The AAAI-17 Workshop on
Computer Poker and Imperfect Information Games
WS-17-06

Spoofing the Limit Order Book: An Agent-Based Model

Xintong Wang, Michael P. Wellman

University of Michigan
xintongw,wellman@umich.edu

Abstract

We present an agent-based model of manipulating prices in financial markets through *spoofing*: submitting spurious orders to mislead other traders. Built around the standard limit-order mechanism, our model captures a complex market environment with combined private and common values, the latter represented by noisy observations of a fundamental time series. We start with zero intelligence traders, who ignore the order book, and introduce a version of *heuristic belief learning* (HBL) strategy that exploits the order book to predict price outcomes. By employing an empirical game-theoretic analysis to derive approximate strategic equilibria, we demonstrate the effectiveness of HBL and the usefulness of order book information in a range of non-spoofing environments. We further show that a market with HBL traders is spoofable, in that a spoofer can qualitatively manipulate prices towards its desired direction. After re-equilibrating games with spoofing, we find spoofing generally hurts market surplus and decreases the proportion of HBL. However, HBL's persistence in most environments with spoofing indicates a consistently spoofable market. Our model provides a way to quantify the effect of spoofing on trading behavior and efficiency, and thus measures the profitability and cost of an important form of market manipulation.

Introduction

Electronic markets have transformed the financial market landscape, with automation of trading scaling of volume and speed across geography and asset classes. Automated traders have unprecedented ability to gather and exploit market information from a broad variety of sources, including transactions and order book information exposed by many market mechanisms. Whereas some of these developments may contribute to improved price discovery and efficiency, they may also introduce new possibilities of disruptive and manipulative practices in financial markets.

Recent years have witnessed several cases of fraud and manipulation in the financial markets, where traders made tremendous profits by deceiving investors or artificially affecting market beliefs. On April 21, 2015, the U.S. Department of Justice charged Navinder Singh Sarao with 22 criminal counts, including fraud and market manipulation (Brush, Schoenberg, and Ring 2015). Prior to the Flash Crash,[1] Sarao allegedly used an automated program to place orders amounting to about $200 million worth of bets that the market would fall, and later replaced or modified those orders 19,000 times before cancellation. The U.S. Commodity Futures Trading Commission (CFTC) concluded that Sarao's manipulative practice was responsible for significant order imbalances. Though recent analysis has cast doubt on the causal role of Sarao on the Flash Crash (Aldrich, Grundfest, and Laughlin 2016), many agree that such manipulation could increase the vulnerability of markets and exacerbate market fluctuations.

The specific form of manipulation we examine in this paper is *spoofing*. Spoofing refers to the practice of submitting large spurious orders to buy or sell some security. The orders are spurious in that the spoofer does not intend for them to execute, but rather to mislead other traders by feigning strong buy or sell interest in the market. Spoof orders may lead other traders to believe that prices may soon rise or fall, thus altering their own behavior in a way that will directly move the price. To profit on its feint, the spoofer can submit a real order on the opposite side of the market and as soon as the real order transacts, cancel all the spoof orders.

In 2010, the Dodd-Frank Wall Street Reform and Consumer Protection Act was signed into federal law, outlawing spoofing as a deceptive practice. In its allegations against Sarao, the CFTC notes that "many market participants, relying on the information contained in the order book, consider the total relative number of bid and ask offers in the order book when making trading decisions". In fact spoofing can be effective only to the extent that traders actually use order book information to make trading decisions. Though regulatory enforcement and detection efforts have been made, spoofing is hard to eliminate due to its adversarial nature and the difficulty of determining the manipulation intent behind placement of orders. By reproducing spoofing in a computational model, the work reported here represents a first step toward developing more robust measures to characterize and prevent spoofing.

In this study, we propose a simple model that captures the

[1]The Flash Crash was a sudden trillion-dollar dip in U.S. stock markets on May 6, 2010, during which stock indexes collapsed and rebounded rapidly.

practice of spoofing in a *continuous double auction* (CDA) market with a single security traded. The CDA is a two-sided mechanism adopted by most financial and commodity markets (Friedman 1993). Under this mechanism, traders can submit orders at any time and whenever an incoming order matches an existing one, they trade at the incumbent order's limit price. We adopt an agent-based modeling approach to simulate the interactions among players with different strategies. The market is populated with multiple background traders and in selected treatments, one spoofer. Background traders are further divided into agents using instances of the *zero intelligence* (ZI) and *heuristic belief learning* (HBL) strategy families. Background traders of either type observe a noisy signal of the current fundamental value at the time they arrive to trade. The HBL strategy further considers information about orders recently submitted to the order book. The spoofer in our model maintains large buy spoof orders at one tick behind the best bid, aiming to manipulate but does not take actions to profit from the manipulation.

We first address the choice of background traders among HBL and ZI strategies, through *empirical game-theoretic analysis*. We demonstrate that in most non-spoofing environments, HBL is adopted in equilibrium and benefits price discovery and social welfare. By executing a spoofer against the found equilibria, we show that spoofing can qualitatively manipulate price given sufficient HBL traders in the market. We finally re-equilibrate games with spoofing and find HBL still exists in some equilibria but with smaller mixture probability. Though the welfare benefits of HBL persist, the presence of spoofing generally decreases market surplus.

