

Empirical Mechanism Design for Optimizing Clearing Interval in Frequent Call Markets

ERIK BRINKMAN, University of Michigan

MICHAEL P. WELLMAN, University of Michigan

Several recent authors have advocated for financial markets to move from continuous clearing to discrete or batched clearing, as a way to defeat the latency arms race: the never-ending quest for small advantages in time to access markets. How frequently should such a modern batch auction clear? We conduct a systematic simulation-based investigation on the relationship between clearing frequency and metrics of market quality, such as allocative efficiency, comparing the performance of discrete and continuous auction mechanisms under empirical equilibrium behavior of all participating traders. In effect we perform empirical mechanism design on frequent batch auctions. We find that in a wide array of environments, equilibrium efficiency is improved for small positive intervals but falls off dramatically when there are too few opportunities to trade. The result is a large range of batch frequencies that are near optimally efficient; this range is wider in thick markets.

CCS Concepts: • **Theory of computation** → *Algorithmic game theory*;

Additional Key Words and Phrases: call market, computational finance, agent-based simulation

1 INTRODUCTION

The defining feature of a *continuous-time* market mechanism is that orders trade as soon as they match. That is, the moment a buy order arrives that is compatible with an incumbent sell order (or vice versa), the trade executes and the two orders are cleared from the market. If two buyers submit orders compatible with an existing sell order, then whichever arrives to the market first is executed. A consequence of the continuous-time property is that there is no lower bound on the speed differential that could be pivotal in deciding the trade. In consequence, financial markets adopting the standard continuous-time mechanism—termed the *continuous double auction* (CDA)—foster a *latency arms race*, where traders may obtain significant benefit by achieving small access-time advantages over their counterparts.

The apparent first-order cost of the latency arms race has led several to advocate switching from continuous to discrete-time clearing mechanisms [3, 22, 23, 28]. In a discrete-time, or *call market*, orders are accumulated (or *batched*) over specified time intervals, with matches determined only at the end of each interval.

The call market mechanism has been known and in substantial use for a long time; in fact typical financial markets open each day with a call market. What is new is the opportunity to run call markets with short clearing intervals, say on the order of one second. The Chicago Stock

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'17, June 26–30, 2017, Cambridge, Massachusetts, USA.

ACM ISBN 978-1-4503-4527-9/17/06...\$15.00.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

DOI: <http://dx.doi.org/10.1145/3033274.3085153>

Exchange recently launched a similar mechanism called a SNAPSM auction,¹ which is roughly a temporary half-second call that occurs in an otherwise standard CDA market. The clearing interval in frequent call markets is intended to be long enough to undermine most of the benefit of tiny speed advantages, yet short enough to avoid imposing serious delays on economic activity. The literature discussing discrete-time markets uses a number of different names, including frequent batch auction [3], periodic auction [19], and frequent call market [28]; we simply refer to them as call markets, with the understanding that the clearing interval will be relatively short.

There are many ways a call market could trigger a clear, such as by volume or time, deterministically or randomly. Random choice of clear time has been proposed as a way to prevent a form of *sniping*, where a fast trader with an edge on information or market access can snipe a stale order. In a call market, this applies only right at the end of the clearing interval, say for a fraction δ of that interval. Randomization prevents such deterministic sniping, but allows the fast trader to snipe with a δ probability of success for the entire time. Thus randomization provides no protection benefit in expectation. It also poses the problem of trusting the mechanism to honestly randomize, which is not easy to audit.

For both of these reasons we focus on deterministic clearing, and the question of how long the interval should be. A specific choice of around one second, which has often been suggested as a reasonable interval, is quite arbitrary. As market designers, we would prefer a principled basis for deciding exactly how much time should elapse between market clears. We investigate this question by examining how the quality of the market—measured as allocative market efficiency—is affected by the clearing interval when traders are strategic and utility maximizing, playing according to an approximate form of Nash equilibrium.

Our model posits finitely many agents trading a single security in a single market, where agents respond to asymmetric private information as they repeatedly enter the market to buy or sell, conditioning on market observations. This model allows us to investigate how the presence of adverse selection, and market parameters that affect it, contribute to the optimal selection of a clearing interval. We do not model agents' response to public information or other important aspects of modern financial markets such as opportunities to invest in better information or faster market access.

We expect the duration of the clearing interval to have certain qualitative effects on a call market. If the clearing interval is sufficiently short, then a frequent call market should behave almost identically to a CDA.² Frequent clearing reduces transaction times at the expense of making hasty matching decisions and thereby potentially inefficient trades. If the clearing interval is sufficiently long, then the behavior approaches a *one-shot double auction* [20, 31]. Infrequent clearing produces more efficient trades at the expense of delayed execution. Waiting for the next clear can be costly, for example when other market actions serially depend on the present one. At some point, the efficiency benefit of longer clears will be outweighed by this time cost. The question then becomes, is there some clearing interval where the benefit of more efficient trades outweighs the loss of immediacy, or is the most efficient call market the one that behaves like a CDA?

Our market model is too complex to analyze using traditional means, so we restrict agents to a finite set of strategies, and use simulation in conjunction with game-theoretic analysis [32] to find approximate equilibria in a CDA market and in call markets with a discrete set of clearing intervals. In effect, we perform *empirical mechanism design* [25] with the goal of maximizing allocative

¹<http://www.chx.com/snap/snap-auction-cycle.html>

²Even at infinite frequency, the behavior is not exactly identical, because the markets price differently. CDAs price at the incumbent order price, whereas call markets price at the midpoint of the clearing range.

efficiency on a much richer model than would be possible to analyze with standard economic analysis.

Our model indicates that frequent batch auctions do not always improve market efficiency, but in markets with low adverse selection, lengthening the clearing interval produces greater efficiency until a point where agent impatience dominates and efficiency falls sharply. We also find that the benefit of switching to a frequent call market is significantly improved in thick markets. These conclusions suggest practical considerations that should be taken when implementing such a mechanism.

2 RELATED WORK

Wah and Wellman [28] and Budish et al. [3] were among the first to propose call markets as a design response to the negative effects of latency arbitrage. Both papers demonstrate that continuous clearing creates arbitrage opportunities that arise at imperceptible frequencies, and that exploiting them reduces market liquidity and efficiency. Budish et al. [3] also roughly calculate what batch interval should be sufficient to prevent the latency arms races for various calibrations of their model. Their results yield a lower bound of one millisecond to one second depending on assumptions about the information arrival process and the magnitude of speed improvement available to high-frequency traders (HFTs). Other studies have also found benefits to frequent call markets. Schnitzlein [21] performed a laboratory experiment comparing a call markets and CDAs and found that the call market was more liquid and better for noise traders, while not harming price efficiency. Baldauf and Mollner [1] compared CDAs, frequent call markets, and delayed CDAs³ in a model with costly research. Both frequent call markets and delayed CDAs performed at the Pareto frontier of price discovery and spread, but traditional CDAs did not. Madhavan [19] considered a model where risk-averse traders start with some initial endowment of both a risky asset and cash and found that, in equilibrium, call markets offered greater price efficiency than CDAs. Follow-up studies have investigated how frequent call markets could coexist with traditional CDAs [18, 27], have outlined practical considerations necessary to implement a frequent call market compliant with current regulations [2], and present some tangential evidence of a smooth tradeoff between immediacy and batching [29]. All these have relegated the question of optimizing call frequency to future work.

