Spoofing and Manipulating Order Books with Learning Algorithms

Álvaro Cartea, Patrick Chang, Gabriel Garcia-Arenas

Oxford-Man Institute
University of Oxford

February 7, 2024

Overview

Motivation:

- There is growing concern about unintended behaviour when decision making is delegated to artificial intelligence algorithms, e.g., algorithmic collusion.
- Regulators and stakeholders (e.g., AFM and OECD) are concerned about algorithms learning to manipulate the market.
- Will algorithms learn to manipulate electronic markets?
- Can we determine when an algorithm will learn to manipulate the book?

Contribution

- Develop an inventory model of the limit order book.
- Derive conditions to test when an algorithm will learn to manipulate the book.
- Results apply to any (generic) learning algorithm.
- Manipulation in our model is unintentional, i.e., happens only when individual actions are sequenced together in a particular order.
- Market conditions in Nasdaq are conducive to algorithms learning to manipulate the book

Limit Order Book

- Limit orders are price-contingent orders to buy or sell an asset.
- Limit orders follow price-time priority, and collectively form the book.

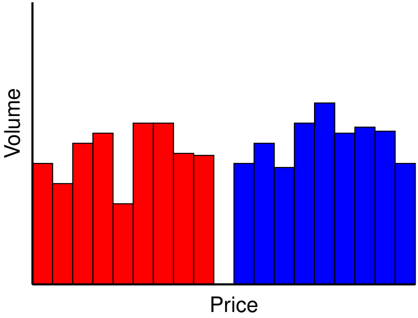

Figure: Limit order book.

Manipulation to sell the asset without crossing spread I

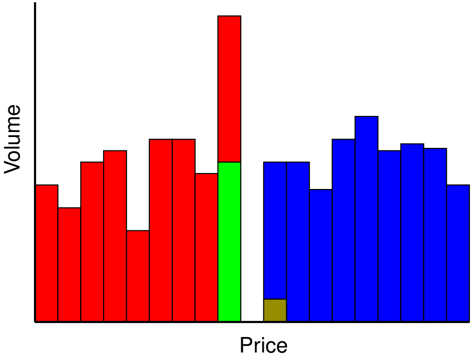

- Quote-based manipulation.
- Limit orders are submitted to both sides of the book — intention is to trade only on one side.

For example, if objective is to sell an asset, then

- submit a large buy limit order that will be cancelled, and
- submit a limit order on the ask that is intended to result in a transaction.

Manipulation to sell the asset without crossing spread II

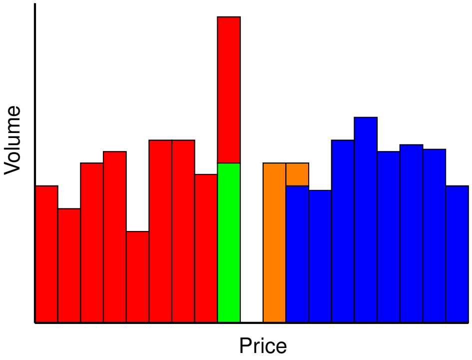

- Increase in buy-pressure is interpreted as an expected increase in the price.
- A buy-heavy book is followed by an increase in the arrival rate of buy market orders that **cross the spread** in anticipation of a price increase.
- These market orders lift the limit sell order that is intended to result in a transaction.

Manipulation to sell the asset without crossing spread III

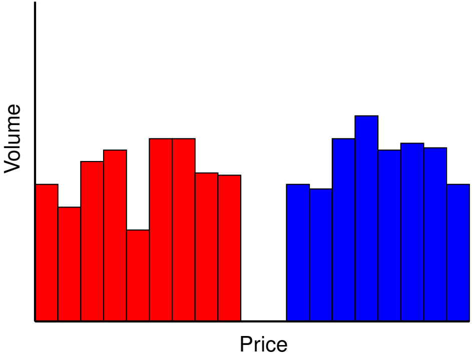

- The manipulative order, i.e., the large limit buy,
  - is cancelled or expires, or
  - is inadvertently filled
- Quote-based manipulation allows one to buy or sell an asset at a more favorable price than was otherwise likely to occur, i.e., not cross the spread.