Related Work and Contributions

Agent-Based Finance

Agent-based modeling (ABM) takes a simulation approach to study complex domains with dynamically interacting decision makers. ABM has been frequently applied to financial markets (LeBaron 2006), for example to study the Flash Crash (Paddrik et al. 2012) or bubbles and crashes in the abstract (LeBaron, Arthur, and Palmer 1999). Often the goal of ABM is to reproduce stylized facts of the financial system (Palit, Phelps, and Ng 2012). Researchers also use ABM to investigate the effects of particular trading practices, such as market making (Wah and Wellman 2015) and latency arbitrage (Li and Das 2016; Wah and Wellman to appear). ABM advocates argue that simulation is particularly well-suited to study financial markets (Bookstaber 2012), as analytic models in this domain typically require extreme stylization for tractability, and pure data-driven approaches cannot answer questions about changing market and agent designs.

Bidding Strategies

There is a substantial literature on autonomous bidding strategies in CDA markets (Wellman 2011). The basic *zero intelligence* (ZI) strategy (Gode and Sunder 1993) submits offers at random offsets from valuation. Despite its simplicity, ZI has been shown surprisingly effective in some cases (Farmer, Patelli, and Zovko 2005). In this study, we adopt an extended and parameterized version of ZI as our representative class of trading strategies that ignore order book information.

Researchers have also extended ZI with adaptive features that exploit observations to tune themselves to market conditions. For example, the zero intelligence plus (ZIP) strategy outperforms ZI by adjusting an agent-specific profit margin based on successful and failed trades (Cliff 1997; 2009). Adaptive Aggressiveness (AA) (Vytelingum, Cliff, and Jennings 2008) adds another level of strategic adaptation, allowing the agent to control its behavior with respect to short and long time scales.

Gjerstad proposed a more direct approach to learning from market observations, termed *GD* in its original version (Gjerstad and Dickhaut 1998) and named *heuristic belief learning* (HBL) in a subsequent generalized form (Gjerstad 2007). The HBL model estimates a heuristic belief function based on market observations over a specific memory length. Variants of HBL (or GD) have featured prominently in the trading agent literature. GDX calculates the belief function in a similar manner, but uses dynamic programming to decide on both the optimal price and time to submit a bid or ask (Tesauro and Bredin 2002). Modified GD further adapts the original GD to markets that support persistent orders (Tesauro and Das 2001).

Our study extends the HBL approach to a more complex financial market environment than addressed in previous studies. We adopt HBL as our representative class of agent strategies that exploit order book information. The extended HBL strategy is well-suited for our study as it considers the full cycle of an order, including the times an order is submitted, accepted, canceled, or rejected. Moreover, HBL can be applied with relatively few tunable strategic parameters, compared to other adaptive strategies in the literature.

Spoofing in Financial Markets

The literature on spoofing and its impact on financial markets is fairly limited. Some empirical research based on historical financial market data has been conducted to understand spoofing. Lee et al. (Lee, Eom, and Park 2013) empirically examine spoofing by analyzing a custom data set, which provides the complete intraday order and trade data associated with identified individual accounts in the Korea Exchange. They found investors strategically spoof the stock market by placing orders with little chance to transact to add imbalance to the order book. They also discovered that spoofing usually targets stocks with high return volatility but low market capitalization and managerial transparency. Wang, similarly, investigates the strategic behavior of spoofing trading orders in the index futures market in Taiwan, including their characteristics, profitability and real-time impacts (Wang 2015). Martinez-Miranda et al. (Martinez-Miranda, McBurney, and Howard 2016) propose a reinforcement learning framework to model spoofing in the context of portfolio growth maximization.

Contributions

Our contributions are threefold. *First*, we adapt the original GD strategy (Gjerstad and Dickhaut 1998) to a complex market environment that supports persistent orders, combined private and fundamental values, noisy observations, stochastic arrivals, and ability to trade multiple units with buy or sell flexibility. *Second*, by employing an empirical game-theoretic analysis to derive approximate strategic equilibria in a range of parametrically different market environments, we demonstrate the effectiveness of the extended HBL and the usefulness of order book information in the absence of spoofing. *Third* and most importantly, we provide the first computational model of spoofing a dynamic financial market, and demonstrate the effectiveness of spoofing against approximate-equilibrium traders in that model. Our model provides a way to quantify the effect of spoofing on trading behavior and efficiency, and thus measures the profitability and cost of an important form of market manipulation.

Market Model

Market Environment

The model employs a CDA mechanism with a single security traded in the market. Prices are fine-grained and take discrete values at integer multiples of the tick size. Time is also fine-grained and discrete and a trading period has a finite horizon T. Agents in the model submit limit orders, which specify the maximum (minimum) price at which they would be willing to buy (sell) together with the number of units to trade.

