Information-based models in centralised and decentralised exchanges: spoofing and private order flow

Gabriel García Arenas
St Hugh's College
University of Oxford

A thesis submitted for the degree of
DPhil in Mathematics

Trinity 2025

Acknowledgements

I am deeply grateful to my advisor, Álvaro Cartea, for all his help and generosity. Álvaro has always offered me outstanding advice. I have learnt a lot, both professionally and personally, from all the inspiring conversations we have had throughout these years of challenging studies.
I also strongly thank my College advisor Tom Sanders for all his good advices over these years.
I would like to thank Peter Howell for his support as Director of Graduate Studies in the Mathematical Institute.
I would like to thank Elisa Alòs for all the encouragement, support and advice given in multiple moments of the DPhil.
I would like to thank the examiners of my viva Leandro Sánchez-Betancourt and Basil Williams for their helpful comments.
A lot of thanks also to Nathanël Berestycki, Miquel Bosch, María Jesús Carro, Laura Costa, Mihai Cucuringu, Ernest Fontich, Àngel Jorba, Luis Ortiz, and José Penalva for their help and advice. I would also like to thank Ramón Caminal and Jordi Galí for their advice.
I would like to thank my colleagues and staff at the Oxford-Man Institute of Quantitative Finance and the Mathematical Institute. In particular, to my coauthors Patrick Chang and Fayçal Drissi, I have enjoyed collaborating with them. Also to my colleagues Álvaro Arroyo and Marcello Monga. To Jen Desmond and Vanessa Wilkins for her work at the Oxford-Man Institute of Quantitative Finance. And to Sandhya Patel for her work in the Academic Administration team at the Mathematical Institute.
I have also benefited greatly from the nice environment that St Hugh's College provided for my studies.
I would like to thank Ramón Areces Foundation for providing partial funding for this DPhil thesis.
Finally, I would like to thank my parents and my brother for always supporting me.

Abstract

In this thesis we focus on information-based models in both traditional exchanges (limit order books) and decentralised ones.

We propose a dynamic model of the limit order book to derive conditions to test if a trading algorithm learns to manipulate the order book. Our results show that as a market maker becomes more tolerant to bearing inventory risk, the learning algorithm will find optimal strategies that manipulate the book more frequently. Manipulation helps to revert inventory to an optimal level and to execute round-trip trades with limit orders at a higher probability than was otherwise likely to occur. Spoofing is a special case of quote-based manipulation where the market maker prefers that the manipulative limit orders are not filled. We use high-frequency data to check our conditions and show that algorithms will learn to manipulate Nasdaq's limit order book. Finally, we extend our model in several directions and we see manipulation can still arise in all of them. In particular, when two market makers use learning algorithms to trade, their algorithms can learn to coordinate their manipulation.

We study how the design of blockchains shapes the interaction between traders in decentralised exchanges (DEXs) and participants in the blockchain security protocol. On blockchains such as Ethereum, traders can route their transactions through either a public memory pool, where transactions are visible and subject to attacks, or private memory pools, where transactions are hidden but can only be executed by specific builders. These features create a fundamental trade-off between execution certainty and protection from attacks. We develop two complementary models to analyse this trade-off. The first model studies the effect of competition among builders (MEV-Boost auction). We show that DEX liquidity depth

governs equilibrium fragmentation: when liquidity is high, traders concentrate in one private pool, whereas scarce liquidity leads to balanced order flow across pools. The second model examines the role of builder composition-specifically, the share of builders able to execute private orders. Here, traders' behaviour and equilibrium fragmentation depend on this composition, with higher shares of public-only builders inducing a shift toward public execution. Together, the models explain how liquidity conditions and builder structure jointly order flow fragmentation.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

This thesis analyzes some information-based models in financial settings involving either traditional centralised exchanges or emerging decentralised ones. In particular, our work focuses on limit order books and memory pools.

Concerns regarding the possibility that algorithms could manipulate financial markets have been manifested by regulators such as the Dutch Authority of Financial Markets, who have the aim of guaranteeing financial markets properly work without any sort of exploitation. One possible manipulation is spoofing, which deliberately aims to provide misleading information to other market participants so as to increase the probability of getting higher profits. This is done by submitting a large amount of limit orders to one side of the limit order book with the intention to cancel them before execution. We devote Chapters 2 and 3 to its study.

Manipulation does not just appear in traditional exchanges such as limit order books, but also on the more modern decentralised ones. Decentralised exchanges (DEXs) exist due to the desire to facilitate transactions avoiding intermediaries such as banks, and have transparency as a relevant feature. However, this attribute enables some market participants to extract rent from others using a variety of manipulative strategies. To fight against this possibility, the appearance of private memory pools has emerged as a potential solution, since it prevents some of these threats from happening. However, they also present some disadvantages fundamentally driven by its lack of transparency, with could contradict to some extent the intrinsic nature and reason of existence of DEXs. In Chapter 4 we study whether market participants can prefer them over public memory pools, and how they behave in equilibrium if there are private pools.

1.2 Contribution

The role of information is key in our spoofing model, where learning algorithms may try to misguide other market participants by conveying misleading information. This can be done by sending a large amount of limit orders and then canceling some of them. In our project of public and private order flow, information also has a prominent role since traders have their own random liquidity needs, which are private information, and take their decisions based on them. Information-based models are standard in the microstructure literature. We focus on models with agents that have some sort of private information (personal inventory level, personal liquidity needs).
In particular, in Chapter 2 we propose a dynamic model of the limit order book to derive conditions to test if a trading algorithm learns to manipulate the order book. We see manipulation can appear under some conditions on both the fill probabilities of execution of limit orders and the transition probabilities between different inventory levels and book regimes.
In Chapter 3, we empirically test the conditions with order book data from Nasdaq. We also study whether manipulation arises under some other scenarios, and we see it can appear in all of them. We consider a finite-horizon setting, the existence of a trend in the fundamental value of the asset, we allow for multiple prices, we allow for multiple filled orders and we also allow for multiple market makers.
In Chapter 4 we study what conditions induce market participants to find optimal to submit their orders to a private memory pool. We consider in particular a model with M traders where each of them decides whether to submit transactions to the public memory pool or to one of two available private memory pools. We see that the split of orders submitted between the two private memory pools in equilibrium ultimately depends on a parameter related to liquidity depth. We also consider a scenario with two traders and one representative private memory pool, where there are multiple builders who are randomly selected to create the next block, and only a fraction $1 - \alpha$ of them have access to the private pool. We see that the value of α affects traders' behavior.

Chapter 2

A Model of Spoofing and Manipulating Order Books with Learning Algorithms

2.1 Introduction

There is growing concern that unintended behavior may arise when decision making is delegated to artificial intelligence algorithms. Recently, the OECD and the Dutch Authority of Financial Markets (AFM) expressed concerns about algorithms learning to manipulate financial markets (see [64],[2]).[1] In this chapter, we derive conditions to test if an algorithm will learn to manipulate the market through manipulative quote-based strategies such as spoofing.

A manipulative quote-based strategy consists of submitting limit orders to both sides of the order book when the objective is either to buy or to sell an asset. If the objective is to buy an asset, the strategy submits a large sell limit order that will be cancelled, and posts a limit order on the bid which is the one intended to result in a transaction. The large ask order is a manipulative order that tilts the order book and creates misleading information about the sell pressure of the asset. Market participants interpret the increase in sell pressure as an expected drop in the price of an asset, so a sell-heavy tilt in the book is followed by an increase in the arrival rate of sell orders that cross the spread in anticipation of a price drop. With the increase in the number of liquidity taking orders, the probability of buying the asset with a

[1] More generally, regulatory bodies around the world are concerned about market manipulation with trading algorithms, and they introduced legislation to address this concern. In the EU, RTS 6 and 7 require firms to test their trading algorithms so they do not behave in an unintended manner or contribute to disorderly trading conditions. In the US, the SEC approved FINRA's rule that requires algorithmic trading developers to register as securities traders, and are therefore subject to the SEC and FINRA rules that govern their trading activities.

limit order is higher than was otherwise likely to occur because market participants will trade on the misleading signal. Similarly, if the objective is to sell an asset, then a manipulative order on the bid creates buy pressure that market participants interpret as an expected increase in the price of an asset, which allows one to sell an asset with a limit order at a higher probability than was otherwise likely to occur.

Quote-based manipulation relies on the change in behavior elicited by a manipulative order, and this change in behavior can be explained with the asymmetric information model of [46] and a non-zero tick size in the order book. The step-function theory of [44] explains that market participants interpret asymmetry in the volumes posted on the best bid and the best ask as good or bad news about the asset. Specifically, when a sell limit order for a large number of shares arrives at a price equal to the existing best offer and there is no increase in the bids at the best bid price, market participants tend to react as if bad news arrived about the asset. Similarly, upon the arrival of a bid for a large number of shares at a price equal to the best bid and there is no increase in the orders at the best offer price, market participants react as if good news arrived about the asset. Therefore, when there is an imbalance between the liquidity posted at the best bid and the best ask quotes, market participants tend to interpret this as a signal to trade in a particular direction, buy or sell, in anticipation of a change in the price of the asset.

We summarize the volume imbalance between limit orders resting on the bid and on the ask sides of the book as buy-heavy, sell-heavy, and neutral. The rates with which market orders, limit orders, and cancellations arrive at the market depend on the tilt of the book; thus, the probability with which limit orders are executed depends on the volume imbalance of the book. Specifically, our empirical results with data from Nasdaq show that the fill probability of a sell limit order is highest (lowest) when the book is buy-heavy (sell-heavy), and the fill probability of a buy limit is highest (lowest) when the book is sell-heavy (buy-heavy). Quote-based manipulation is profitable because traders can manipulate the tilt of the book to buy or to sell an asset with a limit order at a higher probability than was otherwise likely to occur.

To analyze if algorithms can learn to manipulate the book, we develop a dynamic model where the market maker interacts with the limit order book at discrete time intervals for an infinite trading horizon.[2] The market maker is non-myopic and is averse to holding high levels of inventory. Specifically, her objective (i.e., optimality criterion) is to maximize the present value of her expected wealth, while penalizing

[2]We focus on an infinite trading horizon because most learning algorithms are designed for this setting.

exposure to inventory risk. The market maker provides liquidity at the best bid and the best ask prices, and she delegates decision making to a learning algorithm to find an optimal trading strategy.[3] As with most learning algorithms, the market maker's algorithm learns a stationary Markov strategy.[4] Here, the Markov strategy depends on her level of inventory and the state of the limit order book, which is given by its volume imbalance (i.e., tilt of the book). To understand unintended behavior that may emerge, we do not focus on the behavior of a particular learning algorithm. Instead, we analyze the decision framework of learning algorithms, so our results and testable conditions apply to any learning algorithm that finds an optimal stationary Markov strategy.

In our analysis, the market maker does not endow the algorithm with an action that manipulates the order book. Instead, we focus on how an innocuous set of actions leads to manipulation when individual actions are sequenced in a particular order. Unintentional manipulation emerges because the learning algorithm dynamically maximizes the market maker's optimality criterion. Indeed, manipulation in our setting is unintentional, but it is the best course of action when the algorithm learns the optimal strategy. This is different from unintended behavior that arises when an algorithm fails to optimize the optimality criterion. In such cases, the unintended behavior differs on a case-by-case basis and depends on the idiosyncratic assumptions of the learning algorithm.

In our model, manipulation occurs when a large limit order is placed at time t on the side of the book that counters one's objective to buy or to 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 to sell an asset. To derive conditions to test if an algorithm will learn to manipulate the order book, we characterize the optimal stationary Markov strategy as a function of the value of the market maker's inventory aversion parameter for each state of the Markov strategy, i.e., for each pair of inventory level and volume imbalance regime.[5] The optimal strategy manipulates the book when the optimal action in the current state (i.e., inventory and volume imbalance pair) is a manipulative order (i.e., a large limit order in the "wrong" direction that will be cancelled), and the subsequent state prescribes an optimal action of placing

[3]There are several reasons why a market maker would delegate decision making to an algorithm. For example, the rise of high-frequency trading means that delegating decision making to an algorithm is necessary for a market maker to remain competitive.

[4]See [68], [73], and [72] for examples of generic learning algorithms. See also [16, 17], and [1] for examples of learning algorithms that have been studied in the context of algorithmic collusion.

[5]Our characterization follows a similar spirit to the characterization of the optimal order choice in [67].

a limit order on the side of the book that is intended to result in a transaction to complete the manipulative sequence.

Our main result provides sufficient conditions on the limit order book to test if an algorithm can learn to manipulate the order book. If these conditions hold and a trader affects the probability of tilting the book with manipulative orders, then the algorithm will learn to manipulate the book. In particular, we show that as the market maker becomes more tolerant to bearing inventory risk, the learning algorithm is more likely to learn manipulative strategies. The conditions depend only on the parameters of the model, and are applicable to any limit order book, e.g., Euronext, LSE, Nasdaq, NYSE. Our results show that market conditions in Nasdaq are conducive for algorithms to learn optimal strategies that manipulate the order book. In all the stocks we consider, we find that an algorithm will always learn to manipulate the order book for a range of values of the inventory aversion parameter.

One of the consequences associated with quote-based manipulation is that the manipulative order can get "caught out", i.e., the manipulative order inadvertently leads to a transaction. Our model and the learning algorithms account for this possibility. The market maker's decision to manipulate the order book balances the tradeoff between the probability of a fill of the manipulative order, the increase in inventory risk, and profits from round-trip trades due to the manipulative order. Often, when making markets, inventory levels deviate from the preferred inventory position.[6] In our model, the longer and further one deviates away from the preferred inventory position, the more severe is the penalty arising from inventory risk aversion, so the optimal strategy increases the pressure to ensure mean reversion to the preferred level of inventory. Thus, the market maker's strategy balances the trade-off between (i) buying or selling an asset to revert to the preferred inventory position at standard fill probabilities (i.e., without manipulating the book), or (ii) posting a manipulative order to manipulate the fill probabilities through the tilt of the book, which exposes the strategy to deviate further from the preferred inventory position (at least temporarily), but it also exposes the strategy to a round-trip trade. With this trade-off in mind, it is clear that a manipulative strategy becomes dynamically optimal and a more frequent optimal strategy when the market maker is less averse to holding high levels of inventory.

