

Article

Spoofing the Limit Order Book: A Strategic Agent-Based Analysis

Xintong Wang ^{1,†} , Christopher Hoang ², Yevgeniy Vorobeychik ³  and Michael P. Wellman ^{4,*} 

¹ John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA; xintongw@seas.harvard.edu

² Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109, USA; choang@umich.edu

³ Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130, USA; yvorobeychik@wustl.edu

⁴ Lynn A. Conway Collegiate, Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109, USA

* Correspondence: wellman@umich.edu

† This work was completed while Xintong Wang was a research assistant at the University of Michigan.

Abstract: We present an agent-based model of manipulating prices in financial markets through *spoofing*: submitting spurious orders to mislead traders who learn from the order book. Our model captures a complex market environment for a single security, whose common value is given by a dynamic fundamental time series. Agents trade through a limit-order book, based on their private values and noisy observations of the fundamental. We consider background agents following two types of trading strategies: the non-spoofable *zero intelligence* (ZI) that ignores the order book and the manipulable *heuristic belief learning* (HBL) that exploits the order book to predict price outcomes. We conduct empirical game-theoretic analysis upon simulated agent payoffs across parametrically different environments and measure the effect of spoofing on market performance in approximate strategic equilibria. We demonstrate that HBL traders can benefit price discovery and social welfare, but their existence in equilibrium renders a market vulnerable to manipulation: simple spoofing strategies can effectively mislead traders, distort prices and reduce total surplus. Based on this model, we propose to mitigate spoofing from two aspects: (1) mechanism design to disincentivize manipulation; and (2) trading strategy variations to improve the robustness of learning from market information. We evaluate the proposed approaches, taking into account potential strategic responses of agents, and characterize the conditions under which these approaches may deter manipulation and benefit market welfare. Our model provides a way to quantify the effect of spoofing on trading behavior and market efficiency, and thus it can help to evaluate the effectiveness of various market designs and trading strategies in mitigating an important form of market manipulation.

Keywords: market manipulation; agent-based simulation; trading agents; empirical game-theoretic analysis



Citation: Wang, X.; Hoang, C.; Vorobeychik, Y.; Wellman, M.P. Spoofing the Limit Order Book: A Strategic Agent-Based Analysis. *Games* **2021**, *12*, 46. <https://doi.org/10.3390/g12020046>

Academic Editor: Isa Hafalir

Received: 19 February 2021

Accepted: 22 April 2021

Published: 24 May 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Financial exchanges nowadays operate almost entirely electronically, supporting automation of trading and consequential scaling of volume and speed across geography and asset classes. With data and information streaming on an extremely short timescale, often below the limits of human response time, autonomous trading agents directed by algorithms operate on behalf of human traders. Such increasing automation has transformed the financial market landscape from a human decision ecosystem to an algorithmic one, where autonomous agents learn new information, make decisions and interact with each other at an unprecedented speed and complexity. Whereas these developments in market operation and trading technology may contribute to improved efficiency, they also introduce risks, such as the potential for new forms of manipulative practice driven by algorithms.

Market manipulation is defined by the U.S. Securities and Exchange Commission (SEC) as “intentional or willful conduct designed to deceive or defraud investors by controlling or artificially affecting the price of securities, or intentional interference with the free forces of supply and demand”. Although it has long been present, the practice has also evolved in its forms to exploit automated trading and the dissemination of market information offered by many trading platforms [1]. Computer programs are employed to inject deceitful information, as other investors use algorithms to extract information from all possible sources (including the misleading ones) and execute decisions accordingly. On 21 April 2015, nearly five years after the “Flash Crash”—a sudden trillion-dollar dip in U.S. stock markets on 6 May 2010, during which stock indexes collapsed and rebounded rapidly [2], the U.S. Department of Justice charged Navinder Singh Sarao with 22 criminal counts, including fraud and market manipulation. Prior to the Flash Crash, Sarao allegedly used an algorithm to place orders amounting to about \$200 million seemingly betting 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. Although recent analysis has cast doubt on the causal role of Sarao on the Flash Crash [3], 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, *spoofing*, operates through a series of direct trading actions in a market. Traders interact with the market by submitting orders to buy or sell. Orders that do not transact immediately rest in the *order book*, a repository for outstanding orders to trade. At any given time, the order book for a particular security reflects the market’s expressed supply and demand. Spoofing refers to the practice of submitting large *spurious* buy or sell orders with the intent to cancel them before execution. The orders are spurious in that instead of expressing genuine trading intent, they feign a strong buy or sell interest in the market, thus corrupting the order book’s signal on supply and demand. Other traders are then misled by the spoof orders 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 manipulator 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. Figure 1 illustrates an alleged spoofing activity conducted over the course of 0.6 second, demonstrating how quickly and effectively such manipulation behavior can affect the market and profit from the spoofed belief.




Figure 1. Example of alleged spoofing. Source: UK Financial Conduct Authority Final Notice 2013. A series of large out-of-the money manipulation sell orders (red triangles) are first placed to drive the price down and make the buy order accepted (the filled blue triangle). These sell orders are immediately replaced with large buy ones (blue triangles) to push the price up and profit from the sale at a higher price (the filled red triangle).

In 2010, the Dodd–Frank Wall Street Reform and Consumer Protection Act was signed into U.S. law, outlawing spoofing as a deceptive practice. In describing its concern about spoofing, 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. In ideal markets without manipulation, traders may extract useful information from the order book, making more informed decisions over those that neglect such information. A manipulator exploits such learning process, minimizing its own risk in the process. Spoof orders are typically placed at price levels just outside the current best quotes to mislead other investors and withdrawn with high probability before any market movement could trigger a trade [4,5].

We aim to reproduce spoofing in a computational model, as a first step toward developing more robust measures to characterize and prevent spoofing. Figure 2 gives an overview of our agent-based market model. The model implements 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 [6]. Traders can submit limit 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 manipulator who executes the spoofing strategy. Background traders are further divided to follow two types of trading strategies: *zero intelligence* (ZI) that ignores the order book and *heuristic belief learning* (HBL) that learns from the order book to predict price outcomes. Upon each arrival to trade, a background trader receives a noisy observation of the security’s fundamental value. Based on a series of fundamental observations and its private value, a ZI agent computes the limit-order price by shading a random offset from its valuation, and thus it is non-manipulable. An HBL agent, on the other hand, is susceptible to spoofing: it considers information about orders recently submitted to the market, estimates the probability that orders at various prices would be transacted and chooses the optimal price to maximize expected surplus. The manipulator in our model executes a spoofing strategy similar to that illustrated in Figure 1. The spoofer injects and maintains large spurious buy orders at one tick behind the best bid, designed to manipulate the market by misleading others about the level of demand.




Figure 2. An agent-based model of spoofing a CDA market with a single security traded.

We conduct extensive simulation over hundreds of strategy profiles across parametrically different market environments with and without manipulation. The simulation data are used to estimate normal-form game models over the strategies explored in agent-based simulation. From these models, we derive empirical equilibria, where every agent chooses its best response within the set of available strategies to both the market environment and others’ behavior. Studying behavior in (empirical) equilibrium provides robustness to designer choices in agent-based modeling, selecting behaviors based on a rationality crite-

rion [7]. Although the strategies considered in this game model are restricted compared to the true underlying game (where any conceivable trading strategy is possible), imposing a rationality filter on the set considered accounts for strategic response to different settings and ensures that we are evaluating the most relevant configurations of available strategies.

Our fundamental goals of this work are: (1) to reproduce spoofing and understand its impact on market performance (Section 5); and (2) to propose and evaluate variations of market designs (Section 6) and learning-based trading strategies (Section 7) in mitigating manipulation. Below, we overview the structure of the paper and summarize our main contributions and results.

Roadmap

We start by reproducing spoofing in an agent-based model and evaluating its impact on background traders. Section 3 introduces the design of our CDA market model and parameterized trading strategies. Section 4 describes the empirical game-theoretic analysis (EGTA) methodology [7], which we adopt for finding equilibria in games defined by heuristic strategy space and simulated payoff data. Section 5 addresses the choice of background traders among HBL and ZI strategies in markets with and without spoofing. We demonstrate through EGTA that, in a range of non-spoofing environments, HBL is preferred in equilibrium and benefits price discovery and social welfare. However, this renders a market vulnerable to manipulation: by executing a spoofer against the equilibrium profiles, we show that simple spoofing strategies can manipulate prices in a desired direction. After re-equilibrating games with spoofing, we find HBL still persists in equilibria but with smaller mixture probability, suggesting a consistently spoofable market. Although the welfare benefits of HBL remain, the presence of spoofing decreases market surplus.

Building on our computational model of spoofing, we investigate market mechanisms and trading strategies to mitigate spoofing without directly detecting each individual activity. In Section 6, we propose a *cloaking mechanism* to deter spoofing. It adapts the way a standard order book discloses market information by symmetrically concealing a specified number of price levels from the inside of the book. The idea is to make it more difficult for the spoofer to post misleading bids, while not unduly degrading the general usefulness of market information. We characterize market conditions under which such cloaking may mitigate manipulation and benefit market welfare. We further design sophisticated spoofing strategies that probe to reveal cloaked information and demonstrate that the effort and risk of probing exceed the gains.

In Section 7, we explore two variations of the standard HBL strategy to reduce the vulnerability of learning traders to spoofing. The first selectively ignores orders at certain price levels, particularly where spoof orders are likely to be placed. The second considers the full order book, but adjusts its limit order price to correct for bias in decisions based on the learned heuristic beliefs. We evaluate these variations on two criteria: effectiveness in non-manipulated markets and robustness against manipulation.

2. Related Work

2.1. Agent-Based Modeling of Financial Markets

Agent-based modeling (ABM) takes a simulation approach to study complex domains with dynamically interacting decision makers. ABM has been frequently applied to modeling and understanding phenomena in financial markets [8], for example to study the Flash Crash [9] or to replicate the volatility persistence and leptokurtosis characteristic of financial time series [10]. A common goal of agent-based finance studies is to reproduce stylized facts of financial market behavior [11] and to support causal reasoning about market environments and mechanisms. Researchers have also used ABM to investigate the effects of particular trading practices, such as market making [12] and latency arbitrage [13]. ABM advocates argue that simulation is particularly well-suited to study financial markets [14], 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.

2.2. Autonomous Bidding Strategies

There is a substantial literature on autonomous bidding strategies in CDA markets [15]. The basic *zero intelligence* (ZI) strategy [16] submits offers at random offsets from valuation. Despite its simplicity, ZI has been shown surprisingly effective for modeling some cases [17]. In this study, we adopt an extended and parameterized version of ZI to represent 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 [18,19]. Vytelingum et al. [20] introduced another level of strategic adaptation, allowing the agent to control its behavior with respect to short and long time scales. We note that to some extent, the adaptive functions of these strategies are implicitly achieved by the game-theoretic equilibration process which we employ to determine the parametric configurations of the (non-adaptive) trading strategies [21].

Gjerstad proposed a more direct approach to learning from market observations, termed *GD* in its original version [22] and named *heuristic belief learning* (HBL) in a subsequent generalized form [23]. 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. For example, Tesauro and Das [24] adapted the strategy to markets that support persistent orders. Tesauro and Bredin [25] showed how to extend beyond myopic decision making by using dynamic programming to optimize the price and timing of bids.

We adopt HBL as our representative class of agent strategies that exploit order book information. HBL can be applied with relatively few tunable strategic parameters, compared to other adaptive strategies in the literature. We extend HBL to a more complex market environment that supports persistent orders, combined private and fundamental values, noisy observations, stochastic arrivals and the ability to trade multiple units with buy or sell flexibility. The extended HBL strategy considers the full cycle of an order, including the times an order is submitted, accepted, canceled or rejected.

2.3. 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. [26] empirically examined 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 investigated spoofing on the index futures market in Taiwan, identifying strategy characteristics, profitability and real-time impact [27]. Martinez-Miranda et al. [28] implemented spoofing behavior within a reinforcement learning framework to model conditions where such behavior is effective. Tao et al. [29] presented a micro-structural study of spoofing in a static setting, providing conditions under which a market is more likely to admit spoofing behavior as a function of the characteristics of the market.

To our knowledge, we provide the first computational model of spoofing a dynamic financial market and demonstrate the effectiveness of spoofing against approximate-equilibrium traders in this proposed model. Our model provides a way to quantify the effect of spoofing on trading behavior and efficiency, and thus it is a first step in the design of methods to deter or mitigate market manipulation.

2.4. Connections between Spoofing and Adversarial Machine Learning

If we view the market comprised of HBL and ZI agents as mapping a sequence of orders over time—distilled into a collection of features used by the HBL agents—to a market price, spoofing then can be viewed as a variant of *decision-time attacks on machine learning models* [30–33]. Prior work has also used conditional GANs to generate order streams submitted by HBL and ZI agents in market aggregate [34]. These attacks, commonly called *adversarial examples*, have demonstrated vulnerabilities in a broad array of algorithmic models from linear classification [35,36] to deep neural networks [32,37] and across a variety of problem domains, including vision [30,37,38], speech and natural language processing [39,40] and malware detection [41,42].

The spoofing attacks we study here can be viewed as examples of *realizable attacks* that explicitly account for domain constraints [40,42–46]. In our case, for example, manipulation is conducted in the form of submitting and canceling orders, and thus it only indirectly impacts the features extracted by HBL agents. Furthermore, an important aspect of the market context that is central to our analysis, yet rarely considered in prior adversarial learning literature, is market equilibrium behavior; a notable exception is a game-theoretic analysis of adversarial linear regression with multiple learners [47].

Finally, the proposed variations of HBL that increase its robustness to market manipulation are in the spirit of the literature investigating robustness of machine learning to decision-time attacks [42,48,49]. The principle difference is the particular attention we pay to the market structure in balancing robustness and efficacy of learning from order information.

3. Market Model

We present the general structure of the agent-based financial market environment in which we model spoofing. Our model comprises agents trading a single security through a continuous double auction (CDA), the mechanism adopted by most financial markets today. We first describe the market mechanism in Section 3.1. Our model is designed to capture key features of market microstructure (e.g., fundamental shocks and observation noise), supporting a configurable simulator to understand the effect of spoofing under different market conditions. The market is populated with multiple background traders who represent investors in the market, and in selected treatments, a spoofer who seeks trading profit through manipulative action. We specify the valuation model of background traders in Section 3.2 and the two families of background-trader strategies in Section 3.3. In Section 3.4, we discuss the behavior of the spoofing agent.

3.1. Market Mechanism

The market employs a CDA mechanism with a single security traded. Prices are fine-grained and take discrete values at integer multiples of the tick size. Time is also fine-grained and discrete, with trading over 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. Orders are immediately matched as they arrive: if at any time, one agent's maximum price to buy a unit is greater than or equal to another agent's minimum price to sell a unit, a transaction will occur and the agents trade at the price of the incumbent order.

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. On order cancellation or transaction, the market removes the corresponding orders and updates the order book. Agents may use order book information at their own discretion. In Section 6, we investigate how changes made in such order book disclosure may help to mitigate spoofing.

