

Shading and Efficiency in Limit-Order Markets

Erik Brinkman
University of Michigan
erik.brinkman@umich.edu

Michael P. Wellman
University of Michigan
wellman@umich.edu

May 17, 2017

Abstract

Strategic shading of bids is typically viewed as a factor degrading the efficiency of auction outcomes. In continuous double auctions (CDAs), however, strategic bid-shading can actually improve allocative efficiency, as it counteracts the inherent inefficiency of CDAs induced by vagaries in order arrival sequence. This has been demonstrated most clearly in a standard CDA model by Zhan and Friedman (2007). Through a systematic simulation-based game-theoretic study, we show that the phenomenon also exhibits in symmetric mixed-strategy equilibria of Zhan and Friedman's original game, and in equilibria of a richer dynamic model, more representative of financial markets. We find that the presence and degree of the effect can vary significantly, with the greatest benefit of strategic shading in markets with many traders or in the presence of factors such as adverse selection and limited trading opportunity.

1 Introduction

The *continuous double auction* (CDA) (Friedman, 1993) is a simple and well-studied auction mechanism, ubiquitous as the mechanism implementing limit-order markets for financial trading. A *limit order* is an offer to buy or sell a security, specifying the limits of price (maximum for buy or minimum for sell) at which the offeror is willing to trade. CDAs operate continuously, maintaining unmatched limit orders in an *order book*, and executing a trade immediately upon admitting a limit order matching one outstanding in the book. Outstanding orders are matched in order of best offered price. All major stock exchanges, as well as commodity and futures trading institutions, employ some form of CDA for the bulk of their trading activity, and have for a century or more. Given the centrality of this market institution, it is striking that some basic strategic properties remain hazily understood. Financial trading features dynamic interaction with incomplete information (both private and common value elements), in a setting where one's own bid can determine the price of the current transaction, and influence prices of subsequent trades. To date, auction theory has

not successfully tackled this combination of issues in a setting approaching the richness of financial markets.

It has long been understood that CDAs can exhibit *allocative inefficiency* (Gode and Sunder, 1997): at the end of trading the goods may not be held by those who value them the most. Prior studies in our own research group have also found efficiency gaps in continuous trading (Wah and Wellman, 2017). A primary reason, stated simply, is that when agent arrivals to the market are spread over time, the myopic matching procedure of the CDA may produce suboptimal allocations. Biais et al. (2005) survey an extensive literature on reasons for inefficiency, and ways in which concentrating trading in time (as in a call market) can improve efficiency. Modern proposals for online mechanism designs (Blum et al., 2006) aim to maximize efficiency, but redesign of entrenched continuous mechanisms is not generally an option.

Myopic matching leads to inefficiency even with truthful bidding. In actual markets, strategic bidders *shade* their bids from their true values to account for the potential effect of their bids on the transaction price. In a *one-shot* double auction setting, such shading can only degrade allocation quality, from the perfect efficiency that would be achieved with truthful bidders. Strategic financial agents should also be expected to shade their bids, so it bears considering how that affects allocative efficiency in a dynamic limit-order market.

This question was previously addressed within a standard CDA model by Zhan and Friedman (2007), who found that profiles of shading (markup) strategies in a restricted form of Nash equilibrium are highly efficient, and often yield better allocations than truth-telling. We replicate key parts of this prior work and extend it in two major directions. First, we conduct a more comprehensive game-theoretic analysis, evaluating a much larger space of strategy profiles and considering mixed as well as pure strategies in our search for Nash equilibria. Second, in addition to their simple model, we also investigate a richer family of market environments designed to capture key features of financial markets, including private and common valuation elements, significant dynamic structure, and a broader space of agent strategies. Importantly, the common value element of our market model introduces adverse selection, which adds an extra incentive for agents to shade their bids that is not present in independent private value models.

Like Zhan and Friedman, we evaluate profiles of shading strategies through simulation. Analytic solutions of CDA games are intractable beyond simple instances, and all the more out of reach for our richer class of financial trading scenarios. We conduct an extensive empirical game-theoretic analysis, to identify the direction of the effect on efficiency in the region of equilibrium trading behavior. Our results confirm that the efficiency improvements of pure-strategy equilibria found by Zhan and Friedman are also exhibited by mixed-strategy equilibria in their simple model and in more complex financial markets. However, shading equilibria are not always more efficient than their truth-telling counterparts. Only with a large number of agents, meaningfully limited agent arrivals, or large adverse selection do we find outcomes produced by strategic bidders consistently superior to the results of truth-tellers.




Figure 1: How shading can improve efficiency in the scenario of Example 1. Diamonds represent agent valuations, dashed intervals represent inefficient shading, and solid intervals represent efficient shading. When b_1 and s_2 shade enough, they prevent inefficient trades with s_1 and b_2 .

2 Stylized Examples of Shading’s Effect on CDA Equilibrium Efficiency

We develop insight on the effect of strategic shading on CDA efficiency through inspection of some simple CDA bidding scenarios. We measure efficiency as the ratio of the expected total surplus of a profile to the expectation of the maximum surplus possible—also known as the *competitive equilibrium (CE) surplus*. As we show, equilibrium shading can either increase or decrease overall efficiency when compared to truth-telling. In this section, we present an instance that demonstrates *how* shading can benefit efficiency, then we analyze an instance where the benefits are obtained in Perfect Bayesian Equilibrium (PBE), and put forth another instance where PBE shading degrades surplus.

Our first example was previously employed by Wah et al. (2017) to illustrate the allocative inefficiency of CDAs.

Example 1. Consider a market with two buyers and two sellers. The buyers have private values b_1 and b_2 , and sellers have private values s_1 and s_2 , such that values are ordered $b_1 > s_1 > b_2 > s_2$.

Suppose that the agents arrive at the market in order from greatest valuation to least valuation. This sequence is shown in Figure 1. If the agents indeed submit orders at their valuations (the diamonds), then buyer 1 trades with seller 1, yielding surplus $b_1 - s_1$, and buyer 2 trades with seller 2, yielding $b_2 - s_2$. If instead buyer 1 trades with seller 2, the total surplus is $b_1 - s_2$, which is socially optimal. The greedy matching of the CDA in this instance executes trades that preclude efficient allocation. Whether this happens depends on how the limit orders are sequenced. With bids priced at these valuations, a random permutation of limit orders has a two-thirds probability of being suboptimal.

Suppose instead that the agents shade their bids away from their true valuations. For simplicity let us say that they all shade symmetrically, and the

valuations are evenly spaced by $\Delta = b_1 - s_1 = s_1 - b_2 = b_2 - s_2$. If the agents shade more than $\frac{1}{2}\Delta$, then the inefficient trades will not happen. As long as b_1 and s_2 shade less than $\frac{3}{2}\Delta$, then these two still trade, and the allocation is efficient.

The inefficiency of this example is conditioned on the fact that the agents arrived in a particularly bad order, and they cannot reenter to submit new orders. Absent either condition both shading and truth-telling can be efficient. However, this example still elucidates a simple reason for why a little shading promotes efficiency in CDA markets. By definition, competitive equilibrium trades provide relatively high surplus to the agents involved, and are therefore tolerant to some amount of shading. In contrast, inefficient trades generate lower surplus and can be prevented by a moderate amount of shading.

Having demonstrated that non-shading strategies *can* produce inefficient outcomes, and that some shading can restore efficiency, the natural next question is what happens in equilibrium? Strategic agents will clearly shade bids away from their valuations. Will the equilibrium shading levels be sufficient to restore efficiency, or will they perhaps shade too much? Our next example is a stylized situation where PBE is more efficient.

Example 2. *Let there be two buyers and one seller. One buyer, BL , has valuation v_{BL} with $\Pr(v_{BL} = 1) = \Pr(v_{BL} = 2) = 0.5$. The other buyer, BH , has valuation v_{BH} with $\Pr(v_{BH} = 2) = \Pr(v_{BH} = 3) = 0.5$. The sole seller, S , has valuation $v_S = 0$. The agents arrive at the market in a uniform random ordering, and get to observe the order book and their position in the ordering.*

Details of the equilibrium and its derivation are left to Appendix A.1. A simple way to see what happens is to note that when $v_{BL} = 2$, the seller trades with whichever of BL or BH arrives first. When $v_{BL} = 1$, the seller is able to price out the low buyer, and so the seller trades with the high buyer regardless of arrival order. Thus $\frac{1}{4}$ of the time, the seller strategy improves efficiency over truth-telling, and never reduces it. In this example, agent rationality prevents the inefficient trade when the low value buyer arrives early but never prevents efficient trades. It turns out that strategic shading improves efficiency of this market from 0.86 to 0.95. However, equilibrium efficiency is not guaranteed to improve over truth-telling in general, as demonstrated in our next stylized example.