Counter-intuitive to the goal of quote-based manipulation, it is not always bad for the market maker to receive a fill on her manipulative order. Indeed, there are

[6]In our model, the preferred inventory position is zero when the fundamental value of the asset is a martingale.

situations in which the preference is for the manipulative order to be filled with the expectation to unwind the acquired position quicker than otherwise because the strategy tilted the book. In particular, we analyze if the market maker prefers that her manipulative order is filled and we find two driving forces behind the manipulation. One, the manipulation is optimal because it can lead to a manipulative round-trip trade that, in expectation, will be completed faster than otherwise. In these cases, the manipulative order is submitted with the preference for it to be filled and to unwind it immediately with an increased probability due to the tilt in the book caused by the manipulative order. Two, the manipulation is optimal because it increases the chances that the market maker's inventory will revert to a preferred position. In these cases, the preference is that the manipulative order is not inadvertently filled, which is commonly understood as spoofing. That is, the quote-based manipulation we study includes (i) spoofing as a refinement when the preference is that the manipulative order is not inadvertently filled, and (ii) manipulation for a round-trip trade as a refinement when the preference is that the manipulative order is filled. Our project is the first to show that not all quote-based manipulative orders are spoof orders.

Additionally, our model shows that as the quoted spread narrows, learning algorithms will be less likely to manipulate the order book. Specifically, as the quoted spread decreases, the range of values of the inventory aversion parameter where manipulation is optimal decreases. In the limit, when the quoted spread is zero, an algorithm will not learn to manipulate the book because manipulation is suboptimal. Of course, theory shows that the quoted spread is positive even if the tick size is zero.[7] Nonetheless, the insight is that (i) if the profits from using limit orders are negligible and the costs from using market orders are negligible, then it is more efficient to use market orders to revert to the preferred inventory position, and (ii) if the expected profit from the opportunistic round-trip trade, where one leg is a manipulative order, does not outweigh the penalty imposed to manage the inventory risk, then quote-based manipulation is not optimal.

In the literature, traders can attempt to manipulate the market in several ways. Studies focus on information-based manipulation and trade-based manipulation.
Information-based manipulation occurs when the manipulator releases misleading information (see for example [13], [74], [75]), whereas trade-based manipulation occurs when the manipulator buys or sells an asset to effect changes in the price (see for example [3], [4], [34], [35]).

[7]See for example [71], [54], [38], [47], [5], and [46].

Similar to the approach in [54], we propose a model of the market dynamics that is consistent with empirical stylized facts, where the market dynamics do not necessarily derive from principles of individual economic behavior. The purpose of our project is not to explain the underlying economic reasoning of the market dynamics, but to use the market dynamics to derive conditions to test if an algorithm will learn to to manipulate the market.

We focus on the broader quote-based manipulation because existing legislation does not explicitly outlaw spoofing.[8] Spoofing is illegal because the manipulative order (also known as a spoof order) creates misleading information to buy or to sell an asset with a higher probability than was otherwise likely to occur. The manipulative order is what is considered illegal under existing legislation, and the reason why spoofing is illegal. Thus, our results allow us to understand and test for quote-based manipulation, which include spoofing as a particular case. Moreover, our model allows us to make the key distinction between spoofing and other forms of quote-based manipulation based on one's preference of a fill of a manipulative order.

Our results have implications for how a rational market maker should behave within the market dynamics of our model. Indeed, our work is a discrete-time analogue to papers that use stochastic optimal control and continuous-time models to derive algorithmic trading strategies. For example, [33] derive a manipulative quote-based strategy to acquire or liquidate a large position. A key difference is that they explicitly encode a manipulative action, whereas manipulation is optimal but unintentional in our model because the market maker does not endow the algorithm with an action that manipulates the book. Also, their paper is not about learning algorithms.

Most of the literature on spoofing focuses on empirical studies. One exception is [78] who extend the setup of [47] to show that spoofing can occur and sustained in equilibrium. Their model predicts that spoofing slows price discovery, increases bid-ask spreads, and increases the volatility of prices. In our model, manipulative orders hamper price discovery because they erode the information conveyed by volume imbalance.

On the empirical side, [15] study Canadian equity markets. They show that spoofing slows price discovery, and that it increases both the volatility of prices and adverse selection, as predicted in [78], and do not find strong indication that spoofing leads to higher bid-ask spreads.

[8] The exception is the [39] Act, but the Act only applies to the US commodities market.

[58] use a proprietary dataset with trader identification from the Korea Exchange to show that spoofing achieves substantial extra profits and spoofing tends to target stocks with higher return volatility, lower market capitalization, lower price level, and lower managerial transparency. [76] uses data from the Taiwan Futures Exchange to show that market participants spoof the order book in stocks that exhibit high volumes of trading, high volatility, and high prices. [76] also shows that spoofing increases the volume of trading, increases the volatility of prices, and increases the quoted spread. Our empirical results complement their findings because we find that market conditions from Nasdaq are such that algorithms will learn to manipulate the order book. [45] use machine learning to identify and predict market conditions associated with spoofing events in the Moscow Stock Exchange.

Finally, our work is closer to the literature that studies the unintentional effects of algorithms that learn to collude (see for example [16], [17]).[9] Our approach is similar to that in [26] who prove that algorithms can learn to collude. They analyze the equilibria that can be learned and prove convergence to collusive equilibria; whereas in this chapter, we analyze the decision framework where algorithms learn optimal strategies and derive testable conditions to determine if algorithms will learn to manipulate the order book.

The remainder of the chapter proceeds as follows. The next section shows the relationship between volume imbalance and the behavior of market participants. Section 2.3 presents our model and Section 2.4 derives the optimal strategy and testable conditions to determine if a manipulative strategy can be learned. Section 2.5 analyzes the mechanics of quote-based manipulation, its relation to spoofing, and how the parameters of the model affect quote-based manipulation as the optimal strategy. Section 2.6 discusses some regulatory implications. Section 2.7 includes all the proofs required for our results.

2.2 Volume Imbalance and Order Book Activity

This section uses data from Nasdaq for April 2023 to illustrate the relationship between volume imbalance and the activity in the limit order book that is central to quote-based manipulation. For each trading day, we remove the first and last 15 minutes to exclude behavior in the limit order book during the opening and closing auctions.

[9] See also [24, 25], [36], and [40] for studies on algorithmic collusion in financial markets.

Volume imbalance at time t is given by

$$\omega_t = \frac{V_t^b - V_t^a}{V_t^b + V_t^a} \in (-1, 1)\,, \tag{2.1}$$

where $V_t^b, V_t^a > 0$ are the liquidity posted at the best bid and the best ask, respectively, at time t. We follow the definition given in [33] for volume imbalance. Although the choice of only considering volumes at the best bid and best ask is arbitrary, if we considered other possibilities (for example, considering the sum of liquidities posted at the best and second best bid and the sum of liquidities posted at the best and second best ask), we would expect the order of the results in Tables 2.1-2.7 to be the same. For example, for Table 2.1, although we would expect the numbers to change, their relation would still hold in the sense that Buy MO arrival rates (per second) would be highest at the BH regime, and Sell MO arrival rates (per second) would be highest at the SH regime. We would also expect the relations between different book regimes in the rest of tables to still hold. Volume imbalance summarizes the tilt of the limit order book, so when ω_t is close to 1 there is a strong buy pressure and when ω_t is close to -1 there is a strong sell pressure. To simplify our subsequent analysis, we discretise volume imbalance into three regimes: buy-heavy (BH) when $\omega_t \in (1/3, 1)$, neutral (N) when $\omega_t \in [-1/3, 1/3]$, and sell-heavy (SH) when $\omega_t \in (-1, -1/3)$.

Table 2.1: 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
CSCO	0.008	0.025	**0.101**	**0.104**	0.033	0.012
INTC	0.014	0.042	**0.138**	**0.139**	0.036	0.013
MSFT	0.297	0.350	**0.461**	**0.475**	0.352	0.286
TSLA	0.532	0.677	**0.773**	**0.750**	0.635	0.529

Table 2.1 presents the arrival rates (per second) of buy and sell market orders in each volume imbalance regime for the several assets, and Table 2.2 presents the average volume of the market orders that arrive. The arrival rate of buy (sell) market orders is highest when the book is buy-heavy (sell-heavy). However, the average volume of buy (sell) market orders is lowest when the book is buy-heavy (sell-heavy). Nevertheless, the net effect (i.e., arrival rate times average volume) is that there are more buy (sell) transactions when the book is buy-heavy (sell-heavy).

Table 2.2: 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
CSCO	149.88	182.33	**110.30**	**111.74**	161.72	131.67
INTC	212.79	256.93	**143.67**	**134.62**	266.80	227.45
MSFT	72.17	46.41	**20.66**	**21.25**	45.57	70.14
TSLA	132.27	56.33	**26.55**	**25.35**	54.34	120.54

Table 2.3: 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
CSCO	0.567	1.345	0.951	0.882	1.416	0.641
INTC	1.090	2.386	1.829	1.686	2.330	1.106
MSFT	3.631	4.126	4.452	4.830	4.268	3.582
TSLA	3.913	4.628	3.309	3.560	4.784	4.088

Table 2.4: 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
CSCO	228.77	271.07	**307.59**	**319.84**	267.68	234.34
INTC	298.83	372.36	**415.56**	**415.82**	364.61	292.29
MSFT	42.93	44.81	**45.69**	45.25	**45.73**	44.76
TSLA	49.64	55.64	**60.97**	**66.90**	61.38	57.94

Similarly for limit orders, Table 2.3 presents the arrival rates (per second) of buy and sell limit orders in each volume imbalance regime for the assets we consider, and Table 2.4 presents the average volume of the limit orders that arrive. The arrival rates of buy limit orders are higher than those of sell limit orders when the book is buy-heavy, and the arrival rates of sell limit orders are higher than those of buy limit orders when the book is sell-heavy. On the other hand, the average volume of buy (sell) limit orders is largest when the book is buy-heavy (sell-heavy). The net effect (i.e., arrival rate times average volume) is that there are more buy (sell) limit orders than sell (buy) limit orders when the book is buy-heavy (sell-heavy).

Similarly, Table 2.5 presents the arrival rates (per second) of limit buy and limit

Table 2.5: 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
CSCO	0.398	1.021	0.701	0.682	1.054	0.452
INTC	0.818	1.812	1.308	1.193	1.755	0.801
MSFT	3.673	3.561	3.082	3.316	3.756	3.719
TSLA	2.641	3.475	2.824	3.066	3.607	2.777

Table 2.6: 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
CSCO	172.01	281.82	**370.26**	**378.71**	281.96	179.55
INTC	230.90	393.75	**499.81**	**489.77**	392.68	228.14
MSFT	23.17	41.34	**63.64**	**63.67**	41.85	23.55
TSLA	29.17	50.51	**71.24**	**82.05**	54.89	35.21

sell cancellations in each volume imbalance regime for the three assets we consider, and Table 2.6 presents the average volume of limit order cancellations that arrive.

Volume imbalance clearly influences the behavior of market participants. This is consistent with the work of [51], who find that market participants condition their quotation behavior on volume imbalance. These empirical findings are further supported by a survey sent to Dutch algorithmic trading firms, where AFM found that trading algorithms use between 100 and 1000 features, and volume imbalance is one of the key features (see [2]).

Finally, Table 2.7 reports the probability that a limit order on the best bid or best ask is filled within the next five seconds, one second, and half of a second, respectively. We could have considered other time horizons, in particular shorter ones. For example, the roundtrip latency of the HFTs in these LOBs would be a typical horizon to study. However, we are interested in the comparison between the fill probabilities for the different book regimes ω, so that we would expect to get the same relative orders for other time horizons. These fill probabilities account for the effect of time-priority. The details about the estimation procedure are given in the next chapter, in Section 3.2. This link between the fill probability and volume imbalance regimes makes quote-based manipulation viable and profitable. By manipulating the

Table 2.7: 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
CSCO	Ask	0.0674	0.1005	**0.1768**	0.0093	0.0151	**0.0384**	0.0040	0.0060	**0.0198**
	Bid	**0.1851**	0.1034	0.0713	**0.0400**	0.0154	0.0109	**0.0201**	0.0064	0.0048
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
MSFT	Ask	0.5577	0.5838	**0.6286**	0.2033	0.2135	**0.2529**	0.1037	0.1095	**0.1333**
	Bid	**0.6101**	0.5739	0.5605	**0.2339**	0.2081	0.2041	**0.122**	0.1064	0.1048
TLSA	Ask	0.6051	0.6280	**0.6729**	0.2618	0.2751	**0.3222**	0.1459	0.1517	**0.1876**
	Bid	**0.668**	0.6201	0.6047	**0.3137**	0.2660	0.2661	**0.1820**	0.1482	0.1480

volume imbalance and tilting the order book, the probability of buying or selling the asset with a limit order is higher than it was otherwise likely to occur.

2.3 The Model

Here, we present our dynamic model of the limit order book where the market maker provides liquidity at the best bid and best ask. The market maker is non-myopic and interacts with the limit order book at discrete times $t = 0, 1, 2, \cdots, +\infty$. The market maker delegates the decision making process to a learning algorithm that can be trained online or offline.