3.2. Valuation Model

Our valuation model combines individual (private) and fundamental (common) values for a security, following prior computational literature on financial markets [13,50]. The *fundamental value* r_t of the security at time $t \in [0, T]$ changes throughout the trading period according to a mean-reverting stochastic process:

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

The parameter $\kappa \in [0, 1]$ specifies the degree to which the value reverts back to a fundamental mean $\bar{r}$. A process with $\kappa = 0$ corresponds to a martingale Gaussian fundamental, whereas $\kappa = 1$ specifies a process of i.i.d. Gaussian draws around the fundamental mean. A mean-reverting time series of this sort has been empirically observed in financial markets such as foreign exchange and commodity markets [51]. The perturbation u_t captures a systematic random shock upon the fundamental at time t and is normally distributed as $u_t \sim N(0, \sigma_s^2)$, where σ_s^2 represents an environment-specific shock variance. The shock variance governs fluctuations in the fundamental time series and consequently affects the predictability of future price outcomes.

Our time-varying fundamental induces *adverse selection*, a situation where outstanding orders reflect outdated information and thus can be at a disadvantage at the current time. If the fundamental shifts significantly, subsequent agents are more likely to transact with orders on the side opposite to the direction of fundamental change. That is, a positive price shock will tend to trigger transactions with stale sell orders, and a negative shock with stale buys. An agent's exposure to adverse selection in a market is jointly controlled by the fundamental shock variance σ_s^2 , the degree of mean reversion κ and the arrival rate of that agent.

The entries of a background trader follow a Poisson process with an arrival rate λ_a . Upon each entry, the trader observes an agent-and-time-specific noisy fundamental $o_t = r_t + n_t$, where the observation noise n_t is drawn from $n_t \sim N(0, \sigma_n^2)$. Just as in real financial markets, investors will never know the true value of the underlying security, such noisy observations represent each trader's assessment of the security's fundamental value at that time. Given its incomplete information about the fundamental, the agent can potentially benefit by considering market information, which is influenced by and therefore reflects the aggregate observations of other agents. When it arrives, the trader withdraws its previous order (if untransacted) and submits a new single-unit limit order, either to buy or sell as instructed with equal probability.

The *individual (private) value* of a background trader i represents its preferences over holdings of the security:

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

The vector has length $2q_{\max}$, where $q_{\max}$ is the maximum position (long or short) a trader can hold at any time. 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 . Alternatively, θ_i^{q+1} can be understood as the marginal private gain from buying an additional unit given current net position q .

Private values follow the law of diminishing marginal utility observed in many economic settings, as well as commonly assumed for assets trading in financial markets [13]. One natural example is when a trader is aiming for a target investment position. Before reaching this position, the private component is positive for each incremental share, and the more shares the trader needs to buy to reach the optimal position, the higher the marginal utility would be (as the agent is more eager to get close to the optimal position); after a trader passes her desired holding position, she gets negative private value an extra share (e.g., due to cumulative risk), and the more the position deviates from the optimal, the higher the penalty becomes. To capture this 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.

Agent i 's incremental surplus for a trade can be calculated based on its position q before the trade, the value of the fundamental at the end of the trading horizon r_T and the transaction price p :

$$\text{incremental surplus} = \begin{cases} r_T - p + \theta_i^{q+1} & \text{if buying 1 unit,} \\ p - r_T - \theta_i^q & \text{if selling 1 unit.} \end{cases}$$

An agent's total surplus is the sum of the agent's incremental surplus over all transactions. Alternatively, we can also calculate an agent's total surplus by adding its net cash from trading to the *final valuation* of holdings. Specifically, the market's final valuation of trader i with ending holdings H is

$$v_i = \begin{cases} r_T \times H + \sum_{k=1}^{k=H} \theta_i^k & \text{for long position } H > 0, \\ r_T \times H - \sum_{k=H+1}^{k=0} \theta_i^k & \text{for short position } H < 0. \end{cases}$$

We define *background-trader surplus* as the sum of all background agents' surpluses at the end of the trading period T .

3.3. Background Trading Agents

Recall that background traders represent investors with actual preferences for holding long or short positions in the underlying security. The limit-order price submitted by a background trader is jointly decided by its *valuation* and *trading strategy*, which we describe in detail below.

3.3.1. Estimating the Final Fundamental

As holdings of the security are evaluated at the end of a trading period (i.e., $r_T \times H$), a background trader estimates the final fundamental value based on a series of its 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 agent's preceding arrival time. We first update the previous posteriors, $\tilde{r}_{t'}$ and $\tilde{\sigma}_{t'}^2$, by mean reversion for the interval since preceding arrival, denoted $\delta = t - t'$:

$$\tilde{r}_{t'} \leftarrow (1 - (1 - \kappa)^\delta) \bar{r} + (1 - \kappa)^\delta \tilde{r}_{t'} \quad \text{and} \quad \tilde{\sigma}_{t'}^2 \leftarrow (1 - \kappa)^{2\delta} \tilde{\sigma}_{t'}^2 + \frac{1 - (1 - \kappa)^{2\delta}}{1 - (1 - \kappa)^2} \sigma_s^2.$$

The estimates for the current arrive at time t are then given by

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

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:

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

3.3.2. Zero Intelligence (ZI) as a Background Trading Strategy

We consider parameterized trading strategies in the zero intelligence (ZI) family [16]. Background traders who choose to adopt ZI strategies compute limit-order prices solely based on fundamental observations and private values. ZI agents generate bids reflecting a *requested surplus*, determined by a random offset uniformly drawn from $[R_{\min}, R_{\max}]$.

These bids shade from the agent's valuation by this requested surplus. Specifically, a ZI trader i arriving at time t with position q generates a limit price

$$p_i(t) \sim \begin{cases} U[\hat{r}_t + \theta_i^{q+1} - R_{\max}, \hat{r}_t + \theta_i^{q+1} - R_{\min}] & \text{if buying,} \\ U[\hat{r}_t - \theta_i^q + R_{\min}, \hat{r}_t - \theta_i^q + R_{\max}] & \text{if selling.} \end{cases} \quad (3)$$

Our version of ZI further considers the market's current best quotes, and it can choose to immediately trade to get a certain fraction of its requested surplus. This option is governed by a strategic threshold parameter $\eta \in [0, 1]$: 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 submitting a new limit order. Setting η to 1 is equivalent to the strategy without a threshold. Both shading and threshold-taking provide some non-learning ways for ZI agents to strategically adapt to different market environments and improve profitability.

3.3.3. Heuristic Belief Learning (HBL) as a Background Trading Strategy

The second background trading strategy family we consider is heuristic belief learning (HBL). Background traders who choose to adopt HBL go beyond their own observations and private values by also considering order book information. We make a set of changes to adapt the strategy to our dynamic market environment, supporting multiple-unit trading with a flexible buy or sell role.

The strategy is centered on the belief function that a background trader forms on the basis of its observed market data. The agent uses the belief function to estimate the probability that orders at various prices would be accepted in the market and then chooses a limit price that maximizes its expected surplus at current valuation estimates.

Specifically, an HBL agent constructs its belief function based on a dataset $\mathcal{D}$ that records accepted and rejected buy and sell orders during the last L trades. The strategic parameter L represents the agent's memory length, which controls the size of $\mathcal{D}$. Upon an arrival at time t , the HBL agent builds a belief function $f_t(P)$, designed to represent the probability that an order at price P will result in a transaction. Specifically, the belief function is defined for any encountered price P as the following:

$$f_t(P | \mathcal{D}) = \begin{cases} \frac{\text{TBL}_t(P | \mathcal{D}) + \text{AL}_t(P | \mathcal{D})}{\text{TBL}_t(P | \mathcal{D}) + \text{AL}_t(P | \mathcal{D}) + \text{RBG}_t(P | \mathcal{D})} & \text{if buying,} \\ \frac{\text{TAG}_t(P | \mathcal{D}) + \text{BG}_t(P | \mathcal{D})}{\text{TAG}_t(P | \mathcal{D}) + \text{BG}_t(P | \mathcal{D}) + \text{RAL}_t(P | \mathcal{D})} & \text{if selling.} \end{cases} \quad (4)$$

Here, T and R specify *transacted* and *rejected* orders, respectively; A and B represent *asks* and *bids*; and L and G describe orders with prices *less* than or equal to and *greater* than or equal to price P , respectively. For example, $\text{TBL}_t(P | \mathcal{D})$ is the number of transacted bids found in the memory with price less than or equal to P up to time t . An HBL agent updates its dataset $\mathcal{D}$ whenever the market receives new order submissions, transactions or cancellations and computes the statistics in Equation (4) upon each arrival.

Since our market model supports persistent orders and cancellations, the classification of an order as rejected is non-obvious and remains to be defined. To address this, we associate orders with a grace period τ_{gp} and an alive period τ_{al} . We define the grace period as the average time interval per arrival, that is $\tau_{\text{gp}} = 1/\lambda_a$, 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, Equation (4) is defined only at encountered prices; we further extend it over the full price domain by cubic spline interpolation. To speed the computation, we pick knot points and interpolate only between those points.

After formulating the belief function, an agent i with the arrival time t and current holdings q searches for the optimal price $P_i^*(t)$ that maximizes its expected surplus:

$$P_i^*(t) = \begin{cases} \arg \max_P (\hat{r}_t + \theta_i^{q+1} - P) f_t(P | \mathcal{D}) & \text{if buying,} \\ \arg \max_P (P - \theta_i^q - \hat{r}_t) f_t(P | \mathcal{D}) & \text{if selling.} \end{cases} \quad (5)$$

In 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 does not materially affect the overall performance. In Section 7, we explore variations of the HBL strategy to improve its learning robustness in the face of market manipulation.

3.4. The Spoofing Agent

The spoofing agent seeks profits only through manipulating prices. Unlike background traders, the spoofer has no private value for the security. 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 $T_{sp} \in [0, T]$, the spoofing agent submits a buy order at price $BID_{T_{sp}} - 1$ with volume $Q_{sp} \gg 1$. Whenever there is an update on the best bid, the spoofer cancels its original spoof order and submits a new one at price $BID_t - 1$ with the same volume. Since in our model, background traders submit only single-unit orders, they cannot transact with the spoof order, which is always shielded by the order at a higher price $BID_{T_{sp}}$. If that higher-price order gets executed, the spoofer immediately cancels and replaces its spoof orders before another background trader arrives. Here, we assume in effect that the spoofing agent can react infinitely fast, in which case its spoof orders are guaranteed never to transact. Even without making any trade, the spoofer maneuvers HBL traders' pricing beliefs (Equation (4)) via such submissions and cancellations of spoof orders, which ultimately affect their bidding behavior (Equation (5)) and move the price.

By continuously feigning buy interest in the market, this spoofing strategy specifically aims to raise market beliefs. To profit from such manipulation practice, a spoofing agent may first buy some shares of the security, manipulate the market to push prices up and later sell those previously bought shares at higher prices. Other spoofing strategies such as adding sell pressure or alternating between buy and sell pressure can be extended from the basic version.

4. Empirical Game-Theoretic Analysis

To reproduce spoofing and understand its effect, we employ a computational approach that combines agent-based modeling, simulation and equilibrium computation. The point of identifying equilibria of the agent-based model is to focus on the most relevant strategic contexts, where agents are making the best choices among their available strategies, given others' choices. To derive Nash equilibria, we employ empirical game-theoretic analysis (EGTA), a methodology that finds approximate equilibria in games defined by heuristic strategy space and simulated payoff data [7]. We conduct systematic EGTA studies over a range of parametrically defined market environments, all based on the market model described in Section 3.

We model the market as a game with players in two *roles*: N background traders, treated *symmetrically*, and a single spoofer. In most of our games, the spoofing agent, when present, implements a fixed policy so is not considered a strategic player. Symmetry of the background traders means that each has the same set of available strategies (from the ZI and/or HBL families) to choose from, and their payoffs depend on their own strategy and the number of players choosing each of the other strategies (i.e., it does not matter which other-agent plays which other-strategy). For each game, we evaluate a wide variety of *strategy profiles* (i.e., agent-strategy assignments). For each strategy profile, we conduct thousands of simulation runs to account for stochastic effects such as the market

fundamental series, agent arrival patterns and private valuations and get low-variance estimates of the expected payoffs for that profile. Given background trader symmetry, the payoff of a specific strategy in a profile can be taken as the average payoff over all agents playing that strategy in the profile. From the payoff data accumulated from these simulated samples of explored strategy profiles, we induce an *empirical game model*, and from that we derive an approximate Nash equilibrium.

We employ an iterative EGTA process. First, we find candidate equilibria in *subgames*, where a subgame here is defined as a normal-form game over strategy subsets (we note the contrast with the standard notion in extensive-form games of subgame as subtree). We then confirm or refute candidate solutions by examining deviations and incrementally extend subgames, until the termination criteria are satisfied. Below, we describe two key components of the EGTA process we follow: profile search (Section 4.1) and game reduction (Section 4.2).

4.1. Profile Search

We apply EGTA iteratively to guide the profile search over the strategy space. Exploration starts with singleton subgames, and then it incrementally considers each strategy outside the subgame strategy set. Specifically, the singleton subgames are profiles where the same strategy is adopted by all background agents. Starting from this base, we extend evaluation to neighboring profiles with single-agent deviations. Following such a procedure, we systematically explore profiles and incorporate their payoff estimates into the partial payoff matrix corresponding to the empirical game model.

After subgames are completed (all profiles explored for strategy subsets), we compute their equilibria and consider these as candidate solutions of the full game. We attempt to *refute* these candidates by evaluating deviations outside the subgame strategy set, constructing a new subgame when a beneficial deviation is found. If we examine all deviations without refuting, the candidate is *confirmed*. 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 (up to a threshold support size). In this study, we have support sizes (i.e., numbers of strategies played with positive probability) up to five for background agents.

The procedure aims to confirm or refute promising equilibrium candidates found throughout our exploration of the strategy space. By following a fixed procedure and reporting all equilibria found meeting pre-specified criteria, we maintain consistency across the games analyzed. Since it is often not computationally feasible to search the entire profile space, additional distinct equilibria (e.g., with large support size) are possible. We are aware of no systematic biases in the search procedure with respect to properties of interest, and therefore consider it unlikely that we are missing qualitatively important phenomena across the range of environments analyzed. Finally, we note that equilibria identified in empirical games must generally be viewed as provisional, as they are subject to refutation by strategies outside the restricted set considered in the analysis.

4.2. Game Reduction

As the *game size* (i.e., number of possible strategy profiles) grows exponentially in the number of players and strategies, it is computationally prohibitive to directly analyze games with more than a moderate number of players. We therefore apply aggregation methods to approximate a many-player game by a game with fewer players. The specific technique we employ, called *deviation-preserving reduction* (DPR) [52], 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 $(s_2, \dots, s_k)$, is given by the payoff of playing s_1 in the full N -player game when the other $N - 1$ agents are evenly divided among the $k - 1$ strategies $s_2, \dots, s_k$. To facilitate DPR, we choose values for N and k to ensure that the required aggregations come out as integers. For example, in one of

the market environment, we reduce games with 28 background traders to games with four background traders. With one background player deviating to a new strategy, we can reduce the remaining 27 players to three. For games that vary smoothly with the number of other players choosing any particular strategy, we can expect DPR to produce reasonable approximations of the original many-player games with exponential reduction in simulation.