The fundamental value r of the underlying security changes throughout the trading period, according to a mean-reverting stochastic process:

$$r_t = \max\{0, \kappa\bar{r} + (1-\kappa)r_{t-1} + u_t\}; r_0 = \bar{r}. \quad (1)$$

Here r_t denotes the fundamental value of the security at time $t \in [1, T]$. The parameter $\kappa \in [0, 1]$ specifies the degree to which the fundamental reverts back to $\bar{r}$. The perturbation in the fundamental at time t is normally distributed: $u_t \sim N(0, \sigma_s^2)$. A mean-reverting time series of this sort has been empirically observed in financial markets such as foreign exchange and commodity markets and is wildly adopted in related research to investigate the effect of market making and latency arbitrage (Wah and Wellman 2015; Chakraborty and Kearns 2011; Wah and Wellman to appear).

The CDA market maintains a *limit order book* of outstanding orders, and provides information about the book to traders with zero delay. The buy side of the order book starts with BID_t, the highest-price buy order at time t, and extends to lower prices. Similarly, the sell side starts with ASK_t, the lowest-price sell order at time t, and extends to higher prices. When there is an order cancellation or a transaction, the market removes the corresponding orders and updates the order book. Agents may use order book information at their own discretion.

The market is populated with multiple background traders, and in selected treatments, a spoofer. Background traders represent investors with preferences on holding long or short positions in the underlying security. The spoofer seeks trading profits through its price manipulation actions.

The preference of background trader i is defined by its *private value* Θ_i, a vector of length $2q_{\max}$, where $q_{\max}$ is the maximum number of units a trader can be long or short at any time. Private values are subject to diminishing marginal utility and element θ_i^q in the vector specifies the incremental private benefit *foregone* by selling one unit of the security given a current net position of q.

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

Alternatively, θ_i^{q+1} can be understood as the marginal private gain from buying an additional unit given current net position q. To reflect diminishing marginal utility, that is $\theta^{q'} \leq \theta^q$ for all $q' \geq q$, we generate Θ_i from a set of $2q_{\max}$ values drawn independently from $N(0, \sigma_{PV}^2)$, sort elements in descending order, and assign θ_i^q to its respective value in the sorted list.

The entries of a background trader follow a Poisson process with an arrival rate λ_{a}. Upon each entry, a background trader receives a buy or sell signal with equal probability and observes an agent-and-time-specific noisy fundamental $o_t = r_t + n_t$ with the observation noise following $n_t \sim N(0, \sigma_n^2)$. Given its incomplete information about the fundamental, an agent can potentially benefit by considering market information, which is influenced by the aggregate observations of other agents. On each arrival, a background trader withdraws its previous order (if untransacted) and submits a new single-unit limit order. The bid or ask price submitted by a background agent is jointly decided by its *valuation* and *trading strategy*, as discussed in the next section. The spoofing agent, if present, initially arrives at a designated intermediate time $T_{\text{sp}} \in [0, T]$ and executes the manipulation strategy.

Trading Strategies

Estimation of the Final Fundamental. As holdings of the security are evaluated at the end of a trading period, traders estimate the final fundamental value based on their noisy observations. We assume the market environment parameters (mean reversion, shock variance, etc.) are common knowledge for background agents.

Given a new noisy observation o_t, an agent estimates the current fundamental by updating its posterior mean $\tilde{r}_t$ and variance $\tilde{\sigma}_t^2$ in a Bayesian manner. Let t' denote the time of a previous arrival. The updated estimates are given by:

$$\tilde{r}_t = \frac{\sigma_n^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2}\tilde{r}_{t'} + \frac{\tilde{\sigma}_{t'}^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2}o_t \; ; \; \tilde{\sigma}_t^2 = \frac{\sigma_n^2 \tilde{\sigma}_{t'}^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2}.$$

Based on the posterior estimate of $\tilde{r}_t$, the trader computes $\hat{r}_t$, its estimate at time t of the terminal fundamental r_T, by adjusting for mean reversion:

$$\hat{r}_t = \big(1 - (1-\kappa)^{T-t}\big)\bar{r} + (1-\kappa)^{T-t}\tilde{r}_t. \quad (2)$$

ZI as a Background Trading Strategy. ZI traders decide limit-order prices solely based on fundamental observations and private values. The ZI agent submits a bid shaded from its valuation of the security by a random offset, which is uniformly drawn from [$R_{\min}$, $R_{\max}$]. Specifically, a ZI trader i arriving at time t with position q submits a single-unit limit order at price

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

Our version of ZI also takes into account the current quoted price, as governed by a strategic threshold parameter $\eta \in [0, 1]$. Before submitting a new limit order, if the agent could achieve a fraction η of its requested surplus, it would simply take that quote.

HBL as a Background Trading Strategy. HBL agents go beyond their own observations and private values by also considering order book information. The strategy is centered on belief functions that traders form on the basis of observed market data. Agents estimate the probability that orders at various prices would be accepted in the market, and choose a limit price maximizing its own expected surplus at its current valuation estimate.