2.3.1 Setup

Framework We model the decision process of the market maker as a Markov decision process $\mathcal{M} = \langle S, (A_s)_{s\in S}, p, (u_s)_{s\in S}, \delta\rangle$. Let $s \in S$ denote the state, where the set S is finite, and let A_s denote the finite set of actions for the market maker in state s. The state evolves according to the transition function $p : S \times A_s \to \Delta(S)$, where $\Delta(S)$ is the set of probability measures on S. We denote by $p(s'|s, a)$ the probability that the subsequent state is s' given that the current state is s and action a is played. At every time step t, the payoff is given by a utility function $u : S \times A_s \times S \to R$, where $|u(s, a, s')| < \infty$ for all $a \in A_s$ and $s, s' \in S$. The payoff from the utility function $u(s, a, s')$ depends on the transition from state s to state s' under action a. Finally, $\delta \in [0, 1)$ is the parameter with which the market maker discounts the future stream of payoffs.

Strategies The market maker uses an algorithm to learn a time-invariant strategy that depends only on the state s, i.e., we consider stationary Markov strategies. A stationary Markov strategy describes a mapping to a set of probability measures on A_s for each state s, i.e., $\sigma \in \Sigma^{SM} = \prod_{s\in S} \Delta(A_s)$ such that $\sigma : S \to \Delta(A_s)$. Similarly, a stationary pure Markov strategy describes a mapping to an action A_s for each state s, i.e., $\sigma \in \Sigma^{SPM} = \prod_{s\in S} A_s$ such that $\sigma : S \to A_s$. More generally, a strategy is a mapping from the set of all possible histories to a set of probability measures on A_s, i.e., $\sigma \in \Sigma$ such that $\sigma : \mathcal{H} \to \Delta(A_s)$, where $\mathcal{H} = \cup_{t=0}^{\infty} \mathcal{H}_t$ and $\mathcal{H}_t$ satisfies the recursion $\mathcal{H}_t = \mathcal{H}_{t-1} \times S \times \cup_{s\in S} A_s$ with $\mathcal{H}_0 = S$. In general, a strategy need not be time-invariant.

The restriction to stationary Markov strategies is essential for a learning algorithm to find an optimal strategy because it significantly reduces the space of possible strategies. Hence, an algorithm does not need to search for an optimal strategy over the space of all possible (history-dependent) contingency plans. Our focus on stationary Markov strategies is not restrictive because classical results (given below) show that there exists a stationary pure Markov strategy that achieves the same optimality criteria as an optimal strategy from Σ. Indeed, most learning algorithms search for an optimal strategy in the space of stationary Markov strategies.

Optimality Criteria The value of a strategy $\sigma \in \Sigma^{SM}, \Sigma^{SPM}, \Sigma$ that starts in state s is the continuation value of the strategy from state s, i.e., the expected discounted stream of payoffs from implementing strategy σ is given by

$$\upsilon_s(\sigma) = E_\sigma \left[\sum_{t=0}^{\infty} \delta^t \, u(s_t, a_t, s_{t+1}) \,\middle|\, s_0 = s \right] , \tag{2.2}$$

where the expectation in (2.2) is with respect to the strategy σ. That is, actions are sampled from σ and the expectation is taken over $p(s_{t+1}|s_t, a_t)$.

For a fixed value of the discount parameter $\delta \in [0, 1)$, the optimal continuation value is given by $\upsilon_s^* = \sup_{\sigma\in\Sigma} \upsilon_s(\sigma)$. Existence and uniqueness of $\upsilon^* = (\upsilon_s^*)_{s\in S}$ is guaranteed by Theorem 6.2.5a in [68]. Therefore, an optimal strategy $\sigma^* \in \Sigma^{SM}, \Sigma^{SPM}, \Sigma$ exists if $\upsilon_s^* = \upsilon_s(\sigma^*) \geq \upsilon_s(\sigma)$ for all $s \in S$ and all $\sigma \in \Sigma$. Crucially, Theorem 6.2.10 in [68] guarantees that there exists an optimal stationary pure Markov strategy $\sigma^* \in \Sigma^{SPM}$, such that $\upsilon_s(\sigma^*) \geq \upsilon_s(\sigma)$ for all $s \in S$ and all $\sigma \in \Sigma$.

Therefore, with these theoretical guarantees, we ignore all strategies that are history-dependent contingency plans, and for the remainder of the paper, a strategy refers to a stationary Markov strategy, and a pure strategy refers to a stationary pure Markov strategy.

2.3.2 Trading Environment

We present our model of the limit order book. Many of our assumptions are for tractability purposes and conform with the market dynamics described in Section 2.2. We use the midpoint of the bid-ask spread as a proxy for the fundamental value of the asset Z. At each time point, the value of the asset either goes to $Z + \varphi$ with probability $\beta \in (0, 1)$, or goes down to $Z - \varphi$ with probability $1 - \beta$, where $\varphi > 0$ is the tick size. The fundamental value of the asset is a martingale when $\beta = 0.5$, and it drifts up or down when $\beta > 0.5$ or $\beta < 0.5$, respectively.

States The set of states S is the Cartesian product of a set of environmental variables Ω and the inventory of the market maker $\mathcal{Q}$, i.e., $S = \Omega \times \mathcal{Q}$. We restrict the level of inventory to the set $\mathcal{Q} = \{-\bar{q}, ..., 0, ..., \bar{q}\}$, where $\bar{q}$ is some positive integer. The set Ω contains a finite number of environmental variables which are relevant features of the order book and that affect the payoffs the market maker receives. Here, the elements of Ω are the three regimes of volume imbalance in the limit order book, i.e., $\omega \in \Omega = \{BH, N, SH\}$, because quote-based manipulation is the focus of our analysis.

As discussed in Section 2.2, volume imbalance affects the behavior of market participants, and therefore affects the probability with which a limit order is filled. To capture this effect, let $p^b_\omega \in (0, 1)$ and $p^a_\omega \in (0, 1)$ denote the probability that a limit buy and a limit sell order, respectively, are filled in $[t, t + 1)$ for each regime $\omega \in \Omega$.[10] These fill probabilities account for the effect of time-priority in the limit order book and need not sum to unity. Our empirical results in Section 2.2 (see Table 2.7) show that the fill probabilities of bids and offers are similar when the book is neutral (i.e., $p^b_N \approx p^a_N$), the fill probabilities of offers are higher when the book is buy-heavy (i.e., $p^b_{BH} \ll p^a_{BH}$), and the fill probabilities of bids are higher when

[10] In addition to affecting the fill probabilities, volume imbalance has substantial explanatory power in predicting future price movements (see e.g.,[51, 19]). Therefore, submitting a manipulative order may change the volume imbalance regime, which may affect future price movements; however, it should not affect the fundamental value of the asset because manipulative orders contain no real information about the fundamentals. In our model, we do not model how volume imbalance affects future price movements because we use the midpoint of the bid-ask spread as a proxy for the fundamental value of the asset. This prevents unnecessary complications that arise when computing the change in wealth for the payoffs received as a consequence of submitting a manipulative order. Our analysis focuses specifically on how manipulative orders affect the fill probabilities, and how the fill probabilities affect the payoffs at the next period. If manipulative orders were to affect future price movements (temporarily), then quote-based manipulation would become more prevalent in our model because a manipulative order will increase the profits of round-trip trades. In Section 2.5, we show that the profit from round-trip trades is a key factor to determine if quote-based manipulation is optimal.

the book is sell-heavy (i.e., $p^b_{SH} \gg p^a_{SH}$). One could relax the assumption of static probabilities p^a_ω and p^b_ω. They may depend on time ($p^a_\omega = p^a_\omega(t)$ and $p^b_\omega = p^b_\omega(t)$) and may follow some dynamics. However, the key features we want to capture are the comparisons between fill probabilities for different book regimes. In particular, the empirical facts that sell limit orders have the highest probability to be executed in the BH regime and the lowest probability to be executed in the SH regime, and that buy limit orders have the highest probability to be executed in the SH regime and the lowest probability to be executed in the BH regime. These key features are already captured in the model by assuming some static probabilities that ultimately satisfy some inequalities. Otherwise we would have to add some conditions for the time-dependent fill probabilities. We could assume that $p^b_{BH}(t)$ and $p^a_{SH}(t)$ (the two lowest fill probabilities) always lie in an interval (p_1, p_2), $p^b_N(t)$ and $p^a_N(t)$ always lie in an interval (p_3, p_4) and $p^b_{SH}(t)$ and $p^a_{BH}(t)$ (the two highest fill probabilities) always lie in an interval (p_5, p_6) with $0 < p_1 < p_2 < p_3 < p_4 < p_5 < p_6 < 1$.

Actions The market maker does not endow the algorithm with an action that manipulates the order book. Instead, we focus on how an innocuous set of actions leads to manipulation when individual actions are sequenced in a particular order. That is, quote-based manipulation occurs when the set of actions produces unintended manipulative behavior as a consequence of a learning algorithm dynamically optimizing the market maker's optimality criteria.

The set of actions at time t consists of:

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

When a large limit order is submitted, the probability that the manipulative order is filled is low. Recall that when the large limit buy (limit sell) tilts the book to buy-heavy (sell-heavy), Table 2.1 shows that the arrival rate of sell (buy) market orders is very low. Therefore, we assume that at most one unit of the large limit order can be executed. This assumption allows us to validate that a particular sequence of actions has the intention to manipulate; see the analysis in Section 2.5.

We do not model the strategic cancellation of limit orders for the sake of analytical tractability.[11] Rather, we focus on when algorithms learn to create misleading information to buy or to sell an asset with a higher probability than was otherwise likely to occur, which is the key feature that enables quote-based manipulation. Explicitly cancelling a limit order, i.e., sending an instruction to the exchange to cancel an outstanding limit order, is not necessary in Nasdaq if the limit order is submitted with a time-in-force between t and $t+1$ because the limit order expires at time $t+1$.

The set of actions available to the market maker depends on the state $s \in S$. The market maker has access to the full set of actions $A_s = \{LB, LS, LLB, LLS, MB, MS, DN\}$ when $s = (\omega, q)$ for all $\omega \in \Omega$ and for all $q \in \mathcal{Q} \setminus \{\bar{q}, -\bar{q}\}$. At the boundaries of the inventory constraint, the market maker does not buy or does not sell any additional assets. That is, the set of actions reduces to $A_s = \{LS, LLS, MS, DN\}$ when $s = (\omega, \bar{q})$ for all $\omega \in \Omega$, and to $A_s = \{LB, LLB, MB, DN\}$ when $s = (\omega, -\bar{q})$ for all $\omega \in \Omega$.

In this setup, individual actions are not manipulative. Although a large limit order can tilt the book, the action is not manipulative because there is no advantage to be gained from the action alone; rather, a large limit order will only be manipulative when the action following a large limit order aims to profit from the tilt created. For this reason, we focus on a manipulative sequence of actions. We formalize and refine this later in Definition 1.

Transition Dynamics We present the transition dynamics over the set $\mathcal{Q}$ and Ω separately to simplify the presentation, and recall that the set of states is $S = \Omega \times \mathcal{Q}$.

The transition dynamics of the inventory level is intuitive. If an action resulted in a buy transaction, then the level of inventory increases by one unit, i.e., $q_{t+1} = q_t + 1$; if an action resulted in a sell transaction, then the level of inventory decreases by one unit, i.e., $q_{t+1} = q_t - 1$; and if an action resulted in no transaction, then the level of inventory stays the same, i.e., $q_{t+1} = q_t$. The transition probabilities depend on the action taken and the fill probabilities (for limit order submissions).

The transition probabilities of the volume imbalance regime depend on two distinct cases: when a small or no order is submitted, i.e., $a \in \{LB, LS, MB, MS, DN\}$, and

[11] Including cancellation of an order as part of the action prevents us from including an additional state variable that tracks if a market maker has an outstanding order in the book. Additionally, the fill probability between time $[t, t+1)$ and $[t+1, t+2)$ for a limit order submitted at time t is not the same unless the fill probability is Markov.

when a large order is submitted, i.e., $a \in \{LLB, LLS\}$. When a small or no order is submitted, volume imbalance evolves according to some baseline dynamics because unit orders and no orders have little to no impact on the liquidity at the best bid-ask quotes, and hence have little to no impact on the volume imbalance in (2.1).

On the other hand, when a large limit order is submitted, the volume of the buy (sell) limit order is large enough to tilt the volume imbalance regime to the buy-heavy (the sell-heavy) regime with high probability. If we ignore the change in behavior of other market participants, then the volume imbalance regime changes at time t and reverts back before time $t+1$ because of the timing of submission and cancellation of large limit orders posted by the market maker. However, in Section 2.2, we saw that market participants adjust their own liquidity provision. Empirically, we have the following observations: (i) the average volume of limit orders and limit order cancellations is similar in size under each volume imbalance regime, but the arrival rates of limit orders are higher than the arrival rates of limit order cancellations, and (ii) there are more buy (sell) limit orders than sell (buy) limit orders when the book is buy-heavy (sell-heavy). Therefore, the empirical results show that market participants adjust their own liquidity provision to submit more limit orders to the same side of the book as the market maker's large limit order. However, market participants can never react instantaneously and will have a delay when reacting to a large limit order from the market maker. Such delay can occur for a number of reasons, for example, other market participants do not have the monitoring capabilities (i.e., latency, see [9]), or do not have the infrastructure to react immediately.

The effect of the delay and change in liquidity provision is that the volume imbalance regime at time $t+1$ corresponds to the change in regime caused by the market maker's large limit order at time t. When a large limit order is submitted, the volume imbalance regime moves to the buy-heavy (the sell-heavy) regime at the next step with a higher probability than was otherwise likely to occur, i.e., for all $\omega \in \Omega$, $p(BH \,|\, \omega, LLB) = 1 - \kappa > p(BH|\omega, a)$ for any $a \in A_s \setminus \{LLB\}$, and $p(SH \,|\, \omega, LLS) = 1 - \kappa > p(SH|\omega, a)$ for any $a \in A_s \setminus \{LLS\}$, where $\kappa \in [0, 1)$. Moreover, the volume imbalance moves to one of the "wrong" regimes with probability $\kappa/2$. That is, $p(N \,|\, \omega, LLB) = \kappa/2$ and $p(SH \,|\, \omega, LLB) = \kappa/2$ when a large buy limit order is submitted, and $p(N \,|\, \omega, LLS) = \kappa/2$ and $p(BH \,|\, \omega, LLS) = \kappa/2$ when a large sell limit order is submitted.

