

# A Cloaking Mechanism to Mitigate Market Manipulation

Xintong Wang,<sup>1</sup> Yevgeniy Vorobeychik,<sup>2</sup> Michael P. Wellman<sup>1</sup>

<sup>1</sup> University of Michigan, Ann Arbor

<sup>2</sup> Vanderbilt University

xintongw@umich.edu, yevgeniy.vorobeychik@vanderbilt.edu, wellman@umich.edu

## Abstract

We propose a *cloaking mechanism* to deter *spoofing*, a form of manipulation in financial markets. The mechanism works by symmetrically concealing a specified number of price levels from the inside of the order book. To study the effectiveness of cloaking, we simulate markets populated with background traders and an exploiter, who strategically spoofs to profit. The traders follow two representative bidding strategies: the non-spoofable *zero intelligence* and the manipulable *heuristic belief learning*. Through empirical game-theoretic analysis across parametrically different environments, we evaluate surplus accrued by traders, and characterize the conditions under which cloaking mitigates manipulation and benefits market welfare. We further design sophisticated spoofing strategies that probe to reveal cloaked information, and find that the effort and risk exceed the gains.

## 1 Introduction

The automation of trading has transformed today’s financial landscape, with low transaction costs, high convenience, and unprecedented levels of automated trading. With its undeniable virtues, this new marketplace also opens doors to new threats, such as vulnerability to *manipulation*. One recent lawsuit claimed evidence of thousands of manipulation episodes in US Treasury futures observed 2013 and 2014 [Hope, 2015b]. In January 2018, US government agencies filed civil and criminal charges against three major banks for manipulating metals and equities futures [Price, 2018]. New allegations emerge regularly.

We study *spoofing*, a common variety of manipulation, and potential means to deter it. Formally defined in the 2010 Dodd-Frank Act §747 as “bidding or offering with the intent to cancel the bid or offer before execution”, spoof orders, rather than expressing genuine buy or sell intent, are designed to falsely signal demand or supply. In terms of adversarial learning, this can be viewed as a *poisoning attack* [Barreno *et al.*, 2006] targeting the order book. It may lead other investors—those who learn from the order stream—to believe that prices may soon rise or fall, and thus alter their own behavior in a way that will directly move the price.

To design low-risk but effective strategies, manipulators often rely on the instant order book information disclosed by standard market mechanisms [Hope, 2015a; Montgomery, 2016]. 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.

Despite regulatory enforcement and detection efforts, spoofing is hard to catch in each individual case. To identify spoofing, one has to establish the manipulation intent behind cancellations of the placed orders. However, this is not easy, as order cancellation is a legitimate action for many non-manipulative participants, and according to statistics, 95% of NASDAQ limit orders are canceled, with a median order lifetime less than one second [Hautsch and Huang, 2012].

Given the difficulty of detection, we propose a *cloaking mechanism* to deter spoofing. The mechanism extends the traditional continuous double auction (CDA) 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 to post misleading bids, while not unduly degrading the general usefulness of market information. We focus on deterministic cloaking in this study, as a stochastic mechanism may raise issues regarding verification of faithful market operations.

We employ agent-based simulation to evaluate the proposed mechanism under equilibrium settings across a variety of parametrically distinct market environments. The market is populated with multiple background agents and one exploiter, trading a single security. Background traders with private values are further divided into non-spoofable fundamental agents using instances of the *zero intelligence* (ZI) strategy, and *heuristic belief learning* (HBL) agents exploiting the *disclosed* order book to make bidding decisions. Extended from prior literature [Gode and Sunder, 1993; Gjerstad and Dickhaut, 1998], these two broad families of trading strategies have proved robust and competitive, and importantly, are not tailored to either spoofing or cloaking environments. The exploiter profits by first buying the underlying security at low prices and later selling at higher ones. To increase profit, it can try to manipulate the market through spoofing after its original purchase.