The HBL agent's probability estimate is based on observed frequencies of accepted and rejected bids and asks during the last L trades, where L, the agent's *memory length*, is a strategic parameter. On arrival at time t, the HBL agent builds a belief function $f_t(\mathrm{P})$, designed to represent the probability that an order at price P will result in a transaction. Specifically, the belief function is defined for encountered prices P by:

$$f_t(\mathrm{P}) = \begin{cases} \frac{\mathrm{TBL}_t(\mathrm{P})+\mathrm{AL}_t(\mathrm{P})}{\mathrm{TBL}_t(\mathrm{P})+\mathrm{AL}_t(\mathrm{P})+\mathrm{RBG}_t(\mathrm{P})} & \text{if buying} \\ \\ \frac{\mathrm{TAG}_t(\mathrm{P})+\mathrm{BG}_t(\mathrm{P})}{\mathrm{TAG}_t(\mathrm{P})+\mathrm{BG}_t(\mathrm{P})+\mathrm{RAL}_t(\mathrm{P})} & \text{if selling.} \end{cases} \tag{3}$$

Here, *T* and *R* specify *transacted* and *rejected* orders respectively; *A* and *B* represent *asks* and *bids*; *L* and *G* describe orders with prices *less* than or equal to and *greater* than or equal to price P correspondingly. For example, $\mathrm{TBL}_t(\mathrm{P})$ is the number of transacted bids found in memory with price less than or equal to P up to time t. Agents compute the statistics upon each arrival and update their memory whenever the market receives new order submissions, transactions, or cancellations.

Since our market model supports order cancellations and keeps active orders in the order book, the notion of a rejected order is tricky to define. To solve this problem, we introduce a grace period τ_{gp} and an alive period τ_{al} of an order. We define the grace period $\tau_{\text{gp}} = 1/\lambda_a$ and the alive period τ_{al} of an order as the time interval from submission to transaction or withdrawal if it is inactive, or to the current time if active. An order is considered as rejected only if its alive period τ_{al} is longer than τ_{gp}, otherwise it is partially rejected by a fraction of $\tau_{\text{al}}/\tau_{\text{gp}}$.

As the belief function (3) is defined only at encountered prices, we extend it over the full domain by cubic spline interpolation. To speed the computation, we pick knot points and interpolate only between those points.

After formulating the belief function, agent i with an arrival time t and current holdings q searches for the price $\mathrm{P}_i^*(t)$ that maximizes expected surplus:

$$\mathrm{P}_i^*(t) = \begin{cases} \arg\max_p (\hat{r}_t + \theta_i^{q+1} - p) f_t(p) & \text{if buying} \\ \arg\max_p (p - \theta_i^{q} - \hat{r}_t) f_t(p) & \text{if selling.} \end{cases}$$

Under the special cases when there are fewer than L transactions at the beginning of a trading period or when one side of the order book is empty, HBL agents behave the same as ZI agents until enough information is gathered to form the belief function. As those cases are rare, the specific ZI strategy that HBL agents adopt will not affect overall performance.

Spoofing Strategy. We design a simple spoofing strategy which maintains a large volume of buy orders at one tick behind the best bid. Specifically, upon arrival at $T_{\text{sp}} \in [0, T]$, the spoofing agent submits a buy order at price $\mathrm{BID}_{T_{\text{sp}}} - 1$ with volume $Q_{\text{sp}} \gg 1$. Whenever there is an update on the best bid, the spoofer cancels its original spoof order and submits a new one at price $\mathrm{BID}_t - 1$ with the same volume. As background traders submit only single-unit orders, they cannot transact with the spoof order, which is shielded by a higher order at $\mathrm{BID}_{T_{\text{sp}}}$. If that higher order gets executed, the spoofer immediately cancels and replaces its spoof orders before another background trader arrives. We assume in effect that the spoofer can react infinitely fast, in which case its spoof orders are guaranteed never to transact.

By continuously feigning buy interest in the market, this spoofing strategy specifically aims to raise market beliefs. Other spoofing strategies such as adding sell pressure or alternating between buy and sell pressure can be easily constructed by extension from the current version.

Valuation Model

We calculate market surplus as the sum of agents' surpluses at the end of the trading period T. An agent's total surplus is the sum of *cash* paid or gained during trading period and the *final valuation* of holdings. The market's final valuation of trader i with a long position L ($L > 0$) is $r_T \times L + \sum_{k=1}^{k=L} \theta_i^k$ and similarly, the valuation of a trader j with a short position S ($S < 0$) is $r_T \times S - \sum_{k=S+1}^{k=0} \theta_j^k$.

Experiments and Results

Experiments are conducted by simulating the market model described above. We generate data for different games, each of which is comprised of a market environment and a strategy profile specifying the number of background agents playing each strategy. We sample sufficiently large number of runs for each game to account for stochastic effects (market fundamental series, agent arrival patterns, valuations, etc.). Given a specific market environment, we evaluate background-trader performance and the impact of spoofing in empirical Nash equilibrium, where agents have no incentive to deviate to other available strategies, given others' choices.[2]

Market Environment Settings

Our simulations consider nine parametrically distinct environments that differ in market shock variances $\sigma_s^2 \in$

[2]In all of our experiments, outcome features (including agent payoff, surplus, and price discovery) are calculated as the average of 20,000 simulations of games with strategy profiles sampled as the specified equilibrium mixture.

Table 1: Background trading strategies included in empirical game-theoretic analysis.