Example 3. *Let there be two buyers and one seller, all with i.i.d. uniform private values over the unit interval. The agents arrive in a uniform random ordering, and observe the current state of the order book, and the ordering of agents.*

Details of the equilibrium are left to Appendix A.2. The key aspect of this equilibrium is that when the seller arrives last, it always trades with the highest value buyer if it wants to. When buyers truthfully reveal their valuations, this is always maximally efficient. However, When the buyers are strategic they both shade significantly—to extract more surplus from the seller and to outbid each

other. Thus, when the seller arrives last, strategic play significantly decreases the probability of making the efficient trade. When the seller arrives first or second, improvements in trade efficiency roughly balance the reduced number of trades leaving efficiency in these orderings roughly on par with the efficiency of truth-telling. The equilibrium efficiency is 0.77, considerably less than the corresponding truth-telling efficiency of 0.89.

Of course, these are only particular instances; it is easy to construct other simple instances where strategic shading either helps or harms efficiency. It is therefore necessary to understand what aspects of the market cause equilibrium efficiency to be better or worse than truth-telling. In the remainder of this paper we explore the effect of shading in richer scenarios, an expanded form of the model from Zhan and Friedman (2007) and a model more representative of trading situations that arise in financial markets. We overlay game-theoretic reasoning on a systematic simulation-based process, to investigate the impact of strategic bidding on outcomes realized in approximate equilibria, across a range of market environments.

3 Prior Work: Zhan and Friedman

There are many prior studies that attempt to characterize the efficiency of CDAs, however most do so from the perspective of semi-strategic agents, such as human lab participants or fixed algorithms designed without concern for equilibrium (Cason and Friedman, 1996; Gjerstad and Dickhaut, 1998; Rust et al., 1993). In contrast, Zhan and Friedman (2007) address the question with respect to Nash equilibrium shading. In their model, there are N buyers and N sellers with uniform i.i.d. private values for a single unit. The authors considered three different values of N : 4 (thin), 10 (medium), and 100 (thick); and three classes of parameterized shading strategies: Standard, Exponential, and Shift. Strategies are defined for buyers b or sellers s , and map the agent's valuation v to its order price. Let c denote the width of the value distribution.

$$\begin{aligned} \text{standard}_b(v) &= v(1 - m) & \text{exponential}_b(v) &= ve^{-m} & \text{shift}_b(v) &= v - cm \\ \text{standard}_s(v) &= v(1 + m) & \text{exponential}_s(v) &= ve^{+m} & \text{shift}_s(v) &= v + cm \end{aligned}$$

To analyze their model as a normal-form game, Zhan and Friedman instantiated eleven discrete strategies, with shading amounts evenly spaced from 0 to 1, for each class. An agent's strategic choice is the amount of shading conditioned on a global shading class and market thickness. Agents arrive in a random ordering¹ and submit limit orders applying their shading strategies to their private values. In this model strategies do not consider market information, such as price quotes or transaction history; they are functions solely of private value.

Separate from their analysis of equilibrium shading, Zhan and Friedman used this model to explore how uniform non-strategic shading affects market

¹This description is slightly different from the original, but produces identical results.




Figure 2: Moderate shading reduces inefficiency from suboptimal trades (EM-inefficiency) without introducing significant inefficiency from missed optimal trades (IM inefficiency). Significant shading increases efficiency lost from not making optimal trades. The data presented is from a replication of Zhan and Friedman (2007, Figure 6).

efficiency. Their investigation of uniform shading in the thick-market standard-shading scenario follows the intuition from Example 1: that moderate shading mostly precludes inefficient trades. We replicated this experiment with our own implementation, and present the relationship between shading and efficiency in Figure 2 (which faithfully reproduces Zhan and Friedman’s Figure 6).² Let *IM inefficiency* refer to the inefficiency due to agents who would trade in competitive equilibrium do not (missing *intra-marginal* trades), and *EM inefficiency* refer to the inefficiency from agents who would not trade in CE but do (present *extra-marginal* trades). The figure shows that as symmetric fixed shading increases from zero (truth-telling), EM inefficiency significantly decreases, while IM inefficiency remains close to zero. A little before 0.3 shading, both sources of inefficiency are minimized, and then as shading increases more, the IM inefficiency significantly increases. This suggests that the majority of inefficient trades have relatively small margins, and are inhibited by a small amount of shading, whereas the efficient trades have larger margins and so are uninhibited by modest shading.

In addition to analysis of the market environment under uniform shading, Zhan and Friedman also empirically found *role-symmetric pure-strategy* Nash equilibria in the nine variations of this game. In a role-symmetric strategy profile, each player within a role (here buyer or seller) plays the same pure or

²This despite our use of a subtly different definition of expected efficiency, stated in the beginning of Section 2. Section 5 discusses this difference in more detail.

mixed strategy. Conveniently, each of their game instances had a single role-symmetric pure-strategy equilibrium, or a single ε -approximate one for small ε . Zhan and Friedman conclude from the efficiency of these equilibria that CDA equilibrium surplus is close to optimal surplus. However, their results for the thin market—with four buyers and four sellers—indicate that the equilibrium surplus, while high, can be much worse than truth-telling. The issue of non-universal improvement over truth-telling is not discussed much by Zhan and Friedman, but is the focus of our extended investigation. In Section 5 we present the rest of the results of our replication and extension of this work.

4 Empirical Game-Theoretic Analysis

Evaluating the effect of shading on efficiency for all combinations of strategy choices and environments would be infeasible; moreover, the various strategic contexts are not equally relevant. Generally speaking, we are most interested in the effect of shading by rational traders. In other words, we wish to evaluate the impact *in equilibrium*, where the agents are adopting the best strategies, given the setting and other-agent strategy selections.

Since the models we investigate are too complex to support analytical solutions, we compute empirical equilibria using a simulation-based process, known as *empirical game-theoretic analysis* (EGTA) (Wellman, 2016). EGTA employs systematic simulation of strategy profiles to generate samples of profile payoffs in a specified environment. The simulated payoff data is used to induce a normal-form game model, from which we identify role-symmetric mixed-strategy ε -equilibria with respect to the given finite strategy set. A role-symmetric mixed-strategy profile is defined by a distribution over strategies for each role. An ε -equilibrium is a profile such that the regret—the improvement in payoff for unilaterally deviating from the profile when all other agents sample from their role distribution independently—is less than ε .

In the environments we analyze, strategy spaces are continuous and not necessarily compact. We discretized the strategies to make the games amenable to modeling in normal form. For our replication of Zhan and Friedman (2007), we adopted the discretization of the original authors. For our new environments we wanted to choose a set of discrete strategies such that equilibria in the discrete game would tend to exhibit low regret in the continuous game. We started with a rough grid of feasible strategies, and then added a few best responses to initial equilibria, similar to the technique described by Jordan et al. (2010). Analysis of the game over discrete strategies provides no guarantees about regret with respect to the full continuous strategy space, but it does allow us to make precise statements about regret in this restricted normal-form game.

Even with discretized strategies, the number of profiles that would need to be evaluated to directly estimate a complete normal-form game is astronomical— 6×10^{57} for the largest game analyzed here. We adopt two techniques to reduce the number of profiles we sample. First, we employ a heuristic profile search that biases towards low-support equilibria. By sampling only a small num-

ber of strategies played in conjunction we aim to confirm equilibria by sampling payoffs from a small number of profiles. Second, we employ an aggregation technique known as *deviation-preserving reduction* (DPR) (Wiedenbeck and Wellman, 2012) to approximate a game with many players as a reduced game with a smaller number of players. DPR is a heuristic that preserves regret of symmetric pure strategies and approximates the regret of symmetric mixed strategies. The approximation typically works well if the payoff impact of single other agents is small. Since DPR provides no guarantees for the mixture regret in the unreduced game, we gather additional payoff samples from the equilibrium candidates’ mixture distributions and use bootstrapping to compute confidence intervals on regret in the unreduced game. This is a slight variation on the bootstrap regret technique of Wiedenbeck et al. (2014), and has been shown to produce accurate confidence bounds. None of the equilibria we found had large regret upper confidence bounds, so while DPR may have affected which equilibria we found, the statistical evidence suggests that it did not degrade the *quality* of equilibria we found. In addition, in games where we found both pure and mixed-strategy equilibria, the upper confidence regret bound on mixed-strategy equilibria was always close to the confidence bound for pure-strategy equilibria. Since pure strategy equilibria found by DPR are also equilibria in the unreduced game, this evidence suggests that the regret caused by DPR is small compared to the inherent variability of the models.

The heuristics these techniques employ mean that we have no guarantees about the coverage of the whole solution space. In a strict sense, that sacrifice is unavoidable. In infinite games, if the strategy space is not compact or the utility functions are not smooth, then Nash equilibria might not exist (Glicksberg, 1952). Finding *every* Nash equilibrium, even in a finite game, is NP-hard (Gilboa and Zemel, 1989). More pragmatically, sampling every profile from even a relatively small game might be prohibitively expensive, as the number of distinct profiles in a symmetric game grows exponentially. In spite of these difficulties, the combination of our analysis techniques allows us to use game theory to analyze strategic interaction between agents that is much richer than tractable analytic models can provide.