In response, several recent papers have developed stylized analytical models that propose definitions of an optimal clearing interval. Fricke and Gerig [10] studied a financial market model that is very similar to the model used in Garbade and Silber [12]. Like the earlier paper, the authors attempted to determine the optimal clearing frequency based on a version of *liquidity risk*. The authors then calibrated their model with aggregate current market data and estimated that the optimal clearing interval for current S&P 500 securities is somewhere between one and three seconds. Du and Zhu [8] solve for Nash equilibrium in linear strategies of a model with a divisible good and shocks to both common and private information. In this model, faster auctions allow agents to more quickly respond to information, but also cause them to be more aggressive with their demand schedule. A key result is that when information arrives via a Poisson process, then optimal clearing interval can be much longer than the average arrival rate of information, except when there are few traders or a large amount of adverse selection. Their approach is similar to ours, but several aspects of the model, including divisible units, private value shocks, and maintained order books, make it difficult to directly compare results. Haas and Zoican [15] propose a model in which they measure market liquidity as a function of clearing interval in equilibrium. Their market model has two types of participants: HFTs and liquidity agents, where HFTs respond to random information and the liquidity traders are impatient and respond to random private value shocks.

³CDAs with delayed submission but immediate cancels.

Their model parallels Budish et al. [3] until the clearing interval gets long enough, at which point traders’ impatience gives out and longer clearing intervals cease to improve liquidity. The authors find that liquidity may still be worse in these call markets than in traditional CDAs. Each of these models presents a slightly different view on the role and effect of frequent call markets, which is necessary given the complexity of the modern financial system.

3 MARKET MODEL

Our experiments employ a parameterized financial market model, related to the model in Wah et al. [27]. This model features many of the same qualitative phenomena found in today’s financial markets, including adverse selection, multi-unit holdings, and agents with the ability to buy and sell. We model a single security traded in either a CDA or a call market. Agents submit limit orders representing the maximum (minimum) price an agent is willing to buy (sell) a single unit. The market reveals price quotes reflecting the best outstanding orders after the most recent clear; other bids in the market are not visible to agents, and no new information is provided between clears. During a market clear, the mechanism determines the range of prices that would match supply and demand, and executes all compatible trades at a price in this range. If there is an imbalance of supply and demand at the clearing price (due to multiple units with the same price), the ties are broken first by orders that arrived in earlier clears, and then in a random manner. CDAs and call markets differ only in when they clear, and how they set transaction prices. CDAs clear after every market interaction at the price of the incumbent order; call markets clear at discrete intervals at the midpoint of the feasible price range. In this model, prices and times are fine-grained but discrete, and run until a finite horizon T .

3.1 Valuation Model

Each agent has an individual valuation for the security composed of private and common components. The common component is represented as a stochastic *fundamental value*, which determines the common utility of held positions at the end of time. Let f_t denote the fundamental value for the security at time t . The fundamental is generated by a mean-reverting stochastic process:

$$f_t = r\bar{f} + (1-r)f_{t-1} + s_t, \quad f_0 = \bar{f}, \quad s_t \sim \mathcal{N}(0, \sigma_s^2). \quad (1)$$

s_t is a random Gaussian shock at time t . Reversion rate $r \in (0, 1)$ specifies the tendency by which the fundamental reverts back to the mean \bar{f} ; $r = 1$ would correspond to a fundamental process of i.i.d. Gaussian draws around the fundamental mean, while $r = 0$ corresponds to a martingale Gaussian fundamental.

The time-varying fundamental presents agents with an issue of *adverse selection*—trading with agents who have better information—as standing orders reflect outdated information from the time submitted. If the fundamental shifts significantly, subsequently arriving agents are more likely to transact with orders on the side opposite the direction of change. That is, a positive price shock will tend to trigger transactions with stale sell orders, and negative price shocks with stale buys. Strategic agents who leave orders in the order book must shade to account for these events. The degree of adverse selection in our model depends on the fundamental shock variance σ_s^2 and the degree of mean reversion r . With lower mean reversion or higher shock variance the expected shift in the final fundamental estimate after a jump in the fundamental is greater, exposing agents to higher adverse selection.

The private component of an agent’s valuation represents its individual motivation for trading this security, for example needs for consumption or investment, and portfolio management. Differences across agents on these motivations are the sole basis for gains in trade, and naturally arise from

heterogeneous preferences, time in life cycle, risk attitudes, differences in hedging needs, or other reasons. An agent's private value is represented as a vector θ containing the agent-specific marginal utility for acquiring one more unit, relative to the fundamental value—similar to the model of Goettler et al. [14]. The vector has length $2q_{\max}$, where $q_{\max} > 0$ is the maximum number of units the agent can be long or short at any time. Element $\theta_q, q \in [-q_{\max}, q_{\max}] \cap \mathbb{Z}$ is the incremental benefit, over the fundamental, to the agent for gaining one unit of the security given current position q , where positive q indicates a long position. We generate an agent's private value vector by sampling $2q_{\max}$ values from $\mathcal{N}(0, \sigma_{pv}^2)$, sorting the values in descending order, and setting each θ_q respectively. This process ensures agents' valuations exhibit diminishing marginal returns.

At the end of the simulation, an agent's total surplus is the combination of their cash holdings p (from transaction prices), their final position q , and the value of the fundamental at the end of trading f_T .

$$\text{agent surplus} = p + qf_T + \begin{cases} \sum_{i=0}^{q-1} \theta_i & \text{if } q > 0 \\ 0 & \text{if } q = 0 \\ \sum_{i=q}^{-1} -\theta_i & \text{if } q < 0 \end{cases}$$

Since the cash and fundamental terms cancel out in exchange, the total surplus achieved is a result of the differences in transacting agents' private values.

3.2 Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs [7, 11, 31]. We consider trading strategies that are variants of the *Zero Intelligence* (ZI) family [13]. The ZI strategy is exceedingly simple, but often employed in agent-based study of financial markets [4, 17], including recent AI studies of market making [5, 30], because they have been found to generate realistic patterns of market behavior [9]. Though ZI agent instances are typically outperformed by more sophisticated alternatives [24, 26], by using game-theoretic selection to set strategy parameters we can produce highly competitive behavior for a given market environment [34]. In rough effect, game-theoretic equilibration serves the parameter-tuning function accomplished through adaptation and evolutionary search in the *ZI Plus* strategies of [6].

In our market model, agents get information about the fundamental according to an independent geometric process with probability λ of observing the fundamental (f_t) at any specific time. Upon observing the fundamental, agents observe the price quote, withdraw any outstanding orders, and have the opportunity to submit new ones. If multiple agents receive information at the same time, they act in a random order. Agents are assigned on each arrival to either 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 their maximum position (q_{\max}) (either long or short).

An agent assesses its expected valuation at the time of market entry t , using an estimate \hat{f}_t of the terminal fundamental f_T . The estimate is based on the current fundamental, f_t , adjusted to account for mean reversion:

$$\hat{f}_t = (1 - \rho) \bar{f} + \rho f_t, \quad \rho = (1 - r)^{T-t}.$$

The agent then submits a bid *shaded* from this estimate by a random offset—the amount of expected surplus it demands from the trade. The amount of shading is drawn uniformly from the range $[U_{\min}, U_{\max}]$.

The restrictions that agents are assigned to buy or sell randomly, shade in a uniform interval, and shade from the same distribution at every arrival are consistent with the original definition

of ZI [13]. Given such limitations, the ZI equilibria cannot be considered true equilibria of the unrestricted game. Nevertheless, evidence for the efficacy of the ZI strategy set, particularly with parameters tuned by equilibration, suggests that it provides a reasonable basis for evaluating market design choices.