5. Spoofing the Limit Order Book

In this section, we reproduce spoofing in the agent-based market model and study its effect on background trading behavior and market outcomes. We start in Section 5.1 by exploring a range of market environments that can affect the effectiveness of both learning and spoofing. Section 5.2 addresses agents' choices among ZI and HBL strategies in markets without spoofing. This is an important step, as spoofing can be effective only if some fraction of background traders choose to learn from the order book information. Section 5.3 investigates games with spoofing. We first illustrate that a market populated with HBL traders is susceptible to spoofing: a simple spoofing strategy can cause a rise in market prices and a redistribution of surplus between ZI and HBL traders. We finally re-equilibrate the game with spoofing to investigate the impact of spoofing on HBL adoption and market surplus.

5.1. Market Environments

Based on the defined market model, we conduct preliminary explorations over a range of market settings and include the most salient and meaningful ones for our study. We consider nine market environments that differ in fundamental shock, $\sigma_s^2 \in \{10^5, 5 \times 10^5, 10^6\}$, and in observation noise, $\sigma_n^2 \in \{10^3, 10^6, 10^9\}$. Recall that shock variance controls fluctuations in the fundamental time series and observation variance governs the quality of information agents get about the true fundamental. The nine environments cover representative market conditions that can affect an agent's ability and need to learn from market information. For example, when the market shock variance is large, prices fluctuate more and market history may become less predictive; when observation noise is high, agents can glean only limited information from their own observations and may gain more from the market's aggregated order book information. We label the low, medium and high shock variances as $\{LS, MS, HS\}$ and noisy observation variances as $\{LN, MN, HN\}$, respectively. For instance, the label *LSLN* refers to a market with low shock, $\sigma_s^2 = 10^5$, and low observation noise, $\sigma_n^2 = 10^3$.

The global fundamental time series is generated according to Equation (1) with fundamental mean $\bar{r} = 10^5$, mean reversion $\kappa = 0.05$ and specified shock variance σ_s^2 . The minimum tick size is fixed at one. Each trading period lasts $T = 10,000$ time steps. For each environment, we consider markets populated with $N \in \{28, 65\}$ background traders and in selected treatments, a spoofer. 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 $o_t = r_t + n_t$, where $n_t \sim N(0, \sigma_n^2)$. The maximum number of units background traders can hold at any time is $q_{\max} = 10$. Private values are drawn from a Gaussian distribution with zero mean and a variance of $\sigma_{PV}^2 = 5 \times 10^6$. The spoofing agent starts to manipulate at time $T_{sp} = 1000$ by submitting a large buy order at price $BID_{T_{sp}} - 1$ with volume $Q_{sp} = 200$, and later maintains spoofing orders at price $BID_t - 1$ throughout the trading period.

To provide a benchmark for market surplus, we calculate the social optimum—the expected total possible gains from trade, which depends solely on the trader population size and valuation distribution. From 20,000 samples of the joint valuations, we estimate mean social optima of 18,389 and 43,526 for markets with 28 and 65 background traders, respectively. We further calculate the average order book depth (on either buy or sell side) in markets without spoofing. Throughout the trading horizon, the $N = 28$ market has a

relatively thin order book with an average depth of 12 per side, whereas the $N = 65$ market has a thicker one with an average depth of 30.

The background trading strategy set (see Table 1) includes seven versions of ZI and four versions of HBL. Agents are allowed to choose from this restricted set of strategies. We have also explored ZI strategies with larger shading ranges and HBL strategies with longer memory lengths, but they fail to appear in equilibrium in games where they were explored.

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	–	–	–	–	–	–	–	2	3	5	8
$R_{\min}$	0	0	0	0	0	250	250	–	–	–	–
$R_{\max}$	250	500	1000	1000	2000	500	500	–	–	–	–
η	1	1	0.8	1	0.8	0.8	1	–	–	–	–

5.2. Games without Spoofing

Since spoofing targets the order book and can be effective only to the extent traders exploit order book information, we investigate whether background agents adopt the HBL strategy in markets without spoofing. Applying EGTA to the eleven background strategies in Table 1, we found at least one equilibrium for each market environment. Detailed results on equilibrium mixture and outcomes of games without spoofing can be found in Appendix A.2.

Figure 3 (blue circles) depicts the proportion of background traders who choose trading strategies in the HBL family. In most non-spoofing environments, HBL is adopted with positive probability, suggesting that 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 fundamental 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. This is further confirmed in the two exceptions where all agents choose ZI: *HSLN* and *HSMN* with $N = 28$, the environments with fewer traders, high fundamental shocks and at most medium observation noise.




Figure 3. 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.

We further quantify how learning from market information may benefit overall market performance. We conduct EGTA in games where background traders are restricted to strategies in the ZI family (ZI₁–ZI₇ in Table 1). This is tantamount to disallowing learning from order book information. Detailed equilibrium results on games restricted to ZI strategies can be found in Appendix A.4. We compare equilibrium outcomes for each

environment, with and without HBL available to background traders, on two measures: surplus (Figure 4) and price discovery (Figure 5). Recall that we define background-trader surplus as the sum of all background agents' surpluses at time T , the end of trading. Price discovery is defined as the root-mean-squared deviation (RMSD) of the transaction price from the estimate of the true fundamental in Equation (2) over the trading period. It reflects how well transactions reveal the true value of the security. Lower RMSD means better price discovery. We calculate the two measures by averaging the outcomes of 20,000 simulations of games with strategy profiles sampled according to each equilibrium mixture.

Overall, background traders achieve higher surplus (Figure 4) and better price discovery (Figure 5) when the market provides order book information and enables the HBL strategy option. When HBL exists in the equilibrium, we find transactions reveal fundamental estimates well, especially in markets with lower shock and observation variances (i.e., *LSLN*, *LSMN*, *MSLN* and *MSMN*). We also notice small exceptions in scenarios with high observation variance and more background traders (environments *LSHN* and *HSHN* with 65 players) where ZI-only equilibria exhibit higher surplus than equilibria combining HBL and ZI.




Figure 4. Comparisons of background-trader surplus for equilibria 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. The market generally achieves higher surplus when HBL is available.




Figure 5. Comparisons of price discovery 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 where the equilibrium mixture is ZI only, despite the availability of HBL. The market generally achieves better price discovery when HBL is available.

5.3. Games with Spoofing

5.3.1. Comparing across Fixed Strategy Profiles

We examine the effectiveness of our designed spoofing strategy (Section 3.4) by playing a spoofer against each HBL-and-ZI equilibrium found in Section 5.2. We perform controlled comparisons on these games with and without spoofing. As ZI agents are oblivious to spoofing, we ignore the ZI-only equilibria in this analysis. In the paired instances, background agents play identical strategies and are guaranteed to arrive at the same time, receive identical private values and observe the same fundamental values. Therefore, any change in behavior is an effect of spoof orders on HBL traders. For every setting, we simulate 20,000 paired instances, evaluate transaction price differences (Figure 6) and compare surplus attained by HBL and ZI traders. Transaction price difference at a specific time is defined as the most recent transaction price in the run with spoofing minus that of the paired instance without spoofing. Similarly, surplus difference of HBL or ZI is the aggregated surplus obtained in an environment with spoofing minus that of the corresponding environment without spoofing.

Figure 6 shows positive changes in transaction prices across all environments, subsequent to the arrival of a spoofing agent at $T_{sp} = 1000$. This suggests that 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. Although ZI agents do not change their bidding behavior directly, they may transact at higher prices due to the increased bids of HBL traders.




Figure 6. Transaction price differences throughout the trading horizon with and without a spoofer against each HBL-and-ZI equilibrium found in non-spoofing games (Section 5.2). Multiple curves for the same market environment represent different equilibria. The designed spoofing tactic clearly raises market prices when HBL are present. The effect attenuates over time, generally more quickly in the thicker market environments.

Several other interesting findings are revealed by the transaction-price difference series. First, the average price rise caused by spoofing the market with 28 background traders is higher than for $N = 65$. This indicates that a market with fewer background traders can be more susceptible to spoofing, due to the limited pricing information a thin market could aggregate. Second, for markets populated with more HBLs than ZIs in the equilibrium mixture, the transaction price differences tend to increase throughout the trading period. This amplification 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 the HBLs' illusions. Finally, we observe that the spoofing effect tends to attenuate over time: differences in transaction prices first increase, and then stabilize or decrease. This is due to the mean-reverting property of the fundamental series and the way background traders estimate the final fundamental. As time approaches the end of the trading period, background agents rely more on accumulated fundamental observations and get better estimates of the final fundamental value. Therefore, spoofing tends to wear off in the face of accumulated observations and mean reversion.

We further compare background-trader payoffs attained in environments with and without spoofing. We find a redistribution of surplus between HBL and ZI agents: HBL aggregated surplus decreases, while that for ZI increases compared to the non-spoofing baselines. Specifically, across 28-trader market environments, HBL traders suffer an average surplus decrease of 184 across all equilibrium profiles, whereas the ZI traders have an average surplus gain of 19. For the 65-trader markets, the average surplus decrease for HBL traders is 238, while the average increase for ZI is 40. This suggests that the ZI agents benefit from the HBL agents' spoofed beliefs. Since the decreases in HBL surplus are consistently larger than the increases for ZI, the overall market surplus decreases. We leave further discussion of spoofing's impact on market surplus to Section 5.3.2, where background traders can choose other strategies to adjust to the presence of spoofing.

To examine the potential to profit from a successful price manipulation, we extend the spoofing agent with an *exploitation strategy*: buying, (optionally) spoofing to raise the price and then selling. The exploiting spoofer starts by buying when there is a limit sell order with price less than the fundamental mean in the market. It then optionally runs the spoofing trick, or alternatively waits, for 1000 time steps. Finally, the agent sells the previously bought unit (if any) when it finds a limit buy order with price more than fundamental mean. Note that, even without spoofing, this single-unit exploitation strategy is profitable in expectation due to the mean reversion captured by the fundamental process and the reliable arrivals of background traders with private preferences.

Figure 7 compares the exploitation profits with and without spoofing. In controlled experiments, we find that exploitation profits are consistently increased when the spoof action is also deployed. Across 28-trader market environments, the exploiter makes an average profit of 206.1 and 201.8 with and without spoofing, respectively, 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. 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.




Figure 7. Exploitation profits with and without spoofing against each HBL-and-ZI equilibrium found in Section 5.2. Repetitions of the same market environment represent outcomes of multiple equilibria. Exploitation profits are consistently higher when the spoof action is also deployed.

5.3.2. Re-Equilibrating Games with Spoofing

To understand how spoofing changes background-trading behavior, we conduct EGTA again to identify Nash equilibria, allowing background traders to choose any strategy in Table 1, in games with spoofing. Detailed results on equilibrium mixture and outcomes of games with spoofing can be found in Appendix A.3. As indicated in Figure 3 (orange triangles), after re-equilibrating games with spoofing, HBL is generally adopted by a smaller fraction of traders, but it still persists in 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.

We characterize the effect of spoofing on market surplus. Figure 8 compares the total surplus achieved by background traders in equilibrium with and without spoofing. Given the presence of HBL traders, spoofing generally decreases total surplus (as, in Figure 8, most filled orange triangles are below the filled blue circles). However, spoofing has ambiguous effect in the thicker market with large observation variance (environments *LSHN* and *HSN* 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. Finally, we find the welfare effects of HBL strategies persist regardless of spoofing's presence: markets populated with HBL agents in equilibrium achieve higher total surplus than those markets without HBL (as, in Figure 8, the hollow markers are below the filled markers).




Figure 8. 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 equilibria combining HBL and ZI and equilibria with pure ZI are indicated by markers with and without fills, respectively.

5.4. Discussion

Our agent-based model of spoofing aims to capture the essential logic of manipulation through influencing belief about market demand. In our model, the order book reflects aggregate information about the market fundamental, and learning traders can use this to advantage in their bidding strategies. The presence of such learning traders benefits price discovery and social welfare, but it also renders the market vulnerable to manipulation. As we demonstrate, simple spoofing strategies can effectively mislead learning traders, thereby distorting prices and reducing surplus compared to the non-spoofing baseline. Moreover, the persistence of learning traders in equilibrium with manipulation suggests that the elimination of spoofing requires active measures.

We acknowledge several factors that can 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 to evaluate manipulation practices and any preventive or deterrent proposals to mitigate manipulation under strategic settings. In the rest of the paper, we build on this model and conduct comprehensive analysis to investigate the following questions:

- Are there more robust ways for exchanges to disclose order book information (Section 6)?
- Are there strategies by which individual traders can adopt to exploit market information but in less vulnerable ways (Section 7)?

6. A Cloaking Mechanism to Mitigate Spoofing

Despite regulatory enforcement and detection efforts, an individual spoofing episode is hard to catch in high-volume, high-velocity data streams. Legal definitions cannot be

easily translated to computer programs to direct detection, and the lack of datasets with labeled manipulation cases makes training a reliable detector infeasible with supervised machine learning techniques. Based on its definition, to determine that a pattern of activity constitutes spoofing requires establishing the manipulation intent behind submission and cancellation of placed orders. However, this is not easy, as order cancellation is in itself common and legitimate: according to one study, 95% of NASDAQ limit orders are canceled, with a median order lifetime less than one second [53]. Another challenge arises from the adversarial nature of a manipulator who may strategically adapt to evade detection and regulation [54]. Given difficulties in robustly detecting manipulation, we study systematic approaches to deter spoofing, by rendering manipulative practices difficult or uneconomical.

Along these lines, Prewit [55] and Biais and Woolley [56] advocated the imposition of cancellation fees to disincentivize manipulative strategies that rely on frequent cancellations of orders. Others argue that cancellation fees could discourage the beneficial activity of liquidity providers, and, in the event of a market crash, such a policy may lengthen the recovery process [57].

We propose here a *cloaking mechanism* to deter spoofing via the selective disclosure of order book information. The mechanism extends the traditional CDA market with a cloaking parameter K , which specifies the number of price levels to hide symmetrically from inside of the limit order book. The idea is to make it more difficult for the spoofer who relies on the instant order book information to post misleading bids, while not unduly degrading the general usefulness of market information. We focus on deterministic cloaking (i.e., a constant K throughout the trading period), as a stochastic mechanism may raise issues regarding verification of faithful market operations.

We extend our agent-based model of spoofing to support order book cloaking and conduct simulations to evaluate and find the optimal cloaking parameter under strategic settings, where both the learning traders and the spoofer adapt to the new mechanism. Section 6.1 formally defines the cloaking mechanism and describes how we modify the background trading and spoofing strategies accordingly. In Section 6.2, we present an EGTA study conducted to understand agents' strategic responses to the proposed mechanism. Section 6.3 reports results from performing *empirical mechanism design* [58] to set cloaking parameters that maximize efficiency. Finally, in Section 6.4, we explore and evaluate sophisticated spoofing strategies that use probing to reveal cloaked information.

6.1. A Cloaking Market Mechanism

The cloaking mechanism maintains a full limit order book just as the regular CDA market, but discloses only a selective part of the book to traders. Let BID_t^k denote the k th-highest buy price in the book at time t and ASK_t^k the k th-lowest sell price. In a standard order book, at any given time t , the buy side of the book starts with the best bid, BID_t^1 , and extends to lower values; the sell side starts with the best ask, ASK_t^1 , and extends to higher ones. The cloaking mechanism works by symmetrically hiding a deterministic number of price levels K from inside of the order book. Thus, the *disclosed* order book in a cloaking mechanism with parameter K starts with BID_t^{K+1} and ASK_t^{K+1} and extends to lower and higher values, respectively. Upon order submissions, cancellations and transactions, the market updates the full order book and then cloaks the K inside levels. Therefore, an order hidden in the past can be revealed later due to the arrival of new orders at more competitive prices, or it can be hidden throughout its lifetime due to a cancellation. The market discloses all the transaction information at zero delay.