4.1 Profile Search

Even in a game with few players, a large number of strategies can make evaluating every profile through simulation prohibitive. Since it is computationally infeasible to guarantee that we have found every role-symmetric equilibrium, we adopt a heuristic search through profile space in an effort to find low-support equilibria. In order to prevent our results from being too dependent on equilibria found early in the search process, we adopt a set of heuristic criteria that promote thorough coverage of the profile space, while still biasing toward low-support equilibria.

Our search criteria rely on sampling from views of the game where only a

subset of the strategies is considered. We refer to these views as *subgames*.³ We consider a subgame *explored* if we have (i) evaluated every profile in the subgame, (ii) found at least one candidate equilibrium in the subgame, and (iii) for each candidate equilibrium, evaluated all of the one-player deviations to strategies outside the subgame. We consider a profile *evaluated* if we have simulated it enough times to obtain a reliable estimate of expected payoff for each strategy. A *candidate equilibrium* is an equilibrium of a subgame, but whose deviations outside of the subgame may be unevaluated, thus it may or may not be an equilibrium in the full game.

The stopping criteria for our profile search process are as follows:

Criterion 1. *We have found at least one role-symmetric Nash equilibrium.*

Criterion 2. *We have explored every subgame with a single strategy per role (i.e., every strategy in self-play by role).*

Criterion 3. *For every candidate equilibrium found in an explored subgame, we have explored the subgame formed by adding the best-response strategy to the support of that candidate, providing that the subgame size is below a threshold.*

To meet these criteria, we iteratively apply EGTA using a process adapted from the inner loop of the procedure defined by Wellman et al. (2013). Pseudocode for our profile search process is listed as Algorithm 1.⁴

We start with a set of subgames \mathcal{R} we require to be explored per the criteria. This set is initialized to the subgames comprising exactly one strategy per role. For each subgame in this set, we sample all of the profiles in the subgame and then compute Nash equilibria in the induced sample game using replicator dynamics and convex optimization.⁵ For every subgame equilibrium found, we sample every deviating profile with respect to the full-game strategy set. If none of those deviations is beneficial, then we have confirmed an equilibrium. Otherwise, if the deviating subgame is small enough, we add the best-response deviating subgame to the required set of subgames to explore. Other beneficial deviating subgames are kept as a backup to explore in case we have explored all of the required subgames, but found no equilibria.

There are two potential sources of incompleteness in this search process: non-exhaustive coverage of subgames, and failure to find equilibria within a subgame. We could modify Algorithm 1 to ensure finding an equilibrium, by replacing replicator dynamics and convex optimization with a complete equilibrium-finding algorithm, and modifying the definition of brs to include all deviations rather than only the beneficial ones. Neither of these measures has proved necessary for this study, as we successfully found at least one role-symmetric Nash

³Not to be confused with the notion of a subgame in extensive-form specifications.

⁴Our implementation of this procedure is available at <https://github.com/egtaonline/quiesce>.

⁵If neither method finds an equilibrium in a subgame, we gather more profile samples until it does, as failure to find an equilibrium is often a result of sampling noise in payoff estimates. However, this never happened in any game discussed here.

Algorithm 1 Profile Search

```
define br                                ▷ Returns the best-response subgame
define brs                               ▷ Returns all beneficial response subgames
define  $\xi$                                ▷ Subgame size threshold

 $\mathcal{C} \leftarrow \emptyset$                     ▷ Confirmed equilibria
 $\mathcal{R} \leftarrow \{\text{All one-strategy subgames}\}$  ▷ Required subgames
 $\mathcal{S} \leftarrow \emptyset$                     ▷ Explored subgames
 $\mathcal{B} \leftarrow \emptyset$                     ▷ Backup subgames

while  $\mathcal{C}$  is empty or  $\mathcal{R}$  is not empty do
   $S \leftarrow \text{pop}(\mathcal{R})$  or  $\text{pop}(\mathcal{B})$           ▷ Subgame to simulate
  if  $S \not\subseteq S' \forall S' \in \mathcal{S}$  then          ▷ Subgame was not explored
     $\mathcal{S} \leftarrow \mathcal{S} \cup \{S\}$ 
    Sample all profiles in  $S$ 
    for all equilibria  $Q$  of  $S$  found do
      Sample all deviating profiles from  $Q$ 
      if  $|\text{brs}(Q)| = 0$  then
         $\mathcal{C} \leftarrow \mathcal{C} \cup \{Q\}$           ▷ Add confirmed equilibrium
      else if  $|S| \leq \xi$  then
         $\mathcal{R} \leftarrow \mathcal{R} \cup \{\text{br}(Q)\}$ 
         $\mathcal{B} \leftarrow \mathcal{B} \cup \text{brs}(Q) \setminus \{\text{br}(Q)\}$ 
      else
         $\mathcal{B} \leftarrow \mathcal{B} \cup \text{brs}(Q)$ 
      end if
    end for
  end if
end while
```

equilibrium for each game evaluated. We consider this process to exercise reasonable due diligence in equilibrium search, given the practical constraints of analyzing games this large.

5 Replication of Zhan and Friedman

We start our report of results by presenting our replication of Zhan and Friedman (2007), which extends their study in three key ways:

1. The original study separately analyzed games with three classes of discretized strategies. We investigated a fourth category of game where agents can choose any strategy from the original three classes.
2. The original study restricted solutions to pure-strategy ε -Nash equilibria. We broadened consideration to include mixed-strategy equilibria, and accordingly searched more extensively over strategy profiles.
3. We evaluated the found equilibria with many more samples and report statistical confidence on regret.

Our replicated CDA market simulator is identical to that specified by Zhan and Friedman except for one detail. Instead of allowing agents to rebid at the same price, we simply shuffle the agents once and have them submit bids in that order. Aside from subtle effects on time priority, this should produce identical results. We confirm empirically that any differences are negligible.

Also, while the precise aggregate efficiency measure used was unspecified in the original paper, we were able to reproduce the reported efficiency results exactly only by calculating averages over instance efficiencies (i.e., surplus obtained as a fraction of CE surplus for each instance or zero in a no-trade instance). We argue that the proper measure of expected efficiency is the expected surplus over the expected CE surplus. This definition appropriately gives more weight to random instances that allow more surplus and removes the necessity of defining efficiency in a no-trade scenario.

With this simulator, we applied the equilibrium search methodology presented in Section 4. We evaluated a total of twelve scenarios formed by three levels of market thickness combined with four classes of shading strategies. Three of the strategy classes—*Standard*, *Exponential*, and *Shift*—were as introduced by Zhan and Friedman. We added a union shading class—*All*—containing 31 strategies: eleven shading levels from each original class, minus two that are redundant because zero shading corresponds to truth-telling in all three original classes. In order to tractably explore each game, we applied DPR, and reduced the number of players in each role to four. We considered a profile in each of these games evaluated if we sampled it 250,000, 50,000, and 25,000 times respectively for the thin, medium, and thick markets, or an order of magnitude more samples than the previous study. Finally, we set the subgame size limit ξ for equilibrium search to four, meaning we stopped exploring profiles after we

had found at least one equilibrium and all unexplored best response subgames had at least five strategies in support between both roles.

A summary of the experimental results is presented in Figure 3. The equilibrium efficiency between shading classes in a single market thickness varies slightly. This is likely due to a few factors including, the choice of discrete strategies,⁶ the sampling error inherent in random simulation, and the fact that our profile search biases towards low support equilibria. Despite these factors, there is a strong overall trend in the data. As the market gets thinner, the efficiency of truth telling improves, and that of equilibrium shading degrades.

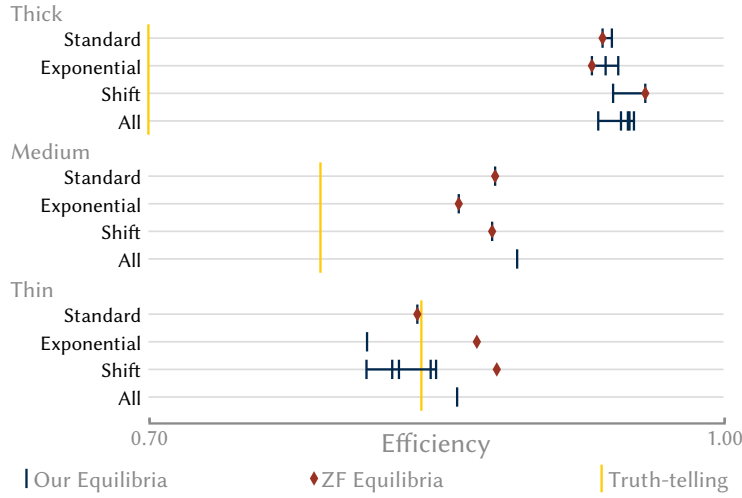


Figure 3: Efficiencies of all found equilibria compared to truth-telling in every scenario. Yellow bars represent the truth-telling efficiency in each market thickness. Blue bars with a line connecting them represent the efficiency of every role-symmetric equilibria found via our profile exploration method. Red diamonds represent the efficiencies of each symmetric pure-strategy equilibrium found by Zhan and Friedman (2007). Efficiency is computed as the expected surplus over the expected CE surplus for a large sample of profiles drawn from the equilibrium distribution. Only one equilibrium was found in each of the Medium environments, likely due to the way DPR handles non-divisible numbers of agents—four does not divide ten. Equilibria with access to All strategies were never subsets of equilibria with access to a restricted set of strategies. Notably, in all *All* equilibria, sellers exclusively play Shift strategies (additive shading), while buyers exclusively play Standard and Exponential strategies (multiplicative shading)—a result of the asymmetry of the formulation of multiplicative shading pointed out by Zhan and Friedman (2007).