We extend the traditional ZI strategy by introducing 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. Thus, an agent’s strategic choice is determined by three parameters: U_{\min} , U_{\max} , and η .

Since the ZI family of trading strategies was conceived for CDA markets, it stands to reason that similarly effective call market strategies would have a different parameterization. To our knowledge, there is no prior literature analyzing strategies for periodic call markets. We propose a simple extension to the ZI class of strategies that allows agents to condition their shading on their location in the clearing interval. Our extension provides the first step in exploring effective call market strategies. An agent that arrives close to a clear has better information about the fundamental than the other agents with orders in the market, and as a result, faces lower adverse-selection. Instead of being parameterized by only U_{\min} and U_{\max} , this strategy class has four shading parameters $U_{\min i}$, $U_{\min f}$, $U_{\max i}$, and $U_{\max f}$, where i and f stand for *initial* and *final* respectively. When an agent arrives, it linearly interpolates its effective U_{\min} and U_{\max} based on its position in the interval. That is, if an agent arrives α fraction of the way through the interval, then its demanded surplus is drawn uniformly from the range $[(1 - \alpha)U_{\min i} + \alpha U_{\min f}, (1 - \alpha)U_{\max i} + \alpha U_{\max f}]$.

3.3 Environments

Characterizing all possible model parameter settings under all possible agent strategies is infeasible and would be largely uninteresting. Many possible parameter settings are degenerate—they provide no opportunities to trade, or insufficient information to trade safely, or fail to present strategic tradeoffs—for example, as fundamental shock variance diminishes, adverse selection vanishes. We obtain an understanding of how different aspects of a market affect optimal mechanism design by analyzing many discrete parameter settings. We refer to each distinct point in parameter space as an *environment*.

Each environment we study shares a base set of parameter settings. These fixed parameters either have an unimportant effect on the simulation, or can be appropriately scaled by settings of the parameters we do vary. The fundamental mean \tilde{f} is 10^9 so that boundary effects at zero never happen in practice. The private value variance σ_{pv}^2 is 5×10^6 and the maximum position q_{\max} is 10 so that the importance of public and private value components are loosely balanced. With $q_{\max} = 10$, there is a 0.0002% chance that an agent’s optimal position is 10, thus agents rarely reach their maximum position. The arrival rate λ is 5×10^{-3} , which gives agents either 5 or 10 opportunities to trade on average.⁴ The remaining parameters are dimensions that we vary in our analysis: the number of agents (market thickness), the simulation length (number of trade opportunities), and the fundamental parameters (mean reversion, shock variance). We explore three different numbers of agents: 25, 58, and 238 respectively referred to as the thin $\downarrow N$, medium $-N$, or thick market $\uparrow N$. We investigate two different simulation lengths: 5000 and 10000 respectively short $\downarrow T$ or long $\uparrow T$. Finally, we consider two different levels of adverse selection: 2×10^{-2} mean reversion with 5×10^6 fundamental shock variance and 2×10^{-4} mean reversion with 5×10^4 fundamental shock variance, respectively $\downarrow A$ or $\uparrow A$. The fundamental parameters were chosen such that the distribution of the

⁴See settings of T in the next paragraph.

final fundamental is roughly equivalent⁵ in the long time horizon setting, but the variance of the difference in final fundamental estimates before and after the initial time step—a loose measure of the benefit for obtaining recent information—are markedly different, roughly 0 in $\downarrow A$ but over 900 in $\uparrow A$. An environment is the combination of these three parameter settings, so describe an environment in shorthand as the combination of the associated symbols. For instance, $\uparrow N \downarrow T \downarrow A$ is the environment with 238 agents and time horizon 5000, with a fundamental that has mean reversion of 0.02 and shock variance of 5×10^6 .

The ZI strategy space defined in Section 3.2 is continuous, not necessarily compact, and not differentiable with respect to utility. We discretize the strategy space in order to make these games amenable to normal-form analysis. We want a set of discrete strategies such that equilibria in the discrete game have low regret in the continuous game. How to guarantee this property is unknown, but we nevertheless select what we consider a reasonable set of discrete strategies by starting with a set of feasible strategies, and then adding best responses to initial equilibria from a significantly larger set. This technique is similar to that described by Jordan et al. [16]. The strategies we ultimately consider for equilibrium analysis are listed in Table 1.

Table 1. Strategies considered for equilibrium analysis. See Section 3.2 for parameter details. The first strategy is *truth-telling*, and the last strategy is effectively *fill-or-kill* where agents demand between 29 and 30 surplus.

U_{\min}	0	0	0	0	0	0	0	0	125	250	500	500	1000	29000
U_{\max}	0	125	250	500	1000	1500	2500	2500	1000	500	1000	1000	2000	30000
η	1	1	1	1	1	0.6	0.4	1	0.4	1	0.4	1	0.4	0.001

In order to test the effect of tailoring strategies to the call mechanism, we also found equilibria in environments where agents have access to a select number of call strategies. In the same way we chose a discrete set of standard strategies, we chose this set of call market strategies by looking for best responses from existing equilibria. Table 2 lists the three strategies agents had additional access to when finding equilibria with call strategies. All of these strategies exhibit decreasing levels of shading as the clear approaches, that is, as agents trade in a call market with better information, they shade less.

4 ANALYSIS

In order to understand how clearing interval affects market efficiency, we first demonstrate the effects when agents are non-strategic—specifically, when agents are truth-telling. We measure efficiency by fraction of optimal social welfare: average total trader surplus divided by the expected maximum possible trader surplus for the corresponding environment. This appropriately gives more weight to sampled instances with higher achievable surplus.

When the clearing interval is sufficiently small, we expect efficiency—as well as most other statistics—to be almost identical for call and CDA markets. The market types still will not be identical due to their different pricing schemes, and the discreteness of time. When the clearing interval is sufficiently large, we expect efficiency to degrade, because traders lack sufficient opportunities to trade. In other words, longer clearing intervals increase the friction due to agents’ impatience for trading. We also expect that at some point between almost-continuous clearing and a single clear, the efficiency of the call market will exceed CDA efficiency. The optimal clearing interval would ideally balance improvement in matching from aggregation over time, and provision to traders of ample opportunities to trade. In the words of Haas and Zoican [15], when the clearing interval is

⁵ $\downarrow A$ has 1.26×10^8 variance versus 1.23×10^8 in $\uparrow A$.

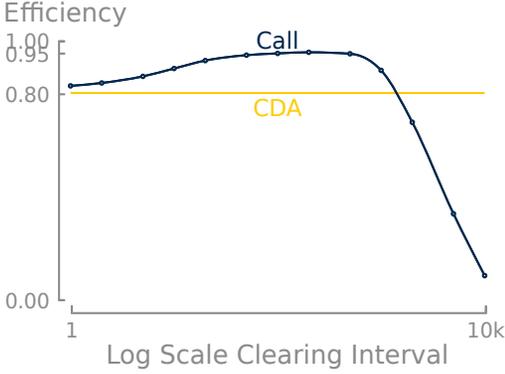


Fig. 1. Call market efficiency in the $-N\uparrow T\downarrow A$ environment as a function of clearing interval when agents are truthful. In this environment, there is a large region of clearing intervals where call market efficiency is near optimal, and significantly higher than CDA market efficiency. Call market efficiency is slightly higher even when clearing every time step because multiple agents can arrive in a single time step and still benefit from batching.