These stylized facts are important for unintentional quote-based manipulation to occur. They create the necessary temporal link between past actions and future actions of the market maker so that quote-based manipulation becomes dynamically optimal as a sequence of actions. Thus, quote-based manipulation emerges as an optimal sequence of actions even if the market maker does not encode quote-based manipulation as a possible action into the learning algorithm.

Utility The market maker is averse to holding inventory and maximizes the present value of her wealth. The wealth $W = X + Z\,q$ of the market maker is the sum of her cash position X and the marked-to-market value of the inventory $Z\,q$, where Z is the fundamental value of the asset that is proxied by the midpoint of the bid-ask spread. To cast the market maker's objective into the optimization problem of a learning algorithm, we write the one-step utility as

$$u(s, a, s') = Y(s, a, s') - \alpha\,(q')^2\,, \tag{2.3}$$

where $q' \in \mathcal{Q}$ is the inventory after action a and $\alpha > 0$ is the inventory aversion parameter.[12] The quadratic penalty ensures that the utility function is concave in the level of inventory. Hence, the inventory aversion parameter α affects the willingness of the market maker to take on larger levels, long or short, of inventory. For example, as the value of α increases, the market maker is more averse to inventory risk so she is less willing to increase the level of inventory, long or short. On the other hand, $Y(s, a, s')$ is the change in wealth as a consequence of action $a \in A_s$ in state $s \in S$. For a value of the tick size $\varphi > 0$, the expected change in wealth from state $= (\omega, q) \in S$ and taking action $a \in A_s$ is given by

$$E\left[Y(s, a, s')\right] = \begin{cases} p^b_\omega\,\vartheta/2 + (2\,\beta - 1)\,(\varphi\,q + p^b_\omega\,\varphi) & \text{for } a = \{LB, LLB\}\,, \\ p^a_\omega\,\vartheta/2 + (2\,\beta - 1)\,(\varphi\,q - p^a_\omega\,\varphi) & \text{for } a = \{LS, LLS\}\,, \\ -\vartheta/2 + (2\,\beta - 1)\,(\varphi\,q + \varphi) & \text{for } a = MB\,, \\ -\vartheta/2 + (2\,\beta - 1)\,(\varphi\,q - \varphi) & \text{for } a = MS\,, \\ (2\,\beta - 1)\,\varphi\,q & \text{for } a = DN\,, \end{cases} \tag{2.4}$$

where the expectation is taken with respect to the fundamental value of the asset Z and the fill probabilities p^b_ω, p^a_ω of limit orders, and $\vartheta > 0$ is the expected quoted spread.

With this one-step utility function, the optimal continuation value is given by

$$\sup_{\sigma \in \Sigma} E_\sigma \left[\sum_{t=0}^{\infty} \delta^t \left(Y(s_t, a_t, s_{t+1}) - \alpha\, q^2_{t+1} \right) \middle|\, s_0 = s \right]\,, \tag{2.5}$$

[12]The units of α are such that (2.5) is in units of wealth. In our model, the value of α is allowed to range between $(0, \infty)$, but in practice, as a rule-of-thumb, the value of α must be several orders of magnitudes smaller than the expected quoted spread, i.e., $\alpha \ll \vartheta$, otherwise the market maker will not be willing to make markets.

which corresponds to maximizing the present value of wealth subject to a running inventory penalty.[13] The optimization problem in (2.5) is similar to the optimization problem posed in [65] where they assume a negative exponential for the one-step utility function. Our choice of (2.3) leads to a clear interpretation of the continuation value, and it also simplifies our analysis in the next section.

The objective to maximize wealth subject to a running inventory penalty is closely related to robustness and ambiguity aversion from [49]. Specifically, [27] show in a related problem that the market maker's objective of maximizing wealth subject to a running inventory penalty is equivalent to a risk-neutral to inventory market maker who is ambiguous to the drift of the fundamental value of the asset.

Finally, the optimality criterion in (2.3) produces behavior consistent with inventory models. The behavior of the optimal strategy depends on the level of inventory, and there is a preferred inventory position (the preferred inventory position is zero when the fundamental price is a martingale), which is consistent with the results of [5].[14] The optimal strategy also prefers to sell if inventory is long and prefers to buy if inventory is short, which is consistent with the results of [71] and [54].

For simplicity, the remainder of the paper assumes that the fundamental value of the asset is a martingale, i.e., $\beta = 0.5$, so that the expected one-step utility from state $s = (\omega, q) \in S$ and taking action $a \in A_s$ is given by

$$\bar{u}(s,a) = \begin{cases} p^b_\omega\,\vartheta/2 - \alpha\, p^b_\omega\,(q+1)^2 - \alpha\,(1-p^b_\omega)\,q^2 & \text{for } a = \{LB, LLB\}\,, \\ p^a_\omega\,\vartheta/2 - \alpha\, p^a_\omega\,(q-1)^2 - \alpha\,(1-p^a_\omega)\,q^2 & \text{for } a = \{LS, LLS\}\,, \\ -\vartheta/2 - \alpha\,(q+1)^2 & \text{for } a = MB\,, \\ -\vartheta/2 - \alpha\,(q-1)^2 & \text{for } a = MS\,, \\ -\alpha\, q^2 & \text{for } a = DN\,. \end{cases} \tag{2.6}$$

2.4 Theory

This section derives conditions to test if quote-based manipulation will occur when decision making is delegated to a learning algorithm. Our analysis focuses on $q \neq 0$, where the overall intention is to buy or to sell the asset so that inventory reverts to the preferred position $q = 0$. We omit the case when $q = 0$ because the overall intention to buy or to sell is unclear.

[13] In (2.5), the payoffs received depends on the realization of the inventory level at the next time step; however, the payoffs are discounted from the start of the period. The timing of the payoffs is constructed to fit within the framework of learning algorithms, where the payoffs are assumed to be immediate. This construction presents no issue in (2.5) because of the expectation operator, and the timing is also consistent with the model of [65].

[14] See also [60] and [53] who examine this preferred inventory position in more detail.

Throughout the project, we maintain the following assumptions on the parameters of our model.

Assumption 1. *The expected quoted spread and the inventory aversion parameter are both greater than zero (i.e., $\vartheta > 0$ and $\alpha > 0$), the fill probabilities $p^b_\omega \in (0,1)$ and $p^a_\omega \in (0,1)$ for all $\omega \in \Omega$, and the market maker is not myopic, i.e., $\delta > 0$.*

2.4.1 Optimal Strategy

The optimal strategy satisfies Bellman's optimality equations, so the optimal action in each state $s = (\omega, q) \in S$ is given by

$$a^* = \underset{a \in A_s}{\arg\max} \left\{ \bar{u}(s,a) + \delta \sum_{\omega' \in \Omega} p(\omega'|\omega, a) \sum_{q' \in \mathcal{Q}} p(q'|q,a)\, \upsilon^*_{\omega',q'} \right\}, \tag{2.7}$$

where υ^* is the optimal continuation value. The optimal action is non-myopic and balances the immediate expected payoff $\bar{u}(s,a)$ with the expected future stream of discounted payoffs. Hence, the optimal action takes into account how the current action will affect the transition to subsequent states s, and therefore accounts for how the current action will affect future actions.

To gain some insight into the optimal strategy, the following lemma shows that the optimal continuation value υ^* decreases in value as the level of inventory deviates further away from zero. Therefore, the optimal continuation value υ^* achieves its maximum value at zero inventory.

Lemma 1. *For $q \geq 0$, the optimal continuation value $\upsilon^*_{\omega,q}$ is non-increasing as q increases, i.e., when $0 \leq q' \leq q$, we have $\upsilon^*_{\omega,q} \leq \upsilon^*_{\omega,q'}$ for all $\omega \in \Omega$. Similarly, for $q \leq 0$, the optimal continuation value $\upsilon^*_{\omega,q}$ is non-decreasing as q increases, i.e., when $q' \leq q \leq 0$, we have $\upsilon^*_{\omega,q'} \leq \upsilon^*_{\omega,q}$ for all $\omega \in \Omega$.*

The result is intuitive because the market maker's aversion to higher levels of inventory and because the fundamental value of the asset is a martingale. One implication is that the preferred level of inventory is zero because it leads to the highest expected stream of discounted payoffs; hence, optimal strategies will induce mean reversion to zero inventory. The concavity of the optimal continuation value $\upsilon^*_{\omega,q}$ as a function of inventory holds for each volume imbalance regime ω. The relationship between the optimal continuation value $\upsilon^*_{\omega,q}$ and the volume imbalance regimes is key for spoofing to be optimal.

The following lemma eliminates suboptimal actions, which reduces the number of actions to consider when solving (2.7).

Lemma 2. *For all volume imbalance regimes $\omega \in \Omega$, the following two statements hold. The actions do nothing (i.e., DN) and market buy order (i.e., MB) are suboptimal if $q > 0$. The actions do nothing (i.e., DN) and market sell order (i.e., MS) are suboptimal if $q < 0$.*

The following proposition characterizes the optimal action for each state $s = (\omega, q)$ when $q \neq 0$. The optimal action is characterized in terms of the value of α, i.e., the willingness to hold inventory. Figure 2.1 shows the optimal action as a function of the value of the inventory aversion parameter α when $q > 0$. For each pair of neighboring actions in Figure 2.1, there is a cutoff value of α such that the optimal strategy is indifferent between the two actions because they yield the same expected stream of discounted payoffs. Hence, for a given value of α, the optimal action is one that maximizes the expected stream of discounted payoffs.

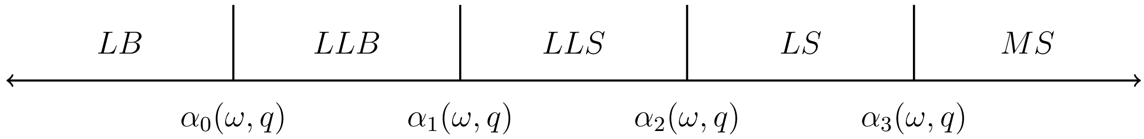

Figure 2.1: Optimal action choice for each state $s = (\omega, q)$ for $q > 0$.

For the remainder of the paper, we denote $x \vee y = \max\{x, y\}$, and $x \wedge y = \min\{x, y\}$.

Proposition 1. *Let $q > 0$ and assume $p^a_{SH} < p^a_N < p^a_{BH}$ holds. Then, for each state $s = (\omega, q)$, there exist cutoff values of the inventory aversion parameter $\alpha_0(\omega, q) \leq \alpha_1(\omega, q) \leq \alpha_2(\omega, q) \leq \alpha_3(\omega, q)$ such that the optimal stationary pure Markov strategy $\sigma^* \in \Sigma^{SPM}$ is given by*

$$\sigma^*(\omega, q) = \begin{cases} LB & \text{if } \alpha \in \big(0, 0 \vee \alpha_0(\omega, q)\big)\,, \\ LLB & \text{if } \alpha \in \big(0 \vee \alpha_0(\omega, q), 0 \vee \alpha_1(\omega, q)\big)\,, \\ LLS & \text{if } \alpha \in \big(0 \vee \alpha_1(\omega, q), 0 \vee \alpha_2(\omega, q)\big)\,, \\ LS & \text{if } \alpha \in \big(0 \vee \alpha_2(\omega, q), 0 \vee \alpha_3(\omega, q)\big)\,, \\ MS & \text{if } \alpha \in \big(0 \vee \alpha_3(\omega, q), +\infty\big)\,. \end{cases} \tag{2.8}$$

Similarly, let $q < 0$ and assume $p^b_{SH} > p^b_N > p^b_{BH}$ holds. Then, for each state $s = (\omega, q)$, there exist cutoff values of the inventory aversion parameter $\alpha_0(\omega, q) \leq \alpha_1(\omega, q) \leq \alpha_2(\omega, q) \leq \alpha_3(\omega, q)$ such that the optimal stationary pure Markov strategy

$\sigma^* \in \Sigma^{SPM}$ *is given by*

$$\sigma^*(\omega, q) = \begin{cases} LS & \text{if } \alpha \in \big(0, 0 \vee \alpha_0(\omega, q)\big)\,, \\ LLS & \text{if } \alpha \in \big(0 \vee \alpha_0(\omega, q), 0 \vee \alpha_1(\omega, q)\big)\,, \\ LLB & \text{if } \alpha \in \big(0 \vee \alpha_1(\omega, q), 0 \vee \alpha_2(\omega, q)\big)\,, \\ LB & \text{if } \alpha \in \big(0 \vee \alpha_2(\omega, q), 0 \vee \alpha_3(\omega, q)\big)\,, \\ MB & \text{if } \alpha \in \big(0 \vee \alpha_3(\omega, q), +\infty\big)\,. \end{cases} \tag{2.9}$$

In general, the cutoff values $\alpha_0(\omega, q)$, $\alpha_1(\omega, q)$, and $\alpha_3(\omega, q)$ are different for positive and negative inventory. However, the cutoff values are the same for positive and negative inventory in the particular case when $p^a_{BH} = p^b_{SH}$, $p^a_{SH} = p^b_{BH}$, and $p^a_N = p^b_N$, and when the transition probability matrix corresponding to the baseline book regime transitions (when small orders or no orders are submitted) is a centrosymmetric matrix (i.e., a matrix that is symmetric about its center).