⁶In particular, any Exponential shading can be achieved by a Standard shading with $m \in \{1 - e^{-m'} \mid m' \in \{0, 0.1, \dots, 1\}\}$ for buyers and $m \in \{e^{m'} - 1 \mid m' \in \{0, 0.1, \dots, 1\}\}$ for sellers, so the essential difference between the classes is really the choice of discrete shading levels.

In addition, our search process failed to find two equilibria from the previous work. However, not all found equilibria are equally important—we prefer equilibria with lower regret, and while we cannot confirm that regret of any equilibrium is low, we can compute confidence bounds on the regret of each equilibrium. Table 2 (Appendix B) lists all of the equilibria, their efficiency, and a 95% bootstrap upper confidence bound on regret. The equilibria from the original study that we did not find have significantly higher—roughly double—the regret bound of our comparable found equilibria.

6 Financial Market Environment

Our extension enriches the CDA scenario to more closely resemble financial markets. Our implementation employs a configurable financial market simulator, related to the system used by Wah et al. (2017). This simulator is not designed to be directly calibrated to financial data, but to exhibit important qualitative characteristics of financial markets, like adverse selection, multi-unit holdings, and agents with the ability to buy and sell. We model a single security traded in a CDA market with a large but finite number of agents. Prices and time are fine-grained but discrete, and simulations run for a finite horizon T . Agents arrive at designated times and submit limit orders to the market. The market maintains and reveals price quotes reflecting the best outstanding orders. Other bids in the order book are not visible to agents.

6.1 Valuation Model

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

$$f_t = \max \left\{ 0, \text{round} \left(\tilde{f}_t \right) \right\}, \quad \tilde{f}_t = r\bar{f} + (1-r)\tilde{f}_{t-1} + \sigma_s s_t, \quad \tilde{f}_0 = \bar{f}. \quad (1)$$

s_t is a standard normal shock at time t , thus σ_s^2 controls the volatility of the fundamental value. Parameter $r \in [0, 1]$ specifies the tendency by which the fundamental reverts back to the mean \bar{f} ; $r = 1$ corresponds 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*, 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. The degree of adverse selection in our model depends on the fundamental shock variance σ_s^2 and the degree of mean reversion r . In particular, higher values of r damp the fundamental variations, and thus reduce agents' exposure to adverse selection.

The private component of agent i 's valuation is a vector θ_i containing the agent-specific marginal utility for acquiring one more unit, relative to the fundamental value—similar to the model of Goettler et al. (2009). 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. Element $\theta_i^q, q \in [-q_{\max}, q_{\max})$ is the incremental benefit, over the fundamental, to agent i for gaining one unit of the security given current position q , where positive q indicates a long position.

We generate θ_i from a set of $2q_{\max}$ values drawn independently from a Gaussian distribution. Let $\hat{\theta} \sim \mathcal{N}(0, \sigma_{pv}^2)$ denote one of these drawn values. To ensure that the valuation reflects diminishing marginal utility, we sort the $2q_{\max}$ $\hat{\theta}$ values in descending order, and set θ_i^q respectively.

Agent i 's incremental surplus for trade j is based on its position q_j before the trade, the value of the fundamental at the end of the trading horizon T , and the transaction price p_j :

$$surplus_{ij} = \begin{cases} f_T - p_j + \theta_i^{q_j} & \text{if buying 1 unit} \\ p_j - f_T - \theta_i^{q_j-1} & \text{if selling 1 unit.} \end{cases}$$

A agent's total surplus is the sum of the agent's surplus over all transactions. Since the price and fundamental terms cancel out in exchange, the total surplus achieved when agent B buys from agent S is $\theta_B^{q_B} - \theta_S^{q_S-1}$, where q^i denotes the pre-trade position of agent i .

6.2 Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs (Das et al., 2001; Friedman and Rust, 1993; Wellman, 2011). In this study, we consider trading strategies that are variants of the *Zero Intelligence* (ZI) family (Gode and Sunder, 1993). The ZI strategy is exceedingly simple, but often employed in agent-based study of financial markets (LeBaron, 2006; Cason and Friedman, 1996), including recent AI studies of market making (Chakraborty et al., 2015; Wah and Wellman, 2017), because they have been found to generate realistic patterns of market behavior (Farmer et al., 2005). Though ZI agent instances are typically outperformed by more sophisticated alternatives (Vytelingum et al., 2008; Tesauro and Das, 2001), by using game-theoretic selection to set ZI strategy parameters we can produce highly competitive behavior for a given market environment (Wright, 2016). In rough effect, game-theoretic equilibration serves the parameter-tuning function accomplished through adaptation and evolutionary search in the *ZI Plus* strategies of Cliff (2009).

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 enter the market, observe the price quote, withdraw any outstanding orders, and have the opportunity to submit new ones. If multiple agents get information at the same time step, they act in a random order.

We consider two cases for order submission. In the first (*single-unit order*) case, 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. This style of order submission is representative the standard ZI strategy. In the second (*full-demand order*), agents submit orders to both buy and sell. Agents consistently submit their total demand, and are not subject to the random direction selection of the single-unit order case. In either case, 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). Since agents withdraw stale orders at arrival, the arrival rate λ serves as a rough proxy for agent urgency. If an agent has a lower arrival rate, then it has fewer opportunities to submit orders, which is particularly constraining in the single-unit order case. With single-unit ordering, agents may find it difficult or impossible to achieve their efficient position levels. In such settings, they have a strong incentive to make each arrival count. In the full-demand order case agents can achieve their efficient position levels in one arrival, thus provides a way to evaluate the effect of reducing urgency.

A ZI 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 ZI 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 $[U_{\min}, U_{\max}]$ —two parameters of the ZI strategy.

We extend the 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. Settings of $\eta < 1$ are often highly advantageous in our simulation environments, suggesting that providing even this simple ability to condition on price quote is an important feature in CDA trading strategy.

The threshold parameter can also be used as a means to submit so-called “fill-or-kill” orders, designed to trade immediately or not at all. Consider the strategy where $U_{\min} = U_{\max} = 10^5$ and $\eta = 10^{-3}$. This strategy takes the outstanding order at the quote—if a surplus of at least 100 is available—or else posts an order at such an unattractive price that it will never transact. In this implementation the unfilled order is not literally killed, but the result is effectively the same. Interestingly, a similar strategy appears in empirical equilibrium in several of our environments.⁷

We restrict our analysis to a discretized finite set of the entire strategy space, shown in Table 1, so we can apply normal-form game analysis techniques. The set of strategies was chosen starting with a roughly exponential grid of

⁷Since the fill-or-kill strategy never leaves reasonable orders, in self-play it never trades. It can however be profitable in mixtures with other strategies.

Table 1: Strategies considered for equilibrium analysis. Agents shade from their true belief by a random offset in $[U_{\min}, U_{\max}]$, the result of this draw is the amount of expected surplus the agent demands from trade. If an agent could get an η fraction of their desired surplus from an outstanding order, they take the order instead. Detailed strategy descriptions are in Section 6.2.

(a) Initial set of strategies at roughly exponentially spaced intervals

U_{\min}	U_{\max}	η
0	0	1
0	50	1
0	125	1
0	250	1
0	500	1
0	1000	1
0	2000	1

(b) Best response strategies, found via iterative process of adding best response to previous equilibria

U_{\min}	U_{\max}	η
125	500	1
2000	4000	1
0	500	0.8
0	2000	0.8
125	1000	0.4
29000	30000	0.001

max shading, then adding strategies that appeared frequently as best responses (among a larger set of heuristic strategies) to equilibria across games. The union of these two sets has thirteen strategies.