Our results show that the proposed cloaking mechanism can effectively mitigate spoofing, but at the expense of a

lower proportion of HBL traders at equilibrium due to loss of order book information. To characterize this tradeoff, we conduct *empirical game-theoretic analysis* [Wellman, 2016] to understand agents’ strategic responses to the proposed mechanism, and perform *empirical mechanism design* [Vorobeychik *et al.*, 2006] to set cloaking parameters that maximize efficiency. Our analyses show for a range of different market environments with moderate shocks, the benefit of cloaking in mitigating spoofing outweighs its costs in information transmission. We also explore more sophisticated spoofing strategies that use probing to reveal cloaked information, but find the cost and risk of such tactics exceed the gains.

## 2 Prior Work

We build on our existing computational model of spoofing [Wang and Wellman, 2017]. In this model, learning traders can benefit price discovery and social welfare, but their existence renders a market vulnerable to manipulation: simple spoofing strategies can effectively mislead traders, distort prices, and reduce total surplus. Interestingly, learning traders persist even with manipulators, which suggests that the elimination of spoofing requires active measures. Along a separate line, Martínez-Miranda *et al.* [2016] implement spoofing within a reinforcement learning framework to model conditions where such behavior is incentivized and effective.

Several works seek to identify spoofing strategy characteristics [Lee *et al.*, 2013], understand its profitability and real-time impact [Wang, 2015] through studying historical market data. Lin [2018] from a legal perspective surveys both traditional and new forms of manipulation and summarizes the challenges of detecting new disruptive practices.

Research on mitigating manipulation is fairly limited. Prewit [2012] and Biais *et al.* [2012] advocate the imposition of cancellation fees to deter manipulative strategies that frequently cancel orders. Others argue that cancellation fees could discourage the beneficial activity of liquidity providers [Copeland and Galai, 1983; Foucault *et al.*, 2003], and in the event of a market crash, such a policy may lengthen the recovery process [Leal and Napoletano, to appear].

## 3 Market Model

### 3.1 Market Mechanism

We extend the CDA spoofing model [Wang and Wellman, 2017] to support order book cloaking. Prices in the market take discrete values at integer multiples of tick size one. Time is also discrete but fine-grained over a finite horizon  $T$ . The fundamental value of the underlying security, denoted by  $r_t$ , changes throughout the trading period according to a mean-reverting stochastic process [Chakraborty and Kearns, 2011; Wah *et al.*, 2017], for  $t \in [0, T]$ :

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

$\kappa \in [0, 1]$  specifies the degree to which the time series reverts back to the fundamental mean  $\bar{r}$ .  $u_t$  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  $\sigma_s^2$  represents an environment-specific shock variance.

Agents trade a single security by submitting limit orders that specify the maximum (minimum) price at which they would be willing to buy (sell) some number of units. The market maintains a *limit order book* of outstanding orders. We use  $BID_t^k$  to denote the  $k$ th-highest bid price in the book at time  $t$ , and  $ASK_t^k$  the  $k$ th-lowest ask price at  $t$ . The cloaking mechanism symmetrically hides a deterministic number of price levels  $K$  from inside of the book. For example, when  $K = 1$ , the mechanism conceals orders at the best quotes, whereas when  $K = 0$ , the market acts as a standard CDA. Thus, the *disclosed* order book in a  $K$ -level cloaking mechanism 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. Agents may learn from the disclosed order book at their own discretion.

### 3.2 Agents in the Market

The market is populated with many background traders and a single exploiter. Background traders with private values represent investors with preferences on longing or shorting in the underlying security, whereas the exploiter without any private value can spoof the market and seek to profit through buying at lower prices and later selling at higher ones.

The position preference of background trader  $i$  is captured by a private value vector  $\Theta_i$  of length  $2q_{\max}$ , where  $q_{\max}$  is the maximum number of units one can be long or short at any time. Element  $\theta_i^{q+1}$  represents the marginal gain from buying an additional unit given current net position  $q$ . We generate  $\Theta_i$  from a set of  $2q_{\max}$  values independently drawn from  $N(0, \sigma_{PV}^2)$ , sort elements to reflect diminishing marginal utility, and assign  $\theta_i^q$  accordingly. The agent’s overall *valuation* for a unit of the security is the sum of fundamental and its private value.