Table 2. Additional strategies considered for equilibrium analysis with call market strategies. See Section 3.2 for parameter details.

$U_{\min i}$	250	29000	39000
$U_{\max i}$	1250	30000	40000
$U_{\min f}$	0	20000	30000
$U_{\max f}$	1000	21000	31000
η	0.4	0.001	0.001

short, the friction due to trader impatience is minimized, when the clearing interval is long, the friction due to imperfect information is minimized, and ideally at some intermediate interval, the two are balanced.

We simulated environment $-N\uparrow T\downarrow A$ ten thousand times for each clearing interval when agents submit truthful bids. Figure 1 shows the results of that simulation, which matches our expectations. The call market is generally more efficient than the CDA, as the clearing interval increases, the agents see the allocation benefit of batched clearing, until the point that infrequent trading opportunities dominate the benefit of batching, and there is a sharp decline in efficiency.

Agents are unlikely to behave truthfully in real markets. Evaluating a market effect on all combinations of strategy choices and environments would be infeasible; moreover, the various strategic contexts are not equally relevant. Given the high level of competition between trading firms and the significant financial incentives, we are most interested in effects produced when agents are being *rational*. In other words, we wish to evaluate the impact *in Nash equilibrium*, where the agents are adopting the best strategies, given the setting and other-agent strategy selections.

Since the model we investigate is too complex to support analytical solutions, we compute empirical equilibria using a simulation-based process, known as *empirical game-theoretic analysis* (EGTA) [32]. EGTA employs systematic simulation of strategy profiles to sample from payoff distributions in a specified environment. The simulated payoff data is used to induce a normal-form game model, from which we identify symmetric mixed-strategy equilibria with respect to the strategies in Table 1. A symmetric mixed-strategy equilibrium is an equilibrium where every agent simultaneously and independently chooses a strategy from the same distribution over all possible strategies prior to play. However, direct simulation of all payoff data is prohibitive given that the number of profiles in the largest games we analyze is over 10^{21} . We employ two techniques to reduce the profile space necessary to identify at least one approximate equilibrium. First, we employ an aggregation technique known as *deviation-preserving reduction* [33] to approximate

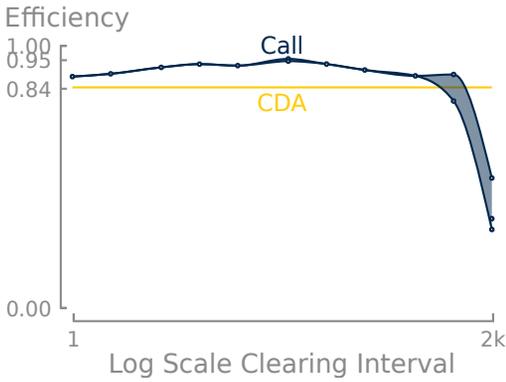
our games with many players as a reduced game with only four players. This reduction heuristic preserves regret of symmetric pure strategies, but only approximates regret of symmetric mixed strategies. The approximation typically works well when the payoff impact of single other agents is small. Second, we use a best-response search to attempt to confirm low-support equilibria, before sampling profiles for large-support equilibria.

Using these techniques we computed approximate empirical equilibria in the environments discussed in Section 3.3. Figure 2 contains plots of clearing interval versus equilibrium efficiency for the six environments with low adverse selection. Since environments may have multiple equilibria, we plot a range of efficiencies for all of the equilibria found at a given clearing interval. In each CDA environment, we found exactly one equilibrium. These methods are not perfect, and occasionally produce *artifacts*—for instance, the dip in Figure 2d—which may be attributable to discretizations imposed in modeling. Despite the few artifacts, the general trends are unmistakable. The trend in efficiency closely matches what was expected and what we see in Figure 1, but to a lesser extent.

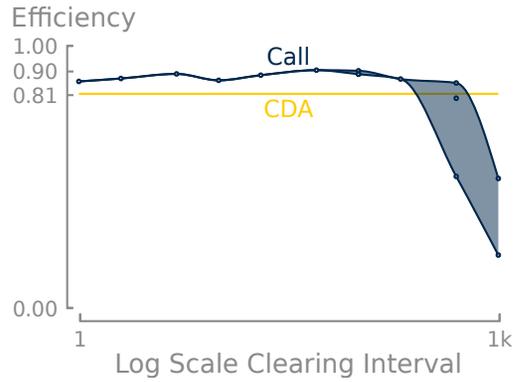
The decrease in equilibrium efficiency improvement of call markets relative to truth-telling is likely due in part to the surprising equilibrium efficiency of CDAs [35]. As clearing interval increases, call markets go from being quite similar to CDAs to being significantly more efficient. In most of the environments the optimal call market mechanism reduces inefficiency by half when compared to the CDA. This trend is stronger as the number of agents increases, which we attribute partially to effective batching that does not happen when the markets are thinner. This suggests that frequent call markets would actually benefit thick markets more than thin, matching the result of Du and Zhu [8]. It also suggests that given an ideal market setting, the specific choice of clearing interval is not important to improve efficiency conditioned on it being short enough to not interfere with the patience of investors.

There are other measures of market quality besides allocative efficiency, and the nature of our model allows us to compute arbitrary market statistics in equilibrium. Price discovery is a general measure of how closely the market represents the true price of a security; we calculate it as the *root-mean-square deviation* (RMSD) of the transaction prices relative to the unbiased estimate of the final fundamental at the time of transaction. A market with *worse* price discovery would have a *larger* RMSD. There is a fundamental tension between common and private information aggregation: as the call interval gets longer, orders in the market represent more outdated fundamental information and may be less accurate, but as the call interval gets longer more orders are present in each clear, and so more private information is aggregated in each transaction. Figure 3 presents the median price discovery RMSD for every low adverse-selection environment. CDAs exhibit far better price discovery in every environment except for thick markets, and as for efficiency, when the clearing interval is long, price discovery significantly degrades. The price discovery improvement in thick markets likely arises because batching large numbers of orders helps to reduce the noise caused by private valuations, while still ensuring that orders represent recent information.

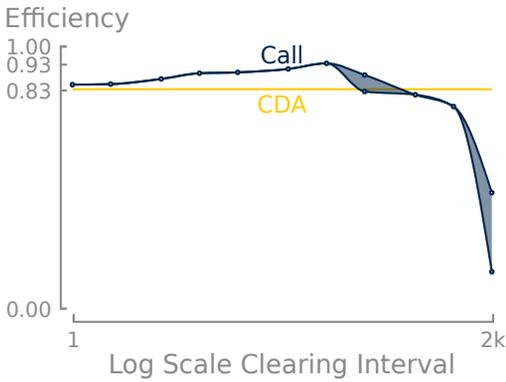
A common measure of liquidity is the market *spread*. The spread is the difference between the highest outstanding buy order and the lowest outstanding sell order. The size of the spread roughly corresponds to the transaction cost for impatient traders, as they must meet the quote to trade immediately. Since our model does not contain a dealer or market maker, there are times when spread is not defined, and we therefore report the median spread over all time. Increasing clearing interval should reduce the spread as agents have more protection from mispriced orders, since transactions happen at the midpoint of all orders instead of the incumbent price, and are therefore willing to price closer to their beliefs. This trend is present in Figure 4, where we show the equilibrium median spread for every low adverse-selection environment. The spread decreases