Example 1 (A Cloaking Mechanism with Parameter K). When $K = 0$, the market acts as a standard CDA, disclosing the full limit order book with zero delay. When $K = 1$, the mechanism conceals orders at the best quotes, that is BID_t^1 and ASK_t^1 . When $K = \infty$, the market does not reveal any part of the book and thus disallows learning from order book information.

Cloaking operates to deter spoofing in two ways. First, it mitigates the effect of spoof orders, pushing them further from the inside of the book. Second, it increases the spoofer's transaction risks, as it cannot as easily monitor the quantity of orders ahead of the spoof. On the other hand, the information hiding also affects the non-manipulative traders, for instance in our model it may degrade the HBL traders' learning capability. To quantify this tradeoff, we start by exploring a range of cloaking parameters, $K \in \{0, 1, 2, 4\}$, which control the amount of information being concealed at any given time. We compare trading behavior and outcomes in markets with cloaking to that of a standard CDA. Among the nine market environments defined in Section 5.1, we consider three representatives that are increasingly challenging for the learning traders: *LSHN* with $\{\sigma_s^2 = 10^5, \sigma_n^2 = 10^9\}$, *MSMN* with $\{\sigma_s^2 = 5 \times 10^5, \sigma_n^2 = 10^6\}$ and *HSLN* with $\{\sigma_s^2 = 10^6, \sigma_n^2 = 10^3\}$. Together with the four cloaking parameters, this gives us a total of 12 market settings, or 24 games with and without spoofing.

The market is populated with 64 background traders and one exploitation agent. Therefore, when adopting DPR to approximate this many-player game, we use simulation data from the (64, 1)-agent environments to estimate reduced (4, 1)-player games, where four players are used to aggregate and represent the background traders. In each game, we consider background trading strategies and spoofing practice similar to those of Section 3, but they are slightly modified to adapt to order book cloaking. Below, we describe changes made to each strategy.

6.1.1. Zero Intelligence

Recall that our ZI strategy uses a threshold parameter $\eta \in [0, 1]$ to immediately transact with an existing order to grasp a portion of desired surplus. That is, if the agent could achieve a fraction η of its requested surplus at the market best quotes, it would simply take that quote rather than posting a limit order for a future transaction. Under a cloaking mechanism, however, ZI may take into account only the current *visible* best quotes that are less competitive compared to the hidden quotes. To adjust to cloaking, we explore a range of more aggressive (smaller) η values to ensure that ZI traders may still transact with incumbent orders to lock a certain fraction of surplus. Besides the seven ZI strategies in Table 1, we further include three ZI strategies with $\eta = 0.4$ (Table 2), which are competitive enough to appear in at least one equilibrium of our explored environments.

Table 2. Additional background trading strategies included in EGTA for cloaking mechanisms.

Strategy	ZI ₈	ZI ₉	ZI ₁₀
$R_{\min}$	0	0	250
$R_{\max}$	1000	2000	500
η	0.4	0.4	0.4

6.1.2. Heuristic Belief Learning

We modify HBL to consider only the *revealed* order book information under the corresponding cloaking markets. Orders at competitive price levels will be missed in the belief function (Equation (4)) if they are hidden throughout order lifetime, or they may be considered with delay if later exposed at visible levels. This reduction in bid information would naturally be expected to degrade HBL's learning effectiveness and thus its trading performance.

6.1.3. Spoofing Strategy

We extend the original spoofing strategy (Section 3.4) to cloaking markets. The strategy includes three stages. At the beginning of a trading period $[0, T_{\text{spoof}}]$, the agent buys by accepting any sell order at price lower than the fundamental mean $\bar{r}$. In a cloaking market, this can be achieved by placing a one-unit limit buy order at price $\bar{r}$ and immediately withdrawing it if does not transact with an existing order.

During the second stage $[T_{\text{spool}}, T_{\text{sell}}]$, the agent submits spoof buy orders at a tick behind the first *visible* bid $\text{BID}_{T_{\text{spool}}}^{K+1} - 1$ with volume $Q_{\text{sp}} \gg 1$. Whenever there is an update on the first visible bid, the spoofer replaces its original spoof with new orders at price $\text{BID}_t^{K+1} - 1$. This spoofing strategy aims to boost price, in the hope that the units purchased in Stage 1 can be later sold at higher prices. In controlled experiments, when the agent is not manipulating, it waits until the selling stage.

During the last stage $[T_{\text{sell}}, T]$, the agent starts to sell the units it previously bought by accepting any buy orders at a price higher than $\bar{r}$. Inverse to the first stage, this operates by placing one-unit limit sell orders at price $\bar{r}$, followed by immediate cancellation if not filled. The agent who also manipulates continues to spoof until all the bought units are sold or the trading period ends. The pure exploitation strategy can be considered as a baseline for the spoofing strategy, allowing us to quantify how much more the agent may profit from spoofing the market.

We refer to the agent who employs the above strategy, whether places spoof orders or not, as an *exploitation agent* or *exploiter*. An exploiter who also spoofs is referred to as a *spoofing agent* or *spoofer*. Note that the spoofing strategy considered here does not face any execution risk on its spoof orders under the assumption it can immediately respond to quote changes. A more sophisticated strategy could *probe* the market to reveal the cloaked bids, and then spoof at a visible price higher than $\text{BID}_t^{K+1} - 1$. We leave discussion of such probing strategies to Section 6.4.

6.2. Tradeoff Faced by Cloaking Mechanisms

We start by separately investigating the impact of cloaking on background traders and on the spoofer. Our first set of games cover the range of cloaking environments without spoofing (i.e., markets populated with background traders and the non-manipulative exploiter). Detailed results on equilibrium mixture and outcomes of cloaking mechanisms without spoofing can be found in Appendix B.1.

Figure 9 displays the HBL adoption rate (i.e., total probability over HBL strategies) at equilibrium across cloaking mechanisms, $K \in \{0, 1, 2, 4\}$. We find that the competitiveness of HBL generally persists when the mechanism hides one or two price levels, but at higher cloaking levels the HBL fraction can drastically decrease. The information loss caused by cloaking weakens HBL's ability to make predictions. The effect is strongest in environments with high fundamental shocks (e.g., *HSLN*), as previous hidden orders can become uninformative or even misleading by the time they are revealed. Such information loss is further confirmed in Figure 10, which compares the price discovery achieved at market equilibrium. We find that in markets where the fundamental shock is relatively high and a higher level of cloaking is adopted, transactions may not reveal fundamental estimates very well. Given the decreasing HBL prevalence and learning effectiveness, background-trader surplus achieved at equilibrium also decreases (as shown in Figure 11b (blue diamonds)).




Figure 9. HBL adoption rate in equilibrium across different cloaking markets without spoofing.




Figure 10. Comparisons of price discovery for equilibrium in each environment with different cloaking.




Figure 11. Equilibrium outcomes in games with and without cloaking. Each marker represents one equilibrium of the environment.

Next, we examine whether cloaking can effectively mitigate manipulation. We perform controlled experiments by letting the exploitation agent also execute the spoofing strategy against each found equilibrium and compare the impact of spoofing under the cloaking mechanism to the standard fully revealed order book ($K = 0$). For every equilibrium, we simulate at least 10,000 paired instances and evaluate their differences on transaction price and agents' payoffs.

From these controlled experiments, we find that cloaking can considerably diminish price distortion caused by spoofing across environments. Recall that we measure price distortion as the transaction price series in a market with spoofing minus that of its paired market without spoofing. Figure 12a demonstrates the case in a specific environment *MSMN*: without cloaking ($K = 0$), transaction prices significantly rise subsequent to the execution of spoofing at $T_{sp} = 1000$, as HBL traders are tricked by the spoof buy orders; in cloaked markets, this price rise is effectively mitigated. Figure 12b further illustrates the surplus change in background traders and the exploiter when it also spoofs. We find the exploiter can robustly profit from manipulating the learning agents in the no-cloaking case. In contrast, partially hiding the order book can significantly reduce spoofing profits and prevent background traders from losing much. These findings indicate the cloaking mechanism may deter or even eliminate the exploiter's incentive to spoof.




Figure 12. The impact of cloaking on spoofing effectiveness. Cloaking mitigates price rise and the decrease in background surplus caused by spoofing

6.3. Finding the Optimal Cloaking

Given the tradeoff between preserving order book informativeness and mitigating manipulation, the question becomes: Under what circumstances do the deterrence benefits of cloaking exceed its efficiency costs? To answer this, we *re-equilibrate* games allowing the exploiter to strategically choose whether to spoof, with background traders able to execute any strategy in Table 1 or Table 2. This allows background traders and the exploitation agent to strategically respond to each other under a certain level of information cloaking. Detailed results on equilibrium mixture and outcomes can be found in Appendix B.2.

Our findings are presented in Figure 11. We compare market outcomes with and without cloaking on two metrics: the probability of spoofing and total background-trader surplus in equilibrium. (Due to the welfare benefits of HBL, equilibria with pure ZIs usually achieve much lower surplus than those with HBLs. For presentation simplicity, we omit all-ZI equilibria from Figure 11b. Environments with such cases are marked with asterisks.) As shown in Figure 11a, the cloaking mechanism effectively decreases the probability of spoofing under most environment settings—completely eliminating spoofing in some cases. Moreover, we find moderate cloaking can preserve the prevalence of HBL at equilibrium, which otherwise would be decreased by spoofing as we saw in Section 5.

This weakened spoofing effect is further confirmed by Figure 11b, which compares the total background-trader surplus achieved in equilibrium under mechanisms with and without cloaking. Without cloaking (i.e., K0 columns), background surplus achieved in equilibrium where the exploiter strategically chooses to spoof (orange triangles) is much lower than the surplus attained when the exploiter is prohibited from spoofing (blue diamonds). We find the decrease in surplus due to spoofing can be considerably mitigated by order book cloaking. As shown in Figure 11b, the vertical distances between the blue diamonds and orange triangles get smaller with $K > 0$. Moreover, we find the benefit of this improved robustness to spoofing can outweigh its associated efficiency costs in markets with moderate fundamental shocks (e.g., LSHN and MSMN). In those environments, background traders in mechanisms that cloak one or two price levels achieve higher surplus than those without cloaking. However, in a market with high shocks (e.g., HSLN), hiding or delaying even a little market information degrades learning to such a degree as to render cloaking counter-productive.

6.4. Probing the Cloaking Mechanism to Spoof

To this point, we have only considered spoofers who are unwilling to risk execution of their spoof orders. A more sophisticated manipulator could *probe* the market, submitting a series of orders at slightly higher prices, in an attempt to reveal the cloaked bids and spoof at a visible price higher than $BID_t^{K+1} - 1$. In this section, we study the feasibility of such probing to the spoofing agent.

We design and evaluate parameterized versions of the spoofing strategy combined with probing. The strategy is governed by two parameters: the step size δ , which controls

probing aggressiveness, and the maximum attempts allowed per time step l , which limits the probing effort.

The spoofer probes by submitting a unit buy order at $BID_t^{K+1} + \delta$, a price inside the visible quotes, in the hopes of exposing BID_t^K . If the probe succeeds, it immediately cancels the probe order, and places a new spoof order at $BID_t^K - 1$, right behind the lowest hidden bid level. If probing fails because the price is too conservative, the spoofer re-probes by raising the price at a decreasing rate (as a function of δ and the attempt number), until a higher price is revealed or the number of probing attempts reaches l . If probing causes a transaction, the spoofer halves the price increment and re-probes. Algorithm 1 describes the detailed probing procedure.

Algorithm 1 Spoofing with probing in a cloaking market with K hidden price levels ($K > 0$)

Input: The probing step size δ and the attempt limit l .

The spoofer's time to place spoof orders T_{spoof} , and its current holding H .

```

1: while  $t \geq T_{\text{spoof}}$  and  $H > 0$  do
2:   if no active spoof orders then
3:      $c \leftarrow 1, \Delta \leftarrow \delta$  ▷ track the number of probing attempts and the price increment
4:     submit a single-unit probe buy order at price  $BID_t^{K+1} + \Delta$ 
5:     while the visible  $BID_t^{K+1}$  remains unchanged and  $c < l$  do
6:        $c \leftarrow c + 1$ 
7:       if the probe buy order gets transacted then
8:          $\Delta \leftarrow \Delta / 2$ 
9:         submit a single-unit probe buy order at price  $BID_t^{K+1} + \Delta$ 
10:      else
11:         $\Delta \leftarrow \Delta + \max\{0.9^{c-1}\delta, 1\}$ 
12:        substitute the probe order with a new one at price  $BID_t^{K+1} + \Delta$ 
13:      end if
14:    end while
15:    submit spoof orders at price  $BID_t^{K+1} - 1$ 
16:    cancel the probe order
17:  else
18:    if spoof orders become hidden then
19:      substitute spoof orders with new ones at price  $BID_t^{K+1} - 1$ 
20:    else if spoof orders are no longer one tick behind  $BID_t^{K+1}$  then
21:      withdraw spoof orders
22:    end if
23:  end if
24: end while

```

Table 3 reports, for cloaking-beneficial environments, the minimum l required for step sizes $\delta \in \{1, 2, 4, 8\}$ to achieve higher payoffs than the equilibrium performance we found for the exploiter in Section 6.3. Multiple rows for the same cloaking parameter correspond to the multiple equilibria found in that market setting. Dashes in the table indicate that an exploiter cannot beat the equilibrium performance with the corresponding δ . We find that, to achieve higher payoffs, the spoofer has to probe with multiple attempts per time step, and conservative probing strategy with smaller δ usually requires more effort. In practice, such frequent cancellation and placement of orders may not be feasible and can largely increase the risk of the associated probing and spoofing intent being identified.

Figure 13 further quantifies the change in exploitation payoff and transaction risk (measured as the number of transactions caused by probing), as we vary the probing step δ and the attempt limit l . As shown Figure 13a, relaxing the maximum number of probing attempts steadily increases the transaction risk, but it does not necessarily improve payoff. Moreover, the spikiness we observe in the exploiter's payoff suggests that the performance is highly sensitive and therefore it would be difficult to find a (δ, l) that robustly maximizes profit. Figure 13b further demonstrates that an exploiter can probe aggressively with larger

step sizes to reduce effort, but it is usually at the cost of a higher transaction risk and consequently a lower payoff. In highly dynamic markets with frequently updated quotes, finding an appropriate δ to successfully probe a cloaking mechanism within a reasonable number of attempts would be quite challenging.

Table 3. Lowest number of probing attempts required to beat equilibrium performance.

Env		(δ, l)			
LSHN	K1	(1, 16)	(2, 9)	–	–
	K2	(1, 8)	(2, 5)	(4, 3)	(8, 3)
	K4	(1, 19)	(2, 3)	–	–
	K4	(1, 10)	(2, 5)	(4, 3)	–
MSMN	K1	(1, 7)	(2, 5)	(4, 4)	(8, 3)
	K1	(1, 7)	(2, 4)	(4, 2)	(8, 1)
	K1	(1, 5)	(2, 3)	(4, 2)	–
	K1	(1, 9)	(2, 4)	(4, 2)	–
	K2	(1, 11)	(2, 3)	(4, 4)	(8, 3)
	K4	(1, 5)	(2, 3)	(4, 3)	(8, 3)




Figure 13. Exploitation payoff and transaction risk as we vary price increment δ and probing limit l .