Volume Imbalance

- Volume imbalance at time t

$$\omega_t = \frac{V_t^b - V_t^a}{V_t^b + V_t^a} \in (-1, 1) \tag{1}$$

  - V_t^b volume at the best bid t
  - V_t^a volume at the best ask
- The book is
  - sell-heavy when $\omega_t \in (-1, -1/3)$
  - neutral when $\omega_t \in [-1/3, 1/3]$
  - buy-heavy when $\omega_t \in (1/3, 1)$

Volume Imbalance — market order type

- More market buys when imbalance is buy-heavy, more market sells when imbalance is sell-heavy.

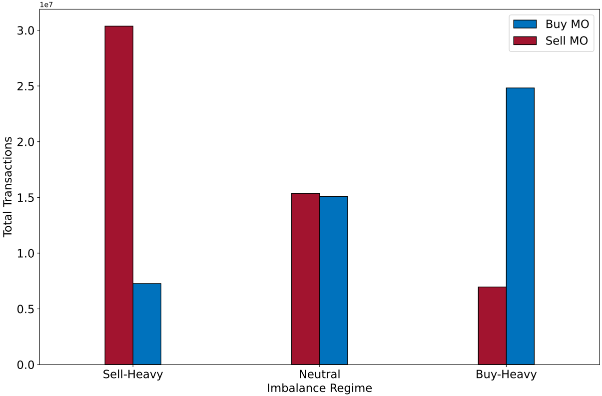

Figure: AAPL, April 2023, Nasdaq.

Volume Imbalance and Order Book Activity — arrival of market orders

- Arrival rate of buys is highest when buy-heavy
- Arrival rate of sells is highest when sell-heavy

Table: Arrival rates of market orders (MOs) for April 2023.

Ticker	Buy MO arrival rates (per second)			Sell MO arrival rates (per second)		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	0.060	0.176	**0.525**	**0.606**	0.179	0.058
AMZN	0.067	0.168	**0.447**	**0.456**	0.167	0.065
INTC	0.014	0.042	**0.138**	**0.139**	0.036	0.013

Volume Imbalance and Fill Probabilities

- Fill probabilities of
  - bids and offers are similar when the book is neutral
  - offers are higher when the book is buy-heavy
  - bids are higher when the book is sell-heavy

Table: Fill probabilities.

Ticker	Side	5 seconds			1 second			0.5 seconds		
		SH	*N*	*BH*	*SH*	*N*	*BH*	*SH*	*N*	*BH*
AAPL	Ask	0.4393	0.4782	**0.5819**	0.1048	0.1286	**0.1910**	0.0449	0.0579	**0.0928**
	Bid	**0.6210**	0.5207	0.4687	**0.2196**	0.1499	0.1180	**0.1121**	0.0697	0.0518
AMZN	Ask	0.4155	0.4651	**0.5669**	0.1008	0.1232	**0.1903**	0.0451	0.0566	**0.0933**
	Bid	**0.5587**	0.4570	0.4201	**0.1767**	0.1228	0.1044	**0.0863**	0.0566	0.0479
INTC	Ask	0.0970	0.1384	**0.2353**	0.0158	0.0222	**0.0561**	0.0071	0.0095	**0.0274**
	Bid	**0.2124**	0.1314	0.1116	**0.0501**	0.0211	0.0161	**0.0251**	0.0089	0.0070

Inventory Model

Basic setup:

- The expected bid-ask spread is $\vartheta > 0$.
- The market maker interacts with the order book at discrete times $t = 0, 1, 2, ..., +\infty$.
- Market maker delegates decision making to a learning algorithm.
- The algorithm has convergence guarantees.
- Midpoint of the bid-ask spread proxies the fundamental value of the asset Z.
- At each time point, the value of the asset either goes up by one-tick $(Z + \varphi)$, or goes down by one-tick $(Z - \varphi)$.