Arrivals of a background trader follow a Poisson process with a rate  $\lambda_a$ . On each entry, the trader observes an agent- and time-specific noisy fundamental  $o_t = r_t + n_t$  with the observation noise following  $n_t \sim N(0, \sigma_n^2)$ . This noisy observation aims to capture different perceptions on the intrinsic value of the underlying security. Given this incomplete information about the fundamental, agents can potentially benefit from considering market information, which is influenced by the aggregate observations of other agents. To react to the new observation, the background trader withdraws its previous order (if untransacted) upon arrival, and submits a new single-unit order either to buy or sell as instructed with equal probability. A background trader’s order price is jointly decided by its valuation and *trading strategy* (see §3.3).

### 3.3 Background Trading Strategies

#### Estimation of the Final Fundamental

As holdings of the security are evaluated at the end of a trading period, background traders estimate the final fundamental value based on their noisy observations. 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 taking ac-

count of mean reversion for the interval since preceding arrival ( $\delta = t - t'$ ):

$$\begin{aligned}\tilde{r}_{t'} &\leftarrow (1 - (1 - \kappa)^\delta)\bar{r} + (1 - \kappa)^\delta \tilde{r}_{t'}; \\ \tilde{\sigma}_{t'}^2 &\leftarrow (1 - \kappa)^{2\delta}\tilde{\sigma}_{t'}^2 + \frac{1 - (1 - \kappa)^{2\delta}}{1 - (1 - \kappa)^2}\sigma_s^2.\end{aligned}$$

The estimates for  $t$  are then given by:

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

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:

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

### ZI Background Trading Strategy

We employ an extended and parameterized version of *zero intelligence* (ZI) as a representative strategy that is non-spoofable, as bids are generated without regard to order book information. The strategy has been widely adopted in agent-based finance due to its simplicity and effectiveness for market modeling [Gode and Sunder, 1993; Farmer *et al.*, 2005].

The ZI agent computes a limit-order price by shading its valuation with a random offset, which is uniformly drawn from  $[R_{\min}, R_{\max}]$ . Specifically, a ZI trader  $i$  arriving at time  $t$  with position  $q$  generates a limit price:

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

The ZI further takes into account the current *visible* quoted price, controlled by a strategic surplus threshold parameter  $\eta \in [0, 1]$ . Before submitting a new limit order, if the agent could achieve a fraction  $\eta$  of its requested surplus by accepting the most competitive visible order, it would take that quote by submitting an order at the same price. However, this may result in a transaction at a better price that is cloaked.

### HBL Background Trading Strategy

We adopt *heuristic belief learning* (HBL) as our representative class of strategies that exploit order book information. Introduced by Gjerstad [1998; 2007], HBL was further adapted to complex market environments that support persistent orders, buy-sell flexibility, and multi-unit trading [Wang and Wellman, 2017]. Both prior works showed the existence of HBL traders can boost price convergence, and notably benefit price discovery and social welfare, compared to markets populated exclusively with non-learning ZI traders. Disclosing aggregate market information to encourage learning can thus be a tool for market designers to promote efficiency.

The strategy is centered on belief functions that traders form on the basis of *observed* market data, and is thus susceptible to order-based manipulation. Specifically, the HBL agent estimates the probability that orders at various prices would be accepted, based on frequencies of transacted and rejected bids and asks during the its *memory length*  $L$ . Based on this estimation, it chooses a limit price that maximizes the expected surplus given current valuation estimates.

Here, we further tailor the HBL strategy to the cloaking mechanism by considering only the *revealed* order book information. Orders at competitive price levels may be ignored in the belief function if they are kept hidden during lifetime; or they can be considered with delays if later exposed in visible levels. This neglect, delay and offsets in estimation can all affect HBL's trading performance.