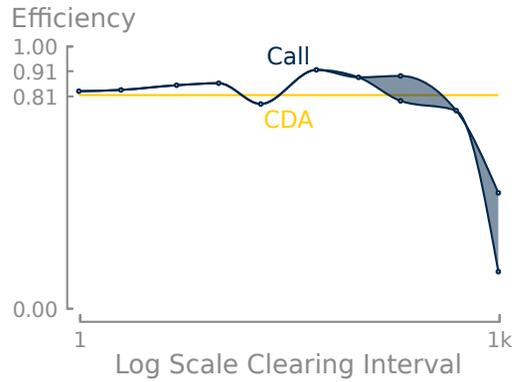
(a) Environment $\uparrow N \uparrow T \downarrow A$



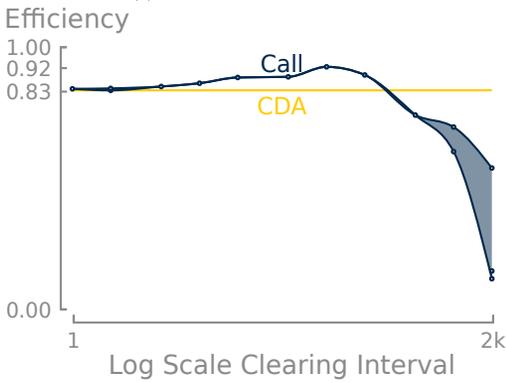
(b) Environment $\uparrow N \downarrow T \downarrow A$



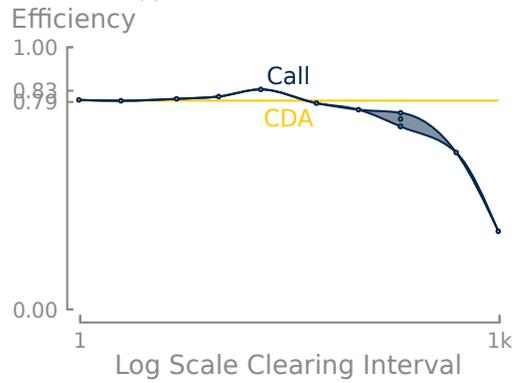
(c) Environment $-N \uparrow T \downarrow A$



(d) Environment $-N \downarrow T \downarrow A$



(e) Environment $\downarrow N \uparrow T \downarrow A$



(f) Environment $\downarrow N \downarrow T \downarrow A$

Fig. 2. Equilibrium efficiency as a function of clearing interval in six environments with low adverse selection. Rows from the top represent environments with 238, 58, and 25 agents respectively. The left column has a time horizon of 10000, while the right has a time horizon of 5000. Games where we found multiple equilibria have their efficiency represented as a range. These plots show the expected tradeoff between the benefits of batching and harm of limited execution opportunities. As the markets become thicker, the range of clearing intervals that are near optimal significantly increases.

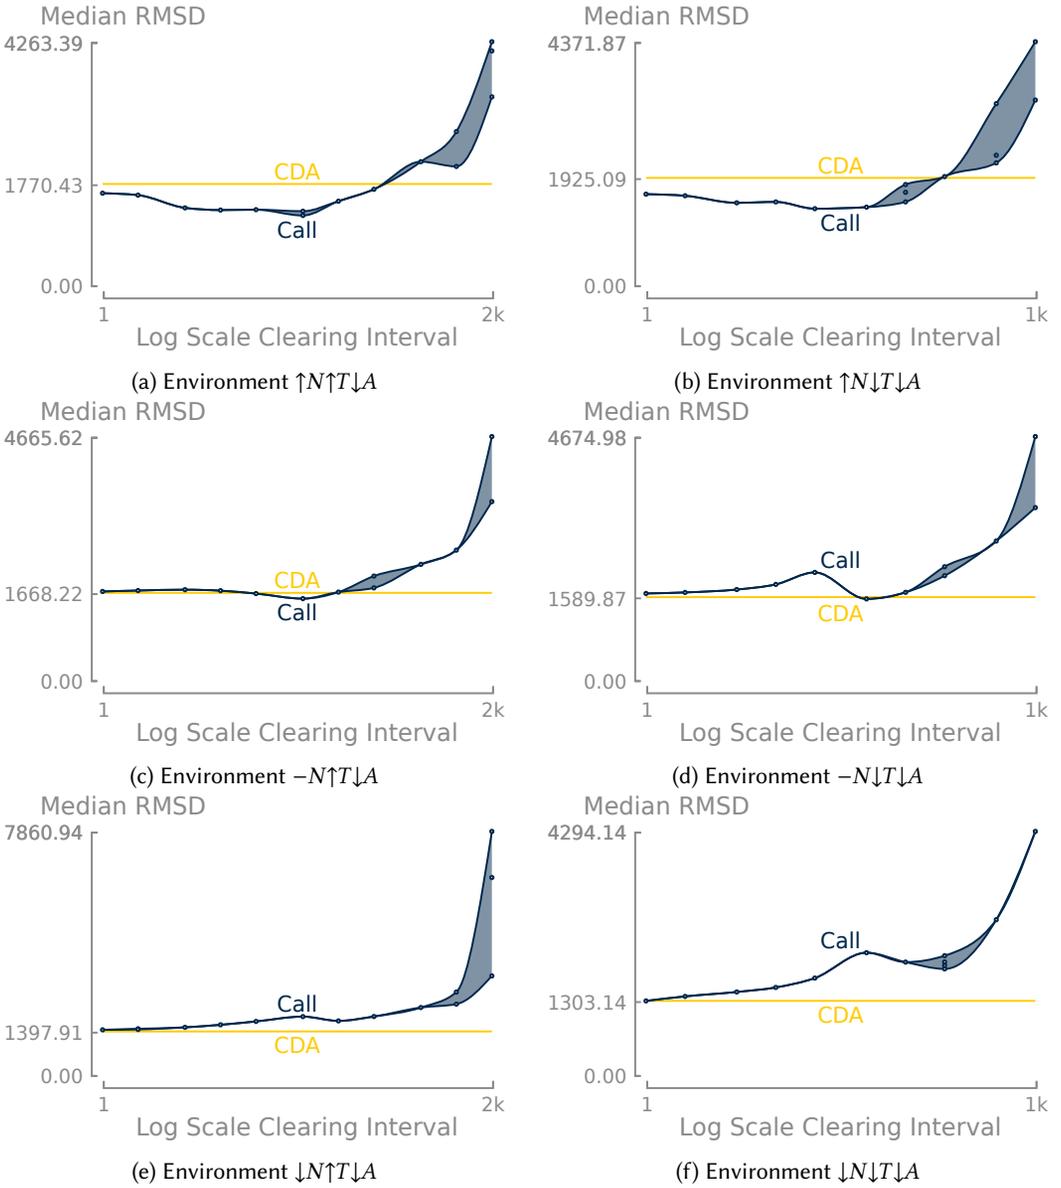


Fig. 3. Equilibrium price discovery as a function of clearing interval in low adverse-selection environments. Rows from the top represent environments with 238, 58, and 25 agents respectively. The left column has a time horizon of 10000, while the right has a time horizon of 5000. Higher RMSD indicates worse price discovery. Call markets improve price discovery only in thick markets with relatively short clearing intervals.

as the clearing interval increases, but the benefit disappears as the clearing interval becomes prohibitively long.

Our results indicate that long clearing intervals can result in low efficiency, poor price discovery, and large spreads, and the transition to this regime of poor market performance can be very fast.