For $q > 0$ and $\alpha = \alpha_0(s)$, the optimal strategy is indifferent between a buy limit order and a large buy limit order in state s; similarly, for $q < 0$ and $\alpha = \alpha_0(s)$, the optimal strategy is indifferent between a sell limit order and a large sell limit order in state s. A similar interpretation applies to the remaining cutoff values $\alpha_1(s)$, $\alpha_2(s)$, and $\alpha_3(s)$. Note that for $q > 0$ ($q < 0$) we can only guarantee MS (MB) is optimal for some range of values of α, by sending $\alpha \to +\infty$. For example, the remaining 4 actions would never be optimal if all the cutoff values were negative. Also, even if all the cutoff values were positive, some actions could never be optimal. This arises in case some cutoff values are equal. For example, for $q > 0$, if $\alpha_0(\omega, q) = \alpha_1(\omega, q)$, then LLB is never optimal. Similarly, for $q > 0$, if $\alpha_1(\omega, q) = \alpha_2(\omega, q)$, then LLS is never optimal and if $\alpha_2(\omega, q) = \alpha_3(\omega, q)$, then LS is never optimal. The conditions in the proposition include the maximum operator because the inventory aversion parameter is strictly positive and there is no guarantee that the cutoff values are positive. The conditions on the fill probabilities are not restrictive because they hold for all assets and all the timescales considered (see Table 2.7).

The proposition is intuitive because the optimal strategy induces mean reversion to zero inventory. For example, when there is less willingness to hold larger values of inventory (i.e., for large values of α), the action preference favors actions that aim to sell the asset and reduce the level of inventory when $q > 0$, and the action preference favors actions that aim to buy the asset and increase the level of inventory when $q < 0$. Similarly, for a fixed value of α, as the level of inventory deviates away from zero, the cutoff value $\alpha_1(\omega, q)$ shifts closer to zero (the value of $\bar{\alpha}_1(\omega, q)$ in (2.10) decreases as the absolute value of q increases). Hence, the action preference favors actions that aim at selling the asset to reduce the level of inventory when $q > 0$, and that aim at buying the asset to increase the level of inventory when $q < 0$.

2.4.2 Quote-Based Manipulation

To derive the conditions to test if algorithms will learn to manipulate the order book, we first define quote-based manipulation in terms of the set of actions available. Recall that an explicit consideration of our model is that the market maker does not endow her algorithm with an individual action that manipulates the book; rather, manipulation occurs in our model when a particular combination of actions is sequenced together. The following definition includes a refinement that also accounts for the market maker's objective to buy or to sell the asset.

Definition 1 (Manipulation)**.** *For $q > 0$, quote-based manipulation occurs if a large buy limit order is placed at time t and tilts the book to buy-heavy regime, and it is followed at time $t + 1$ by a unit or large sell limit order. Similarly, for $q < 0$, quote-based manipulation occurs if a large sell limit order is placed at time t and tilts the book to sell-heavy regime, and it is followed at time $t + 1$ by a unit or large buy limit order.*

The sequence of actions in Definition 1 is manipulative because the optimal strategy intends to revert to the preferred inventory position of zero, so sending a large buy limit order when the market maker is long or a large sell limit order when the market maker is short counters this objective. Therefore, if these sequences arise in optimality, then the market maker must be profiting through this manipulative order by tilting the book in her favor for future actions.

One consequence of quote-based manipulation is that the manipulative order may get caught out and lead to a transaction. We distinguish this possibility into two cases. Type I manipulation occurs when the manipulative order is not caught out, i.e., it is not filled. Specifically, type I manipulation is the sequence initiated by LLB in state $s = (\omega, q)$ and followed by LS or LLS in state $s' = (BH, q)$ when $q > 0$, or the sequence initiated by LLS in state $s = (\omega, q)$ and followed by LB or LLB in state $s' = (SH, q)$ when $q < 0$. On the other hand, type II manipulation occurs when the manipulative order is caught out, i.e., it gets filled. Specifically, type II manipulation is the sequence initiated by LLB in state $s = (\omega, q)$ and followed by LS or LLS in state $s' = (BH, q + 1)$ when $q > 0$, or the sequence initiated by LLS in state $s = (\omega, q)$ and followed by LB or LLB in state $s' = (SH, q - 1)$ when $q < 0$.

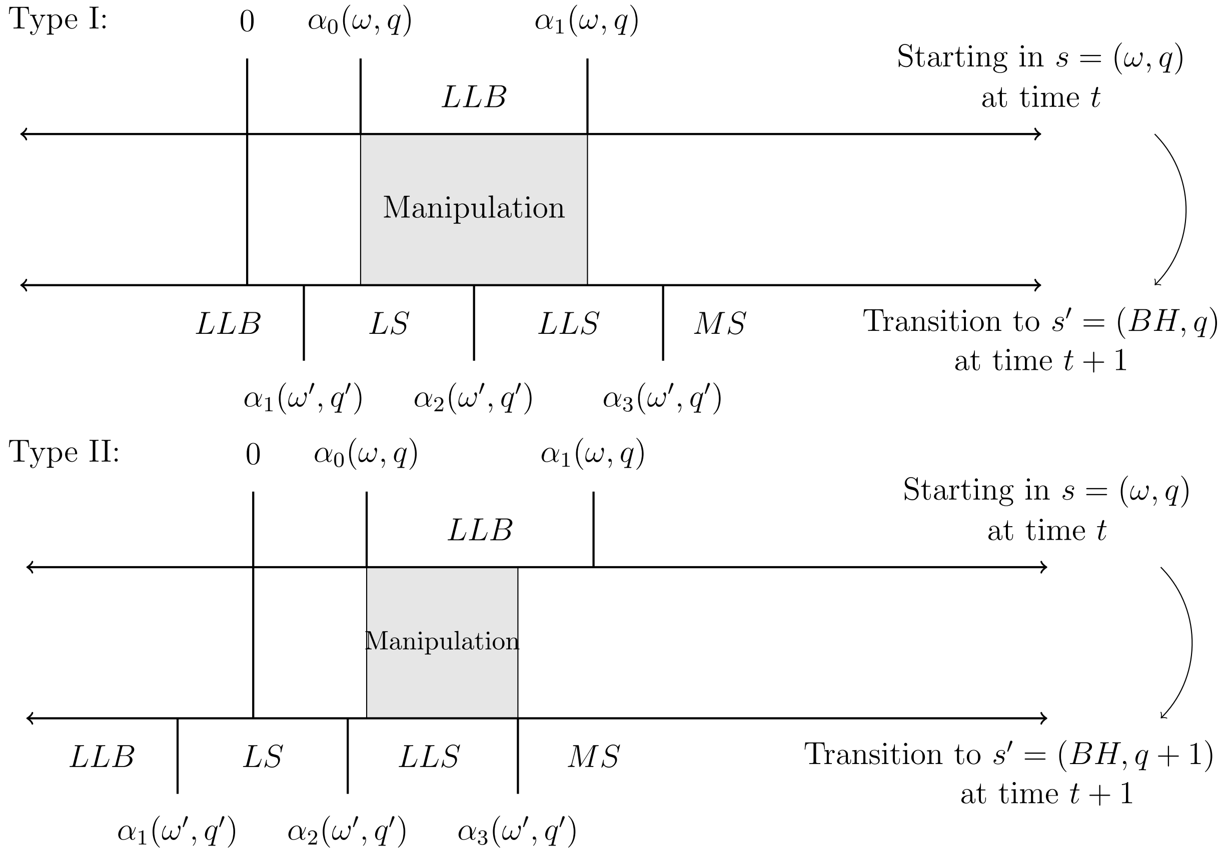

Figure 2.2: Quote-based manipulation is optimal when the value of the inventory aversion parameter α lies within the shaded region for $q > 0$. The top panel describes type I manipulation, and the bottom panel describes type II manipulation.

Figure 2.2 illustrates the two types of manipulation when $q > 0$. Here, manipulation occurs if the value of the inventory aversion parameter α is within the shaded region. We say that manipulation occurs in state $s = (\omega, q)$ if the optimal action is to submit a large buy limit order in state $s = (\omega, q)$, and when the market transitions to the subsequent state s' the optimal strategy prescribes to submit either a sell limit order or a large sell limit order.

We formalize the definition of quote-based manipulation in the context of an optimal strategy. We adopt the standard convention that the interval $(x, y) = \emptyset$ if $x \geq y$, and recall that $x \vee y = \max\{x, y\}$, and $x \wedge y = \min\{x, y\}$.

Definition 2 (Manipulative Strategy)**.** *If there exists a state $s = (\omega, q)$ such that*

(i) $$\emptyset \neq I_1(s) = \begin{cases} \big(0 \vee \alpha_0(\omega, q) \vee \alpha_1(BH, q), \big(0 \vee \alpha_1(\omega, q)\big) \wedge \big(0 \vee \alpha_3(BH, q)\big)\big) & \text{if } q > 0\,, \\ \big(0 \vee \alpha_0(\omega, q) \vee \alpha_1(SH, q), \big(0 \vee \alpha_1(\omega, q)\big) \wedge \big(0 \vee \alpha_3(SH, q)\big)\big) & \text{if } q < 0\,, \end{cases}$$

or

(ii) $\emptyset \neq I_2(s) = \begin{cases} \big(0 \vee \alpha_0(\omega, q) \vee \alpha_1(BH, q+1), \big(0 \vee \alpha_1(\omega, q)\big) \wedge \big(0 \vee \alpha_3(BH, q+1)\big)\big) & \text{if } q > 0\,, \\ \big(0 \vee \alpha_0(\omega, q) \vee \alpha_1(SH, q-1), \big(0 \vee \alpha_1(\omega, q)\big) \wedge \big(0 \vee \alpha_3(SH, q-1)\big)\big) & \text{if } q < 0\,. \end{cases}$

Then, if the value of the inventory aversion parameter $\alpha \in I_1(s)$, the optimal strategy is a manipulative strategy, where type I manipulation occurs in states s where condition (i) is satisfied. Similarly, if the value of the inventory aversion parameter $\alpha \in I_2(s)$, the optimal strategy is a manipulative strategy, where type II manipulation occurs in states s where condition (ii) is satisfied.

The intervals $I_1(s)$ and $I_2(s)$ describe conditions for quote-based manipulation to be dynamically optimal as a sequence of actions. Specifically, if $\alpha \in I_1(s)$ or $\alpha \in I_2(s)$, then manipulation occurs in state s; if $\alpha \in I_1(s)$, then manipulation occurs in state s when the manipulative order is not filled; if $\alpha \in I_2(s)$, then manipulation occurs in state s when the manipulative order is filled; and if $\alpha \in I_1(s) \cap I_2(s)$, then manipulation occurs in state s regardless of whether the manipulative order was filled or not.

2.4.3 Testable Conditions

To obtain testable conditions that apply to a wide variety of learning algorithms, we make the following assumption.

Assumption 2. *The learning algorithm used by the market maker learns an optimal stationary pure Markov strategy $\sigma^* \in \Sigma^{SPM}$.*

This assumption is not restrictive and allows us to analyze the framework where algorithms learn optimal strategies so that the testable conditions we derive apply to generic learning algorithms. The premise and objective of designing a learning algorithm is to learn an optimal stationary pure Markov strategy. Indeed, the most popular offline learning algorithms (such as policy iteration and value iteration) and online learning algorithms (such as Q-learning and SARSA) satisfy this assumption (see for example [68], [72]).[15]

Clearly, when Assumption 2 and the conditions for manipulation in Definition 2 hold, the algorithm will learn a manipulative strategy. The issue is that the intervals $I_1(s)$ and $I_2(s)$ may be empty. Additionally, the intervals depend on the cutoff values, which in turn, depend on the parameters of the model and on the optimal continuation value v^*. The following theorem provides sufficient conditions to determine if

[15] See [77] for a convergence proof of Q-learning and [69] for a convergence proof of SARSA. See also [48], [63], and [70] for applications of learning algorithms in financial markets that satisfy this assumption.

algorithms will learn manipulative strategies. These conditions depend only on the fill probabilities of the limit orders.

Theorem 1. *Let $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If the conditions*

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

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

hold, if the transition probabilities associated with large limit orders are such that

$$\begin{aligned} p(BH\,|\,\omega, LLB) = 1-\kappa > p_{BH|\omega}\,,\; p(N\,|\,\omega, LLB) = \frac{\kappa}{2} < p_{N|\omega}\,,\; p(SH\,|\,\omega, LLB) = \frac{\kappa}{2} < p_{SH|\omega}\,, \\ p(SH\,|\,\omega, LLS) = 1-\kappa > p_{SH|\omega}\,,\; p(N\,|\,\omega, LLS) = \frac{\kappa}{2} < p_{N|\omega}\,,\; p(BH\,|\,\omega, LLS) = \frac{\kappa}{2} < p_{BH|\omega}\,, \end{aligned} \tag{C3}$$

hold for all $\omega \in \Omega$, and if

$$\begin{aligned} p^a_{BH} - p^b_{SH} &< \min\left\{ (p^b_{SH} - \max\{p^a_N, p^b_N\}) \frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}}\,,\, (p^b_{SH} - \max\{p^a_N, p^b_N\}) \frac{p_{N|N} - \frac{\kappa}{2}}{p_{BH|N} - \frac{\kappa}{2}} \right\} \\ p^b_{SH} - p^a_{BH} &< \min\left\{ (p^a_{BH} - \max\{p^a_N, p^b_N\}) \frac{p_{N|SH} - \frac{\kappa}{2}}{p_{SH|SH} - \frac{\kappa}{2}}\,,\, (p^a_{BH} - \max\{p^a_N, p^b_N\}) \frac{p_{N|N} - \frac{\kappa}{2}}{p_{SH|N} - \frac{\kappa}{2}} \right\} \end{aligned} \tag{C4}$$

hold, then $I_1(s) \neq \emptyset$ and $I_2(s) \neq \emptyset$ for all $s \in S$ such that (i) $s = (SH, q > 0)$, and (ii) $s = (BH, q < 0)$.
Moreover,

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

In short, if the fill probabilities and the transition probabilities satisfy certain conditions, then there are values of the inventory aversion parameter α where an algorithm will learn a manipulative strategy in which type I and/or type II manipulation occurs in state s. Whether both types or only one type is learned depends on the value of the inventory aversion parameter α and if the intervals $I_1(s)$ and $I_2(s)$ overlap.