States and Actions

States $\mathcal{S} = \mathcal{Q} \times \Omega$:

- $\mathcal{Q} = \{-\bar{q}, ..., 0, ..., \bar{q}\}$ is the level of inventory.
- $\Omega = \{BH, N, SH\}$ is the three regimes of volume imbalance.
- $p^b_\omega \in (0, 1)$ and $p^a_\omega \in (0, 1)$ are the fill probabilities of a limit buy and a limit sell order being filled between $[t, t+1)$ in each regime $\omega \in \Omega$.

Actions at time t:

- Submit a buy limit order (*LB*) on the best bid or a sell limit order (*LS*) on the best offer
  - If order is not executed between $[t, t+1)$, then it is cancelled before start of $t+1$
  - LB and LS are for one unit of the asset
- Submit a large buy limit order (*LLB*) on the best bid or a large sell limit order (*LLS*) on the best offer and cancel order before start of $t+1$
- Submit a market order to buy (*MB*) or to sell (*MS*) one unit of the asset.
- Do nothing (*DN*).

Data and other assumptions

- Empirically we find
  - Average volume of limit orders and limit order cancellations are similar in size under each volume imbalance regime.
  - Arrival rates of limit orders are higher than the arrival rates of limit order cancellations.
  - More buy (sell) limit orders than sell (buy) limit orders when the book is buy-heavy (sell-heavy).
- Market participants cannot react instantaneously, so have a delay.
- We assume
  - $p(BH|\omega, LLB) = 1$ and $p(SH|\omega, LLS) = 1$ for all $\omega \in \Omega$.
  - changes in fill probabilities come into effect at time $t + 1$.

Thus, manipulation can occur even if the market maker does not explicitly encode the manipulation as a possible action into the learning algorithm.

Utility I

- The objective of the market maker is to maximize the present value of her wealth subject to a running inventory penalty.
- The total wealth $X + Z\,q$ of the market maker is the sum of her cash position X and the marked-to-market value of the inventory $Z\,q$.

One-step utility:

$$u(\mathbf{s}, a, \mathbf{s}') = Y(\mathbf{s}, a, \mathbf{s}') - \alpha\,(q')^2 \tag{2}$$

$\implies$ optimization problem of learning algorithm is

$$\sup_{\sigma \in \Sigma} \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \delta^t \left(Y(\mathbf{s}_t, a_t, \mathbf{s}_{t+1}) - \alpha\, q_{t+1}^2 \right) \middle| \mathbf{s}_0 = \mathbf{s} \right] \tag{3}$$

Optimal Strategy when $q > 0$

Let $q > 0$. Then, for each state $\mathbf{s} = (\omega, q)$, there exist cutoff values of the inventory aversion parameter $\alpha_0(\omega, q) < \alpha_1(\omega, q) < \alpha_2(\omega, q) < \alpha_3(\omega, q)$ such that the optimal stationary pure Markov strategy $\sigma^* \in \Sigma^{SPM}$ is given by

$$\sigma^*(\omega, q) = \begin{cases} LB & \text{if } \alpha \in \left(0, \alpha_0(\omega, q)\right), \\ LLB & \text{if } \alpha \in \left(\alpha_0(\omega, q), \alpha_1(\omega, q)\right), \\ LLS & \text{if } \alpha \in \left(\alpha_1(\omega, q), \alpha_2(\omega, q)\right), \\ LS & \text{if } \alpha \in \left(\alpha_2(\omega, q), \alpha_3(\omega, q)\right), \\ MS & \text{if } \alpha \in \left(\alpha_3(\omega, q), +\infty\right). \end{cases} \tag{4}$$

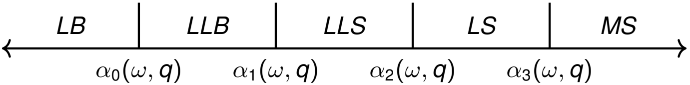

Figure: Optimal action choice for each state $\mathbf{s} = (\omega, q)$ for $q > 0$.