Strategy	ZI_1	ZI_2	ZI_3	ZI_4	ZI_5	ZI_6	ZI_7	HBL_1	HBL_2	HBL_3	HBL_4
L	NA	NA	NA	NA	NA	NA	NA	2	3	5	8
R_{min}	0	0	0	0	0	250	250	250	250	250	250
R_{max}	250	500	1000	1000	2000	500	500	500	500	500	500
η	1	1	0.8	1	0.8	0.8	1	1	1	1	1

$\{10^5, 5 \times 10^5, 10^6\}$ and noisy observation variances $\sigma_n^2 \in \{10^3, 10^6, 10^9\}$. A higher shock variance means larger intrinsic fluctuations in the fundamental time series, whereas a higher observation variance implies agents receive less information of the true fundamental. Explorations regarding the significance of different environment parameters suggest market shock and observation noise are the most relevant and sensitive ones to our study. We label the three low, medium and high shock variances as $\{\mathrm{A}, \mathrm{B}, \mathrm{C}\}$ and noisy observation variances as $\{1, 2, 3\}$ respectively to describe the nine environments. For example, A1 represents a market with low shock $\sigma_s^2 = 10^5$, and low observation noise $\sigma_n^2 = 10^3$.

In all environments, we further consider markets with $N \in \{28, 65\}$ background traders and in selected treatments, a spoofer. The global fundamental time series is generated according to (1) with a fundamental mean $\bar{r} = 10^5$, a mean-reverting parameter $\kappa = 0.05$ and a specific shock variance σ_s^2. Market has a minimum tick size of one and each trading period lasts $T = 10,000$ time steps. Background traders arrive at the market according to a Poisson distribution with a rate $\lambda_a = 0.005$ and upon each arrival, the trader observes a noisy fundamental o_t. The maximum number of units background traders can hold at any time is $q_{\text{max}} = 10$. Private values are drawn from a Gaussian distribution with zero mean and a variance of $\sigma_{\text{PV}}^2 = 5 \times 10^6$. The spoofing agent initially arrives at time $\mathrm{T}_{\text{sp}} = 1000$, submits a large buy order at price $\mathrm{BID}_{\mathrm{T}_{\text{sp}}} - 1$ with volume $Q_{\text{sp}} = 200$ and later maintains spoofing orders at price $\mathrm{BID}_t - 1$.

The background trading strategy set (Table 1) includes seven versions of ZI and four versions of HBL.[3] Agents are allowed to choose from this restricted set of strategies to maximize their payoffs.

EGTA Process

Each market game is defined by a market environment and multiple players partitioned into two roles, background traders and a spoofer. We are interested in agents' strategic choices in Nash equilibrium of a game. The payoff of a specific strategy is the average of payoffs of all agents playing that strategy, and thus only depends on the number of agents playing each strategy, not on individual mappings.

As game size grows exponentially in the number of players and strategies, it is computationally prohibitive to analyze games with this many traders. We therefore apply aggregation to approximate the many-player games as games with fewer players. The specific technique we employ, called deviation-preserving reduction (DPR) (Wiedenbeck and Wellman 2012), defines reduced-game payoffs in terms of payoffs in the full game as follows. Consider an N-player symmetric game, which we want to reduce to a k-player game. The payoff for playing strategy s_1 in the reduced game, with other agents playing strategies $(s_2, \ldots, s_k)$, is given by the payoff of playing s_1 in the full N-player game when the other $N-1$ agents are evenly divided by among strategies $s_2, \ldots, s_k$. To facilitate DPR, we choose values for N to ensure that the required aggregations come out as integers. Specifically, in this study we reduce the market environments with 28 (65) background traders and a spoofer to games with four (five) background traders and one spoofer. With one background player deviating to a new strategy, we can reduce the remaining 27 (64) players to three (four).

To find the Nash equilibrium of a game, we run simulations to get payoffs for background strategies and conduct EGTA based on those payoffs. Exploration starts with games where all players in a role adopt the same strategy, and spreads to other strategies by single-agent deviations. Equilibria found in each subgame are considered as candidates of the full game. We can refute these candidates by finding a beneficial deviation outside the subgame strategy set, or confirm by examining all deviations without refuting. We continue to refine the empirical subgame with additional strategies and corresponding simulations until at least one equilibrium is confirmed and all non-confirmed candidates are refuted.

Games without Spoofing

Since spoofing targets the order book and can be effective only to the extent traders exploit order book information, we first investigate whether background agents adopt the HBL strategy at equilibrium in games without spoofing. Applying EGTA to the eleven background strategies in Table 1, we found at least one equilibrium for each market environment.[4]

As indicated in Figure 4, HBL is adopted with positive probability by background traders in most non-spoofing environments. That is, in the absence of spoofing, investors generally have incentives to make bidding decisions based on order book information. We find that HBL is robust and widely preferred in markets with more traders, low fun-

[3]We also considered ZI strategies with larger shading ranges and HBL strategies with longer memory lengths, but they fail to appear in equilibrium.