6.3 Market Environment Parameter Settings

We focus our analysis on two basic market *settings*, where we measure the impact of varying arrival rate, mean reversion, and market thickness. Restricting our analysis to these two settings allows us to investigate qualitative effects without having to do an infeasible exhaustive grid search over the space of parameter configurations. Both environments have sixty six agents and a time horizon T of 60000. The fundamental mean \bar{f} is 10^7 , sufficiently large to ensure that the probability of the fundamental reaching zero is negligible. The fundamental shock variance is 10^6 with mean reversion 0.05, making the a priori final fundamental variance roughly 10^7 . Private value variance is 5×10^6 so that the relative importance of private value and common value are close. The fundamental, private value, and shading parameters are the only parameters that affect valuation, and thus only their relative values dictate equilibria.

The *multi-position* setting most closely matches the environment for which ZI agents were originally defined (Gode and Sunder, 1993). Agents have a maximum absolute position $q_{\max} = 10$ (hence the name multi-position), and follow the single-unit order scheme. As a result, the agents never have more than one outstanding single-unit order in the market at a time. Agents arrive with a rate λ of 10^{-4} giving them six bidding opportunities in expectation, enough to reach their optimal position, which rarely exceeds four.

The second setting we call the *single-position* setting. In this setting, agents

have a maximum position $q_{\max} = 1$, and follow the full-demand order scheme.⁸ This setting lifts the restriction that agents have exactly one outstanding single-unit order at a time, but agents will never hold an absolute position greater than a single unit (hence the name single-position). Agents arrive with a rate λ of 5×10^{-5} giving them three entries in expectation, even though agents need only to arrive once to achieve their optimal position.

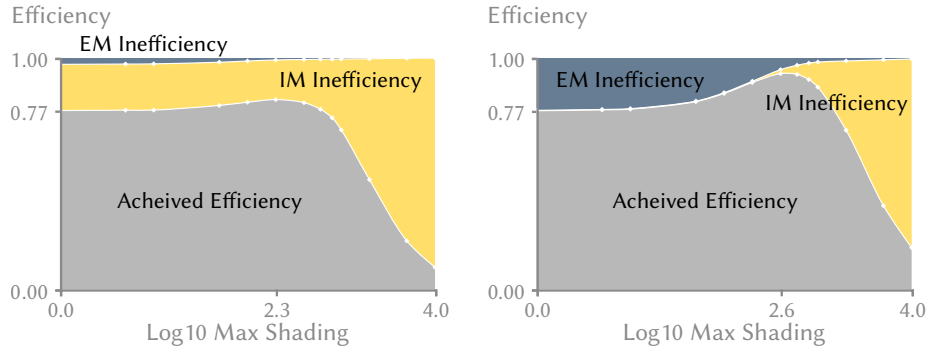
6.4 Uniform Shading Analysis

One would hope that the intuition behind the tradeoff of EM and IM inefficiency in the Zhan and Friedman model—seen in Figure 2—would carry over to this financial market model as well. However, the situation is more complicated when agents have multiple opportunities to trade, for multiple units, and are not restricted to only buying or selling. Figure 4a is a plot of EM and IM inefficiency for the baseline multi-position setting. Unlike the standard CDA model, EM and IM inefficiency do not provide clear clues as to *how* shading affects efficiency in this model. We hypothesize that the definitions of EM and IM inefficiency are muddled when agents have the opportunity to trade counter to their CE position, that is, an agent who *should* buy a unit might opportunistically sell one due to advantageous information. Figure 4b is an identical plot of EM and IM efficiency for the baseline multi-position setting, except half of the agents were assigned to be buyers, the other half sellers, allowed to trade only in their assigned directions. When we apply this structure to the market, we regain the illustrative breakdown of efficiency. It is clear from these plots that the ability for agents to trade in both directions adds a significant level of complexity to analyzing the effect of shading on efficiency in CDAs. We leave a more appropriate breakdown of efficiency when agents can trade in both directions as an open problem, and instead focus on the effects of shading in equilibrium.

7 Financial Market Equilibrium Analysis

Using the model described in Section 6, and the methodology to find approximate role-symmetric Nash equilibria described in Section 4, we can investigate the effects of equilibrium shading and truth-telling on financial market efficiency. To measure efficiency in this model, we calculate an agent’s competitive equilibrium position constrained by its number of arrivals. This is necessary since agents may not arrive at all. Due to the large number of agents in the simulation, we used DPR to reduce the effective number of agents to six. We considered a profile evaluated if we sampled it 10,000 times; at this level, the standard error of profile payoffs was sufficiently small. We set the subgame size limit $\xi = 3$, meaning we stopped exploring profiles after we had found an equilibrium, and all best response subgames had support over at least four strategies.

⁸We chose $q_{\max} = 1$ for full-demand ordering in part because generalizing the ZI strategy to multiple units in the same direction is not obvious, particularly without adding parameters.



(a) EM and IM inefficiency in the financial market model. Unlike the standard CDA model, this decomposition does not present a clear picture of *how* shading affects efficiency.

(b) EM and IM inefficiency in the financial market model when agents are assigned to be buyers or sellers. With this restriction, the results are remarkably similar to Figure 2.

Figure 4: Uniform shading analysis of the financial market model.

In both the single- and multi-position setting we varied three key parameters from their baseline value: number of agents (6, 36, 66, 216), arrival rate (0.8, 1, 2, 5, and 10 times baseline), and fundamental mean reversion (0, 0.1, 0.2, 0.4, and 1 times baseline). Increasing the number of agents in a simulation increases the market thickness. Increasing the arrival rate increases agents' access to relevant information, decreases the amount of time potentially stale order sit in the market, and gives agents more opportunities to trade. Finally, increasing the mean reversion exposes agents to less adverse selection due to the smaller impact of a shock on the final fundamental price. In each setting, we analyzed twelve distinct environments corresponding to different levels of these three parameters. With no mean reversion and other settings at baseline, the adverse selection is severe enough to preclude all trading. To restore some profitability in this environment, we compensate by also reducing the shock variance σ_s^2 to 100.

Figure 5 compares the efficiency achieved in equilibrium with that produced by truth-telling agents for each parameter variation. Both the single- and multi-position setting show identical trends for each variable. As for the simple CDA model, increasing the number of agents increases the efficiency with equilibrium shading. Unlike in the simple CDA, truth-telling efficiency for the financial model also increases with market thickness. This is probably because agents in our model can retrade, making initial extra-marginal trades less deleterious to final efficiency. This result illustrates that the relative effect of strategic shading depends not just on the inherent efficiency of the setting, as increasing the number of agents improves the truth-telling efficiency, and also the benefit of equilibrium shading. Increasing either arrival rate or mean reversion has the opposite trend, the benefit to equilibrium shading decreases, alongside an increase

in truth-telling efficiency. This is to be expected, as a limited number of trader arrivals means little time to correct for bad trades and the increase in adverse selection from low mean reversion tends to generate more inefficient trades. Shading ameliorates inefficient trades, yielding an improvement in efficiency in both environments.

8 Conclusions

Shading in continuous double auctions presents a tradeoff between preventing low efficiency extra-marginal trades and inhibiting efficient intra-marginal trades. We employed a simulation-based approach to analyze the effect of *strategic* bid shading on the efficiency of simple CDA markets and richer financial markets. We confirmed the results from Zhan and Friedman (2007) using a more complete equilibrium search. In both market models that we investigated, we consistently found that strategic bid shading helps efficiency when there are more agents in the market. In the financial market, we also observe a benefit to strategic shading when there is a large amount of adverse selection due to a noisy common valuation or a high level of urgency due to limited trading opportunities. Our results strengthen the claims made by Zhan and Friedman by broadening the search for Nash-equilibria and extending the model to environments with dynamic elements.

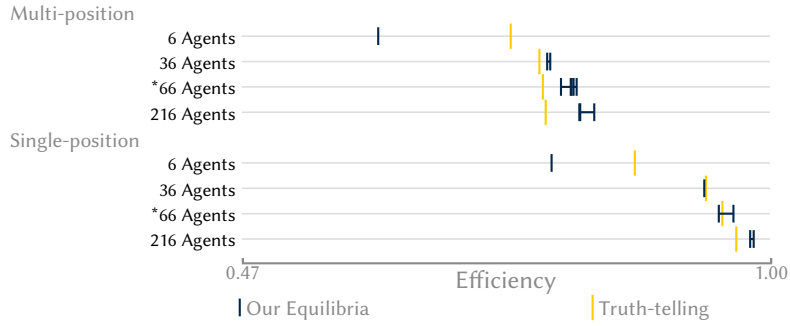
More generally, this phenomenon is highly germane to design of rules and regulations surrounding CDA markets. For example, measures aimed at promoting true value revelation—following the typical intuition of mechanism designers—may be counter-productive to mechanism design goals in this setting. Our analysis also underscores the need for accounting for strategic behavior when comparing CDAs to alternative mechanisms, particularly those (e.g., call markets) that may not exhibit this phenomenon.