### 3.4 Exploitation and Spoofing Strategy

The exploitation strategy includes three stages. At the beginning of a trading period  $[0, T_{\text{spoofer}}]$ , the exploiter accepts any sell order at price lower than the fundamental mean  $\bar{r}$ .

During the second stage  $[T_{\text{spoofer}}, T_{\text{sell}}]$ , the exploiter, if also manipulating, submits spoof buy orders at a tick behind the first *visible* bid  $\text{BID}_{T_{\text{spoofer}}}^{K+1} - 1$  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  $\text{BID}_t^{K+1} - 1$ . By maintaining a large volume of buy orders a tick behind, the exploiter is spoofing at the most competitive price level that is not subject to any transaction risk. This simple spoofing strategy specifically aims to boost price, in the hope that the units purchased earlier can be later sold at higher prices.

During the last stage  $[T_{\text{sell}}, T]$ , the exploiter starts to sell the units it previously purchased by accepting any buy orders at a price higher than  $\bar{r}$ . The exploiter who also manipulates continues to spoof until  $T$  or all the bought units are sold.

### 3.5 Surplus

A background trader's surplus is its net cash from trading plus the *final valuation* of holdings at  $T$ , whereas an exploiter's payoff is its gain or loss from trading. The market's final valuation of background trader  $i$  with final holdings  $H$  is  $r_T H + \sum_{k=1}^{k=H} \theta_i^k$  for long position  $H > 0$ , or alternatively,  $r_T H - \sum_{k=H+1} \theta_i^k$  for short position  $H < 0$ .

## 4 Empirical Game-theoretic Analysis

To evaluate the proposed cloaking mechanism and understand potential strategic responses from market participants, we conduct agent-based simulations and game-theoretic analysis of the model described in §3.<sup>1</sup>

### 4.1 Market Environment Settings

We consider three environments varying in market shock  $\sigma_s^2$  and observation noise  $\sigma_n^2$ . **LSHN** represents a market with *low* shock and *high* observation noise  $\{10^5, 10^9\}$ , **MMSN** a market with *medium* shock and *medium* observation noise  $\{5 \times 10^5, 10^6\}$ , and **HSLN** a market with *high* shock and *low* observation noise  $\{10^6, 10^3\}$ . Shock variance governs fluctuations in the fundamental time series, and observation variance the quality of information agents get about the true fundamental. Intuitively, low shocks increase the predictability of future price outcomes and high observation noise limits what an agent can glean from its own information, and thus may encourage exploiting market information. For each environment, we consider cloaking mechanisms with  $K \in$

<sup>1</sup>Detailed equilibrium outcomes and simulation results are posted at <http://hdl.handle.net/2027.42/143507>.

Strategy	ZI <sub>1</sub>	ZI <sub>2</sub>	ZI <sub>3</sub>	ZI <sub>4</sub>	ZI <sub>5</sub>	ZI <sub>6</sub>	ZI <sub>7</sub>	ZI <sub>8</sub>	ZI <sub>9</sub>	HBL <sub>1</sub>	HBL <sub>2</sub>	HBL <sub>3</sub>	HBL <sub>4</sub>
$L$	-	-	-	-	-	-	-	-	-	2	3	5	8
$R_{\min}$	0	0	0	0	0	0	250	250	250	250	250	250	250
$R_{\max}$	1000	1000	1000	2000	2000	2000	500	500	500	500	500	500	500
$\eta$	0.4	0.8	1	0.4	0.8	1	0.4	0.8	1	1	1	1	1

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

$\{1, 2, 4\}$ , and compare market performance with cloaking to that of a standard CDA ( $K = 0$ ). 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 exploiter. The global fundamental time series is generated according to (1) with fundamental mean  $\bar{r} = 10^5$ , mean reversion  $\kappa = 0.05$ . Each trading period lasts  $T = 10,000$  time steps. Background traders arrive in the market according to a Poisson distribution with a rate  $\lambda_a = 0.005$  and the maximum number of units background traders can hold at any time is  $q_{\max} = 10$ . Private value variance is  $\sigma_{pV}^2 = 5 \times 10^6$ .