The parameters of the model, such as the fill probabilities and the transition probabilities, may change if market participants adapt their behavior to the behavior of the algorithm. This would result in a sequence of changes to the parameter values, which may or may not lead to an equilibrium. Nonetheless, one can check if conditions (C1)–(C4) hold whether we are in an equilibrium or not.

The role of the conditions in Theorem 1 is intuitive. The conditions for the fill probabilities of the ask and of the bid are so that the ordering of action preferences in Proposition 1 hold. For $q > 0$, condition (C1) ensures that buy limit orders are not optimal when the book is buy-heavy, while condition (C2) ensures that there are values of the inventory aversion parameter α for which it is optimal to initiate the manipulative sequence with a large buy limit order in sell-heavy. Similarly, for $q < 0$, condition (C2) ensures that sell limit orders are not optimal when the book is sell-heavy, while condition (C1) ensures that there are values of the inventory aversion parameter α for which it is optimal to initiate the manipulative sequence with a large sell limit order in buy-heavy. Condition (C3) ensures that submission of buy (sell) large limit orders make the volume imbalance regime to move to the buy-heavy (the sell-heavy) regime at the next step with a higher probability than was otherwise likely to happen.

Condition (C4) is a formal condition to describe that the values of the fill probabilities p^a_{BH} and p^b_{SH} are similar, i.e., $p^a_{BH} \approx p^b_{SH}$. This condition simplifies the analysis to determine if the manipulative sequence is initiated in the neutral regime with a large buy limit order for $q > 0$, or if the manipulative sequence is initiated with a large sell limit order for $q < 0$. On the other hand, the conditions $(p^b_N - p^a_N) > \frac{\delta(1-\kappa-\frac{\kappa}{2})}{1+\delta(1-\kappa-\frac{\kappa}{2})}(p^b_{SH} - p^a_{BH})$ or $(p^a_N - p^b_N) > \frac{\delta(1-\kappa-\frac{\kappa}{2})}{1+\delta(1-\kappa-\frac{\kappa}{2})}(p^a_{BH} - p^b_{SH})$ allow one to determine if manipulation occurs when inventory is long or short in the neutral regime. These conditions are such that one inequality will always hold, but never both or neither.

One can think of the fill probabilities in the neutral regime as the short-term incentives associated with the immediate payoffs, and the fill probabilities in the heavy regimes as the long-term incentives associated with discounted future payoffs. The intuition behind these conditions is clear. If the signs of $p^b_N - p^a_N$ and $p^b_{SH} - p^a_{BH}$ are different, then the short-term and long-term incentives align. On the other hand, if the signs of $p^b_N - p^a_N$ and $p^b_{SH} - p^a_{BH}$ are the same, then the short-term and long-term incentives are not aligned, so the determining factor is the tradeoff between short-term and long-term incentives.

For example, for manipulation to occur in the neutral regime for $q > 0$, it must be that there are values of the inventory aversion parameter α for which it is optimal to initiate the manipulative sequence with a large buy limit order. If $p^b_N > p^a_N$ and $p^b_{SH} < p^a_{BH}$, then the signs of $p^b_N - p^a_N$ and $p^b_{SH} - p^a_{BH}$ are different, so the short-term and long-term incentives align because the immediate payoff from a buy limit order is greater than that of a sell limit order, and the future payoffs from being in buy-heavy (proxied with p^a_{BH}) are better than the future payoffs from being in sell-heavy (proxied with p^b_{SH}). Thus, it is clear that there are values of the inventory aversion parameter α for which it is optimal to initiate the manipulative sequence with a large buy limit order because the incentives align. On the other hand, if $p^b_N > p^a_N$ and $p^b_{SH} > p^a_{BH}$, then the signs of $p^b_N - p^a_N$ and $p^b_{SH} - p^a_{BH}$ are the same, so the short-term and long-term incentives are not aligned. In this case, if the immediate payoff outweighs the discounted future payoffs, then there are values of the inventory aversion parameter α for which it is optimal to initiate the manipulative sequence with a large buy limit order.

Theorem 1 gives sufficient (but not necessary) conditions for the intervals $I_1(s)$ and $I_2(s)$ to exist. Hence, the theorem provides conditions to test when a manipulative strategy could be optimal; however, failure to satisfy the conditions does not mean that a manipulative strategy cannot be optimal. Moreover, the theorem does not specify the values of the inventory aversion parameter α for which a manipulative strategy is optimal. To narrow the search for values of the inventory aversion parameter α where an algorithm will learn to manipulate the book, we derive an interval $I'(s)$ that contains both $I_1(s)$ and $I_2(s)$, which uses the following upper bound.

Lemma 3. *Let $m \in (0,1)$ be the minimum element of the transition probability matrix corresponding to the baseline book regime transitions (when no large limit orders are submitted). Then for all $\omega, \omega' \in \Omega$ and $q \in \mathcal{Q}$, we have $\upsilon^*_{\omega,q} - \upsilon^*_{\omega',q} \leq \vartheta/m$.*

With this lemma, the following proposition characterizes the interval $I'(s)$ that contains both $I_1(s)$ and $I_2(s)$, and does not depend on the optimal continuation value υ^*.

Proposition 2. *Let $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. For all $s = (\omega, q)$ when $q \neq 0$, the interval $I'(s) = \big(0, \bar{\alpha}_1(\omega, q)\big)$ is such that $I_1(s) \subset I'(s)$ and $I_2(s) \subset I'(s)$.*

The interval $I'(s)$ is characterized in terms of the upper bound of the cutoff value $\alpha_1(\omega, q)$ given by $\alpha_1(\omega, q) \leq \bar{\alpha}_1(\omega, q)$, where

$$\bar{\alpha}_1(\omega, q) = \begin{cases} \Big[\, \big(p^b_\omega - p^a_\omega\big)\,\vartheta/2 + \vartheta/m \Big] \Big[p^a_\omega\,(2\,q - 1) + p^b_\omega\,(2\,q + 1) \Big]^{-1} & \text{if } q > 0\,, \\ \Big[\, \big(p^a_\omega - p^b_\omega\big)\,\vartheta/2 + \vartheta/m \Big] \Big[-\, p^a_\omega\,(2\,q - 1) - p^b_\omega\,(2\,q + 1) \Big]^{-1} & \text{if } q < 0\,, \end{cases} \tag{2.10}$$

for all $\omega \in \Omega$. The value of $\bar{\alpha}_1(\omega, q)$ is strictly positive, depends only on parameters of the model, and is easy to compute. The interval $I'(s)$ narrows the search for values of the inventory aversion parameter α for which an algorithm will learn a manipulative strategy.

Indeed, $\alpha \in I'(s)$ is a necessary but not a sufficient condition for an algorithm to learn a manipulative strategy in state s. Specifically, not all values of $\alpha \in I'(s)$ will lead to manipulation in state s because not all values of $\alpha \in I'(s)$ lie within $I_1(s)$ or $I_2(s)$. Nonetheless, if $\alpha \notin I'(s)$, then an algorithm cannot learn a strategy that manipulates the order book in state s. In the next section, we use this condition to analyze how parameters of the model affect an algorithm's ability to learn to manipulate the order book.

2.5 Understanding Manipulation and Spoofing

This section analyzes the mechanics of what makes quote-based manipulation dynamically optimal and how parameters of the model affect the optimality of manipulative strategies.

2.5.1 Workings of Manipulation

Mechanics of Manipulation Figure 2.3 illustrates the mechanics behind manipulation in optimality with data from AMZN and CSCO at 0.5 second decision intervals. The model parameters are estimated with the dataset from Section 2.2; the estimation procedure is discussed in the next chapter, in Section 3.2. With a discount factor of $\delta = 0.95$, we solve for the optimal strategy with the policy iteration algorithm and solve for the optimal continuation values $v_{\omega,q}$ for each volume imbalance regime as a function of the inventory level. The inventory aversion parameter is $\alpha = 10^{-5}$.

For each volume imbalance regime, there is a gravitational pull towards zero inventory because $v_{\omega,q}$ achieves its maximum at $q = 0$. However, the optimal continuation values $v_{\omega,q}$ as a function of inventory q differs based on the volume imbalance regimes, i.e., there is asymmetry in the volume imbalance regimes. This asymmetry is key for quote-based manipulation to arise. Specifically, consider AMZN at $\omega = SH$ with $q = 1$, and focus only on the expected discounted future payoffs. If the market maker tries to revert to zero inventory with a sell limit order, then the expected discounted future payoffs is an expectation across the optimal continuation values $v_{\omega,q}$ for all regimes and for inventory levels $q = 0$ and $q = 1$. On the other hand, if the market maker uses a manipulative order (a large buy limit order), then the expected

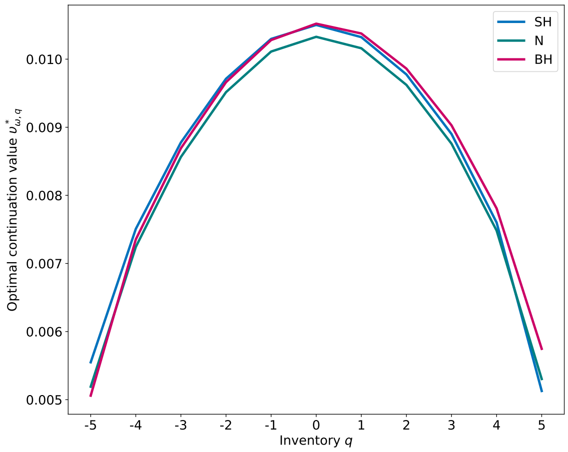

(a) AMZN

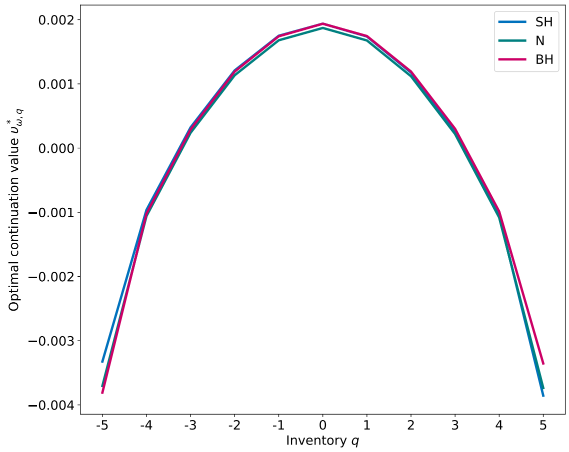

(b) CSCO

Figure 2.3: Optimal continuation values $v_{\omega,q}$.

discounted future payoffs is an expectation across the optimal continuation values $v_{\omega,q}$ for buy-heavy and for inventory levels $q = 1$ and $q = 2$. This difference in the expected discounted future payoffs induced by the asymmetry in the volume imbalance regimes is a key factor that makes a manipulative strategy dynamically optimal. Specifically, if this difference outweighs the gravitational pull to zero inventory and the immediate payoffs, then a manipulative strategy becomes dynamically optimal.

The concavity of these curves decreases as the value of the inventory aversion parameter α decreases. A decrease in the concavity increases the difference in the expected discounted future payoffs induced by the asymmetry; hence, quote-based manipulation is more likely to become dynamically optimal as the market maker becomes more tolerant to bearing inventory risk.

Spoofing and Fill Preferences Counter-intuitive to the motivation to manipulate the book, we find that getting caught out with a manipulative order is not always suboptimal. Indeed, there are situations in which manipulation occurs with the preference for the manipulative order to get filled. To show this, we analyze if the market maker prefers that her manipulative order is caught out, or if she prefers that her manipulative order does not get caught out. We compute

$$\begin{aligned} v_s(LLB, \text{filled}) &= E_{\sigma^*}\left[\sum_{t=0}^{\infty} \delta^t\, u(s_t, a_t, s_{t+1}) \,\middle|\, s_0 = s, a_0 = LLB, q_1 = q_0 + 1\right], \\ v_s(LLB, \text{not filled}) &= E_{\sigma^*}\left[\sum_{t=0}^{\infty} \delta^t\, u(s_t, a_t, s_{t+1}) \,\middle|\, s_0 = s, a_0 = LLB, q_1 = q_0\right], \end{aligned} \tag{2.11}$$

which corresponds to the expected stream of discounted payoffs conditional on a manipulative order getting filled or not for $q > 0$. If $v_s(LLB, \text{filled}) > v_s(LLB, \text{not filled})$, then the market maker prefers that her manipulative order is caught out. Conversely, if $v_s(LLB, \text{filled}) < v_s(LLB, \text{not filled})$, then the market prefers that her manipulative order does not get caught out.

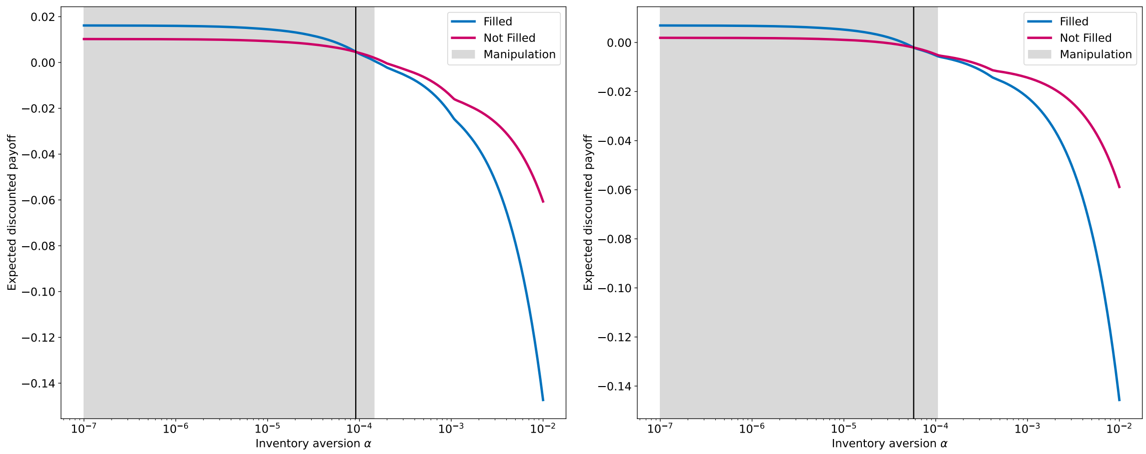

(a) AMZN (b) CSCO

Figure 2.4: Expected stream of discounted payoffs when the manipulative order is filled or not for $s = (SH, q = 2)$.