Manipulation I

Manipulation

- occurs if a **large** limit order is placed at time t on the side of the book that counters one's objective to buy or sell an asset, and the following action at time $t+1$ is to place a limit order on the side of the book that aligns with one's objective to buy or sell an asset.
- When $q > 0$, we want to revert to $q = 0$, so manipulation occurs if the sequence is initiated by *LLB* and followed by *LS* or *LLS*.

Manipulation II

A manipulative sequence when the manipulative order is not filled.

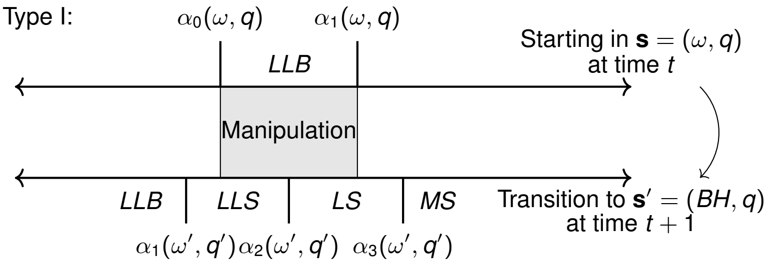

Shaded area is $I_1(\mathbf{s}) = \Big(\max\{\alpha_0(\omega, q), \alpha_1(BH, q)\}, \min\{\alpha_1(\omega, q), \alpha_3(BH, q)\}\Big) \neq \emptyset$.

Manipulation III

A manipulative sequence when the manipulative order is caught out.

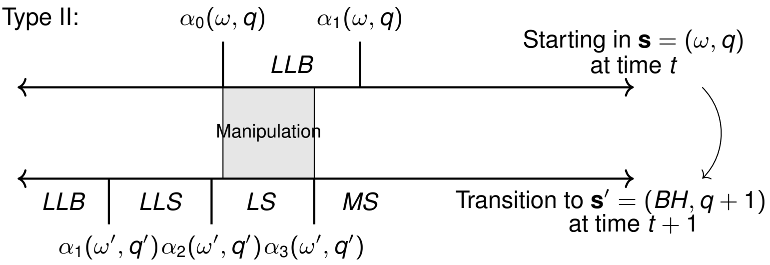

Shaded area is $I_2(\mathbf{s}) = \Big(\max\{\alpha_0(\omega, q+1), \alpha_1(BH, q)\}, \min\{\alpha_1(\omega, q), \alpha_3(BH, q+1)\}\Big) \neq \emptyset$.

Manipulation IV

Theorem

If the conditions

$$p^b_{BH} < p^a_{BH} \tag{C1}$$
$$p^a_{SH} < p^b_{SH} \tag{C2}$$

hold, then $I_1(\mathbf{s}) \neq \emptyset$ and $I_2(\mathbf{s}) \neq \emptyset$ for all $\mathbf{s} \in \mathcal{S}$ such that (i) $\mathbf{s} = (SH, q > 0)$, (ii) $\mathbf{s} = (BH, q < 0)$, and (iii) $\mathbf{s} = (N, q)$ for either $q > 0$ or $q < 0$.

In other words, there exist values of α such that the algorithm will learn to manipulate the book.

Manipulation V

Theorem

Let (C1) and (C2) hold, and let

$$
\begin{aligned}
p^a_{BH} - p^b_{SH} &< \min\left\{(p^b_{SH} - p^b_N)\frac{p_{N|BH}}{p_{BH|BH}},\ (p^b_{SH} - p^b_N)\frac{p_{N|N}}{p_{BH|N}}\right\} \\
p^b_{SH} - p^a_{BH} &< \min\left\{(p^a_{BH} - p^a_N)\frac{p_{N|SH}}{p_{SH|SH}},\ (p^a_{BH} - p^a_N)\frac{p_{N|N}}{p_{SH|N}}\right\}
\end{aligned} \tag{C3}
$$

hold.