Table 1 specifies our background trading strategy set, comprising nine versions of ZI and four of HBL.<sup>2</sup> Background agents choose from this restricted set to maximize payoff. Following the spoof strategy described in §3.4, the exploiter stops buying and starts to manipulate at  $T_{\text{spoof}} = 1000$  by submitting orders with volume  $Q_{\text{sp}} = 200$ . If following a pure exploitation strategy, it waits until  $T_{\text{sell}} = 2000$  and starts to sell units purchased earlier at prices above  $\bar{r}$ .

## 4.2 EGTA Process

To identify equilibria under different market settings, we employ empirical game-theoretic analysis (EGTA), a methodology for finding equilibria in games defined by heuristic strategy space and simulated payoff data [Wellman, 2016]. It takes an iterative process to find candidate equilibria in subgames, incrementally add strategies and confirm or refute candidate solutions by examining deviations, until quiescence. This methodology has been adopted by a variety of multi-agent system studies, especially under complex market-based scenarios where applying standard analytic means is hard [Wah *et al.*, 2016; Brinkman and Wellman, 2017].

We model the market as a role-symmetric game, which is defined by an environment and a set of players representing two roles: background traders and one exploiter. As game size grows exponentially in players and strategies, we apply *deviation-preserving reduction* (DPR) [Wiedenbeck and Wellman, 2012] to approximate many-player games as fewer-player games through aggregation. DPR preserves payoffs from single-player deviations and has been shown to produce good approximations in several settings.

To facilitate DPR, we choose 64 background traders in this study to ensure that the required aggregations come out as integers. Specifically, we reduce markets with 64 background traders and one exploiter to games with four background traders and a single exploiter; with one background player deviating to a new strategy, the remaining 63 players

<sup>2</sup>We also explored ZI strategies with different shading ranges and HBL strategies with longer memory lengths, but they fail to appear in equilibrium.

can be further reduced to three. To account for stochastic features, such as market fundamental series, agent arrival patterns and private valuations, we sample at least 20,000 simulation runs for a specified strategy profile of each game to reduce sampling error.

## 4.3 Tradeoff Faced by Cloaking Mechanisms

We start by separately investigating the impact of cloaking on background traders, and on exploitation with spoofing. Our first set of games cover the range of cloaking environments without manipulation (i.e., the exploiter is non-spoofing).

Fig. 1a displays the HBL adoption rate (i.e., total probability over HBL strategies) at equilibrium across cloaking mechanisms in the three environments. We find 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, as previous hidden orders can become uninformative or even misleading by the time they are revealed. Given the decreasing HBL prevalence and effectiveness, background surplus achieved at equilibrium also decreases, as we see in Fig. 2b (blue diamonds).

Next, we examine whether the cloaking mechanism can effectively mitigate manipulation. We perform controlled experiments by letting the exploiter also execute the spoofing strategy against each found equilibrium, and compare the impact of spoofing under the cloaking mechanisms to that under a standard CDA. For every equilibrium, we simulate at least 10,000 paired instances and evaluate their differences on transaction price and agents’ payoffs. Transaction price difference measures the extent of price distortion, and is defined as the most recent transaction price of a game with spoofing minus that of the paired game without spoofing. Similarly, surplus difference is calculated by comparing profits obtained in a game with spoofing and that of its pair without spoofing. In each paired instance, background agents play the same strategies, and experience identical arrival times, private values, and noisy observations of the fundamental. Therefore, all changes in bidding behavior and outcome are caused by the spoof orders.

Empirical results from controlled experiments show that cloaking can considerably diminish price distortion caused by spoofing across environments. Fig. 1b demonstrates a specific but representative environment MSMN: in the standard CDA ( $K = 0$ ), transaction prices significantly rise subsequent to the execution of spoofing at  $T_{\text{sp}} = 1000$ , as HBL traders are tricked by the spoof buy orders; in cloaked markets, this price rise is effectively mitigated. Fig. 1c further illustrates the surplus redistribution between background traders and the ex-




Figure 1: The impact of cloaking on HBL adoption rate and spoofing respectively.