[4]Details of the HBL adoption rates and market surpluses of all found equilibria in games with and without spoofing are available in an online appendix at https://goo.gl/nRfF4L

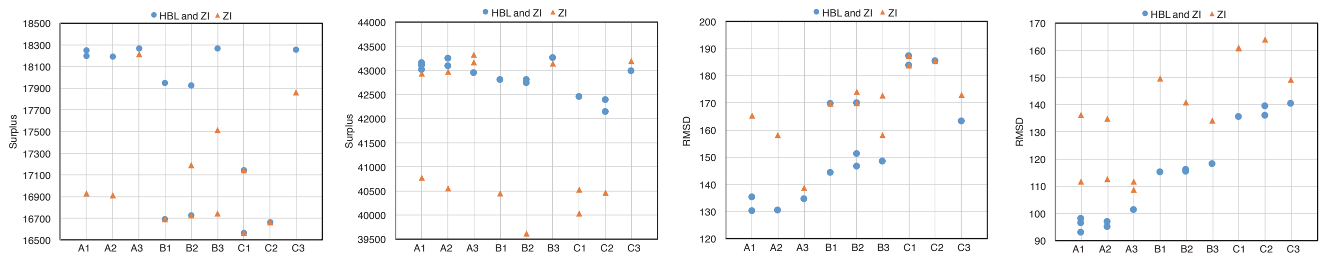

(a) N = 28 (b) N = 65 (c) N = 28 (d) N = 65

Figure 1: Comparisons of market surplus (Figures 1a and 1b) and price discovery (Figures 1c and 1d) for equilibrium in each environment, with and without the HBL strategies available to background traders. Blue circles represent equilibrium outcomes when agents can choose both HBL and ZI strategies; orange triangles represent equilibrium outcomes when agents are restricted to ZI strategies. Overlapped markers are outcomes from the same equilibrium mixture, despite the availability of HBL.

damental shocks, and high observation noise. Intuitively, a larger population size implies a thick order book with more learnable aggregated data; low shocks in fundamental time series increase the predictability of future price outcomes; and high observation noise limits what an agent can glean about the true fundamental from its own information. The two exceptions (environments C1 and C2 with 28 background traders) where all agents choose ZI can be explained by the environments' small population size, large fundamental shocks, and relatively small observation noise.

We further conduct EGTA in games where background traders are restricted to strategies in the ZI family ($\mathrm{ZI}_1 - \mathrm{ZI}_7$ in Table 1). This is tantamount to disallowing learning from order book information. To understand the effect of order book disclosure on market performance, we compare equilibrium outcomes for each environment, with and without the HBL strategy set available to background traders, on two measures: market surplus and price discovery (Figure 1). Price discovery reflects how well transactions reveal the true value of the security; it is defined as the root-mean-squared deviation (RMSD) of the transaction price from the estimate of the true fundamental (as calculated by (2)) over the trading period. Lower RMSD indicates better price discovery.

Overall in our experiments, background traders achieve higher surplus (Figures 1a and 1b) and better price discovery (Figures 1c and 1d) when the market provides order book information and the HBL strategy option. When the equilibrium includes HBL, we find transactions reveal fundamental estimates well, especially in markets with lower shock and observation variances (as in Figures 1c and 1d, where blue circles at lower left have low RMSDs). We also notice small exceptions in scenarios with high observation variance and more background traders (environment A3 and C3 with 65 players) where ZI-only equilibria exhibit higher surplus than equilibria combining HBL and ZI.

Games with Spoofing

Spoofing the HBLs. To examine the effectiveness of the designed spoofing strategy, we play a spoofer against each equilibrium found in games without spoofing and perform controlled experiments upon a pair of games with and without spoofing. In the paired games, background agents are guaranteed to arrive at the same time, receive identical private values, and observe the same fundamental values. Therefore, any change in HBL's bidding behavior is caused by the spoof orders. For every paired games, we run 20,000 simulations for each and compare their transaction price differences (Figure 2), and surplus differences attained by HBL and ZI traders respectively (Figure 3). Transaction price difference at a specific time is defined as the most recent transaction price of a game with spoofing minus that of the paired game without spoofing. Similarly, surplus difference of HBL or ZI is the aggregated surplus obtained in a game with spoofing minus that of the paired game without spoofing.

Figure 2 shows consistent positive changes in transaction prices since the arrival of a spoofing agent at $t = 1000$ across all environments.[5] This suggests HBL traders are tricked by the spoof buy orders: they believe the underlying security should be worth more and therefore submit or accept limit orders at higher prices. Though ZI agents do not change their bidding behavior directly, they may be passively affected and make transactions at higher prices.

Several other interesting findings are revealed by the transaction price difference series (Figure 2). *First*, the average price rise caused by spoofing in market with 28 background traders is higher than that of the 65-background-trader market. This indicates markets with less background traders are generally more susceptible to spoofing, possibly due to the limited pricing information a thin market could aggregate. *Second*, for markets populated with more HBLs than ZIs, the transaction price differences increase throughout the trading period. This exacerbated spoofing effect can be explained by HBLs consistently submitting orders at higher prices and confirming each other's spoofed belief. However, for markets with more ZIs, the spoofing effect diminishes as ZIs who do not change their limit-order pricing can partly correct HBLs' illusions. *Third*, throughout the trading horizon, differences in transaction prices first increase and then stabilize or decrease as time approaches the end of a trading period. This suggests as time approaches $\mathrm{T} = 10,000$, HBL agents may better estimate the termi-

[5]As spoofing has no impact on pure ZI populations, we display transaction price differences only for equilibria with HBL.