1. If $(p^b_N - p^a_N) > \frac{\delta}{1+\delta}(p^b_{SH} - p^a_{BH})$ holds, then $I_1(\mathbf{s}) \neq \emptyset$ and $I_2(\mathbf{s}) \neq \emptyset$ for all states $\mathbf{s} = (N, q > 0)$.
2. If $(p^a_N - p^b_N) > \frac{\delta}{1+\delta}(p^a_{BH} - p^b_{SH})$ holds, then $I_1(\mathbf{s}) \neq \emptyset$ and $I_2(\mathbf{s}) \neq \emptyset$ for all states $\mathbf{s} = (N, q < 0)$.

Nasdaq

Table: Testable conditions.

Ticker	5 seconds			1 second			0.5 seconds		
	(C1), (C2)	(C3)	Side	(C1), (C2)	(C3)	Side	(C1), (C2)	(C3)	Side
AAPL	✓	✓	$q > 0$	✓	✓	$q > 0$	✓	✓	$q > 0$
AMZN	✓	✓	$q < 0$	✓	✓	$q > 0$	✓	✓	$q > 0$
INTC	✓	✓	$q > 0$	✓	✓	$q > 0$	✓	✓	$q > 0$

Manipulation and ϑ

Proposition

If $\vartheta \to 0$, then the algorithm will not learn to manipulate the order book for any state $\mathbf{s} = (\omega, q)$ where $q \neq 0$.

- If quoted spread is zero, manipulating the book does not provide any advantages over market orders (limit orders do not obtain better prices than market orders)
- manipulation is not optimal because inadvertent fills are penalised

Spoofing and Fill Preferences

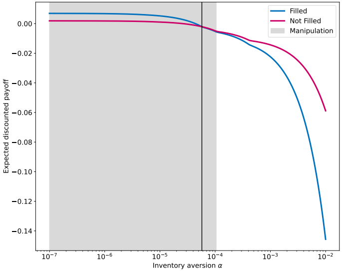

Figure: Fill preferences

In our model,
quote-based manipulation contains:

- Spoofing is the manipulative sequence + preference for the manipulative order not to get caught out.
- Manipulation for round-trip trade is the manipulative sequence + preference for the manipulative order to get filled.

Multiple Market Makers

- In the offline learning setting, the algorithms either coordinate (i.e., market makers ride the manipulative sequences of each other) or mis-coordinate (i.e., market makers send large opposing orders that cancel each other out) depending on their initial inventory.

- In the online learning setting, the algorithms learn to coordinate.
  - If the market makers start with zero inventory, then they coordinate by riding the sequences of each other.
  - If the market makers start with the same level of inventory or with opposing levels of inventory, then they coordinate by allowing one to ride the other's sequences to avoid their large limit orders cancelling each other out.

Legal Implications

EU:

- Article 12(2)(c) of Regulation (EU) No 596/2014 makes manipulation illegal.
- RTS 6 and 7 (part of MiFID II) require firms to test their trading algorithms so they do not behave in an unintended manner or contribute to disorderly trading conditions.

US (securities):

- Section 9(a)(2) of the Securities Exchange Act of 1934 makes manipulation illegal.
- FINRA's rule requires algorithmic trading developers to register as securities traders, and are therefore subject to the SEC and FINRA rules that govern their trading activities.

US (commodities):

- Dodd–Frank Act of 2010 defines spoofing as bidding or offering with the intent to cancel the bid or offer before execution. Spoofing is illegal under the Act.
- Our definition captures the spirit of DF but is broader in scope, e.g., expired orders.
- Narrow focus of DF misses other forms of quote-based manipulation.

References I

Amihud, Y. and Mendelson, H. (1986).
Asset pricing and the bid-ask spread.
Journal of Financial Economics, 17(2):223–249.

