviating to the CALL from an all-CDA RSNE. Note that the FAST and SLOW trader NE regrets in environment IV are overlaid as they are nearly identical.

threshold $\eta < 1$, regardless of environment or trader speed. Because environment agents arrive even more frequently than FAST traders, any player faces significant adverse selection if alone with the environment agents in the CDA market. Adopting a lower η decreases the tendency to leave standing orders, thus avoiding some of the pick-off risk.

5. RELATED WORK

There are only a few isolated examples of call markets in today's financial markets, most of which clear on a semi-frequent basis. The Taiwan Stock Exchange matches orders by call auction, with clears occurring every 60 to 90 seconds, depending on trading activity [26]. From mid-1998 to 2002, the Taiwan Futures Exchange employed a periodic call market to match orders; the clearing interval in the auction was incrementally reduced from 30 seconds to 20 and then 10, before finally being eliminated in favor of a predominantly continuous market mechanism [43]. More recently, both the London Stock Exchange and the NYSE have announced plans to introduce a midday batch auction in hopes of encouraging institutional investors to trade large blocks of shares on their exchanges [21, 39]. An intraday call auction has been standard for the past 15 years on Xetra, an electronic trading system for securities operated by Deutsche Börse [4]. Outside the equities space, batching to prevent exploitation by fast traders is currently in place on several foreign-exchange platforms. EBS, one of the largest currency trading platforms, has introduced a so-called latency floor, in which orders are batched in randomized clearing intervals (of lengths ranging from one to three milliseconds) in an effort to curb the advantages of super-fast traders [9]. The competing ParFX platform applies a randomized delay of 20 to 80 milliseconds to all order elements, and Thomson Reuters is currently trialling randomization of order execution on its foreign-exchange platform [9].

Several prior works have argued for frequent call markets as a means to end the latency arms race. Budish et al. [6] show that frequent batch auctions can potentially eliminate the latency arms race by reducing the value of very small speed advantages. Using millisecond-level exchange data, they demonstrate the breakdown of correlation between securities at high frequency, arguing that this phenomenon creates arbitrage opportunities that can be

exploited by the fastest traders. They develop a simple theoretical model of a single security traded on a continuous limit order book and show that in equilibrium, HFT profits come from fundamental investors via wider spreads. In a complementary analysis, Budish et al. [5] discuss the implementation details of frequent batch auctions in today's regulatory environment. Farmer and Skouras [13] likewise advocate frequent sealed-bid auctions as a means to end the technological arms race, suggesting that clear times be randomized. McPartland [29] proposes matching orders every half-second and switching to a cardinal time-weighted pro rata trade allocation formula to eliminate the advantage of speed in tie-breaking. This work also recommends randomization of the trade match algorithm, that is, matching orders to trade at a random time within each fixed-length clearing interval. Additional variants of randomized frequent call markets to deter HFT sniping have been suggested by Sellberg [37] and ISN [22].

Others have focused not on the role of call markets in mitigating the harmful effects of HFT, but on the difference in market quality offered in a discrete-time versus a continuous market. Pancs [30] compares three models—a dark pool, a continuous market, and a periodic call auction—focusing on both allocative efficiency and informational efficiency (which is high when observed transactions reveal traders' private information). This study finds that the periodic auction is more allocatively efficient than the continuous protocol when the demand for immediacy is low. Pellizzari and Dal Forno [31] use an agent-based model to compare the efficiency of a call auction (clearing only once), a continuous double auction, and a dealership. They find that the dealer market is the most efficient market structure of the three, offering the lowest volatility and the highest perceived gains by traders.

Baldauf and Mollner [2] develop a model of order anticipation to examine the impact of exchange competition on the spreads faced by investors. They study selective delay, an alternative trading mechanism in which cancellation orders are processed immediately but all other order types have a small delay, showing that selective delay reduces adverse selection by allowing liquidity providers to cancel stale quotes before being sniped by HFTs. In the specific setting of their work, they demonstrate that selective delay leads to the same outcome as a frequent batch auction. In another study, Baldauf and Mollner [1] consider a setting in which selective delay and frequent batch auctions result in different equilibrium outcomes. They show that a frequent batch auction in this case results in wider spreads than both selective delay and a continuous market.

Another relevant question is the frequency of clearing in a periodic call market, with some prior work suggesting that more frequent trading leads to increased volatility [25, 43]. Fricke and Gerig [15] argue that the optimal speed at which a security clears is related to volatility, trading intensity, and correlation of the security's value with other securities. They estimate that a range of 0.2 to 0.9 seconds is optimal. Du and Zhu [11] study the effects of trading speed on overall welfare via a series of uniform-price double auctions held at discrete time intervals. They find that the optimal trading frequency varies depending on trader speed: fast traders prefer a higher trading frequency, whereas slow traders prefer a lower frequency (and consequently thicker) market.