We have explored other more aggressive probing strategies, where the spoofer probes to expose multiple hidden levels and spoofs at even higher prices. To accomplish that, the spoofer is forced to keep at least one order in the cloaked levels to guarantee that its spoof orders are *visible*. However, according to our experiments, such aggressive probing strategies fail to beat the equilibrium performance, as orders kept in hidden levels are often accepted by background traders due to adverse selection. Those transactions tend to accumulate the spoofer's position, and consequently they impose losses at the end of the trading period.

6.5. Discussion

Our cloaking mechanism offers a systematic approach to disincentivizing spoofing. We conduct EGTA to understand agents' strategic responses to the proposed mechanism and evaluate the effectiveness and robustness of cloaking. Experimental results demonstrate that cloaking the order book can significantly diminish the efficacy of spoofing, but at the loss of useful information for the learning traders. With the goal of balancing this tradeoff to maximize background-trader surplus, we perform empirical mechanism design to choose the optimal cloaking across parametrically distinct environments. We find that, in markets with moderate shocks, the benefit of cloaking in mitigating spoofing can outweigh its efficiency cost, whereas, in markets with large fundamental fluctuations, hiding even a little order book information can largely degrade learning efficiency and render the cloaking mechanism counter-productive. By further exploring sophisticated spoofing strategies that probe to reveal cloaked information, we observe that associated effort and risk generally exceeds the gains and that finding reliably profitable probing regiments

is quite difficult. We conclude that the proposed cloaking mechanism cannot be easily circumvented by probing.

7. Learning-Based Trading Strategies under the Presence of Market Manipulation

We next consider how individual traders may construct strategies that are more robust to manipulation. In realistic market scenarios, traders are aware of potential manipulation but unable to reliably detect spoofing behavior in real time. In the absence of manipulation, traders submit orders that reflect their private observations and preferences, and thus learning from others' actions enables more informed decisions. Indeed, as shown above, learning as implemented by HBL agents is effective in a realistic market model and provides benefits to the learning agent as well as to market efficiency. HBL is vulnerable to spoofing, however, and agents adopting such learning are harmed by spoofing compared to non-learning strategies that are oblivious to spoofers and thus non-manipulable. The question we investigate in this section is whether learning-based strategies can be designed to be similarly robust to spoofing. We seek strategies by which individual traders can learn from market information, but in less vulnerable ways.

We treat the original HBL described in Section 3.3.3 as a baseline strategy and propose two variations that aim to reasonably trade off learning effectiveness in non-manipulated markets for robustness against manipulation. The first variation works by selectively ignoring orders at certain price levels, particularly where spoof orders are likely to be placed. The second variation considers the full order book, but has the flexibility to adjust the offer price by a stochastic offset. The adjustment serves to correct biases in learned price beliefs either caused by manipulation or the intrinsic limitation built in the belief function. We formally define the two variations in Section 7.1, and then evaluate the proposed strategies in terms of the effectiveness in non-manipulated markets and robustness against manipulation in Section 7.2.1.

We adopt the standard CDA market mechanism as described in Section 3.1. The market is populated with 64 background traders and one profitable exploiter. Background traders can choose from a select set of strategies that covers ZI, original HBL and the two proposed variations of HBL. The exploiter follows the three-stage exploitation strategy specified in Section 6.1 and executes spoofing in selected treatments. As in our study of cloaking mechanisms, we consider three representative market settings for our experiments, namely *LSHN*, *MSMN* and *HSLN*.

7.1. Two Variations of HBL

7.1.1. HBL with Selective Price Level Blocking

Our first HBL variation is inspired by the success of our cloaking mechanism. It takes advantage of the common placement of spoof orders closely behind the market best quotes. Instead of including all observed trading activities in its memory to construct the belief function just as the standard HBL, the idea is to neglect limit orders at a specified price level when assembling the dataset $\mathcal{D}$ to learn from. We extend standard HBL with a blocking parameter χ , which specifies the index of a single price level to ignore symmetrically from inside of the limit order book. For example, when $\chi = 1$, the trading agent constructs a dataset, $\mathcal{D} \setminus O_{\chi=1}$, by considering only orders strictly outside the best bid and ask. The goal of this additional strategic parameter is to exclude price levels where spoof orders are likely to appear. However, ignoring orders may come at the cost of less effective learning, especially when information that conveys true insight is blocked from the belief function.

7.1.2. HBL with Price Offsets

Our second HBL variation considers all orders in its memory, but translates the target price $P_i^*(t)$ derived by surplus maximization in Equation (5) with an offset uniformly drawn from $[R_{\min}, R_{\max}]$. Specifically, a background trader i who arrives the market at

time t with the optimized price $P_i^*(t)$ submits a limit order for a single unit of the security at price

$$p_i(t) \sim \begin{cases} U[P_i^*(t) - R_{\max}, P_i^*(t) - R_{\min}] & \text{if buying,} \\ U[P_i^*(t) + R_{\min}, P_i^*(t) + R_{\max}] & \text{if selling.} \end{cases} \quad (6)$$

A positive offset can be viewed as a hedge against misleading information, effectively shading the bid to compensate for manipulation risk. A negative offset increases the probability of near-term transaction, which may have benefits in reducing exposure to future spoofing. Offsets (positive or negative) may also serve a useful correction function even when manipulation is absent. In particular, negative offsets may compensate for the myopic nature of HBL optimization Equation (5), which considers only the current bid, ignoring subsequent market arrivals and opportunities to trade additional units. Our design here is in line with prior literature [24,25] that refines the original HBL to become more competitive.

7.2. Empirical Evaluation

7.2.1. Standard HBL

We start with our baseline market environments where background traders are restricted to choose from the standard HBL strategies and five parametrically different ZI strategies in Table 4. Figure 18 (dark grey columns) verifies what is observed in Section 5 within this restrictive set of background-trading strategies: (1) the learning-based trading strategy is more widely preferred in environments where fundamental shock is low and observation noise is high (e.g., *LSHN* is the most learning-friendly environment); and (2) the presence of spoofing generally hurts the learning-based strategy and reduces background-trader surplus. Detailed equilibrium outcomes can be found in Appendix C.1. We next evaluate the two HBL variations.

Table 4. A set of basic background trading strategies that we include to compare to the two HBL variations.

Strategy	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	HBL ₁	HBL ₂
L	-	-	-	-	-	2	5
$R_{\min}$	0	0	0	0	0	-	-
$R_{\max}$	1000	1000	1000	500	250	-	-
η	0.4	0.8	1	0.8	0.8	-	-

7.2.2. HBL with Selective Price Level Blocking

Learning traders who choose to ignore certain orders face a natural tradeoff between losing useful information and correctly blocking spoof orders to avoid manipulation. We first examine, under *non-spoofing* environments, how learning effectiveness may be compromised by excluding orders at each price level. Starting with the equilibrium strategy profile of each non-spoofing market environment found in Section 7.2.1 (we arbitrarily select one if there are multiple equilibria found in a certain environment), we perform controlled experiments by letting background traders who adopt the standard HBL strategy ignore orders from a selected price level throughout the trading period. Table 5 compares the payoffs obtained by HBL in its standard form and variations that, respectively, block orders at the first, second and third price level in the order book. We find that, consistently across market settings, HBL agents benefit the most by learning from market best bids and asks and can achieve fairly similar performance even when orders at a selected level beyond the market best quotes are ignored.

In response to the HBL variation that ignores price levels, we extend the exploiter to be able to place spoof orders behind a chosen price level, denoted by ψ . For example, when $\psi = 2$, the exploiter injects spoof orders at one tick behind the second-best bid. We start with the same set of equilibrium strategy profiles and conduct controlled experiments to

evaluate how injecting spoof orders at different levels can change the manipulation effect, even when learning traders are considering the full order book (i.e., adopting standard HBL). We measure the effectiveness of each spoofing strategy by profits from trade as well as the price deviation caused by spoof orders. Experimental results (Table 5) show that the exploiter benefits the most by placing spoof orders behind the best bid (i.e., $\psi = 1$) and moving spoof orders to less competitive levels reduces exploitation profit. We further confirm this weakened manipulation effect in Figure 14, which showcases market price deviations caused by different spoofing strategies in the MSMN environment. We find the price rise diminishes as spoof orders are placed further away from the best bid.

Table 5. Average payoffs of learning-based background traders and the exploiter, as they deviate from the equilibrium strategy profiles found in Section 7.2.1. We deviate either background traders or the exploiter to its corresponding strategy variation. We refer to the exploiter who spoofs as SP and the one who only executes trades as EXP. Asterisks denote statistical significance at the 1% level for the paired t-test in payoffs compared to the standard HBL (*), $SP_{K=1}$ (*) and EXP (**).

Env	HBL	$HBL_{\chi=1}$	$HBL_{\chi=2}$	$HBL_{\chi=3}$	$SP_{\psi=1}$	$SP_{\psi=2}$	$SP_{\psi=3}$	EXP
LSHN	658	650 *	658	658	525	494 *,**	488 *	483 *
MSMN	655	645 *	655	655	356	312 *	299 *	295 *
HSLN	649	641 *	649	649	295	264 *	268 *,**	253 *




Figure 14. Price deviations caused by spoof orders placed behind different price levels in the order book.

Although our exploration of possible spoofing strategies here is limited, the results suggest that spoof orders near the market quotes tend to maximize manipulation effect. In response, HBL traders who adapt to the presence of spoofing may naturally block orders around such levels. Figure 15 shows that, when blocking the correct level, HBL traders can significantly increase their payoffs and reduce the amount the exploiter could profit via manipulation. This mitigated manipulation effect is further verified by the dashed blue line in Figure 14, which shows price deviations close to zero. Price differences are not strictly zero before spoofing (time 1000), as traders who adopt $HBL_{\chi=2}$ consistently block orders throughout the trading period.

Given these beneficial payoff deviations, in the final set of experiments, we conduct EGTA to find approximate Nash equilibria in games where background traders may choose trading strategies from the ZI family and HBLs that block a selected price level (any strategy in Table 4 or Table 6). Detailed equilibrium results can be found in Appendix C.2. As shown in Figure 18 (light grey columns), we find that: (1) adding the blocking strategic parameter does not affect the competitiveness of learning-based strategies with respect to ZI (HBL adoption rates in equilibrium remain in similar ranges as those of markets where only the standard HBL strategy is provided); and (2) the extended order blocking ability improves the learning robustness of HBL traders (compared to surplus decreases caused

by manipulation in markets where background agents are restricted to the standard HBL, background-trader surpluses are no longer significantly reduced when agents can strategically block orders in the face of manipulation). In other words, background traders who learn from market information but also strategically ignore orders can achieve robustness against manipulation and retain comparable effectiveness in non-manipulated markets.




Figure 15. Correctly blocking spoof orders increases background-trader surplus and decreases manipulation profits.

Table 6. A set of first HBL variation with different price level blocking parameters.

Strategy	HBL ₃	HBL ₄	HBL ₅	HBL ₆
L	2	2	5	5
χ	1	2	1	2

7.2.3. HBL with Price Offsets

Our second HBL variation relies on a price adjustment rather than information selection to adapt to different market conditions. We start by exploring a set of price offset intervals $[R_{\min}, R_{\max}]$, ranging from positive values that understate the learned offer prices (e.g., similar to price shading) to negative values that adjust prices to become more competitive. As in Section 7.2.2, we conduct controlled experiments starting from equilibrium profiles found in Section 7.2.1 and then deviating from standard HBL to allow price offsets. Figure 16 shows for the MSMN non-spoofing environment how HBL surplus and number of transactions vary in markets where HBL traders adopt different offset intervals. (HBL with positive offset usually generates much lower payoff. For presentation simplicity, we cropped the surplus decrease at -400 in Figure 16.) We find adjusting learned prices with a range of negative offsets can be generally beneficial in our setting where agents have reentry opportunities. It increases HBL payoff and facilitates transactions, thus improving overall price convergence in markets.

To test the effectiveness of spoofing against the new HBL variation, we further have the $SP_{\psi=1}$ spoof in markets where the learning background traders, respectively, adopt the standard HBL, $HBL_{[-10,0]}$, $HBL_{[-20,0]}$, $HBL_{[-40,0]}$ and $HBL_{[-80,0]}$. Figure 17 compares market price deviations caused by spoof orders in those markets. We find that, although all markets experience initial price rise as a result of misled pricing beliefs, the spoofing effect tends to wear off faster in markets where HBL traders adopt negative price offsets. This may be because negative offsets promote near-term transaction: as more transactions happen, HBL traders can glean true information from the transaction prices to construct more accurate belief functions, and the $SP_{\psi=1}$ places spoof orders at lower prices due to the widened bid-ask spreads. Indeed, we find that markets populated with the standard HBL, $HBL_{[-10,0]}$, $HBL_{[-20,0]}$ and $HBL_{[-40,0]}$, respectively, have average *spoof-order* prices of 99,972, 99,951, 99,945 and 99,950.




Figure 16. Average HBL surplus differences and total number of transactions in non-spoofing markets where HBL traders use different price offsets.






Figure 17. Market price deviations caused by spoofing in markets where HBL traders use different price offsets.

Finally, we conduct EGTA in games with and without spoofing to find Nash equilibria where background traders can choose from ZI strategies and HBL variations that adjust learned prices with certain offsets (Tables 4 and 7). Detailed equilibrium results can be found in Appendix C.3. Equilibrium results (Figure 18 white columns) show that the extended price offsets tend to largely improve HBL's profitability and background-trader surpluses, in markets both with and without manipulation. Such price adjustments can especially help learning traders to better adapt to high shock environments where prices are less predictable from past observations. However, the extended offsets may not directly address manipulation and improve learning robustness against spoofing.



(a) HBL (and its variations) adoption rates in equilibrium.

Figure 18. Cont.



(b) Background-trader surpluses achieved in equilibrium.

Figure 18. Total background-trader surpluses and HBL strategy adoption rates achieved at equilibria across different market settings. For each market environment, we compare four settings where background traders are, respectively, provided with the standard HBL strategy (dark grey), HBL with selective price blocking (light grey), HBL with price offsets (white) and HBL that combines the two variations (striped). Each marker specifies one equilibrium outcome in markets with spoofing (orange) and without spoofing (blue).

Table 7. A set of second HBL variations with different price offsets.

Strategy	HBL ₇	HBL ₈	HBL ₉	HBL ₁₀	HBL ₁₁	HBL ₁₂	HBL ₁₃	HBL ₁₄
L	2	2	2	2	5	5	5	5
$R_{\min}$	−10	−20	−40	−80	−10	−20	−40	−80
$R_{\max}$	0	0	0	0	0	0	0	0

7.3. Combine Order Blocking and Price Offsets

We observe that HBL with price offsets is overall competitive across different market settings, but its performance still degrades in markets with spoofing (refer to Figure 18, white columns). Since the second HBL variation demonstrates a general improvement in both settings with and without manipulation, we augment this variation with price level blocking to reduce vulnerability to spoofing. Specifically, we extend the background trading strategy set in Table 6 with three strategies: $HBL_{[-10,0]}^{K=2}$, $HBL_{[-20,0]}^{K=2}$ and $HBL_{[-40,0]}^{K=2}$, which appear to be competitive in our preliminary explorations. We conduct EGTA in a similar manner across market environments with and without spoofing (detailed results can be found in Appendix C.4). Equilibrium outcomes (Figure 18, striped columns) show that: (1) compared to markets where only the standard and the price-blocking HBL are provided, HBL that combines the two variations is more widely preferred and can help to increase overall background-trader surplus in equilibrium; and (2) across all environments, background-trader surpluses in markets with and without spoofing fall roughly into the same ranges. These suggest that, by combining the two proposed variations, HBL traders can enjoy both improved competitiveness and robustness against manipulation.