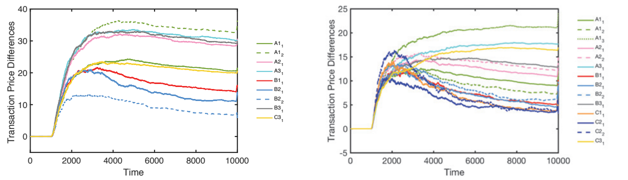

(a) N = 28 (b) N = 65

Figure 2: Transaction price differences throughout the trading horizon with and without a spoofer against each HBL-and-ZI equilibrium. Multiple lines of the same environment represent different equilibria.

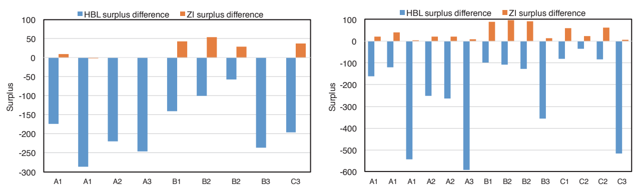

(a) N = 28 (b) N = 65

Figure 3: Background trader surplus differences in markets with and without a spoofer against each HBL-and-ZI equilibrium. Repetitions of the same market environment represent outcomes of multiple equilibria.

nal value from observations given the mean-reverting adjustment and thus are less spoofed.

Figure 3 demonstrates a redistribution of surplus between HBL and ZI agents when we include a spoofer: HBL's aggregated surplus decreases, while ZI's total surplus increases compared to those of the non-spoofing games. This implies that ZI can take advantages of HBL's spoofed beliefs to profit more. Since the decreases in HBL's surplus are consistently larger than the increases of ZI's, the overall market surplus decreases. However, we leave the discussion of spoofing's impact on market surplus to the next section, where background traders can choose other strategies to adjust to spoofing. We also find that markets with a spoofer against background traders at equilibrium have statistically significantly higher RMSDs, which affirms the notion that spoofing, as a deceptive practice, hurts price discovery.

To examine the potential to profit from a successful price manipulation, we extend the spoofing agent with an exploitation strategy: buying, then (optionally) spoofing to raise the price, then selling. It starts by buying when it finds a limit sell order with price less than the fundamental mean. It then optionally runs the spoofing trick, or alternatively waits, for 1000 time steps. Finally, the exploiter sells when it finds a limit buy order with price more than fundamental mean. Note that even without the spoof, this exploitation strategy is profitable in expectation due to the mean reversion, and the reliable arrival of background traders willing to sell at prices better than the fundamental.

In controlled experiments, we find that exploitation profits are consistently increased when the spoof is also deployed. Specifically, across 28-trader market environments, the exploiter makes an average profit of 206.1 and 201.8 with and without spoofing, and the increases in profit range from 1.2 to 11.5. For the 65-trader market, the average profits of this exploitation strategy with and without spoofing are 50.5 and 46.3 respectively, with the increases in profit varying from 1.7 to 9.4 across environments.[6]

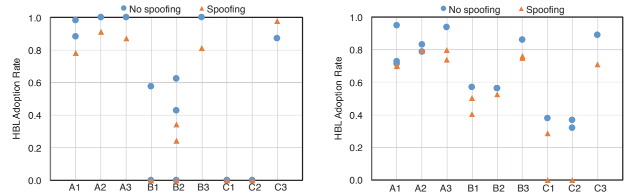

(a) N = 28 (b) N = 65

Figure 4: HBL adoption rates at equilibria in games with and without spoofing. Each blue (orange) marker specifies the HBL proportion at one equilibrium found in a specific game environment without (with) spoofing.

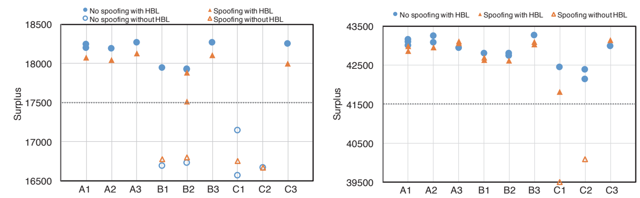

(a) N = 28 (b) N = 65

Figure 5: Total surplus achieved at equilibria in games with and without spoofing. Each blue (orange) marker specifies the surplus at one equilibrium found in a specific game environment without (with) spoofing. Surplus achieved at equilibrium combining HBL and ZI and equilibria with pure ZI are indicated by markers with and without fills respectively.

Re-equilibrating Games with Spoofing. To understand how spoofing affects background-trader interactions, we conduct EGTA again to find Nash equilibrium in games with spoofing, where background traders can choose any strategy in Table 1. As indicated in Figure 4, after re-equilibrating games with spoofing, HBL is generally adopted by a smaller fraction of traders, but may still persist at equilibrium in most market environments. HBL's existence after re-equilibration indicates a consistently spoofable market: the designed spoofing tactic fails to eliminate HBL agents and in turn, the persistence of HBL may incentivize a spoofer to continue effectively manipulating the market.