9 Acknowledgements

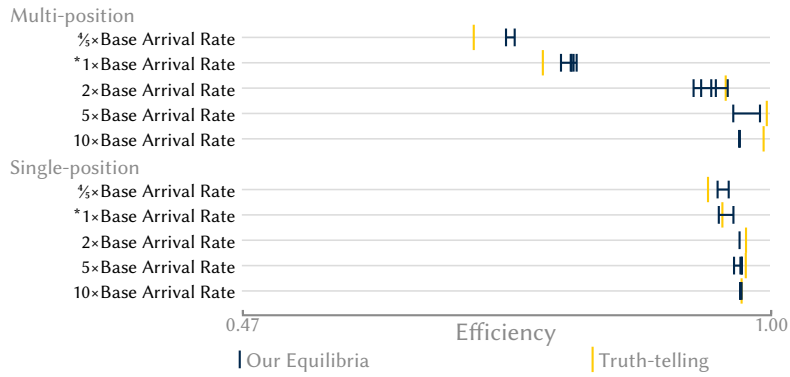
This work was supported by the National Science Foundation. An early version of this paper was presented at the Algorithmic Game Theory Workshop at IJCAI 2016.

References

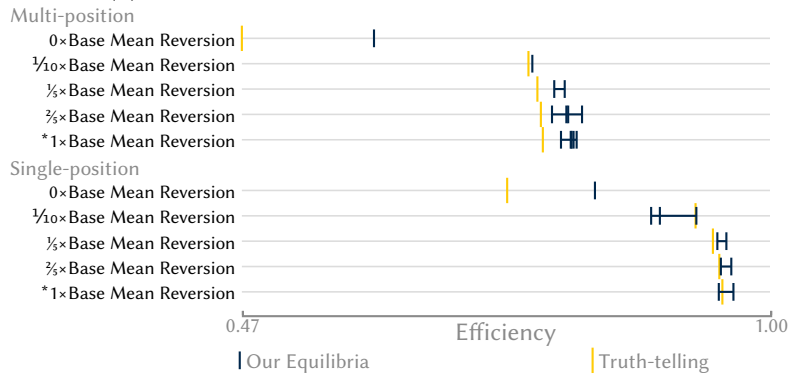
- Bruno Biais, Larry Glosten, and Chester Spatt. Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264, 2005.
- Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. *Journal of the ACM*, 53:845–879, 2006.



(a) Equilibrium efficiency as a function of number of agents.



(b) Equilibrium efficiency as a function of arrival rate.



(c) Equilibrium efficiency as a function of mean reversion.

Figure 5: Fraction of optimal surplus for truth-telling agents (yellow), and in equilibrium (blue) for each of the environments. Blue ranges correspond to multiple found equilibria. Baseline settings are marked with an asterisk. Thicker markets, fewer arrivals, and less mean reversion lead to more efficient CDA equilibria.

- Timothy N. Cason and Daniel Friedman. Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20(8):1307–1337, 1996.
- Mithun Chakraborty, Sanmay Das, and Justin Peabody. Price evolution in a continuous double auction prediction market with a scoring-rule based market maker. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 835–841, 2015.
- Dave Cliff. ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computation*, 13:3–18, 2009.
- Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauero. Agent-human interactions in the continuous double auction. In *Seventeenth International Joint Conference on Artificial Intelligence*, pages 1169–1178, 2001.
- J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102:2254–2259, 2005.
- Daniel Friedman. The double auction market institution: A survey. In Friedman and Rust (1993), pages 3–25.
- Daniel Friedman and John Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1993.
- Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, 1998.
- Irving L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, pages 119–137, 1993.
- Dhananjay K. Gode and Shyam Sunder. What makes markets allocationally efficient? *Quarterly Journal of Economics*, 112(2):603–630, 1997.
- Ronald L. Goettler, Christine A. Parlour, and Uday Rajan. Informed traders and limit order markets. *Journal of Financial Economics*, 93(1):67–87, 2009.
- Patrick R. Jordan, L. Julian Schvartzman, and Michael P. Wellman. Strategy exploration in empirical games. In *Ninth International Conference on Autonomous Agents and Multiagent Systems*, pages 1131–1138, 2010.

- Blake LeBaron. Agent-based computational finance. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Agent-Based Computational Economics*. Elsevier, 2006.
- John Rust, John Miller, and Richard Palmer. Behavior of trading automata in a computerized double auction market. In Friedman and Rust (1993), pages 155–198.
- Gerald Tesauro and Rajarshi Das. High-performance bidding agents for the continuous double auction. In *Third ACM Conference on Electronic Commerce*, pages 855–872, 2001.
- Perukrishnen Vytelingum, Dave Cliff, and Nicholas R. Jennings. Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172:1700–1729, 2008.
- Elaine Wah and Michael P. Wellman. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance*, 2017.
- Elaine Wah, Mason D. Wright, and Michael P. Wellman. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 2017.
- Michael P. Wellman. *Trading Agents*. Morgan & Claypool Publishers, 2011.
- Michael P. Wellman. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.
- Michael P. Wellman, Tae Hyung Kim, and Quang Duong. Analyzing incentives for protocol compliance in complex domains: A case study of introduction-based routing. In *Twelfth Workshop on the Economics of Information Security*, 2013.
- Bryce Wiedenbeck and Michael P. Wellman. Scaling simulation-based game analysis through deviation-preserving reduction. In *Eleventh International Conference on Autonomous Agents and Multi-Agent Systems*, pages 931–938, 2012.
- Bryce Wiedenbeck, Ben-Alexander Cassell, and Michael P. Wellman. Bootstrap statistics for empirical games. In *Thirteenth International Conference on Autonomous Agents and Multiagent Systems*, pages 597–604, 2014.
- Mason Wright. Using reinforcement learning to validate empirical game-theoretic analysis: A continuous double auction study. 2016.
- Wenjie Zhan and Daniel Friedman. Markups in double auction markets. *Journal of Economic Dynamics and Control*, 31(9):2984–3005, 2007.

Appendices

A Stylized Examples: Equilibrium Analysis

A.1 Example 2

Let there be two buyers and one seller. One buyer, BL , has valuation v_{BL} with $\Pr(v_{BL} = 1) = \Pr(v_{BL} = 2) = 0.5$. The other buyer, BH , has valuation v_{BH} with $\Pr(v_{BH} = 2) = \Pr(v_{BH} = 3) = 0.5$. The sole seller, S , has valuation $v_S = 0$. The agents arrive at the market in a uniform random ordering, and get to observe the order book and their position in the ordering.

In the optimal outcome BH trades with S , for expected welfare (total surplus) 2.5.

The truth-telling outcome is straightforward. If everyone bids their valuations, then when S arrives first, it will trade with the first buyer to enter the market (welfare = $\frac{1}{2}\mathbb{E}[v_{BL}] + \frac{1}{2}\mathbb{E}[v_{BH}] = 2$). When S arrives second, it will also trade with the first buyer to arrive (welfare = 2); when S arrives last, it will trade with the best order in the market (welfare = 2.5). The expected welfare under truth-telling is $\frac{1}{3}(2 + 2 + 2.5) \approx 2.167$.

To characterize strategic behavior in this example, we adopt perfect Bayesian equilibrium as our solution concept. We also assume for this construction that players break ties by accepting indifferent trades; this assumption does not affect welfare, but allows us to avoid discussing the multitude of qualitatively similar equilibria due to indifference. If S enters first, it bids 2. This offer will always trade, since at least one and possibly both of the subsequently arriving buyers have value at least this high.

If a buyer arrives first, the PBE behavior of that buyer can be described by the following cases:

1. BL arrives first, $v_{BL} = 1$. BL offers any $x < 1$.
2. BL arrives first, $v_{BL} = 2$. BL offers $\frac{4}{3}$.
3. BH arrives first, $v_{BH} = 2$. BH offers $\frac{4}{3}$.
4. BH arrives first, $v_{BH} = 3$. BH offers $\frac{4}{3}$.

If the seller arrives second, it will reject any offer less than $\frac{4}{3}$, and ask for 2, otherwise it will accept. If another buyer arrives second, it will bid ϵ over the current bid if it can profit from doing so. The last agent will take the best offer it can.

First we show that when S arrives second, it is best responding under its presumption that the third agent is a price-taker. In Case 1, the third agent is BH (and S infers this), in which case offering to sell at 2 maximizes surplus. Cases 2–4 form an information set for S , in which the third agent is BL with probability $\frac{2}{3}$. If S were to reject the first agent's bid, its optimal offer would be 2, yielding an expected surplus of $(\frac{2}{3})(\frac{1}{2})2 + (\frac{1}{3})2 = \frac{4}{3}$. So it may as well

accept the first agent's bid. Given S 's strategy, no first-round buyer bid less than $\frac{4}{3}$ would suffice. The only advantage from a greater bid would be to BH with $v_{BH} = 3$ if could prevent being outbid by BL with $v_{BL} = 2$. If BH bid 2 in this instance, then it would make a guaranteed profit of 1, but by bidding $\frac{4}{3}$ it makes an expected profit of $\frac{5}{4} = \frac{5}{3} \cdot \frac{3}{4}$, due to being outbid 25% of the time. As a result, no first-round buyer benefits from bidding more than $\frac{4}{3}$. The strategies are therefore in PBE.