Much of the empirical work in the call auction literature examines the effects of discrete-time trading through natural experiments. For example, Kalay et al. [24] analyze the move of stocks on the Tel Aviv Stock Exchange from discrete-time trading to continuous trading. They argue that investors prefer stocks that trade continuously, based on observed losses in volume in stocks that trade by call auction. Webb et al. [43] examine the effect of the decision of TAIFEX, at the time a periodic call auction, to match

the trading hours of the Singapore Exchange (SGX), a continuous market, finding that this switch led to a statistically significant reduction in volatility on the SGX. They attribute these results to better price formation in the discrete-time market.

6. CONCLUSIONS

We examined strategic market choice in four environments with both FAST and SLOW traders who must decide between two market mechanisms: a frequent call market and a continuous double auction. We modeled this interaction as a game of market selection. We employed empirical simulation methods to compare the market type selected in equilibrium, the trading gains accrued, and the regret of deviating from equilibrium. We also analyzed best-response patterns in order to characterize the frequent call market's basin of attraction in multiple environments.

This study offers the first analysis of adoption of frequent call markets, framed as a question of strategic market choice. Our findings demonstrate that in equilibrium, SLOW-trader welfare is generally higher in the discrete-time market—further evidence that frequent call markets offer both increased gains from trade as well as protection from speed-advantaged HFTs capable of picking off resting orders. We also find strong evidence of a predator-prey interaction between FAST and SLOW traders. The FAST traders chase the SLOW traders into either market, whereas the SLOW traders flee their pursuers for the protection and efficiency gains of the frequent call market, as long as the CALL is sufficiently thick.

Overall, our results demonstrate that a frequent call market functions as an attractor for SLOW traders, as FAST traders are willing to follow the SLOW traders to either market. The predators (e.g., the HFT real-world counterparts to FAST traders) will always pursue their prey (e.g., institutional and retail investors), but in a frequent call market, the SLOW traders will be better protected from adverse selection and sniping. This suggests that frequent call markets in the wild could attract sufficient volume for viability, while deterring the wasteful pursuit of tiny latency advantages.

Several limitations should be taken into account in evaluating our results. We employ a simulation-based methodology for deriving our payoff estimates, and whereas we effectively reduce sampling error though the collection of large numbers of observations, there are limitations to the DPR approximations we use to compute equilibria (as evidenced by the observed negative NE-regret values). As we are not able to exhaustively search the entire strategy space, our equilibria are subject to refutation by other strategies, and additional qualitatively distinct equilibria are always possible.

Our trader strategy set is fairly limited, with both FAST and SLOW traders employing the same set of strategies. One particularly unrealistic restriction is that traders cannot alter their market choice once it has been made. Interesting extensions might include strategies that permit learning or adaptive selection of the market mechanism, or formulating an iterated form of our market choice game. Similarly, broader exploration of environments with, for instance, zero mean reversion in the fundamental, slower environment agents, as well as different clearing frequencies, could provide further insight on the relative attractiveness of alternative market mechanisms.

Acknowledgments

References

[1] M. Baldauf and J. Mollner. High-frequency trade and market performance, December 2014.

[2] M. Baldauf and J. Mollner. Trading in fragmented markets, May 2015.

[3] B. Biais, L. Glosten, and C. Spatt. Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264, 2005.

[4] M. Budimir. The Xetra intraday auction: Growing potential for strong price discovery. Technical report, Deutsche Börse, 2014.

[5] E. Budish, P. Cramton, and J. Shim. Implementation details for frequent batch auctions: Slowing down markets to the blink of an eye. *The American Economic Review*, 104(5): 418–424, 2014.

[6] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. Technical report, Booth School of Business, University of Chicago, 2015.

[7] L. Cardella, J. Hao, I. Kalcheva, and Y.-Y. Ma. Computerization of the equity, foreign exchange, derivatives, and fixed-income markets. *Financial Review*, 49(2): 231–243, 2014.

[8] B.-A. Cassell and M. P. Wellman. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent-Based Simulation XIII*, volume 7838 of *Lecture Notes in Artificial Intelligence*. Springer, 2013.

[9] J. Clark. Thomson Reuters to trial randomization on FX matching platform. *Euromoney*, Mar. 2014.

[10] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *17th International Joint Conference on Artificial Intelligence*, pages 1169–1176, 2001.

[11] S. Du and H. Zhu. Welfare and optimal trading frequency in dynamic double auctions. Technical report, National Bureau of Economic Research, 2014.

[12] N. Economides and R. A. Schwartz. Electronic call market trading. *Journal of Portfolio Management*, 21(3):10–18, 1995.

[13] J. D. Farmer and S. Skouras. Review of the benefits of a continuous market vs. randomised stop auctions and of alternative priority rules (policy options 7 and 12), 2012.

[14] D. Featherstone. Frequent batch auctions. Technical report, Optiver, December 2014.

[15] D. Fricke and A. Gerig. Too fast or too slow? Determining the optimal speed of financial markets, January 2015.