Figure 2.4 illustrates this preference. The shaded region denotes the values of the inventory aversion parameter where manipulation is optimal, i.e., $I_1(s) \cup I_2(s)$. The vertical black line denotes the cutoff value of α for which the preference switches. Within the shaded area, when α is to the left of the vertical line, the market maker prefers that her manipulative order is caught out; and to the right of the vertical line, the market prefers that her manipulative order does not get caught out.

A market maker who prefers a fill of her manipulative order may seem counter-intuitive when the motivation to manipulate the book is to manage inventory risk and revert to the preferred inventory position. However, managing inventory risk is only one part of the optimization problem to determine if manipulation is optimal. The market maker may prefer that her manipulative order is filled because the manipulative order increases the probability to complete a round-trip trade; i.e., the additional profit from the round-trip trade outweighs the costs to manage inventory risk.

Based on the fill preferences, we further refine quote-based manipulation into two forms: manipulation for a round-trip trade where manipulation occurs with the preference that the manipulative order is filled, and spoofing where manipulation occurs with the preference that the manipulative order is not filled. This refinement is different from type I and type II manipulation, which determines the optimality of the manipulation sequence. The refinement of fill preferences allow us to determine the intention behind a market makers manipulative order, which allows us to make the distinction between spoofing and manipulation for a round-trip trade.

One might argue that manipulation for a round-trip trade is not market manipulation but rather a byproduct of making markets. In our model, we assume (based on empirical results) that at most one unit of the large limit order can be executed. We make this assumption so that the market maker has the same expected one-step utility if she sends a limit order for one unit. Therefore, if she were to make markets without manipulating the book, then she would not use a large limit order.

2.5.2 Model Parameters and Manipulation

Building on the insights behind the driving forces of quote-based manipulation, we formalize how parameters of the model affect an algorithm's ability to learn to manipulate the order book. The results rely on Proposition 2 so that if $\alpha \notin I'(s)$, then an algorithm cannot learn a strategy that manipulates the order book in state s.

Inventory Aversion. The following proposition shows that if the market maker is sufficiently averse to holding larger levels of inventory, then their algorithm will not learn to manipulate the order book.

Proposition 3. *Let Assumption 2 hold and let $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If the market maker's inventory aversion parameter is such that*

$$\alpha > \max_{\omega \in \Omega,\, q \in \{-1,1\}} \{\bar{\alpha}_1(\omega, q)\}\,, \tag{2.12}$$

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

For a fixed volume imbalance regime ω, the upper bound of the cutoff $\bar{\alpha}_1(\omega, q)$ decreases monotonically as the absolute value of q increases. When the value of the inventory aversion parameter α satisfies (2.12), then $\alpha \notin I'(s)$ for all states $s = (\omega, q)$ where $q \neq 0$, and the algorithm will not learn a manipulative strategy for all $q \neq 0$.

Conversely, if the value of the inventory aversion parameter does not satisfy the inequality in (2.12), then $\alpha \in I'(s)$ for some states $s = (\omega, q)$ where $q \neq 0$. Here, there is a possibility that an algorithm will learn a manipulative strategy, but it is not guaranteed.

The result is intuitive if we consider the factors that make a manipulative strategy dynamically optimal. When initiating the manipulative sequence with a large limit order, there is a possibility that the large limit order is filled. Therefore, if the market maker is sufficiently averse to holding larger levels of inventory, then the cost associated with a manipulative order getting filled is too high for manipulation to be optimal; thus, the algorithm will not learn to manipulate the order book.

Quoted Spread. The following proposition shows how the expected quoted spread affects an algorithm's ability to learn to manipulate the order book.

Proposition 4. *Let Assumption 2 hold and let $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If the expected quoted spread $\vartheta \to 0$, then the algorithm will not learn to manipulate the order book for any state $s = (\omega, q)$ where $q \neq 0$.*

The result follows because $\bar{\alpha}_1(\omega, q) \to 0$ for all $\omega \in \Omega$ and $q \neq 0$ as $\vartheta \to 0$. This ensures that for $\alpha > 0$, we have $\alpha \notin I'(s)$ for all states $s = (\omega, q)$ where $q \neq 0$, so the algorithm will not learn a manipulative strategy.

Theory shows that the quoted bid-ask spread will not be zero even if the tick size is zero (i.e., $\varphi = 0$) because market makers must recover inventory costs (e.g., [71, 54]), and losses due to asymmetric information (e.g., [38, 47, 46]). Nonetheless, the proposition demonstrates the relationship between the expected quoted spread and quote-based manipulation. As the expected quoted spread decreases, the range of values of the inventory aversion parameter for which manipulation is optimal decreases. Conversely, as the expected quoted spread increases, the range of values of the inventory aversion parameter for which manipulation is optimal increases.

The result is also intuitive when one analyzes the factors that make manipulation optimal. First, as the expected quoted spread decreases, the gains from using limit orders and the costs from using market orders become negligible, so it is more efficient to revert to the preferred inventory position using market orders because they guarantee execution. Therefore, the uncertainty of a limit order execution and the possibility that a manipulative order is caught out makes a manipulative strategy suboptimal. Finally, as the expected quoted spread decreases, the expected profit from round-trip trades also decreases. The decrease in the profits does not outweigh the costs required to manage the inventory risk, and hence a manipulative strategy is suboptimal.

2.6 Discussion

Our analysis focuses on when an algorithm learns to create misleading information to buy or to sell an asset with a higher probability than was otherwise likely to occur, and spoofing is a special case when the preference is for the manipulative order not to be filled. In both types of manipulation, the manipulative step is to submit a large quantity of limit orders to mislead other market participants who react to the misleading signal, so the manipulator benefits from this manipulation. Indeed, this manipulative step is consistent with what is considered illegal in Article 12(2)(c) of Regulation (EU) No 596/2014 and Section 9(a)(2) of the Securities Exchange Act of 1934.[16,17] Therefore, our results can help identify limit order books in both the European Union and US securities exchanges where quote-based manipulation is likely to occur.

In this chapter, we define spoofing as a manipulative sequence in Definition 1 with the preference for the manipulative order not to be filled. Another definition of spoofing is given by the [39] Act, which defines spoofing as bidding or offering with the intent to cancel the bid or offer before execution. Our definition encapsulates the spirit of the Dodd–Frank definition by including the preference not to get caught out with the manipulative order, while also having greater reach.[18] For example, manipulative orders with a time-in-force achieve the same effect as cancellations, but the Dodd–Frank definition will not cover this case because there are no cancellations.

[16] Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02014R0596-20210101 and https://www.govinfo.gov/content/pkg/COMPS-1885/pdf/COMPS-1885.pdf, respectively.

[17] It is worth noting that Section 9(a)(2) of the Securities Exchange Act of 1934 refers to trade-based manipulation, whereas spoofing is a form of quote-based manipulation. However, existing case law roughly accomplishes the goal of making spoofing illegal, but it lacks clarity and makes errors in its reach.

[18] Our requirement to have the preference not to get caught with the manipulative order is restrictive, but it is no more restrictive than proving intent, which has its own set of issues. For example, it is the main reason behind why tacit collusion is legal [see pp. 339–340 of 50]. Nevertheless, showing a preference is not required to outlaw spoofing in securities exchanges.

Another shortfall of the Dodd–Frank definition is its specific focus on spoofing, which is only a particular case of quote-based manipulation. This narrow focus fails to prohibit other forms of quote-based manipulation.

Our model predicts simple manipulative sequences that are easy to identify. Regulators could monitor the probabilities p^a_ω and p^b_ω (following for instance our estimation procedure given in the next Chapter) across different LOBs to try to flag the marketplaces where manipulation is more likely to appear. Intuitively, higher p^b_{SH} would make action LLB to be more likely to appear at book regime SH, and similarly higher p^a_{BH} would make action LLS to be more likely to appear at book regime BH. However, in practice, identifying and detecting manipulative sequences is not straightforward. For example, market fragmentation allows for cross market manipulation, so our sequences need not appear within the same order book. Moreover, the interaction of multiple market makers makes it more difficult to detect these sequences because the market makers can coordinate and ride the manipulative sequences of each other. Nonetheless, a straightforward mechanical artifact of quote-based manipulation is that the volatility of the microprice (volume weighted midprice) increases as quote-based manipulation increases; how one establishes a counterfactual baseline without quote-based manipulation remains unclear. Finally, although we focused on quote-based manipulation, our approach can be used to understand analytically other forms of unintended market manipulation from learning algorithms such as layering or quote stuffing.

2.7 Proofs

Define the action value as

$$v_s(a) = E_{\sigma^*}\left[\sum_{t=0}^{\infty} \delta^t\, u(s_t, a_t, s_{t+1}) \,\middle|\, {}_0 =, a_0 = a\right] \tag{2.13}$$

$$= \bar{u}(s,a) + \delta \sum_{\omega' \in \Omega} p(\omega'|\omega, a) \sum_{q' \in \mathcal{Q}} p(q'|q, a)\, v^*_{\omega', q'}\,, \tag{2.14}$$

which is the expected stream of discounted payoffs from playing action a in state s then playing optimally thereafter according to an optimal strategy σ^*.

Proof of Lemma 1. Order the set $\mathcal{Q}$ so that $\mathcal{Q} = \{-\bar{q}, ..., -2, -1, 0, 1, 2, ..., \bar{q}\}$. To keep notation simple, let $A_s = A'$ denote the set of all actions for all $s \in S$. The result continues to hold without this simplification (see [68], pp. 108).[19] We have the following observations:

(i) For $q \geq 0$, $\bar{u}(\omega, q, a)$ is non-increasing in q; i.e., $\bar{u}(\omega, q, a) \leq \bar{u}(\omega, q', a)$ for all $\omega \in \Omega$ and for all $a \in A'$ when $0 \leq q' \leq q$;

(ii) For $q \leq 0$, $\bar{u}(\omega, q, a)$ is non-decreasing in q; i.e., $\bar{u}(\omega, q', a) \leq \bar{u}(\omega, q, a)$ for all $\omega \in \Omega$ and for all $a \in A'$ when $q' \leq q \leq 0$; and

(iii) $\sum_{j=k}^{\infty} p(j|q, a)$ is non-decreasing in q for all $k \in \mathcal{Q}$ and for all $a \in A'$, where $p(j|q, a)$ is the transition probability from inventory level q to inventory level j under action a.[20]

Consider the finite-horizon version of our model up to period N. Let $\upsilon_s^*(t)$ denote the optimal continuation value in state s at period t and adopt the convention that the terminal payoff $\bar{u}_N(\omega, q) = \bar{u}(\omega, q, DN)$.

Claim 1. *For $t = 0, 1, 2, ..., N$, we have that*

$$\upsilon_{\omega,q}^*(t) = \max_{a \in A'} \left\{ \bar{u}(\omega, q, a) + \delta \sum_{\omega'} p(\omega'|\omega, a) \sum_{j=0}^{\infty} p(j|q, a)\, \upsilon_{\omega',j}^*(t+1) \right\}$$

is non-increasing in q when $q \geq 0$, so $\upsilon_{\omega,q}^*(t) \leq \upsilon_{\omega,q'}^*(t)$ for all $\omega \in \Omega$ and for all $t = 0, 1, 2, ..., N$ when $0 \leq q' \leq q$. Similarly, when $q \leq 0$, $\upsilon_{\omega,q}^*(t)$ is non-decreasing in q, i.e., $\upsilon_{\omega,q'}^*(t) \leq \upsilon_{\omega,q}^*(t)$ for all $\omega \in \Omega$ and for all $t = 0, 1, 2, ..., N$ when $q' \leq q \leq 0$.

Proof of Claim 1. The claim follows from a straightforward modification of Proposition 4.7.3 in [68] using backwards induction. Consider the case $q \geq 0$. First, the result holds for $t = N$ from observation (i) because $\upsilon_{\omega,q}^*(N) = \bar{u}_N(\omega, q)$. Next, assume (for induction) that when $0 \leq q' \leq q$, we have $\upsilon_{\omega,q}^*(t) \leq \upsilon_{\omega,q'}^*(t)$ for all $\omega \in \Omega$ and for all $t = n+1, ..., N$. By Proposition 4.4.3 in [68], there exists $a^* \in A'$ so that

$$\upsilon_{\omega,q}^*(t) = \bar{u}(\omega, q, a^*) + \delta \sum_{\omega'} p(\omega'|\omega, a^*) \sum_{j=0}^{\infty} p(j|q, a^*)\, \upsilon_{\omega',j}^*(t+1)\,.$$

Let $0 \leq q' \leq q$. Use observations (i) and (iii), the induction hypothesis, and Lemma 4.7.2 in [68] to write

$$\upsilon_{\omega,q}^*(t) \leq \bar{u}(\omega, q', a^*) + \delta \sum_{\omega'} p(\omega'|\omega, a^*) \sum_{j=0}^{\infty} p(j|q', a^*)\, \upsilon_{\omega',j}^*(t+1)$$

[19] Indeed, the three conditions laid out on pp. 108 are satisfied with an appropriate adjustment to the ordering of the set $\mathcal{Q}$ and the proof follows with some minor adjustments.