When $v_{BL} = 2$, the seller trades with whichever of BL or BH arrives first. When $v_{BL} = 1$, the low buyer's shading precludes the trade, and so the seller trades with the high buyer regardless of arrival order. Thus, $\frac{3}{4}$ of the time the seller sells to BH and the other $\frac{1}{4}$ the seller sells to BL when it has a valuation of 2. The expected welfare in this equilibrium is $2.375 = \frac{3}{4} \cdot 2.5 + \frac{1}{4} \cdot 2$. Strategic shading significantly improves the efficiency of this market from 0.86 ($2.167 / 2.5$) to 0.95 ($2.375 / 2.5$).

A.2 Example 3

Let there be two buyers and one seller, all with i.i.d. uniform private values over the unit interval. The agents arrive in a uniform random ordering, and observe the current state of the order book, and the ordering of agents. Since the buyers are ex ante identical, we refer to the first buyer as buyer 1 (b_1) and the second buyer as buyer 2 (b_2).

The social optimum occurs when the seller trades with the largest buyer. If we let \bar{v}_b represent the largest buyer valuation, then

$$\text{Optimal Welfare} = \mathbb{E}[\bar{v}_b - v_s \mid v_s < \bar{v}_b] = \frac{1}{4}.$$

If agents tell the truth, and the seller arrives last, it trades with the maximum of the buyer valuations, and so the welfare is the same as the optimal social welfare $\frac{1}{4}$. In the other two instances the seller has a chance to trade with the first buyer before trading with the second. Therefore

$$\begin{aligned} \text{Welfare}_{\text{seller not last}} &= \mathbb{E}[v_{b_1} - v_s \mid v_s \leq v_{b_1}] + \mathbb{E}[v_{b_2} - v_s \mid v_s > v_{b_1}, v_s \leq v_{b_2}] \\ &= \frac{5}{24}. \end{aligned}$$

The efficiency of truth-telling is the expected welfare for truth-telling divided by the maximum social welfare, or $\frac{8}{9} = 4 \left(\frac{2}{3} \cdot \frac{5}{24} + \frac{1}{3} \cdot \frac{1}{4} \right)$.

The Perfect Bayesian Equilibrium (PBE) solutions can be calculated using backward induction. To calculate the efficiency in PBE, we first consider the case when the seller arrives first. In this case, both buyers will accept any offer below their valuation. The only strategic decision is the bid the seller should make. Using similar notation to the social optimum, the seller's expected profit for a bid s is

$$\text{Profit}_{\text{seller}} = \mathbb{E}[s - v_s \mid s < \bar{v}_b] = (s - v_s)(1 - s^2).$$

This profit is maximized at $s^* = \frac{1}{3}(v_s + \sqrt{v_s^2 + 3})$. The welfare calculation is the same as in the truthful case, except that the conditions are in terms of s^* instead of v_s . The social welfare when the seller arrives first is therefore $\sqrt{3}/9$.

When the seller arrives in the middle, the last buyer takes the seller's bid if it exists and the buyer can profit. The seller has a choice between taking the existing bid or placing a new one, attempting to get more surplus from the second bidder. The seller makes profit

$$\text{Profit}_s = \begin{cases} b_1 - v_s & \text{take order} \\ \mathbb{E}[s - v_s \mid s \leq v_{b_2}] & \text{place order } s. \end{cases}$$

The optimal bid in the later case is $s^* = \frac{1}{2}(1 + v_s)$, which implies that the seller takes the existing order if $v_s \leq 2\sqrt{b_1} - 1$, otherwise it places an ask at s^* . A placed ask will always be greater than the existing bid, that is, it never transacts with the old bid.

The first buyer's profit conditioned on this information is

$$\text{Profit}_{b_1} = \mathbb{E}[v_{b_1} - b_1 \mid v_s \leq 2\sqrt{b_1} - 1] = (v_{b_1} - b_1)(2\sqrt{b_1} - 1).$$

Which is maximized when $b_1^* = \frac{1}{18}(1 + 6v_{b_1} + \sqrt{12v_{b_1} + 1})$, conditioned on $v_{b_1} \geq \frac{1}{4}$. The resulting social welfare conditioned on these strategies is $\frac{13}{5184}(47 + 8\sqrt{13})$.

When the seller arrives last, the first buyer knows it can win only by bidding over the second buyer's valuation, thus

$$\text{Profit}_{b_1} = \mathbb{E}[v_{b_1} - b_1 \mid v_{b_2} < b_1, v_s \leq b_1] = b_1^2(v_{b_1} - b_1).$$

The optimal bid is $b_1^* = \frac{2}{3}v_{b_1}$. The second bidder will over bid the existing order if it can profit, but might bid more than epsilon over if it can extract more expected profit from the seller. The profit for placing a new order is

$$\text{Profit}_{b_2} = \mathbb{E}[v_{b_2} - b_2 \mid v_s < b_2, b_1 < b_2] = b_2(v_{b_2} - b_2).$$

This profit is maximized at $b_2^* = \frac{1}{2}v_{b_2}$ as long as $b_2^* \geq b_1$. Buyer 2's optimal strategy is to bid $\min\{v_{b_2}, \max\{b_2^*, b_1 + \epsilon\}\}$. The social welfare of this permutation is $\frac{19}{96}$, making the expected social welfare in PBE roughly 0.19.

The efficiency in PBE is the average of each permutation, which is approximately 0.77, significantly less than the corresponding truth-telling efficiency of approximately 0.89.

B Equilibrium Tables

Table 2: Equilibria found in CDA games, by Zhan and Friedman (2007) or our replication. Where multiple equilibria were found, they are given numbers to differentiate. Equilibria numbered “ZF” correspond to equilibria found in the previous work, but not this one, while equilibria labeled with a † were also identified in the previous work. In the “All” shading class agents can play any strategy from standard, exponential, or shift. To differentiate these strategies in the “All” class, Standard shading strategies have a “T” suffix, Exponential strategies have an “E” suffix, and Shift strategies have an “S” suffix. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Density	Shading Class	Num	Buyers’ Shade	Buyers’ Prob (%)	Sellers’ Shade	Sellers’ Prob (%)	Efficiency	95% Regret
Thick	Truthful	–	0.0	100.0	0.0	100.0	0.697	–
		Standard	1†	0.4	100.0	0.3	100.0	0.936
	2		0.4	100.0	0.2	69.4	0.941	0.780
						0.4	30.6	
	Exponential	1†	0.5	100.0	0.3	100.0	0.931	0.344
		2	0.5	73.9	0.2	100.0	0.944	0.273
				0.6	26.1			
	3	0.5	100.0	0.3	61.3	0.938	0.091	
					38.7			
	Shift	1†	0.4	100.0	0.4	100.0	0.959	0.524
		2	0.4	100.0	0.5	87.5	0.942	0.194
					0.3	12.5		
	All	1	0.4T	100.0	0.2S	100.0	0.950	0.416
			0.4E	100.0	0.3S	86.9	0.953	0.559
		3	0.5E	56.0	0.3S	100.0	0.934	0.355
0.4T						27.9		
4		0.3T	16.1	0.4E	63.0	0.3S	100.0	0.946
	0.5E				24.2			
5	0.4T	100.0	0.2S	95.9	0.950	0.631		
				0.3S	4.1			
Medium	Truthful	–	0.0	100.0	0.0	100.0	0.787	–
	Standard	1†	0.3	100.0	0.3	100.0	0.879	0.108
	Exponential	1†	0.4	100.0	0.3	100.0	0.860	0.197
	Shift	1†	0.4	100.0	0.4	100.0	0.878	0.274
	All	1	0.3T	100.0	0.3S	100.0	0.891	0.336

Table 2: Equilibria found in CDA games, by Zhan and Friedman (2007) or our replication. Where multiple equilibria were found, they are given numbers to differentiate. Equilibria numbered “ZF” correspond to equilibria found in the previous work, but not this one, while equilibria labeled with a † were also identified in the previous work. In the “All” shading class agents can play any strategy from standard, exponential, or shift. To differentiate these strategies in the “All” class, Standard shading strategies have a “T” suffix, Exponential strategies have an “E” suffix, and Shift strategies have an “S” suffix. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Density	Shading Class	Num	Buyers' Shade	Buyers' Prob (%)	Sellers' Shade	Sellers' Prob (%)	Efficiency	95% Regret	
Thin	Truthful	–	0.0	100.0	0.0	100.0	0.841	–	
	Standard	1†	0.3	100.0	0.3	100.0	0.838	0.065	
	Exponential	ZF	0.3	100.0	0.2	100.0	0.870	0.391	
		1	0.4	100.0	0.3	100.0	0.812	0.009	
	Shift	ZF	0.3	100.0	0.3	100.0	0.880	0.241	
		1	0.3	100.0	0.4	100.0	0.848	0.033	
		2	0.4	100.0	0.3	93.1	0.845	0.032	
						0.4	6.9		
		3	0.4	56.3	0.4	100.0	0.825	0.076	
			0.3	43.7					
	All	4	0.4	90.2	0.4	97.7	0.812	0.092	
			0.3	9.8	0.3	2.3			
		5	0.4	88.2	0.4	57.5	0.829	0.059	
			0.3	11.8	0.3	42.5			
1		0.3E	100.0	0.3S	100.0	0.859	0.034		