8. Conclusions

In this paper, we construct a computational model of spoofing: the tactic of manipulating market prices by targeting the order book. We design an HBL strategy that uses order book information to make pricing decisions. Since HBL traders use the order book, they are potentially vulnerable to spoofing attacks, and we confirm this in simulation analysis. We demonstrate that, in the absence of spoofing, HBL is generally adopted in equilibrium and benefits price discovery and social welfare. Although 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 that 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.

We further propose a cloaking mechanism to deter spoofing. The mechanism discloses a partially cloaked order book by symmetrically concealing a deterministic number of price levels from the inside. Our results demonstrate that the proposed cloaking mechanism can significantly diminish the efficacy of spoofing, but at the cost of a reduced HBL proportion and surplus in equilibrium. With the goal of maximizing background trader surplus, we perform empirical game-theoretic analysis across parametrically different mechanisms and environments, and find in markets with moderate shocks, the benefit of cloaking in mitigating spoofing outweighs its efficiency cost. By further exploring sophisticated spoofing strategies that probe to reveal cloaked information, we demonstrate the associated effort and risk exceed the gains, and verified that the proposed cloaking mechanism cannot be circumvented.

Two strategy variations based on the standard HBL strategy are explored. The first variation considers common characteristics of spoofing activities and works by offering agents the flexibility to neglect limit orders at a specified price level when assembling a dataset to learn from. The second variation learns from full order book information and later adjusts the target price derived from surplus maximization with a random offset to correct any biases in the learning process. Our analysis shows that the first HBL variation offers learning traders a way to strategically block orders to improve robustness against spoofing, while achieving similar competitiveness in non-manipulated markets. Our second HBL variation exhibits a general improvement over baseline HBL, in markets both with and without manipulation. Further explorations suggest that traders can enjoy both improved profitability and robustness by combining the two HBL variations.

Author Contributions: Spoofing market model conceptualization, methodology, and game-theoretic analysis, X.W. and M.P.W.; implementation of the spoofing simulator and experiments, X.W.; the cloaking mechanism conceptualization and analysis, X.W., Y.V. and M.P.W.; variations of learning trading strategies conceptualization, implementation, and result analysis, X.W., C.H. and M.P.W.; writing—original draft preparation, X.W.; writing—review and editing, C.H., Y.V. and M.P.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported in part by grant IIS-1421391 and CAREER Award IIS-1905558 from the US National Science Foundation.

Acknowledgments: We thank Uday Rajan and Erik Brinkman for helpful advice and discussions. This work extends and supersedes a series of conference papers: “Spoofing the Limit Order Book: An Agent-Based Model” presented at the 16th International Conference on Autonomous Agents and Multiagent Systems [59], “A Cloaking Mechanism to Mitigate Market Manipulation” presented at the 27th International Joint Conference on Artificial Intelligence [60], and “Learning-Based Trading Strategies in the Face of Market Manipulation” presented at the 1st ACM International Conference on AI in Finance [61].

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Detailed Equilibrium Results for Spoofing the Limit Order Book

Appendix A.1. Summary and Comparison of Markets without and with Spoofing

Table A1. Background-trader surplus and HBL proportion in equilibrium of markets without spoofing. Each row describes one Nash equilibrium found in a game (rounded to the nearest integer). Surpluses marked with asterisks indicate statistically significantly higher surpluses than those achieved in their corresponding markets with spoofing (see Table A3). $N = 28$ without spoofing.

Env	Surplus	95% CI	HBL%
LSLN	18,198 *	[18,127, 18,269]	88
LSLN	18,246 *	[18,177, 18,315]	98
LSMN	18,189 *	[18,116, 18,262]	100
LSHN	18,265 *	[18,200, 18,330]	100
MSLN	17,947 *	[17,851, 18,043]	58
MSLN	16,693 *	[16,613, 16,773]	0
MSMN	17,923 *	[17,829, 18,017]	62
MSMN	17,927 *	[17,839, 18,015]	43
MSMN	16,726	[16,638, 16,814]	0
MSHN	18,266 *	[18,188, 18,344]	100
HSLN	16,565	[16,485, 16,645]	0
HSLN	17,143 *	[17,055, 17,231]	0
HSMN	16,667	[16,565, 16,769]	0
HSHN	18,253 *	[18,179, 18,327]	87

Table A2. Similar to Table A1 above. Surpluses marked with asterisks indicate statistically significantly higher surpluses than those achieved in their corresponding markets with spoofing (see Table A4). $N = 65$ without spoofing.

Env	Surplus	95% CI	HBL%
LSLN	43,157 *	[43,016, 43,298]	71
LSLN	43,102 *	[42,980, 43,224]	73
LSLN	43,010 *	[42,885, 43,135]	95
LSMN	43,249 *	[43,106, 43,392]	83
LSMN	43,086 *	[42,964, 43,208]	79
LSHN	42,946	[42,817, 43,075]	94
MSLN	42,804 *	[42,647, 42,961]	57
MSMN	42,807 *	[42,652, 42,962]	56
MSMN	42,745 *	[42,610, 42,880]	56
MSHN	43,265 *	[43,128, 43,402]	86
HSLN	42,455*	[42,316, 42,594]	37
HSMN	42,383 *	[42,248, 42,518]	37
HSMN	42,144 *	[41,999, 42,289]	32
HSHN	42,981	[42,867, 43,095]	89

Table A3. Background-trader surplus and HBL proportion in equilibrium of markets with spoofing. Each row describes one Nash equilibrium found in a game (rounded to the nearest integer). $N = 28$ with spoofing.

Env	Surplus	95% CI	HBL%
LSLN	18,076	[18,000, 18,152]	78
LSMN	18,040	[17,971, 18,109]	91
LSHN	18,125	[18,054, 18,196]	87
MSLN	16,774	[16,707, 16,841]	0
MSMN	17,883	[17,795, 17,971]	34
MSMN	17,517	[17,429, 17,605]	24
MSMN	16,796	[16,708, 16,884]	0
MSHN	18,108	[18,032, 18,184]	81
HSLN	16,749	[16,680, 16,818]	0
HSMN	16,667	[16,565, 16,769]	0
HSHN	17,999	[17,923, 18,075]	97

Table A4. Similar to Table A3 above, but with $N = 65$ with spoofing

Env	Surplus	95% CI	HBL%
LSLN	42,868	[42,741, 42,995]	70
LSLN	42,993	[42,850, 43,136]	70
LSMN	42,961	[42,812, 43,110]	80
LSHN	43,061	[42,943, 43,179]	80
LSHN	43,103	[42,983, 43,223]	74
MSLN	42,639	[42,508, 42,770]	41
MSLN	42,698	[42,549, 42,847]	50
MSMN	42,624	[42,477, 42,771]	52
MSHN	43,038	[42,887, 43,189]	75
MSHN	43,101	[42,946, 43,256]	76
HSLN	41,815	[41,664, 41,966]	29
HSLN	39,502	[39,398, 39,606]	0
HSMN	40,091	[39,968, 40,214]	0
HSHN	43,143	[43,012, 43,274]	71

Appendix A.2. Markets without Spoofing

Table A5. Equilibria for games without spoofing, $N = 28$, calculated from the four-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus, HBL adoption rate and the equilibrium mixture probabilities of strategies included.

Env	Surplus	HBL	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSLN	18,198	0.88	0	0.12	0	0	0	0	0	0.88	0	0	0
LSLN	18,246	0.98	0	0	0	0	0.02	0	0	0	0.92	0.06	0
LSMN	18,189	1	0	0	0	0	0	0	0	0.82	0	0.18	0
LSHN	18,265	1	0	0	0	0	0	0	0	1	0	0	0
MSLN	17,947	0.58	0	0	0	0.40	0.02	0	0	0	0.40	0.18	0
MSLN	16,693	0	0	0	0	0	0	0.74	0.26	0	0	0	0
MSMN	17,923	0.62	0	0	0	0	0.38	0	0	0.44	0.18	0	0
MSMN	17,927	0.43	0	0.04	0.53	0	0	0	0	0.43	0	0	0
MSMN	16,726	0	0	0	0	0	0	0.80	0.20	0	0	0	0
MSHN	18,266	1	0	0	0	0	0	0	0	0.74	0.24	0	0.02

Table A5. Cont.

Env	Surplus	HBL	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	HBL ₁	HBL ₂	HBL ₃	HBL ₄
HSLN	16,565	0	0	0	0	0	0	0.73	0.27	0	0	0	0
HSLN	17,143	0	0	0	0.53	0	0	0	0.47	0	0	0	0
HSMN	16,667	0	0	0	0	0	0	1	0	0	0	0	0
HSHN	18,253	0.87	0	0	0.13	0	0	0	0	0.84	0	0	0.03

Table A6. Equilibria for games without spoofing, $N = 65$, calculated from the five-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus, HBL adoption rate and the equilibrium mixture probabilities of strategies included.

Env	Surplus	HBL	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSLN	43,157	0.71	0	0.29	0	0	0	0	0	0.59	0	0.12	0
LSLN	43,102	0.73	0	0	0.27	0	0	0	0	0.73	0	0	0
LSLN	43,010	0.95	0	0	0.05	0	0	0	0	0.22	0	0.73	0
LSMN	43,249	0.83	0	0	0.17	0	0	0	0	0.57	0	0.26	0
LSMN	43,086	0.79	0	0.05	0.16	0	0	0	0	0	0.79	0	0
LSHN	42,946	0.94	0	0.04	0.02	0	0	0	0	0.75	0.19	0	0
MSLN	42,804	0.57	0	0	0.43	0	0	0	0	0.31	0.26	0	0
MSMN	42,807	0.56	0	0	0.44	0	0	0	0	0.31	0.25	0	0
MSMN	42,745	0.56	0.01	0	0	0.43	0	0	0	0	0.56	0	0
MSHN	43,265	0.86	0	0.06	0	0.08	0	0	0	0.67	0.19	0	0
HSLN	42,455	0.37	0	0	0.63	0	0	0	0	0	0.18	0.19	0
HSMN	42,383	0.37	0	0	0.63	0	0	0	0	0.26	0	0.11	0
HSMN	42,144	0.32	0	0	0	0.54	0.14	0	0	0	0	0.32	0
HSHN	42,981	0.89	0	0.08	0	0	0	0.03	0	0.89	0	0	0

Appendix A.3. Markets with Spoofing

Table A7. Equilibria for games with spoofing, $N = 28$, calculated from the four-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus, HBL adoption rate and the equilibrium mixture probabilities of strategies included.

Env	Surplus	HBL	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSLN	18,076	0.78	0	0	0.22	0	0	0	0	0.78	0	0	0
LSMN	18,040	0.91	0	0	0	0.09	0	0	0	0.91	0	0	0
LSHN	18,125	0.87	0	0	0.13	0	0	0	0	0.87	0	0	0
MSLN	16,774	0	0	0	0	0	0	1	0	0	0	0	0
MSMN	17,883	0.34	0	0	0.11	0.55	0	0	0	0	0.34	0	0
MSMN	17,517	0.24	0	0	0.54	0	0	0	0.21	0.24	0	0	0
MSMN	16,796	0	0	0	0	0	0	1	0	0	0	0	0
MSHN	18,108	0.81	0	0	0.12	0.07	0	0	0	0.81	0	0	0
HSLN	16,749	0	0	0	0	0	0.04	0.96	0	0	0	0	0
HSMN	16,667	0	0	0	0	0	0	1	0	0	0	0	0
HSHN	17,999	0.97	0	0	0	0.03	0	0	0	0.75	0	0.22	0

Table A8. Equilibria for games with spoofing, $N = 65$, calculated from the five-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus, HBL adoption rate and the equilibrium mixture probabilities of strategies included.

Env	Surplus	HBL	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSLN	42,868	0.70	0.21	0	0.09	0	0	0	0	0.70	0	0	0
LSLN	42,993	0.70	0	0.30	0	0	0	0	0	0.54	0.16	0	0
LSMN	42,961	0.80	0	0	0.20	0	0	0	0	0.51	0.29	0	0
LSHN	43,061	0.80	0	0	0.20	0	0	0	0	0.80	0	0	0
LSHN	43,103	0.74	0	0.26	0	0	0	0	0	0.74	0	0	0
MSLN	42,639	0.41	0	0	0	0.59	0	0	0	0	0.41	0	0
MSLN	42,698	0.50	0	0	0.50	0	0	0	0	0.32	0	0.18	0
MSMN	42,624	0.52	0	0	0.48	0	0	0	0	0	0.38	0.14	0
MSHN	43,038	0.75	0	0.25	0	0	0	0	0	0.48	0.27	0	0
MSHN	43,101	0.76	0	0.24	0	0	0	0	0	0.41	0.35	0	0
HSLN	41,815	0.29	0	0	0.50	0	0	0.21	0	0	0.29	0	0
HSLN	39,502	0	0	0	0	0	0	0	1	0	0	0	0
HSMN	40,091	0	0	0	0	0	0	0.77	0.23	0	0	0	0
HSHN	43,143	0.71	0.10	0	0.19	0	0	0	0	0.71	0	0	0

Appendix A.4. Markets with Background Agents Restricted to ZI Strategies

Table A9. Equilibria for games where agents are restricted to ZI strategies, $N = 28$, calculated from the four-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus and the equilibrium mixture probabilities of strategies included.

Env	Surplus	95% CI	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇
LSLN	16,929	[16,862, 16,996]	0	0	0	0	0	0.98	0.02
LSMN	16,914	[16,841, 16,987]	0	0	0	0	0	0.89	0.11
LSHN	18,213	[18,137, 18,289]	0.22	0.78	0	0	0	0	0
MSLN	16,693	[16,613, 16,773]	0	0	0	0	0	0.74	0.26
MSMN	17,192	[17,106, 17,278]	0	0	0.42	0	0	0	0.58
MSMN	16,726	[16,638, 16,814]	0	0	0	0	0	0.80	0.20
MSHN	16,746	[16,675, 16,817]	0	0	0.09	0	0	0	0.91
MSHN	17,516	[17,438, 17,594]	0.38	0	0	0	0	0.62	0
HSLN	16,565	[16,485, 16,645]	0	0	0	0	0	0.73	0.27
HSLN	17,143	[17,055, 17,231]	0	0	0.53	0	0	0	0.47
HSMN	16,667	[16,565, 16,769]	0	0	0	0	0	1	0
HSHN	17,861	[17,787, 17,935]	0.31	0.39	0	0	0	0.30	0

Table A10. Equilibria for games where agents are restricted to ZI strategies, $N = 65$, calculated from the five-player DPR approximation. Each row of the table describes one equilibrium found with its corresponding surplus and the equilibrium mixture probabilities of strategies included.