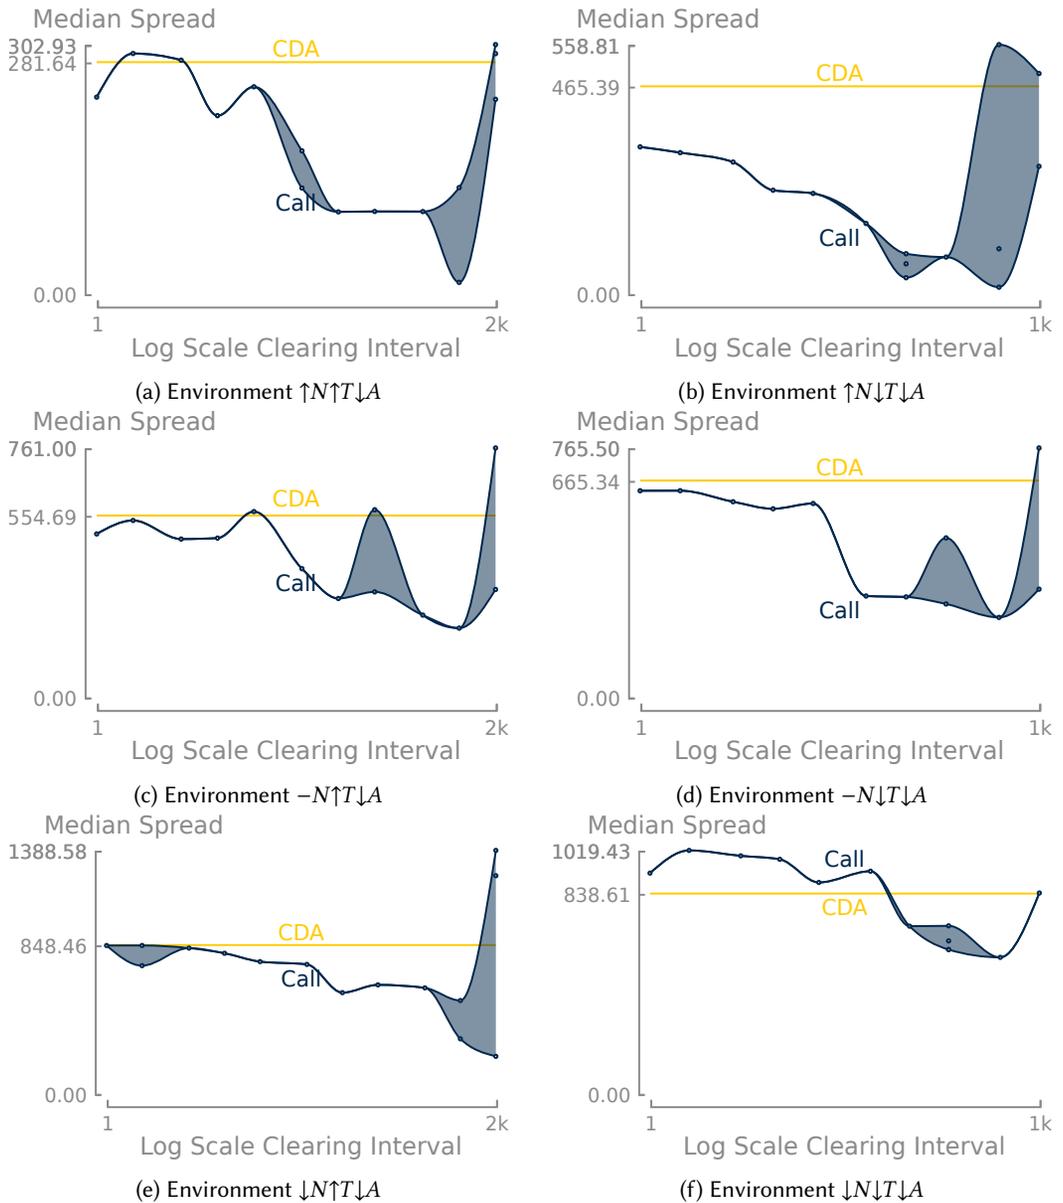
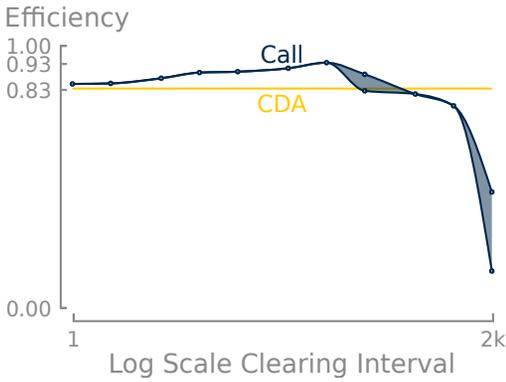
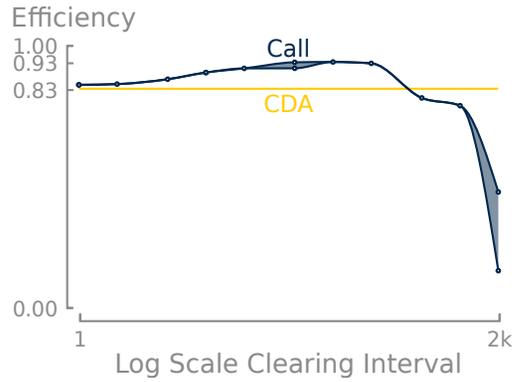


Fig. 4. Equilibrium median spreads as a function of clearing interval in low adverse-selection environments. Rows from the top are for environments with 238, 58, and 25 agents respectively; columns from the left represent time horizons of 10000 and 5000 respectively. Following intuition, longer clearing intervals decrease the spread; however, excessively long clearing intervals increase the spread.

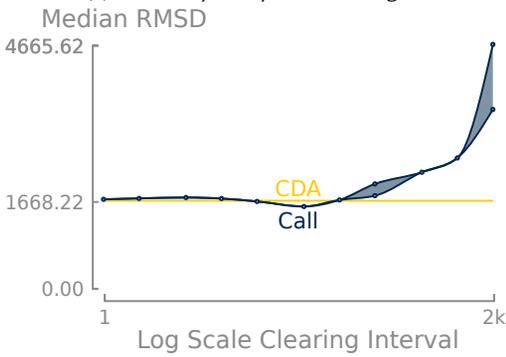
Our results also indicate that appropriate clearing interval selection in thick markets with low adverse selection can significantly improve performance over CDAs. However, these equilibria contain only strategies intended for use in CDAs. How robust are these results when agents can condition their strategy on their location in the clearing interval? Figure 5 shows the difference