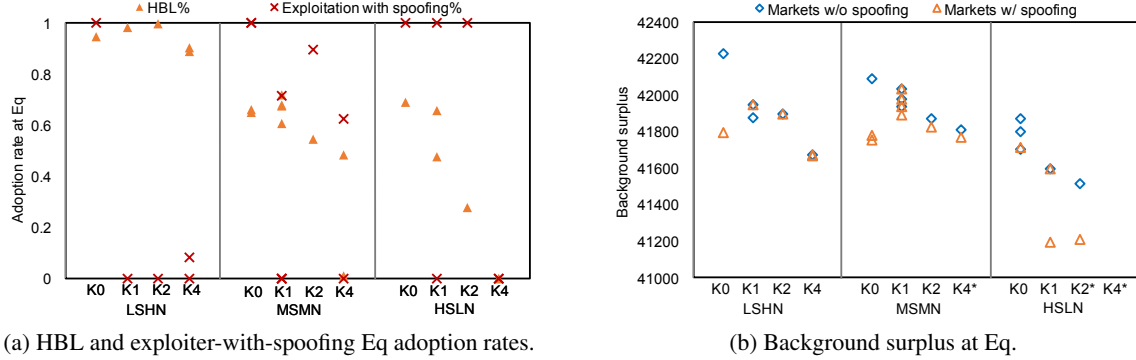


Figure 2: Equilibrium outcomes in games with and without cloaking. Each marker represents one equilibrium of the environment.

exploiter when it also spoofs. We find the exploiter can robustly profit from learning agents’ manipulated beliefs in standard CDAs. 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.

#### 4.4 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 where the exploiter can strategically choose to spoof, and background traders can execute any strategy in Table 1. This allows traders to adapt to spoofing and spoofers to traders, at given levels of order book cloaking.

Our findings are presented in Fig. 2.<sup>3</sup> We compare equilibrium outcomes of the cloaking mechanisms to those of standard CDAs primarily on two metrics: the probability of spoofing and background-trader surplus in equilibrium. As shown in Fig. 2a, 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.

<sup>3</sup>Equilibria with only ZIs usually achieve much lower surplus than those with HBLs. For presentation simplicity, we omit all-ZI equilibria from the chart. Environments with such cases are marked with asterisks.

This weakened spoofing effect is further confirmed in Fig. 2b, which compares the total background-trader surplus achieved in equilibrium under mechanisms with and without cloaking. Results show that under standard CDAs, 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). Favorably, we find the decrease in surplus due to spoofing can be considerably mitigated by order book cloaking. As shown in Fig. 2b, the vertical distances between the blue diamonds and orange triangles get smaller with  $K > 0$ . More importantly, we find the benefit of this improved robustness to spoofing can outweigh its associated efficiency costs in markets with moderate shocks (LSHN and MSMN). In those environments, background traders in mechanisms that cloak one or two price levels achieve higher surplus than those in standard CDAs. However, in a market with high shocks (HSLN), hiding or delaying even a little market information degrades learning to such a degree as to render cloaking counter-productive.

#### 4.5 Smarter Spoofing Strategies

To this point we have considered only spoofers unwilling to take risk of execution on their spoof orders. A more sophisticated manipulator can *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 value of such probing to the spoofing agent.

We design and evaluate parameterized versions of the spoofing strategy combined with probing. The strategy is

Env		$(\delta, l)$			
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)

Table 2: Least number of probes required by the smart spoofing strategy to beat equilibrium performance.

governed by two parameters: the step size  $\delta$ , which controls the probing aggressiveness, and the maximum attempts allowed per time step  $l$ , which limits 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, iteratively at a decreasing rate (as a function of  $\delta$  and the attempt number), until a higher price is displayed or the number of probing attempts reaches  $l$ . If probing causes a transaction, the spoofer halves the price increment and re-probes.