Cartea, A., Jaimungal, S., and Wang, Y. (2020).
Spoofing and Price Manipulation in Order-Driven Markets.
Applied Mathematical Finance, 27(1-2):67–98.

Fox, M. B., Glosten, L. R., and Guan, S. S. (2021).
Spoofing and its Regulation.
Colum. Bus. L. Rev., page 1244.

Glosten, L. R. (1994).
Is the Electronic Open Limit Order Book Inevitable?
The Journal of Finance, 49(4):1127–1161.

Harris, L. E. and Panchapagesan, V. (2005).
The Information Content of the Limit Order Book: Evidence from NYSE Specialist Trading Decisions.
Journal of Financial Markets, 8(1):25–67.

References II

Ho, T. and Stoll, H. R. (1981).
Optimal dealer pricing under transactions and return uncertainty.
Journal of Financial Economics, 9(1):47–73.

Stoll, H. R. (1978).
The Supply of Dealer Services in Securities Markets.
The Journal of Finance, 33(4):1133–1151.

Williams, B. and Skrzypacz, A. (2021).
Spoofing in Equilibrium.
Available at SSRN 3742327.

Volume of market orders

- Volume of buys is lowest when buy-heavy
- Volume of sells is lowest when sell-heavy

Table: Average volume of market orders (MOs) for April 2023.

Ticker	Buy MO average volume			Sell MO average volume		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	145.58	111.16	**62.27**	**64.29**	103.11	135.29
AMZN	205.95	108.59	**61.68**	**60.92**	107.78	181.50
INTC	212.79	256.93	**143.67**	**134.62**	266.80	227.45

Volume Imbalance and Order Book Activity — arrival of limit orders

Table: Arrival rates of limit orders (LOs) for April 2023.

Ticker	Buy LO arrival rates (per second)			Sell LO arrival rates (per second)		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	4.245	7.000	4.129	4.285	6.970	4.303
AMZN	4.367	7.346	4.381	4.637	7.600	4.213
INTC	1.090	2.386	1.829	1.686	2.330	1.106

Volume of limit orders

- Volume of buys is largest when buy-heavy
- Volume of sells is largest when sell-heavy

Table: Average volume of limit orders (LOs) for April 2023.

Ticker	Buy LO average volume			Sell LO average volume		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	98.19	109.51	**112.32**	**115.30**	109.50	97.40
AMZN	87.95	96.32	**101.95**	**101.53**	95.44	87.83
INTC	298.83	372.36	**415.56**	**415.82**	364.61	292.29

Volume Imbalance and Order Book Activity — arrival of cancellations

Table: Arrival rates of limit order cancellations for April 2023.

Ticker	Arrival rates of limit buy cancellation (per second)			Arrival rates of limit sell cancellation (per second)		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	3.092	6.037	3.785	3.946	6.154	3.334
AMZN	3.594	6.464	3.631	3.885	6.734	3.529
INTC	0.818	1.812	1.308	1.193	1.755	0.801

Volume of cancellations

- Volume of limit buy cancellations is largest when buy-heavy
- Volume of limit sell cancellations is largest when sell-heavy

Table: Average volume of limit order cancellations for April 2023.

Ticker	Average volume of limit buy cancellations			Average volume of limit sell cancellations		
	SH	*N*	*BH*	*SH*	*N*	*BH*
AAPL	92.25	109.29	**112.92**	**116.77**	111.55	91.26
AMZN	78.53	95.69	**106.34**	**103.48**	94.28	79.87
INTC	230.90	393.75	**499.81**	**489.77**	392.68	228.14

Expected one-step utility:

$$\bar{u}(\mathbf{s}, a) = \begin{cases} p_\omega^b\,\vartheta/2 - \alpha\,p_\omega^b\,(q+1)^2 - \alpha\,(1-p_\omega^b)\,q^2 & \text{for } a = \{LB, LLB\}\,, \\ p_\omega^a\,\vartheta/2 - \alpha\,p_\omega^a\,(q-1)^2 - \alpha\,(1-p_\omega^a)\,q^2 & \text{for } a = \{LS, LLS\}\,, \\ -\vartheta/2 - \alpha\,(q+1)^2 & \text{for } a = MB\,, \\ -\vartheta/2 - \alpha\,(q-1)^2 & \text{for } a = MS\,, \\ -\alpha\,q^2 & \text{for } a = DN\,. \end{cases} \tag{5}$$