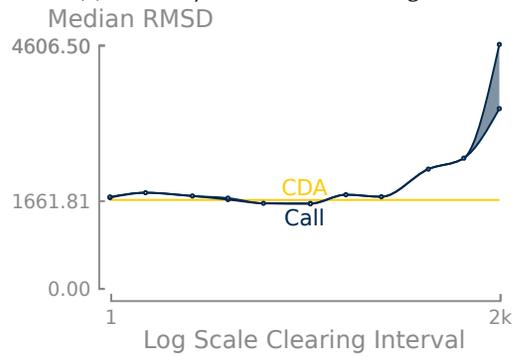
(a) Efficiency: Only CDA strategies



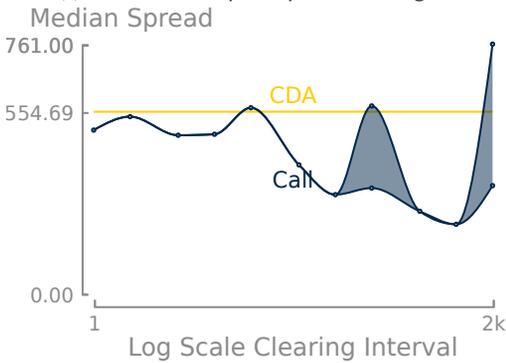
(b) Efficiency: CDA and call strategies



(c) Price Discovery: Only CDA strategies



(d) Price Discovery: CDA and call strategies



(e) Spread: Only CDA strategies



(f) Spread: CDA and call strategies

Fig. 5. How equilibrium statistics change when agents have access to call market specific strategies in environment $-N \uparrow T \downarrow A$. The left column contains equilibria when agents have access only to CDA strategies; the right column contains equilibria when agents additionally have access to call market specific strategies. Call market strategies have negligible effect on the qualitative trends in market quality.

between equilibrium market statistics when call market agents have access to call market strategies. 82% of equilibria and 90% of games had equilibria where at least one call market strategy was in support, which indicates that it can be effective to develop customized strategies for call markets.

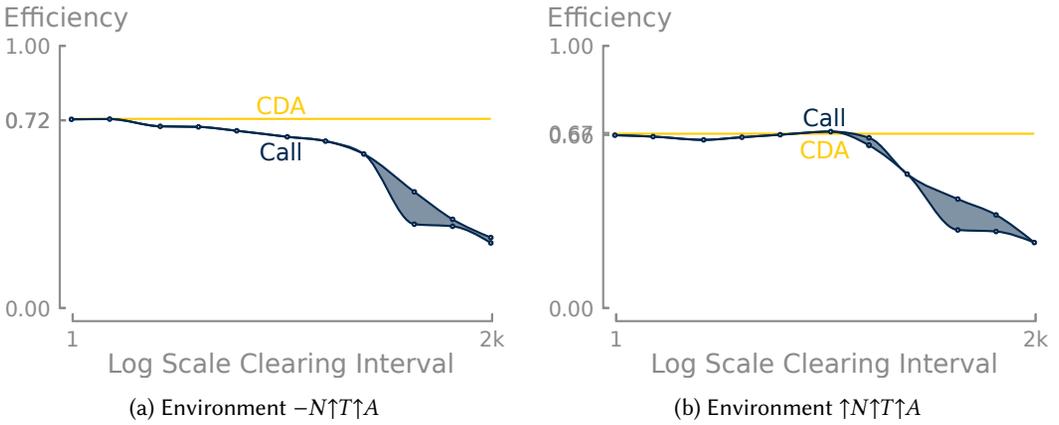


Fig. 6. Equilibrium efficiencies in two environments with high adverse selection. Even in a thick market (Figure b), short clearing intervals are on par with a CDA, but never surpass it.

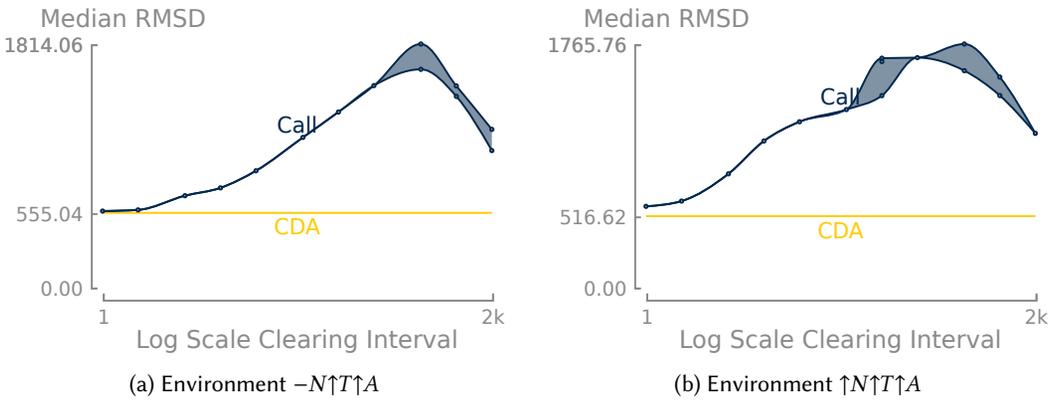


Fig. 7. Equilibrium price discovery with high adverse selection. Unlike with low adverse selection, price discovery in call markets with high adverse selection is consistently worse.

However, despite incorporating some call market strategies, qualitative market performance in equilibrium is essentially the same.

Our analysis to this point has considered only environments with low adverse-selection. Figures 6–8 show the same set of experiments for varying numbers of agents with a long time horizon and large adverse selection. We find that in high adverse-selection environments, the optimally efficient clearing interval approaches a CDA. Even when the number of agents is increased, short clearing intervals match, but do not exceed CDA efficiency. Thick markets do not restore price discovery as they do with low adverse selection, but the spread is generally tighter.

5 CONCLUSION

We employed empirical mechanism design to investigate the effects of clearing interval on market statistics. Given the significant difference between equilibrium trading behavior in a call market and in a CDA, we believe analyses of strategic interaction, like the results presented here, are

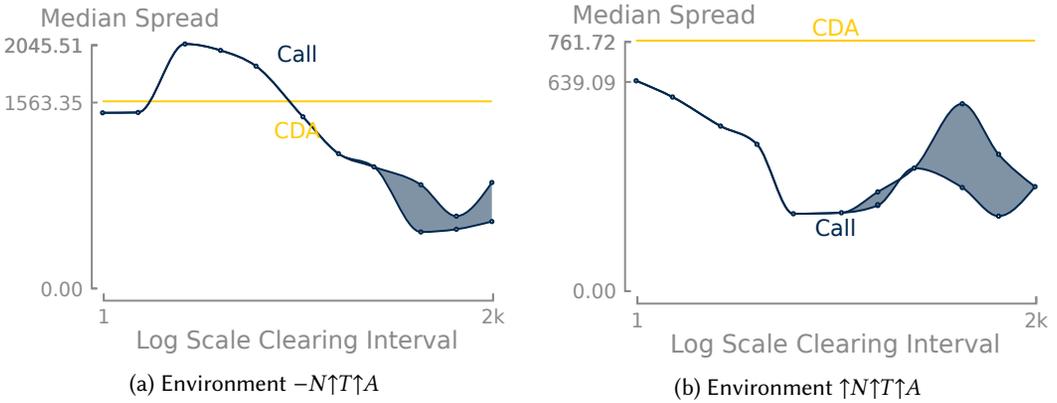


Fig. 8. Equilibrium spread with high adverse selection. Spreads are reduced at large clearing intervals, even with high adverse selection.

necessary to characterize expected behavior in call markets where little historic data exists. Our analysis has produced a number of unique insights.

In a first exploration of agent strategies geared toward call markets, we find evidence suggesting that even simple consideration of the call interval can provide gains. However, while these strategies appear in equilibrium, they do not qualitatively change equilibrium market quality in our experiments. Are the lack of qualitative differences due simply to the first-order nature of the call strategies, or does effective call market trading not need a significantly different strategy space? We leave this question for future work.

More to the point, there are a number of practical revelations about frequent call markets that we demonstrate. Environments with high adverse selection may be served more efficiently by CDAs than by frequent call markets. Otherwise, call markets tend to have some clearing interval region where they operate more efficiently, and this region is significantly wider in thick markets. Thick markets also universally saw improved spreads and price discovery than their thinner counterparts. These results roughly parallel those of Du and Zhu [8], who also found long optimal clearing intervals in thick markets with low adverse selection. Our results suggest that contrary to what might seem natural, the markets best suited to frequent call markets are thick with low adverse selection.