Table 2 reports, for cloaking-beneficial environments, the minimum  $l$  required for the corresponding step size  $\delta \in \{1, 2, 4, 8\}$  to achieve statistically significantly higher payoffs than the equilibrium performance of an exploiter in §4.4.4. We find in order to achieve higher payoffs, the spoofer has to probe with multiple attempts each time, and conservative probing strategy with smaller  $\delta$  usually requires more effort. However, in practice, such frequent cancellations and placements of orders can largely increase the risk of the associated spoofing intent being identified.

Fig. 3 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  $\delta$  and the attempt limit  $l$ . As we see from Fig. 3a, extra probing attempts steadily increase the transaction risk, but do not necessarily improve payoff. Moreover, the spikiness of the exploiter’s payoff indicates optimizing  $(\delta, l)$  to maximize profit is a challenging task. Fig. 3b further demonstrates that an exploiter can probe aggressively with larger step sizes to reduce effort, but usually at the cost of a higher transaction risk, and consequently a lower payoff. In highly dynamic markets with frequently updated quotes, it seems that finding an appropriate  $\delta$  to successfully probe a cloaking mechanism in a reasonable number of attempts would be challenging.

We have also explored more aggressive probing strategies, where the spoofer probes to expose multiple hidden levels and spoof at even higher prices. To accomplish that, the spoofer

<sup>4</sup>Dashes in the table indicate that an exploiter cannot beat the equilibrium performance with the corresponding  $\delta$ .




Figure 3: Exploitation payoff and transaction risk with varying price increment  $\delta$  and probing limit  $l$ .

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 usually accepted by background traders due to adverse selection. Those transactions will typically cause an accumulation in position, and consequently a negative payoff at the end of the trading period.

## 5 Conclusion

We proposed a cloaking mechanism to deter spoofing, a manipulative tactic that targets the order book. The mechanism discloses a partially cloaked order book by symmetrically concealing a deterministic number of price levels from the inside. We adopted an agent-based simulation approach to model such cloaking markets populated with multiple background traders and a single exploiter. Background traders can strategically choose from a set of ZI and HBL strategies, and the exploiter can spoof to maximize payoff. We conducted empirical game theoretic analysis to understand agents’ strategic responses to the proposed mechanism, and evaluate the effectiveness and robustness of cloaking.

Our results demonstrate 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 performed empirical game-theoretic analysis across parametrically different mechanisms and environments, and found 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 demonstrated the associated effort and risk exceed the gains, and verified that the proposed cloaking mechanism cannot be easily circumvented.

We acknowledge several aspects that may affect of our equilibrium analysis, including sampling error, reduced-game approximation and restricted bidding strategy coverage. Despite these limitations that are inherent in any complex modeling effort, we believe our proposal and analysis of the cloaking mechanism can serve as a constructive basis to study other methods deterring spoofing (or similar forms of manipulation), and identify practical considerations that should be regarded when making regulatory decisions.

## Acknowledgments

We thank the anonymous reviewers for constructive feedback. This work was supported in part by the US National Science

Foundation (IIS-1649972, IIS-1741190) and Army Research Office (W911NF-16-1-0069).