Finally, we investigate the effect of spoofing on market surplus. Figure 5 compares the total surplus achieved by background traders in equilibrium with and without spoofing. It reveals several interesting findings. *First*, given the presence of HBL traders, spoofing generally decreases total

[6] Statistical tests show all increases in profit are significantly larger than zero. Regardless of spoofing, the exploitation strategy profits more in the thinner market due to the greater variance in transaction prices.

surplus (as in Figure 5, most filled orange triangles are below the filled blue circles). However, spoofing has ambiguous effect in the thicker market with large observation variance (environment A3 and C3 with 65 background agents). This may be because noise and spoofing simultaneously hurt the prediction accuracy of the HBL agents and therefore shift agents to other competitive ZI strategies with higher payoffs. *Second*, we find the welfare effects of HBL strategies persist regardless of spoofing's presence: markets populated with HBL agents in equilibrium generally achieve higher total surplus than those markets without HBL (as in Figure 5, the hollow markers are below the filled markers).

Conclusion

We constructed a computational model of spoofing market prices by targeting the order book. To do so, we design an HBL strategy that uses order book information to make pricing decisions. Since HBL traders use the order book, they are spoofable, which we confirmed in simulation analysis. We demonstrate that in the absence of spoofing, HBL is generally adopted in equilibrium and benefits price discovery and social welfare. Though the presence of spoofing decreases the HBL proportion in background traders, HBL's persistence in equilibrium indicates a robustly spoofable market. By comparing equilibrium outcomes with and without spoofing, we find spoofing tends to decrease market surplus. Comparisons across parametrically different environments reveal factors that may influence the adoption of HBL and the impact of spoofing. Our agent-based model aims to capture the complex essence of real-world financial markets and the strategic interactions among investors.

We acknowledge several factors that limit the accuracy of our equilibrium analysis in individual game instances; these include sampling error, reduced-game approximation, and restricted strategy coverage. Despite such limitations (inherent in any complex modeling effort), we believe the model offers a constructive basis for other researchers, regulators, and policymakers to better evaluate spoofing and understand its interplay with other trading strategies.

Acknowledgement

This work was supported in part by NSF grant IIS-1421391.

References

Aldrich, E. M.; Grundfest, J.; and Laughlin, G. 2016. The Flash Crash: A new deconstruction. *Available at SSRN 2721922*.

Bookstaber, R. 2012. Using agent-based models for analyzing threats to financial stability. Working Paper 3, Office of Financial Research.

Brush, S.; Schoenberg, T.; and Ring, S. 2015. How a mystery trader with an algorithm may have caused the Flash Crash. *Bloomberg News*.

Chakraborty, T., and Kearns, M. 2011. Market making and mean reversion. In *12th ACM Conference on Electronic Commerce*, 307–314.

Cliff, D. 1997. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report, Hewlett-Packard Labs.

Cliff, D. 2009. ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computation* 13(1):3–18.

Farmer, J. D.; Patelli, P.; and Zovko, I. I. 2005. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences* 102(6):2254–2259.

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

Gjerstad, S., and Dickhaut, J. 1998. Price formation in double auctions. *Games and Economic Behavior* 22(1):1–29.

Gjerstad, S. 2007. The competitive market paradox. *Journal of Economic Dynamics and Control* 31(5):1753–1780.

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

LeBaron, B.; Arthur, W. B.; and Palmer, R. 1999. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23(9):1487–1516.

LeBaron, B. 2006. Agent-based computational finance. In Tesfatsion, L., and Judd, K. L., eds., *Handbook of Computational Economics*, volume 2. Elsevier. 1187–1233.

Lee, E. J.; Eom, K. S.; and Park, K. S. 2013. Microstructure-based manipulation: Strategic behavior and performance of spoofing traders. *Journal of Financial Markets* 16(2):227–252.

Li, Z., and Das, S. 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 Multiagent Systems*, 50–58.

Martinez-Miranda, E.; McBurney, P.; and Howard, M. 2016. Learning unfair trading: A market manipulation analysis from the reinforcement learning perspective. In *IEEE International Conference on Evolving and Adaptive Intelligent Systems*.

Paddrik, M.; Hayes, Jr., R.; Todd, A.; Yang, S.; Beling, P.; and Scherer, W. 2012. An agent based model of the E-Mini S&P 500 applied to Flash Crash analysis. In *IEEE Conference on Computational Intelligence for Financial Engineering and Economics*, 1–8.

Palit, I.; Phelps, S.; and Ng, W. L. 2012. Can a zero-intelligence plus model explain the stylized facts of financial time series data? In *11th International Conference on Autonomous Agents and Multiagent Systems*, 653–660.

Tesauro, G., and Bredin, J. L. 2002. Strategic sequential bidding in auctions using dynamic programming. In *First International Joint Conference on Autonomous Agents and Multiagent Systems*, 591–598.

Tesauro, G., and Das, R. 2001. High-performance bidding agents for the continuous double auction. In *Third ACM Conference on Electronic Commerce*, 206–209.

Vytelingum, P.; Cliff, D.; and Jennings, N. R. 2008. Strategic bidding in continuous double auctions. *Artificial Intelligence* 172(14):1700–1729.

Wah, E., and Wellman, M. P. 2015. Welfare effects of market making in continuous double auctions. In *14th International Conference on Autonomous Agents and Multiagent Systems*, 57–66.

Wah, E., and Wellman, M. P. to appear. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance*.

Wang, Y. 2015. Strategic spoofing order trading by different types of investors in the futures markets. *Wall Street Journal*.

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

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