More generally, our methodology adapts readily to other metrics of interest, other market designs, and other agent behaviors, and provides a unique perspective on complex problems where models will always be stylized. For example, analyzing the effectiveness of random or volume-based clearing and other market-quality metrics can be easily factored into this model.

ACKNOWLEDGMENTS

Aaron Pollack contributed to the implementation and exploration of call-market strategies. Anonymous referees provided helpful suggestions for improving our exposition. This work was supported in part by grants from the National Science Foundation (IIS-1421391) and Future of Life Institute.

REFERENCES

- [1] Markus Baldauf and Joshua Mollner. 2015. *High-Frequency Trade and Market Performance*. Technical Report. University of British Columbia and Northwestern University. Available at SSRN, 2674767.

- [2] Eric Budish, Peter Cramton, and John Shim. 2014. Implementation Details for Frequent Batch Auctions: Slowing Down Markets to the Blink of an Eye. *American Economic Review* 104, 5 (2014), 418–424.
- [3] Eric Budish, Peter Cramton, and John Shim. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics* 130, 4 (2015), 1547–1621.
- [4] Timothy N. Cason and Daniel Friedman. 1996. Price Formation in Double Auction Markets. *Journal of Economic Dynamics and Control* 20, 8 (1996), 1307–1337.
- [5] Mithun Chakraborty, Sanmay Das, and Justin Peabody. 2015. Price Evolution in a Continuous Double Auction Prediction Market with a Scoring-Rule Based Market Maker. In *29th AAAI Conference on Artificial Intelligence*. 835–841.
- [6] Dave Cliff. 2009. ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computation* 13 (2009), 3–18.
- [7] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. 2001. Agent-Human Interactions in the Continuous Double Auction. In *Seventeenth International Joint Conference on Artificial Intelligence*. 1169–1178.
- [8] Songzi Du and Haoxiang Zhu. 2017. What is the Optimal Trading Frequency in Financial Markets? *Review of Economic Studies* (2017).
- [9] J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. 2005. The Predictive Power of Zero Intelligence in Financial Markets. *Proceedings of the US National Academy of Sciences* 102 (2005), 2254–2259.
- [10] Daniel Fricke and Austin Gerig. 2016. *Too Fast or Too Slow? Determining the Optimal Speed of Financial Markets*. Technical Report. Available at SSRN, 2363114.
- [11] Daniel Friedman and John Rust (Eds.). 1993. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley.
- [12] Kenneth D. Garbade and William L. Silber. 1979. Structural Organization of Secondary Markets: Clearing Frequency, Dealer Activity and Liquidity Risk. *Journal of Finance* 34, 3 (1979), 577–593.
- [13] Dhananjay K. Gode and Shyam Sunder. 1993. Allocative Efficiency of Markets with Zero-intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy* (1993), 119–137.
- [14] Ronald L. Goettler, Christine A. Parlour, and Uday Rajan. 2009. Informed Traders and Limit Order Markets. *Journal of Financial Economics* 93, 1 (2009), 67–87.
- [15] Marlene Haas and Marius A. Zoican. 2016. Beyond the Frequency Wall: Speed and Liquidity on Batch Auction Markets. *Paris Finance Meeting EUROFIDAI* (2016).
- [16] Patrick R. Jordan, L. Julian Schvartzman, and Michael P. Wellman. 2010. Strategy Exploration in Empirical Games. In *9th International Conference on Autonomous Agents and Multiagent Systems*. 1131–1138.
- [17] Blake LeBaron. 2006. Agent-Based Computational Finance. In *Handbook of Agent-Based Computational Economics*, Leigh Tesfatsion and Kenneth L. Judd (Eds.). Elsevier.
- [18] Zhuoshu Li and Sanmay Das. 2016. An Agent-Based Model of Competition between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs?. In *15th International Conference on Autonomous Agents and Multi-Agent Systems*. 50–58.
- [19] Ananth Madhavan. 1992. Trading mechanisms in securities markets. *Journal of Finance* 47, 2 (1992), 607–641.
- [20] Mark A. Satterthwaite and Steven R. Williams. 1989. Bilateral Trade with the Sealed Bid k -Double Auction: Existence and Efficiency. *Journal of Economic Theory* 48, 1 (1989), 107–133.
- [21] Charles R. Schnitzlein. 1996. Call and Continuous trading Mechanisms under Asymmetric Information: An Experimental Investigation. *Journal of Finance* 51, 2 (1996), 613–636.
- [22] Robert A. Schwartz and Liuren Wu. 2013. Equity Trading in the Fast Lane: The Staccato Alternative. *Journal of Portfolio Management* 39 (2013), 3–6. Issue 3.
- [23] Chris Sparrow. 2012. The Failure of Continuous Markets. *Journal of Trading* 7, 2 (2012), 44–47.
- [24] Gerald Tesauro and Rajarshi Das. 2001. High-Performance Bidding Agents for the Continuous Double Auction. In *Third ACM Conference on Electronic Commerce*. 855–872.
- [25] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P. Wellman. 2006. Empirical Mechanism Design: Methods, with Application to a Supply Chain Scenario. In *Seventh ACM Conference on Electronic Commerce*. 306–315.
- [26] Perukrishnen Vytelingum, Dave Cliff, and Nicholas R. Jennings. 2008. Strategic Bidding in Continuous Double Auctions. *Artificial Intelligence* 172 (2008), 1700–1729.
- [27] Elaine Wah, Dylan R. Hurd, and Michael P. Wellman. 2015. Strategic Market Choice: Frequent call Markets vs. Continuous Double Auctions for Fast and Slow Traders. In *Third EAI Conference on Auctions, Market Mechanisms and Their Applications*.
- [28] Elaine Wah and Michael P. Wellman. 2013. Latency Arbitrage, Market Fragmentation, and Efficiency: A Two-Market Model. In *14th ACM Conference on Electronic Commerce*. 855–872.
- [29] Elaine Wah and Michael P. Wellman. 2016. Latency Arbitrage in Fragmented Markets: A Strategic Agent-Based Analysis. *Algorithmic Finance* 5 (2016), 69–93.

- [30] Elaine Wah, Mason D. Wright, and Michael P. Wellman. 2017. Welfare Effects of Market Making in Continuous Double Auctions. *Journal of Artificial Intelligence Research* (2017).
- [31] Michael P. Wellman. 2011. *Trading Agents*. Morgan & Claypool Publishers.
- [32] Michael P. Wellman. 2016. Putting the Agent in Agent-Based Modeling. *Autonomous Agents and Multi-Agent Systems* 30 (2016), 1175–1189.
- [33] Bryce Wiedenbeck and Michael P. Wellman. 2012. Scaling simulation-based game analysis through deviation-preserving reduction. In *Eleventh International Conference on Autonomous Agents and Multi-Agent Systems*. 931–938.
- [34] Mason Wright. 2016. Using Reinforcement Learning to Validate Empirical Game-Theoretic Analysis: A Continuous Double Auction Study. (2016). arXiv:[1604.06710](https://arxiv.org/abs/1604.06710)
- [35] Wenjie Zhan and Daniel Friedman. 2007. Markups in Double Auction Markets. *Journal of Economic Dynamics and Control* 31, 9 (2007), 2984–3005.