## References

- [Barreno *et al.*, 2006] Marco Barreno, Blaine Nelson, Russell Sears, Anthony D. Joseph, and J. Doug Tygar. Can machine learning be secure? In *ACM Symposium on Information, Computer and Communications Security*, pages 16–25, 2006.
- [Biais and Woolley, 2012] Bruno Biais and Paul Woolley. High frequency trading. Technical report, Toulouse University, 2012.
- [Brinkman and Wellman, 2017] Erik Brinkman and Michael P. Wellman. Empirical mechanism design for optimizing clearing interval in frequent call markets. In *18th ACM Conference on Economics and Computation*, pages 205–221, 2017.
- [Chakraborty and Kearns, 2011] T. Chakraborty and M. Kearns. Market making and mean reversion. *11th ACM Conference on Electronic Commerce*, pages 307–314, 2011.
- [Copeland and Galai, 1983] Thomas E. Copeland and Dan Galai. Information effects on the bid-ask spread. *Journal of Finance*, 38(5):1457–1469, 1983.
- [Farmer *et al.*, 2005] J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102:2254–2259, 2005.
- [Foucault *et al.*, 2003] Thierry Foucault, Ailsa Röell, and Patrik Sandås. Market making with costly monitoring: An analysis of the SOES controversy. *Review of Financial Studies*, 16(2):345–384, 2003.
- [Gjerstad and Dickhaut, 1998] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- [Gjerstad, 2007] Steven Gjerstad. The competitive market paradox. *Journal of Economic Dynamics and Control*, 31:1753–1780, 2007.
- [Gode and Sunder, 1993] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, pages 119–137, 1993.
- [Hautsch and Huang, 2012] Nikolaus Hautsch and Ruihong Huang. Limit order flow, market impact, and optimal order sizes: Evidence from NASDAQ TotalView-ITCH data. In Frederic Abergel, Jean-Philippe Bouchaud, Thierry Foucault, Charles-Albert Lehalle, and Mathieu Rosenbaum, editors, *Market Microstructure: Confronting Many Viewpoints*. Wiley, 2012.
- [Hope, 2015a] Bradley Hope. How ‘spoofing’ traders dupe markets. *Wall Street Journal*, 2015.
- [Hope, 2015b] Bradley Hope. Was ‘John Doe’ manipulating Treasury futures? New lawsuit says yes. *Wall Street Journal MoneyBeat*, 2015.
- [Leal and Napoletano, to appear] Sandrine Jacob Leal and Mauro Napoletano. Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading. *Journal of Economic Behavior and Organization*, to appear.
- [Lee *et al.*, 2013] Eun Jung Lee, Kyong Shik Eom, and Kyung Suh Park. Microstructure-based manipulation: Strategic behavior and performance of spoofing traders. *Journal of Financial Markets*, 16(2):227–252, 2013.
- [Martínez-Miranda *et al.*, 2016] Enrique Martínez-Miranda, Peter McBurney, and Matthew Howard. Learning unfair trading: A market manipulation analysis from the reinforcement learning perspective. In *IEEE International Conference on Evolving and Adaptive Intelligent Systems*, pages 103–109, 2016.
- [Montgomery, 2016] John Montgomery. Spoofing, market manipulation, and the limit-order book. Technical report, Navigant Economics, 2016.
- [Prewit, 2012] Matt Prewit. High-frequency trading: Should regulators do more. *Michigan Telecommunications and Technology Law Review*, 19:131–161, 2012.
- [Price, 2018] Michelle Price. U.S. authorities charge three banks, eight individuals in futures ‘spoofing’ probe. *Reuters Business News*, 2018.
- [Vorobeychik *et al.*, 2006] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P. Wellman. Empirical mechanism design: Methods, with application to a supply-chain scenario. In *7th ACM Conference on Electronic Commerce*, pages 306–315, 2006.
- [Wah *et al.*, 2016] Elaine Wah, Sébastien Lahaie, and David M. Pennock. An empirical game-theoretic analysis of price discovery in prediction markets. In *26th International Joint Conference on Artificial Intelligence*, pages 510–516, 2016.
- [Wah *et al.*, 2017] Elaine Wah, Mason Wright, and Michael P. Wellman. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 59:613–650, 2017.
- [Wang and Wellman, 2017] Xintong Wang and Michael P. Wellman. Spoofing the limit order book: An agent-based model. In *16th International Conference on Autonomous Agents and Multiagent Systems*, pages 651–659, 2017.
- [Wang, 2015] Yun-Yi Wang. Strategic spoofing order trading by different types of investors in the futures markets. *Wall Street Journal*, 2015.
- [Wellman, 2016] Michael P. Wellman. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.
- [Wiedenbeck and Wellman, 2012] Bryce Wiedenbeck and Michael P. Wellman. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pages 931–938, 2012.