Behaviour from (2) is consistent with inventory models:

- Behaviour depends on the level of inventory, and there is a preferred inventory position, e.g., Amihud and Mendelson (1986).
- Prefer to sell if inventory is long and prefer to buy if inventory is short, e.g., Stoll (1978) and Ho and Stoll (1981).

Non-Deterministic Transitions

Theorem

Let (C1) and (C2) hold. If the transition probabilities associated with large limit orders are such that

$$
\begin{aligned}
& p(BH \mid \omega, LLB) = 1 - \kappa > p_{BH|\omega}\,,\ p(N \mid \omega, LLB) = \frac{\kappa}{2} < p_{N|\omega}\,,\ p(SH \mid \omega, LLB) = \frac{\kappa}{2} < p_{SH|\omega}\,, \\
& p(SH \mid \omega, LLS) = 1 - \kappa > p_{SH|\omega}\,,\ p(N \mid \omega, LLS) = \frac{\kappa}{2} < p_{N|\omega}\,,\ p(BH \mid \omega, LLS) = \frac{\kappa}{2} < p_{BH|\omega}\,,
\end{aligned} \tag{C4}
$$

hold for all $\omega \in \Omega$. Then $I_1(\mathbf{s}) \neq \emptyset$ and $I_2(\mathbf{s}) \neq \emptyset$ for all $\mathbf{s} \in \mathcal{S}$ such that (i) $\mathbf{s} = (SH, q > 0)$, (ii) $\mathbf{s} = (BH, q < 0)$, and (iii) $\mathbf{s} = (N, q)$ for either $q > 0$ or $q < 0$.

Multiple Market Makers I

Table: AMZN: Average number of manipulative sequences over 50 trading intervals. Agent 1 uses $\alpha = 10^{-4}$, agent 2 uses $\alpha = 10^{-5}$.

Setup	Decision Interval Δt	Zero inventory		Same inventory		Opposing inventory	
		Agent 1 $q = 0$	Agent 2 $q = 0$	Agent 1 $q = 4$	Agent 2 $q = 4$	Agent 1 $q = 4$	Agent 2 $q = -4$
Baseline	5 seconds	24.87	20.87	20.79	25.93	21.97	22.11
	1 second	25.22	14.92	14.78	29.25	18.52	18.77
	0.5 seconds	27.03	14.51	14.52	32.42	17.37	14.45
Offline	5 seconds	24.92	26.29	21.01	22.65	20.85	22.52
	1 second	27.01	29.62	17.12	19.02	17.46	19.40
	0.5 seconds	30.71	32.76	16.20	18.32	22.05	18.27
Online	5 seconds	24.40	25.89	20.47	22.12	20.41	22.04
	1 second	22.49	29.16	12.69	19.26	11.98	18.16
	0.5 seconds	21.27	32.13	1.21	15.20	1.12	14.43

Multiple Market Makers II

Table: AMZN: Average manipulation statistics.

Setup	Δt	Mismatching manipulative orders			Single manipulative order		
		Zero inv.	Same inv.	Opposing inv.	Zero inv.	Same inv.	Opposing inv.
Offline	5s	0.1554%	0.2388%	0.4408%	13.46	18.42	18.75
	1s	0.1215%	1.3516%	0.0054%	22.47	22.01	29.52
	0.5s	0%	0%	0%	19.07	18.71	34.65
Online	5s	0.2256%	0.3738%	0.4949%	19.39	21.58	21.92
	1s	0.4190%	1.8937%	2.3348%	24.25	26.66	23.93
	0.5s	0.7635%	0%	0%	25.25	25.96	25.69