Table 3: Equilibria found in multi-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered “T” correspond truth-telling profiles. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Environment	Num	U_{\min}	U_{\max}	η	Prob (%)	Efficiency	95% Regret
$4/5 \times$ Base Arrival Rate	T	0	0	1	100.0	0.702	–
	1	0	500	0.8	100.0	0.743	35.096
	2	125	500	1	80.6	0.734	6.936
6 Agents		0	500	0.8	19.4		
	T	0	0	1	100.0	0.739	–
	1	125	1000	0.4	94.3	0.606	17.519
36 Agents		125	500	1	3.5		
		0	500	1	2.2		
	T	0	0	1	100.0	0.768	–
	1	125	500	1	64.3	0.776	8.426
		0	500	0.8	25.3		
		0	500	1	10.4		
$0 \times$ Base Mean Reversion		0	500	1	67.5	0.779	24.002
		125	500	1	32.5		
	T	0	0	1	100.0	0.469	–
$1/10 \times$ Base Mean Reversion	1	125	1000	0.4	55.2	0.602	36.613
		29000	30000	0.001	44.8		
	T	0	0	1	100.0	0.757	–
$1/5 \times$ Base Mean Reversion	1	125	1000	0.4	49.9	0.761	98.385
		125	500	1	46.6		
		0	500	1	3.5		
$1/5 \times$ Base Mean Reversion	T	0	0	1	100.0	0.766	–
	1	125	500	1	97.7	0.783	37.485
		0	500	0.8	2.3		
	2	0	500	1	63.0	0.793	47.580
		0	250	1	26.3		
$2/5 \times$ Base Mean Reversion		125	1000	0.4	10.6		
	T	0	0	1	100.0	0.769	–
	1	0	250	1	70.9	0.811	60.549
		0	500	1	29.1		
	2	0	500	0.8	97.5	0.797	40.140
		125	1000	0.4	2.5		
	3	125	500	1	82.1	0.781	44.436
		125	1000	0.4	14.8		
		0	500	0.8	3.1		
	4	0	500	0.8	45.1	0.795	19.879
	125	500	1	44.1			
	0	125	1	10.7			

Table 3: Equilibria found in multi-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered “T” correspond truth-telling profiles. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Environment	Num	U_{\min}	U_{\max}	η	Prob (%)	Efficiency	95% Regret	
Baseline	T	0	0	1	100.0	0.771	–	
	1	0	500	0.8	80.7	0.802	31.807	
		0	250	1	19.3			
	2	0	500	0.8	90.6	0.800	31.260	
		0	250	1	8.5			
		125	500	1	0.9			
	3	125	500	1	43.3	0.800	12.652	
		0	250	1	29.6			
		0	500	0.8	27.1			
	4	0	500	1	55.7	0.805	12.196	
		0	250	1	23.4			
		0	125	1	20.8			
	5	125	500	1	67.7	0.790	15.456	
		0	500	0.8	21.4			
		0	500	1	10.9			
2×Base Arrival Rate	T	0	0	1	100.0	0.955	–	
	1	0	250	1	94.7	0.957	15.694	
		125	1000	0.4	5.3			
	2	0	500	1	87.3	0.941	32.729	
		29000	30000	0.001	12.7			
	3	125	500	1	75.9	0.923	52.861	
		0	500	1	24.1			
	4	0	500	0.8	83.6	0.945	31.820	
		0	50	1	8.7			
		0	125	1	7.7			
	5	125	500	1	50.1	0.931	26.311	
		0	500	0.8	36.7			
		0	500	1	13.1			
	5×Base Arrival Rate	T	0	0	1	100.0	0.997	–
		1	0	500	0.8	90.5	0.990	62.958
0			50	1	9.5			
2		125	500	1	92.9	0.963	31.901	
		0	1000	1	7.1			

Table 3: Equilibria found in multi-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered “T” correspond truth-telling profiles. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Environment	Num	U_{\min}	U_{\max}	η	Prob (%)	Efficiency	95% Regret
10×Base Arrival Rate	T	0	0	1	100.0	0.993	–
	1	125	500	1	92.1	0.970	15.152
		0	1000	1	7.9		
		125	500	1	91.9	0.969	16.657
	2	0	2000	0.8	7.6		
		0	250	1	0.5		
216 Agents	T	0	0	1	100.0	0.774	–
	1	0	250	1	78.1	0.823	19.134
		0	500	1	21.9		
		0	500	0.8	98.1	0.809	38.010
	2	125	1000	0.4	1.9		
		0	500	0.8	45.2	0.808	3.904
		125	500	1	31.0		
	3	0	500	1	23.8		

Table 4: Equilibria found in single-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered “T” correspond truth-telling profiles. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Environment	Num	U_{\min}	U_{\max}	η	Prob (%)	Efficiency	95% Regret
$4/5 \times$ Base Arrival Rate	T	0	0	1	100.0	0.938	–
	1	125	500	1	100.0	0.958	3.992
	2	125	1000	0.4	68.3	0.947	12.242
6 Agents		0	1000	1	31.7		
	T	0	0	1	100.0	0.864	–
	1	0	2000	0.8	49.1	0.780	5.656
		125	1000	0.4	45.7		
36 Agents		29000	30000	0.001	5.2		
	T	0	0	1	100.0	0.935	–
	1	125	1000	0.4	90.9	0.934	8.409
$0 \times$ Base Mean Reversion		125	500	1	9.1		
	T	0	0	1	100.0	0.736	–
	1	0	2000	0.8	66.6	0.824	34.781
$1/10 \times$ Base Mean Reversion		29000	30000	0.001	33.4		
	T	0	0	1	100.0	0.925	–
	1	0	2000	0.8	100.0	0.880	40.943
	2	0	2000	0.8	84.9	0.889	25.701
$1/5 \times$ Base Mean Reversion		125	1000	0.4	15.1		
	3	125	1000	0.4	83.2	0.926	6.577
		0	2000	0.8	16.8		
	T	0	0	1	100.0	0.943	–
	1	125	1000	0.4	77.1	0.947	9.790
$2/5 \times$ Base Mean Reversion		0	1000	1	22.9		
	2	125	500	1	71.8	0.956	51.101
		0	500	0.8	28.2		
	T	0	0	1	100.0	0.949	–
Baseline	1	125	1000	0.4	55.1	0.951	12.401
		0	1000	1	44.9		
	2	125	500	1	61.5	0.961	9.872
		0	500	0.8	33.7		
		0	500	1	4.7		
$2 \times$ Base Arrival Rate	T	0	0	1	100.0	0.952	–
	1	0	1000	1	100.0	0.948	8.419
	2	125	500	1	77.3	0.963	13.534
$2 \times$ Base Arrival Rate		0	500	1	22.7		
	T	0	0	1	100.0	0.976	–
	1	125	500	1	58.2	0.969	11.489
	125	1000	0.4	41.8			

Table 4: Equilibria found in single-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered “T” correspond truth-telling profiles. “95% Regret” is the bootstrapped 95% upper confidence interval on regret. Equilibria with support for more than a single strategy use additional rows for each strategy.

Environment	Num	U_{\min}	U_{\max}	η	Prob (%)	Efficiency	95% Regret
5×Base Arrival Rate	T	0	0	1	100.0	0.976	–
	1	0	2000	1	90.8	0.964	18.835
		125	500	1	9.0		
		0	1000	1	0.2		
	2	125	1000	0.4	65.6	0.972	7.141
		0	1000	1	34.4		
		125	1000	0.4	60.0	0.972	7.927
	3	125	500	1	22.3		
		0	1000	1	17.7		
		125	1000	0.4	71.3	0.970	7.431
	4	0	2000	0.8	15.4		
		125	500	1	13.2		
125		500	1	13.2			
10×Base Arrival Rate	T	0	0	1	100.0	0.971	–
	1	0	1000	1	68.4	0.971	13.149
		0	2000	1	31.6		
	2	125	1000	0.4	75.3	0.970	11.784
		0	2000	1	24.7		
	3	125	1000	0.4	77.9	0.970	13.335
	0	2000	0.8	22.1			
216 Agents	T	0	0	1	100.0	0.966	–
	1	125	500	1	100.0	0.984	6.245
		0	500	1	89.5	0.980	13.563
	2	125	1000	0.4	9.8		
		125	500	1	0.7		