Env	Surplus	95% CI	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇
LSLN	42,938	[42,840, 43,036]	0	1	0	0	0	0	0
LSLN	40,779	[40,677, 40,881]	0	0	0	0	0	1	0
LSMN	42,972	[42,870, 43,074]	0	0.97	0	0	0	0	0.03
LSMN	40,557	[40,439, 40,675]	0	0	0	0	0	0.83	0.17
LSHN	43,327	[43,192, 43,462]	0.44	0.56	0	0	0	0	0
LSHN	43,173	[43,063, 43,283]	0.11	0.89	0	0	0	0	0

Table A10. Cont.

Env	Surplus	95% CI	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇
MSLN	40444	[40342, 40546]	0	0	0	0	0	1	0
MSMN	39,622	[39518, 39726]	0	0	0	0	0	1	0
MSHN	43,140	[43017, 43263]	0	0.73	0.27	0	0	0	0
HSLN	40,523	[40398, 40648]	0	0	0.28	0	0	0	0.72
HSLN	40,038	[39903, 40173]	0	0	0	0	0	0.60	0.40
HSMN	40,458	[40327, 40589]	0	0	0	0.08	0	0.73	0.19
HSHN	43,197	[43087, 43307]	0	0.88	0	0.12	0	0	0

Appendix B. Detailed Equilibrium Results for a Cloaking Mechanism to Mitigate Spoofing

Table A11. Background trading strategies included in EGTA for cloaking mechanisms.

Strategy	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	ZI ₈	ZI ₉	HBL ₁	HBL ₂	HBL ₃	HBL ₄
<i>L</i>	-	-	-	-	-	-	-	-	-	2	3	5	8
<i>R</i> _{min}	0	0	0	0	0	0	250	250	250	250	250	250	250
<i>R</i> _{max}	1000	1000	1000	2000	2000	2000	500	500	500	500	500	500	500
<i>η</i>	0.4	0.8	1	0.4	0.8	1	0.4	0.8	1	1	1	1	1

Appendix B.1. Cloaking Markets without Spoofing

Table A12. Equilibria for games where the exploiter does not spoof. Each row of the table describes one equilibrium found with its corresponding background surplus, total surplus and HBL adoption rate. The results reported are based on at least 20,000 simulation runs.

Env	K	95% CI Background Surplus	95% CI Total Surplus	HBL Fraction
LSHN	K0	[42,121, 42,329]	[42,548, 42,694]	1.00
LSHN	K1	[41,848, 42,048]	[42,254, 42,396]	0.98
LSHN	K1	[41,769, 41,977]	[42,264, 42,406]	0.92
LSHN	K2	[41,788, 42,000]	[42,205, 42,347]	0.997
LSHN	K4	[41,572, 41,772]	[42,046, 42,188]	0.89
MSMN	K0	[41,958, 42,220]	[42,274, 42,388]	0.67
MSMN	K1	[41,902, 42,164]	[42,210, 42,324]	0.67
MSMN	K1	[41,849, 42,107]	[42,170, 42,284]	0.60
MSMN	K1	[41,801, 42,067]	[42,167, 42,281]	0.68
MSMN	K2	[41,742, 42,000]	[42,123, 42,237]	0.66
MSMN	K4	[41,693, 41,924]	[42,116, 42,230]	0.47
MSMN	K4	[38,809, 39,025]	[39,367, 39,485]	0.012
HSLN	K0	[41,529, 41,871]	[41,974, 42,088]	0.59
HSLN	K0	[41,698, 42,040]	[42,102, 42,216]	0.67
HSLN	K0	[41,625, 41,973]	[42,021, 42,135]	0.67
HSLN	K1	[41,417, 41,769]	[41,869, 41,983]	0.66
HSLN	K2	[41,377, 41,655]	[41,776, 41,890]	0.38
HSLN	K2	[39,728, 39,972]	[40,484, 40,594]	0
HSLN	K2	[38,691, 38,965]	[39,419, 39,537]	0
HSLN	K4	[39,557, 39,803]	[40,256, 40,374]	0
HSLN	K4	[39,558, 39,804]	[40,290, 40,408]	0

Table A13. Detailed equilibria for games where the exploiter does not spoof. Each row of the table describes one equilibrium found with its corresponding mixture of background strategies.

Env	K	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	ZI ₈	ZI ₉	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSHN	K0	0	0	0	0	0	0	0	0	0	0.70	0	0	0.30
LSHN	K1	0	0	0.02	0	0	0	0	0	0	0.73	0	0	0.25
LSHN	K1	0.05	0	0.03	0	0	0	0	0	0	0	0.88	0	0.04
LSHN	K2	0	0	0	0	0	0	0.003	0	0	0	0.856	0	0.141
LSHN	K4	0	0	0.11	0	0	0	0	0	0	0.29	0.60	0	0
MSMN	K0	0	0	0.33	0	0	0	0	0	0	0.51	0.16	0	0
MSMN	K1	0.11	0.01	0.21	0	0	0	0	0	0	0	0.14	0.53	0
MSMN	K1	0	0.20	0.20	0	0	0	0	0	0	0.20	0.15	0.25	0
MSMN	K1	0	0	0.32	0	0	0	0	0	0	0.14	0.39	0.03	0.12
MSMN	K2	0	0.15	0	0.19	0	0	0	0	0	0.11	0.40	0.15	0
MSMN	K4	0.20	0	0.33	0	0	0	0	0	0	0.40	0	0.07	0
MSMN	K4	0	0	0	0	0	0	0.674	0.312	0.002	0	0	0.012	0
HSLN	K0	0	0	0.12	0	0	0	0.29	0	0	0.49	0.10	0	0
HSLN	K0	0	0	0	0.33	0	0	0	0	0	0.66	0	0	0.01
HSLN	K0	0	0	0	0.19	0.14	0	0	0	0	0	0.50	0.17	0
HSLN	K1	0.05	0	0	0.29	0	0	0	0	0	0	0.09	0.57	0
HSLN	K2	0.27	0.35	0	0	0	0	0	0	0	0.08	0	0.30	0
HSLN	K2	0.03	0.29	0.13	0	0	0	0	0	0.55	0	0	0	0
HSLN	K2	0	0	0	0	0	0	0.25	0.34	0.41	0	0	0	0
HSLN	K4	0	0.35	0	0	0	0	0.65	0	0	0	0	0	0
HSLN	K4	0	0.36	0	0	0	0	0.64	0	0	0	0	0	0

*Appendix B.2. Cloaking Markets with Spoofing***Table A14.** Equilibria for games where the exploiter strategically chooses to spoof. Each row of the table describes one equilibrium found with its corresponding background surplus, total surplus, HBL and spoofing adoption rate. The results reported are based on at least 20,000 simulation runs.

Env	K	95% CI Background Surplus	95% CI Total Surplus	HBL Fraction	Spoofing Fraction
LSHN	K0	[41,693, 41,893]	[42,243, 42,389]	0.95	1.00
LSHN	K1	[41,848, 42,048]	[42,254, 42,396]	0.98	0.00
LSHN	K2	[41,788, 42,000]	[42,205, 42,347]	0.997	0.00
LSHN	K4	[41,564, 41,764]	[42,010, 42,152]	0.90	0.08
LSHN	K4	[41,572, 41,772]	[42,046, 42,188]	0.89	0.00
MSMN	K0	[41,652, 41,902]	[42,151, 42,265]	0.65	1.00
MSMN	K0	[41,622, 41,884]	[42,106, 42,220]	0.66	1.00
MSMN	K1	[41,902, 42,164]	[42,210, 42,324]	0.67	0.00
MSMN	K1	[41,849, 42,107]	[42,170, 42,284]	0.60	0.00
MSMN	K1	[41,801, 42,067]	[42,167, 42,281]	0.68	0.00
MSMN	K1	[41,749, 42,031]	[42,146, 42,260]	0.72	0.71
MSMN	K2	[41,700, 41,946]	[42,109, 42,223]	0.54	0.90
MSMN	K4	[41,655, 41,883]	[42,111, 42,225]	0.48	0.62
MSMN	K4	[38,809, 39,025]	[39,367, 39,485]	0.012	0.00
HSLN	K0	[41,538, 41,882]	[42,047, 42,161]	0.69	1.00
HSLN	K1	[41,417, 41,769]	[41,869, 41,983]	0.66	0.00
HSLN	K1	[41,039, 41,345]	[41,593, 41,707]	0.48	1.00
HSLN	K2	[41,080, 41,342]	[41,719, 41,833]	0.28	1.00
HSLN	K4	[39,557, 39,803]	[40,256, 40,374]	0	0.00
HSLN	K4	[39,558, 39,804]	[40,290, 40,408]	0	0.00

Table A15. Detailed Equilibria for games where the exploiter strategically chooses to spoof. Each row of the table describes one equilibrium found with its corresponding mixture of background strategies.

Env	K	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	ZI ₆	ZI ₇	ZI ₈	ZI ₉	HBL ₁	HBL ₂	HBL ₃	HBL ₄
LSHN	K0	0	0	0	0.05	0	0	0	0	0	0.95	0	0	0
LSHN	K1	0	0	0.02	0	0	0	0	0	0	0.73	0	0	0.25
LSHN	K2	0	0	0	0	0	0	0.003	0	0	0	0.856	0	0.141
LSHN	K4	0	0	0.10	0	0	0	0	0	0	0.38	0.52	0	0
LSHN	K4	0	0	0.11	0	0	0	0	0	0	0.29	0.60	0	0
MSMN	K0	0	0.19	0	0.16	0	0	0	0	0	0.65	0	0	0
MSMN	K0	0	0.20	0	0	0	0	0.14	0	0	0.61	0.05	0	0
MSMN	K1	0.11	0.01	0.21	0	0	0	0	0	0	0	0.14	0.53	0
MSMN	K1	0	0.20	0.20	0	0	0	0	0	0	0.20	0.15	0.25	0
MSMN	K1	0	0	0.32	0	0	0	0	0	0	0.14	0.39	0.03	0.12
MSMN	K1	0.28	0	0	0	0	0	0	0	0	0	0.51	0.21	0
MSMN	K2	0	0	0.46	0	0	0	0	0	0	0.35	0	0.19	0
MSMN	K4	0	0.52	0	0	0	0	0	0	0	0.46	0	0	0.02
MSMN	K4	0	0	0	0	0	0	0.674	0.312	0.002	0	0	0.012	0
HSLN	K0	0	0	0	0.31	0	0	0	0	0	0.69	0	0	0
HSLN	K1	0.05	0	0	0.29	0	0	0	0	0	0	0.09	0.57	0
HSLN	K1	0	0	0	0.33	0	0	0.19	0	0	0.08	0.40	0	0
HSLN	K2	0	0.72	0	0	0	0	0	0	0	0	0	0.28	0
HSLN	K4	0	0.35	0	0	0	0	0.65	0	0	0	0	0	0
HSLN	K4	0	0.36	0	0	0	0	0.64	0	0	0	0	0	0

Appendix C. Detailed Equilibrium Results for Learning-Based Trading Strategies under the Presence of Market Manipulation

Table A16. Background trading strategies used in EGTA for HBL variations. HBL_{n-L} in tables below means HBL_n with memory length of *L*.

Strategy	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	HBL ₁	HBL ₂	HBL ₃	HBL ₄	HBL ₅	HBL ₆	HBL ₇	HBL ₈	HBL ₉	HBL ₁₀
χ	NA	NA	NA	NA	NA	NA	1	2	NA	NA	NA	NA	2	2	2
$R_{\min}$	0	0	0	0	0	0	0	0	-10	-20	-40	-80	-10	-20	-40
$R_{\max}$	1000	1000	1000	500	250	0	0	0	0	0	0	0	0	0	0
η	0.4	0.8	1	0.8	0.8	1	1	1	1	1	1	1	1	1	1

Appendix C.1. Standard HBL

Table A17. Equilibria for games where the learning-based trading strategy set is restricted to standard HBL. Each row describes an equilibrium found for the game described by the Env column, detailing the adoption rate of each strategy considered and the corresponding background surplus. The equilibrium strategy profiles with checkmarks in the “Baseline” column indicates those used as baseline strategy profiles for controlled experiments.

Env	Baseline	ZI ₁	ZI ₂	ZI ₃	ZI ₄	ZI ₅	HBL ₁₋₂	HBL ₁₋₅	95% CI Background Surplus
LSHN	✓	0	0	0	0	0	0	1.00	[42,050, 42,142]
		0	0	0	0	0	1.00	0	[41,609, 41,703]
LSHN—Spoof	✓	0	0	0	0	0	0	1.00	[41,641, 41,733]
		0	0	0	0	0	1.00	0	[41,300, 41,393]
MSMN	✓	0.01	0	0	0	0	0	0.99	[41,820, 41,978]
		0	0	0	0	0.25	0.75	0	[41,779, 41,965]
		0	0	0.27	0	0	0.73	0	[41,693, 41,866]
MSMN—Spoof	✓	0	0	0.37	0	0	0.63	0	[41,493, 41,669]
		0.22	0	0	0	0	0.78	0	[41,702, 41,876]
		0	0.27	0	0	0	0.73	0	[41,642, 41,814]

Appendix C.3. HBL with Price Offsets

Table A19. Equilibria for games where the learning-based trading strategy set is comprised of standard HBL and HBL with price offsets. Each row describes an equilibrium found for the game described by the Env column, detailing the adoption rate of each strategy considered and the corresponding background surplus.

Env	ZI ₁	ZI ₂	ZI ₃	ZI ₄	HBL ₄₋₂	HBL ₅₋₂	HBL ₆₋₂	HBL ₄₋₅	HBL ₅₋₅	HBL ₆₋₅	95% CI Background Surplus
LSHN	0	0	0	0	0	0	0	0	1.00	0	[41,518, 42,562]
	0	0	0	0	0	0	0	1.00	0	0	[41,512, 42,556]
	0	0	0	0	0	0	1.00	0	0	0	[42,420, 42,507]
	0	0	0	0	0	1.00	0	0	0	0	[42,551, 42,640]
	0	0	0	0	1.00	0	0	0	0	0	[42,551, 42,639]
LSHN—Spoof	0	0	0	0	0	0	0	1.00	0	0	[42,430, 42,474]
	0	0	0	0	0	0	1.00	0	0	0	[42,406, 42,492]
	0	0	0	0	0	1.00	0	0	0	0	[42,527, 42,614]
	0	0	0	0	1.00	0	0	0	0	0	[42,516, 42,603]
MSMN	0	0	0	0	0	0	1.00	0	0	0	[42,085, 42,229]
	0.03	0	0	0	0	0.97	0	0	0	0	[42,227, 42,383]
	0	0	0	0	1.00	0	0	0	0	0	[42,219, 42,366]
	0	0	0.21	0	0	0	0	0	0	0.79	[41,702, 41,787]
	0.10	0	0	0	0	0	0	0	0.90	0	[42,058, 42,142]
MSMN—Spoof	0	0	0	0.08	0	0	0	0	0	0.92	[41,912, 41,991]
	0	0	0	0	0	0	1.00	0	0	0	[42,054, 42,197]
	0	0.16	0	0	0	0.84	0	0	0	0	[41,951, 42,119]
	0	0.13	0	0	0	0.87	0	0	0	0	[42,021, 42,185]
HSLN	0.12	0	0	0	0	0.88	0	0	0	0	[42,140, 42,255]
HSLN—Spoof	0.16	0	0	0	0	0	0	0	0	0.84	[41,782, 41,893]
	0	0.37	0	0	0	0	0	0	0.63	0	[41,457, 41,578]

Appendix C.4. HBL Price Offsets and Price Level Blocking

Table A20. Equilibria for games where the learning-based trading strategy set is comprised of standard HBL, HBL with price level blocking, HBL with price offsets and HBL with both price offsets and price level blocking (HBL₁ and HBL₂ are not shown because they do not appear in any equilibrium). Each row describes an equilibrium found for the game described by the Env column, detailing the adoption rate of each strategy considered and the corresponding background surplus.