[16] D. Friedman. The double auction market institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 3–25. Addison-Wesley, 1993.

[17] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101 (1):119–137, 1993.

[18] D. K. Gode and S. Sunder. What makes markets allocationally efficient? *Quarterly Journal of Economics*, 112(2):603–630, 1997.

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

[20] M. A. Goldstein, P. Kumar, and F. C. Graves. Computerized and high-frequency trading. *The Financial Review*, 49: 177–202, 2014.

[21] B. Hope. NYSE Group planning midday auction for its stock markets. *The Wall Street Journal*, Jan. 2015.

[22] ISN. Toward a fairer and more efficient share market, Feb. 2013.

[23] P. R. Jordan, M. P. Wellman, and G. Balakrishnan. Strategy and mechanism lessons from the first Ad Auctions Trading Agent Competition. In *11th ACM Conference on Electronic Commerce*, pages 287–296, New York, NY, USA, 2010. ACM.

[24] A. Kalay, L. Wei, and A. Wohl. Continuous trading or call auctions: Revealed preferences of investors at the Tel Aviv Stock Exchange. *The Journal of Finance*, 57(1):523–542, 2002.

[25] L. H. P. Lang and Y. T. Lee. Performance of various transaction frequencies under call markets: The case of Taiwan. *Pacific-Basin Finance Journal*, 7:23–39, 1999.

[26] Y.-T. Lee, Y.-J. Liu, R. Roll, and A. Subrahmanyam. Order imbalances and market efficiency: Evidence from the Taiwan Stock Exchange. *Journal of Financial and Quantitative Analysis*, 39:327–341, June 2004.

[27] M. Lewis. *Flash Boys: A Wall Street Revolt*. W. W. Norton & Company, 2014.

[28] A. Madhavan. Trading mechanisms in securities markets. *The Journal of Finance*, 47(2):607–641, 1992.

[29] J. McPartland. Recommendations for equitable allocation of trades in high frequency trading environments. Technical report, Federal Reserve Bank of Chicago, 2013.

[30] R. Pancs. Comparing market structures: Allocative and informational efficiencies of continuous trading, periodic auctions, and dark pools, February 2013.

[31] P. Pellizzari and A. Dal Forno. A comparison of different trading protocols in an agent-based market. *Journal of Economic Interaction and Coordination*, 2:27–43, 2007.

[32] N. Popper. Flood of errant trades is a black eye for Wall Street. *The New York Times*, Aug. 2012.

[33] S. Rosov. Are frequent batch auctions a solution to HFT latency arbitrage? `http://blogs.cfainstitute.org/marketintegrity/2014/11/10/are-frequent-batch-auctions-a-solution-to-hft-latency-arbitrage`, 2014.

[34] D. K. Ross. Synchronized frequent batch auctions: A rebuttal. `http://www.pdqats.com/portfolio/synchronized-frequent-batch-auctions-a-rebuttal`, 2014.

[35] E. Schneiderman. Remarks on high-frequency trading & insider trading 2.0, Mar. 2014. New York Law School Panel on "Insider Trading 2.0—A New Initiative to Crack Down on Predatory Practices".

[36] R. A. Schwartz and L. Wu. Equity trading in the fast lane: The staccato alternative. *Journal of Portfolio Management*, 39:3–6, 2013.

[37] L.-I. Sellberg. Using adaptive micro auctions to provide efficient price discovery when access in terms of latency is differentiated among market participants, 2010.

[38] C. Sparrow. The failure of continuous markets. *Journal of Trading*, 7(2):44–47, 2012.

[39] P. Stafford. LSE to launch midday stock auction. *Financial Times*, November 2014.

[40] X. Vives. *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press, 2010.

[41] E. Wah and M. P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *14th ACM Conference on Electronic Commerce*, pages 855–872, 2013.

[42] E. Wah, M. Barr, U. Rajan, and M. P. Wellman. Public Comment in response to the U.S. Commodity Futures Trading Commission Concept Release on Risk Controls and System Safeguards for Automated Trading Environments. `http://comments.cftc.gov/PublicComments/ViewComment.aspx?id=59450`, December 2013.

[43] R. I. Webb, J. Muthuswamy, and R. Segara. Market microstructure effects on volatility at the TAIFEX. *Journal of Futures Markets*, 27:1219–1243, 2007.

[44] M. P. Wellman. Methods for empirical game-theoretic analysis (extended abstract). In *21st National Conference on Artificial Intelligence*, pages 1552–1555, 2006.

[45] M. P. Wellman. Countering high-frequency trading. `http://ai.eecs.umich.edu/people/wellman/?p=40`, July 2009.

[46] M. P. Wellman. *Trading Agents*. Morgan & Claypool, 2011.

[47] M. J. White. Enhancing our equity market structure, June 2014. Sandler O'Neill & Partners, L.P. Global Exchange and Brokerage Conference.

[48] B. Wiedenbeck and M. P. Wellman. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pages 931–938, 2012.