Env	ZI ₁	ZI ₂	ZI ₃	HBL ₃₋₂	HBL ₄₋₂	HBL ₅₋₂	HBL ₆₋₋₂	HBL ₈₋₂	HBL ₉₋₂	HBL ₁₀₋₂	HBL ₄₋₅	HBL ₅₋₅	HBL ₆₋₅	HBL ₈₋₅	HBL ₉₋₅	HBL ₁₀₋₅	95% CI
LSHN	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.00	0	[42,520, 42,565]
	0	0	0	0	0	0	0	0	0	0	0	0	0	1.00	0	0	[42,511, 42,556]
	0	0	0	0	0	0	0	0	0	0	0	1.00	0	0	0	0	[41,518, 42,562]
	0	0	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	[41,512, 42,556]
	0	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	[42,423, 42,509]
	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	[42,550, 42,638]
	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	[42,555, 42,642]
	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	[42,420, 42,507]
	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	0	[42,551, 42,640]
	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	0	0	[42,551, 42,639]
	0	0	0	0	0	0	0	0	0	0	0	0	0	1.00	0	0	[42,511, 42,556]
	0	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	[42,422, 42,509]
	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	[42,551, 42,639]
	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	[42,554, 42,641]
	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	[42,406, 42,492]
	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	0	[42,527, 42,614]
0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	0	0	[42,516, 42,603]	
MSMN	0.22	0	0	0	0	0	0	0	0	0	0	0	0	0	0.78	0	[41,877, 41,967]
	0.11	0	0	0	0	0	0	0	0.89	0	0	0	0	0	0	0	[42,123, 42,291]
	0	0	0.20	0	0	0	0	0	0	0.80	0	0	0	0	0	0	[41,755, 41,921]
	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	[42,085, 42,229]
	0.03	0	0	0	0	0.97	0	0	0	0	0	0	0	0	0	0	[42,227, 42,383]
	0	0	0	0	1.00	0	0	0	0	0	0	0	0	0	0	0	[42,219, 42,366]
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.00	[42,060, 42,133]
	0	0	0	0	0	0	0	0	1.00	0	0	0	0	0	0	0	[42,246, 42,395]
	0	0.13	0	0	0.87	0	0	0	0	0	0	0	0	0	0	0	[42,021, 42,185]
0.25	0	0	0.75	0	0	0	0	0	0	0	0	0	0	0	0	[41,808, 41,988]	
HSLN	0.23	0	0	0	0	0	0	0	0.77	0	0	0	0	0	0	0	[41,845, 42,085]
	0.24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.76	[41,705, 41,824]
	0	0.37	0	0	0	0	0	0	0	0	0	0.63	0	0	0	0	[41,457, 41,578]
	0.29	0	0	0	0	0	0	0.71	0	0	0	0	0	0	0	0	[41,754, 41,997]
	0	0.32	0	0	0	0	0	0	0.68	0	0	0	0	0	0	0	[41,639, 41,877]

References

1. Lin, T.C.W. The new market manipulation. *Emory Law J.* **2015**, *66*, 1253–1314.
2. Kirilenko, A.A.; Kyle, A.S.; Samadi, M.; Tuzun, T. The Flash Crash: High frequency trading in an electronic market. *J. Financ.* **2017**, *72*, 967–998.
3. Aldrich, E.M.; Grundfest, J.; Laughlin, G. The Flash Crash: A New Deconstruction. 2017. Available online: <http://dx.doi.org/10.2139/ssrn.2721922> (accessed on 10 May 2017).
4. Hope, B. How ‘Spoofing’ traders dupe markets. *Wall Str. J.* **2015**. Available online: <https://www.google.com.hk/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKEwib2uakquHwAhXDBKYKHfjXAvAQFjACegQIAxAD&url=https%3A%2F%2Fwww.wsj.com%2Farticles%2Fhow-spoofing-traders-dupe-markets-1424662202&usg=AOvVaw0E29hp4F8GCnGgut00lkzm> (accessed on 22 February 2015).
5. Montgomery, J.D. Spoofing, Market Manipulation, and the Limit-Order Book. 2016. Available online: <http://dx.doi.org/10.2139/ssrn.2780579> (accessed on 3 May 2016).
6. Friedman, D. The double auction market institution: A survey. In *The Double Auction Market: Institutions, Theories, and Evidence*; Addison-Wesley: Boston, MA, USA, 1993; pp. 3–25.
7. Wellman, M.P. Putting the agent in agent-based modeling. *Auton. Agents Multi Agent Syst.* **2016**, *30*, 1175–1189.
8. LeBaron, B. Agent-based computational finance. In *Handbook of Computational Economics*, 1st ed.; Tesfatsion, L., Judd, K.L., Eds.; Elsevier: Amsterdam, The Netherlands, 2006; Volume 2, Chapter 24, pp. 1187–1233.
9. Paddrik, M.; Hayes, R.; and Steve Yang, A.T.; Beling, P.; Scherer, W. An agent based model of the E-Mini S&P 500 applied to Flash Crash analysis. In Proceedings of the IEEE Conference on Computational Intelligence for Financial Engineering and Economics, New York, NY, USA, 29–30 March 2012; pp. 1–8.
10. LeBaron, B.; Arthur, W.B.; Palmer, R. Time series properties of an artificial stock market. *J. Econ. Dyn. Control* **1999**, *23*, 1487–1516.
11. Palit, I.; Phelps, S.; Ng, W.L. Can a zero-intelligence plus model explain the stylized facts of financial time series data? In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, 4–8 June 2012; pp. 653–660.
12. Wah, E.; Wright, M.; Wellman, M.P. Welfare effects of market making in continuous double auctions. *J. Artif. Intell. Res.* **2017**, *59*, 613–650.
13. Wah, E.; Wellman, M.P. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Financ.* **2016**, *5*, 69–93.
14. Bookstaber, R. *Using Agent-Based Models for Analyzing Threats to Financial Stability*; Working paper; Office of Financial Research: Washington, DC, USA, 2012.
15. Wellman, M.P. *Trading Agents*; Morgan & Claypool: San Rafael, CA, USA, 2011.
16. Gode, D.K.; Sunder, S. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *J. Political Econ.* **1993**, *101*, 119–137.
17. Farmer, J.D.; Patelli, P.; Zovko, I.I. The predictive power of zero intelligence in financial markets. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 2254–2259.
18. Cliff, D. *Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments*; Technical report; Hewlett-Packard Labs: Palo Alto, CA, USA, 1997.
19. Cliff, D. ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Trans. Evol. Comput.* **2009**, *13*, 3–18.
20. Vytelingum, P.; Cliff, D.; Jennings, N.R. Strategic bidding in continuous double auctions. *Artif. Intell.* **2008**, *172*, 1700–1729.
21. Wright, M.; Wellman, M.P. Evaluating the stability of non-adaptive trading in continuous double auctions. In Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems, Stockholm, Sweden, 10–15 July 2018; pp. 614–622.
22. Gjerstad, S.; Dickhaut, J. Price formation in double auctions. *Games Econ. Behav.* **1998**, *22*, 1–29.
23. Gjerstad, S. The competitive market paradox. *J. Econ. Dyn. Control* **2007**, *31*, 1753–1780.
24. Tesouro, G.; Das, R. High-performance bidding agents for the continuous double auction. In Proceedings of the 3rd ACM Conference on Electronic Commerce, Tampa, FL, USA, 14–17 October 2001; pp. 206–209.
25. Tesouro, G.; Bredin, J.L. Strategic sequential bidding in auctions using dynamic programming. In Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 2, Bologna, Italy, 15–19 July 2002; pp. 591–598.
26. Lee, E.J.; Eom, K.S.; Park, K.S. Microstructure-based manipulation: Strategic behavior and performance of spoofing traders. *J. Financ. Mark.* **2013**, *16*, 227–252.
27. Wang, Y.Y. Strategic spoofing order trading by different types of investors in Taiwan Index futures market. *J. Financ. Stud.* **2019**, *27*, 65.
28. Martínez-Miranda, E.; McBurney, P.; Howard, M. Learning unfair trading: A market manipulation analysis from the reinforcement learning perspective. In Proceedings of the IEEE International Conference on Evolving and Adaptive Intelligent Systems, Natal, Brazil, 23–25 May 2016; pp. 103–109.
29. Tao, X.; Day, A.; Ling, L.; Drapeau, S. On detecting spoofing strategies in high frequency trading. *arXiv* **2020**, arXiv:2009.14818.

30. Athalye, A.; Carlini, N.; Wagner, D. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 274–283.
31. Goodfellow, I.; Shlens, J.; Szegedy, C. Explaining and harnessing adversarial examples. In Proceedings of the International Conference on Learning Representations, San Diego, CA, USA, 7–9 May 2015.
32. Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; Fergus, R. Intriguing properties of neural networks. In Proceedings of the International Conference on Learning Representations, Banff, AB, Canada, 14–16 April 2014.
33. Vorobeychik, Y.; Kantarcioglu, M. *Adversarial Machine Learning*; Morgan & Claypool: San Rafael, CA, USA, 2018.
34. Li, J.; Wang, X.; Lin, Y.; Sinha, A.; Wellman, M.P. Generating realistic stock market order streams. In Proceedings of 34th AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; pp. 727–734.
35. Dalvi, N.; Domingos, P.; Mausam, S.; Sanghai, S.; Verma, D. Adversarial classification. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 22–25 August 2004; pp. 99–108.
36. Lowd, D.; Meek, C. Adversarial learning. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, IL, USA, 21–24 August 2005.
37. Carlini, N.; Wagner, D.A. Towards evaluating the robustness of neural networks. *IEEE Symp. Secur. Priv.* **2017**, pp. 39–57.
38. Papernot, N.; McDaniel, P.; Jha, S.; Fredrikson, M.; Celik, Z.B.; Swami, A. The limitations of deep learning in adversarial settings. In Proceedings of the IEEE European Symposium on Security and Privacy (EuroS&P), Saarbruecken, Germany, 21–24 March 2016; pp. 372–387.
39. Alzantot, M.; Sharma, Y.; Elgohary, A.; Ho, B.J.; Srivastava, M.; Chang, K.W. Generating natural language adversarial examples. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, 31 October–4 November 2018; pp. 2890–2896.
40. Carlini, N.; Mishra, P.; Vaidya, T.; Zhang, Y.; Sherr, M.; Shields, C.; Wagner, D.; Zhou, W. Hidden voice commands. In Proceedings of the 25th USENIX Security Symposium (USENIX Security 16), Austin, TX, USA, 10–12 August 2016.
41. Grosse, K.; Papernot, N.; Manoharan, P.; Backes, M.; McDaniel, P.D. Adversarial examples for malware detection. In *European Symposium on Research in Computer Security*; Foley, S.N., Gollmann, D., Sneekenes, E., Eds.; Springer International Publishing: New York City, NY, USA, 2017; pp. 62–79.
42. Tong, L.; Li, B.; Hajaj, C.; Xiao, C.; Zhang, N.; Vorobeychik, Y. Improving robustness of ML classifiers against realizable evasion attacks using conserved features. In Proceedings of the 28th USENIX Security Symposium (USENIX Security 19), Santa Clara, CA, USA, 14–16 August 2019; pp. 285–302.
43. Bolloor, A.; He, X.; Gill, C.; Vorobeychik, Y.; Zhang, X. Simple physical adversarial examples against end-to-end autonomous driving models. In Proceedings of the IEEE International Conference on Embedded Software and Systems, Las Vegas, NV, USA, 2–3 June 2019.
44. Eykholt, K.; Evtimov, I.; Fernandes, E.; Li, B.; Rahmati, A.; Xiao, C.; Prakash, A.; Kohno, T.; Song, D.X. Robust physical-world attacks on deep learning visual classification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 1625–1634.
45. Sharif, M.; Bhagavatula, S.; Bauer, L.; Reiter, M.K. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, 24–28 October 2016; pp. 1528–1540.
46. Wu, T.; Tong, L.; Vorobeychik, Y. Defending against physically realizable attacks on image classification. In Proceedings of the International Conference on Learning Representation, Virtual Conference, 26 April–1 May 2020.
47. Tong, L.; Yu, S.; Alfeld, S.; Vorobeychik, Y. Adversarial regression with multiple learners. In Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 4946–4954.
48. Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; Vladu, A. Towards deep learning models resistant to adversarial attacks. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
49. Wong, E.; Kolter, J.Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018.
50. Goettler, R.L.; Parlour, C.A.; Rajan, U. Informed traders and limit order markets. *J. Financ. Econ.* **2009**, *93*, 67–87.
51. Chakraborty, T.; Kearns, M. Market making and mean reversion. In Proceedings of the 12th ACM Conference on Electronic Commerce, San Jose, CA, USA, 5–9 June 2011; pp. 307–314.
52. Wiedenbeck, B.; Wellman, M.P. Scaling simulation-based game analysis through deviation-preserving reduction. In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia Spain, 4–8 June 2012; pp. 931–938.
53. Hautsch, N.; Huang, R. Limit order flow, market impact, and optimal order sizes: Evidence from NASDAQ TotalView-ITCH data. In *Market Microstructure: Confronting Many Viewpoints*; Abergel, F., Bouchaud, J.P., Foucault, T., Lehalle, C.A., Rosenbaum, M., Eds.; Wiley: Hoboken, NJ, USA, 2012.
54. Wang, X.; Wellman, M.P. Market manipulation: An adversarial learning framework for detection and evasion. In Proceedings of the 29th International Joint Conference on Artificial Intelligence, Virtual Conference, 7–15 January 2020; pp. 4626–4632.
55. Prewit, M. High-frequency trading: Should regulators do more. *Mich. Telecommun. Technol. Law Rev.* **2012**, *19*, 131–161.
56. Biais, B.; Woolley, P. *High Frequency Trading*; Technical report; Toulouse University: Toulouse, France, 2012.

-
57. Leal, S.J.; Napoletano, M. Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading. *J. Econ. Behav. Organ.* **2019**, *157*, 15–41.
 58. Vorobeychik, Y.; Kiekintveld, C.; Wellman, M.P. Empirical mechanism design: Methods, with application to a supply-chain scenario. In Proceedings of the 7th ACM Conference on Electronic Commerce, Ann Arbor, MI, USA, 11–15 June 2006; pp. 306–315.
 59. Wang, X.; Wellman, M.P. Spoofing the limit order book: An agent-based model. In Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems, Sao Paulo, Brazil, 8–12 May 2017; pp. 651–659.
 60. Wang, X.; Vorobeychik, Y.; Wellman, M.P. A cloaking mechanism to mitigate market manipulation. In Proceedings of the 27th International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 10–15 July 2018; pp. 541–547.
 61. Wang, X.; Hoang, C.; Wellman, M.P. Learning-based trading strategies in the face of market manipulation. In Proceedings of the 1st ACM International Conference on AI in Finance, Virtual Conference, 15–16 October 2020.