[20] We adopt the convention that $p(j|q, a) = 0$ when $j \notin \mathcal{Q}$.

$$\leq \max_{a \in A'} \left\{ \bar{u}(\omega, q', a) + \delta \sum_{\omega'} p(\omega'|\omega, a) \sum_{j=0}^{\infty} p(j|q', a)\, \upsilon^*_{\omega',j}(t+1) \right\} = \upsilon^*_{\omega,q'}(t)\,.$$

When $q \leq 0$, similar calculations hold with observations (ii) and (iii). Thus, the claim follows. □

Finally, the pointwise limit (as $N \to \infty$) of non-increasing functions is non-increasing, and the pointwise limit (as $N \to \infty$) of non-decreasing functions is non-decreasing. Hence, $\upsilon^*_{\omega,q}(t)$ is non-increasing in q for all t when $q \geq 0$, and $\upsilon^*_{\omega,q}(t)$ is non-decreasing in q for all t when $q \leq 0$, so the lemma follows. □

Proof of Lemma 2. The optimal action a^* in state $s = (\omega, q)$ maximizes the action value in (2.13). To show that an action a is not optimal in state s, we show that the action is strictly dominated by another action a', i.e., $\upsilon_s(a) < \upsilon_s(a')$ for state s.

For $q > 0$, we first show that submitting a sell limit order LS always dominates do nothing DN. Specifically, we have

$$\begin{aligned}
\upsilon_s(LS) = p^a_\omega \left[\frac{\vartheta}{2} - \alpha\,(q-1)^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q-1} \right] &+ (1-p^a_\omega) \left[-\alpha\, q^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} \right] \\
> -\alpha\, q^2 + p^a_\omega\, \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q-1} &+ (1-p^a_\omega)\, \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} \\
\geq -\alpha\, q^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} &= \upsilon_s(DN)\,,
\end{aligned}$$

where the last inequality follows from Lemma 1. Next, do nothing DN always dominates a buy market order MB because

$$\begin{aligned}
\upsilon_s(DN) &= -\alpha\, q^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} \\
&> -\frac{\vartheta}{2} - \alpha\,(q+1)^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q+1} = \upsilon_s(MB)\,,
\end{aligned}$$

where the inequality follows from Lemma 1. Thus, LS dominates DN, which dominates MB. Therefore, both DN and MB are not optimal for $q > 0$.

The same reasoning and calculations hold for $q < 0$. Thus, the result follows because a buy limit order always dominates do nothing, and do nothing always dominates a sell market order. □ □

Proof of Proposition 1. For $q > 0$, Lemma 2 ensures that the optimal action is never DN or MB. We consider the ten pairwise comparisons from the set of actions $\{MS, LS, LLS, LLB, LB\}$. For each (ω, q), there exists a unique value $\alpha_{a,a'}(\omega, q)$

where two action values $\upsilon_{\omega,q}(a)$ and $\upsilon_{\omega,q}(a')$ as a function of α intersect because the action values are linear in α.

To compare the pairwise actions excluding LS versus LLS and LB versus LLB, we choose a large value of α so that it is optimal to revert to zero inventory as fast as possible and then stop making markets at zero inventory. That is, MS is the optimal action for all pairs $(\omega, q > 0)$ and DN is the optimal action for all pairs $(\omega, q = 0)$, so it is easy to see that for $q > 0$ and all $\omega \in \Omega$:

1. There exists $\alpha_{LS,MS}(\omega, q)$ such that MS is preferred to LS if and only if $\alpha > \alpha_{LS,MS}(\omega, q)$.
2. There exists $\alpha_{LLS,MS}(\omega, q)$ such that MS is preferred to LLS if and only if $\alpha > \alpha_{LLS,MS}(\omega, q)$.
3. There exists $\alpha_{LLB,MS}(\omega, q)$ such that MS is preferred to LLB if and only if $\alpha > \alpha_{LLB,MS}(\omega, q)$.
4. There exists $\alpha_{LB,MS}(\omega, q)$ such that MS is preferred to LB if and only if $\alpha > \alpha_{LB,MS}(\omega, q)$.
5. There exists $\alpha_{LLB,LS}(\omega, q)$ such that LS is preferred to LLB if and only if $\alpha > \alpha_{LLB,LS}(\omega, q)$.
6. There exists $\alpha_{LB,LS}(\omega, q)$ such that LS is preferred to LB if and only if $\alpha > \alpha_{LB,LS}(\omega, q)$.
7. There exists $\alpha_{LLB,LLS}(\omega, q)$ such that LLS is preferred to LLB if and only if $\alpha > \alpha_{LLB,LLS}(\omega, q)$.
8. There exists $\alpha_{LB,LLS}(\omega, q)$ such that LLS is preferred to LB if and only if $\alpha > \alpha_{LB,LLS}(\omega, q)$.

For the final two comparisons, we consider a reduced action set without MS. We choose a large value of α so that it is optimal to revert to zero inventory as fast as possible and then stop making markets at zero inventory. For a large enough value of α, we have $\upsilon^*_{BH,q} > \upsilon^*_{N,q} > \upsilon^*_{SH,q}$ and $\upsilon^*_{BH,q-1} > \upsilon^*_{N,q-1} > \upsilon^*_{SH,q-1}$ because the probability of selling a limit order is highest (lowest) in BH (SH). The action values of LLS and LS (excluding the one-step utility) are

$$\upsilon_{\omega,q}(LLS) = p^a_\omega \left((1-\kappa)\, \upsilon^*_{SH,q-1} + \frac{\kappa}{2}\, \upsilon^*_{N,q-1} + \frac{\kappa}{2}\, \upsilon^*_{BH,q-1} \right)$$

$$+(1-p^a_\omega)\ \Big((1-\kappa)\,\upsilon^*_{SH,q}+\frac{\kappa}{2}\,\upsilon^*_{N,q}+\frac{\kappa}{2}\,\upsilon^*_{BH,q}\Big)$$

$$\upsilon_{\omega,q}(LS)=p^a_\omega\ \big(p_{SH|\omega}\,\upsilon^*_{SH,q-1}+p_{N|\omega}\,\upsilon^*_{N,q-1}+p_{BH|\omega}\,\upsilon^*_{BH,q-1}\big)$$

$$+(1-p^a_\omega)\ \big(p_{SH|\omega}\,\upsilon^*_{SH,q}+p_{N|\omega}\,\upsilon^*_{N,q}+p_{BH|\omega}\,\upsilon^*_{BH,q}\big)\ ,$$

respectively. Therefore, $\upsilon_{\omega,q}(LLS)<\upsilon_{\omega,q}(LS)$ if and only if

$$(1-\kappa-p_{SH|\omega})\,(p^a_\omega\,\upsilon^*_{SH,q-1}+(1-p^a_\omega)\,\upsilon^*_{SH,q})$$

$$<\ p^a_\omega\ \Big((p_{N|\omega}-\frac{\kappa}{2})\,\upsilon^*_{N,q-1}+(p_{BH|\omega}-\frac{\kappa}{2})\,\upsilon^*_{BH,q-1}\Big)$$

$$+(1-p^a_\omega)\ \Big((p_{N|\omega}-\frac{\kappa}{2})\,\upsilon^*_{N,q}+(p_{BH|\omega}-\frac{\kappa}{2})\,\upsilon^*_{BH,q}\Big)\ ,$$

which follows because

$$\Big(p_{N|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{N,q-1}+\Big(p_{BH|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{BH,q-1}>(p_{N|\omega}+p_{BH|\omega}-\kappa)\,\upsilon^*_{N,q-1}$$

$$=(1-\kappa-p_{SH|\omega})\,\upsilon^*_{N,q-1}>(1-\kappa-p_{SH|\omega})\,\upsilon^*_{SH,q-1}$$

and

$$\Big(p_{N|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{N,q}+\Big(p_{BH|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{BH,q}>(p_{N|\omega}+p_{BH|\omega}-\kappa)\,\upsilon^*_{N,q}$$

$$=(1-\kappa-p_{SH|\omega})\,\upsilon^*_{N,q}>(1-\kappa-p_{SH|\omega})\,\upsilon^*_{SH,q}\,.$$

Hence, LS is preferred to LLS for a large enough value of α. Next, we compare LLB and LB. The action values of LLB and LB (excluding the one-step utility) are

$$\upsilon_{\omega,q}(LLB)=p^b_\omega\ \Big((1-\kappa)\,\upsilon^*_{BH,q+1}+\frac{\kappa}{2}\,\upsilon^*_{N,q+1}+\frac{\kappa}{2}\,\upsilon^*_{SH,q+1}\Big)$$

$$+(1-p^b_\omega)\ \Big((1-\kappa)\,\upsilon^*_{BH,q}+\frac{\kappa}{2}\,\upsilon^*_{N,q}+\frac{\kappa}{2}\,\upsilon^*_{SH,q}\Big)$$

$$\upsilon_{\omega,q}(LB)=p^b_\omega\ \big(p_{BH|\omega}\,\upsilon^*_{BH,q+1}+p_{N|\omega}\,\upsilon^*_{N,q+1}+p_{SH|\omega}\,\upsilon^*_{SH,q+1}\big)$$

$$+(1-p^b_\omega)\ \big(p_{BH|\omega}\,\upsilon^*_{BH,q}+p_{N|\omega}\,\upsilon^*_{N,q}+p_{SH|\omega}\,\upsilon^*_{SH,q}\big)\ ,$$

respectively. Thus, $\upsilon_{\omega,q}(LLB)>\upsilon_{\omega,q}(LB)$ if and only if

$$(1-\kappa-p_{BH|\omega})\,(p^b_\omega\,\upsilon^*_{BH,q+1}+(1-p^b_\omega)\,\upsilon^*_{BH,q})$$

$$>p^b_\omega\ \Big((p_{N|\omega}-\frac{\kappa}{2})\,\upsilon^*_{N,q+1}+(p_{SH|\omega}-\frac{\kappa}{2})\,\upsilon^*_{SH,q+1}\Big)$$

$$+(1-p^b_\omega)\ \Big((p_{N|\omega}-\frac{\kappa}{2})\,\upsilon^*_{N,q}+(p_{SH|\omega}-\frac{\kappa}{2})\,\upsilon^*_{SH,q}\Big)\ ,$$

which follows because

$$\Big(p_{N|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{N,q+1}+\Big(p_{SH|\omega}-\frac{\kappa}{2}\Big)\ \upsilon^*_{SH,q+1}<(p_{N|\omega}+p_{SH|\omega}-\kappa)\,\upsilon^*_{N,q+1}$$

$$= (1 - \kappa - p_{BH|\omega})\, \upsilon^*_{N,q+1} < (1 - \kappa - p_{BH|\omega})\, \upsilon^*_{BH,q+1}$$

and

$$\left(p_{N|\omega} - \frac{\kappa}{2}\right) \upsilon^*_{N,q} + \left(p_{SH|\omega} - \frac{\kappa}{2}\right) \upsilon^*_{SH,q} < (p_{N|\omega} + p_{SH|\omega} - \kappa)\, \upsilon^*_{N,q}$$
$$= (1 - \kappa - p_{BH|\omega})\, \upsilon^*_{N,q} < (1 - \kappa - p_{BH|\omega})\, \upsilon^*_{BH,q} \,.$$

With the ten pairwise preferences for $q > 0$, we have that Figure 2.1 is the only (non-contradictory) ordering of action preferences and cutoffs. The same reasoning holds for $q < 0$. □ □

Lemma 4. *Let $0 < q' \leq q$. We have that $\upsilon^*_{\omega,q} - \upsilon^*_{\omega,q'}$ converges to zero from below as $\alpha \to 0$ for all $\omega \in \Omega$. Similarly, $\upsilon^*_{\omega,q'} - \upsilon^*_{\omega,q}$ converges to zero from above as $\alpha \to 0$.*

Proof of Lemma 4. The result follows from adapting the proof of Lemma 1 to show that if $\alpha = 0$ then $\upsilon^*_{\omega,q} - \upsilon^*_{\omega,q'} = 0$. First, observe that if $\alpha = 0$, then $\bar{u}(\omega, q, a)$ is constant in q when $q \geq 0$ for all $\omega \in \Omega$ and for all $a \in A'$. Following the proof of Lemma 1, we have that the optimal continuation value $\upsilon^*_{\omega,q}$ is both non-increasing and non-decreasing in q when $q \geq 0$. Therefore, $\upsilon^*_{\omega,q} - \upsilon^*_{\omega,q'} = 0$ for all $q, q' \geq 0$, and hence, the result follows as a consequence of Lemma 1. □ □

Lemma 5. *For $q > 0$, there exists $\alpha > 0$ such that $MS \prec LS$.*

Proof of Lemma 5. The action values of LS and MS are given by

$$\upsilon_s(LS) = p^a_\omega \left[\frac{\vartheta}{2} - \alpha\,(q-1)^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q-1}\right] + (1 - p^a_\omega) \left[-\alpha\, q^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q}\right],$$

$$\upsilon_s(MS) = -\frac{\vartheta}{2} - \alpha\,(q-1)^2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q-1}\,,$$

and their difference is

$$\upsilon_s(LS) - \upsilon_s(MS) = (1 + p^a_\omega)\, \vartheta/2 - \alpha\,(1 - p^a_\omega)\,(2q-1) + \delta\,(1 - p^a_\omega) \sum_{\omega'} p_{\omega'|\omega} \left(\upsilon^*_{\omega,q} - \upsilon^*_{\omega,q-1}\right)\,.$$

From Lemma 4, we have that $\upsilon_s(LS) - \upsilon_s(MS) \to (1 + p^a_\omega)\, \vartheta/2 > 0$ as $\alpha \to 0$. Thus, the claim follows because there are values of $\alpha > 0$ for which $MS \prec LS$. □

Lemma 6. *Let $\kappa = 0$, $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If condition* (C1) *holds, then action LS dominates action LB at state (BH, q) for $q > 0$. Moreover, action LLB cannot be optimal at state (BH, q) for $q > 0$. Similarly, if condition* (C2) *holds, then action LB dominates action LS at state (SH, q) for $q < 0$. Moreover, action LLS cannot be optimal at state (SH, q) for $q < 0$.*

Proof of Lemma 6. We focus on $q > 0$. Since $p^a_{BH} > p^b_{BH}$ and lower levels of inventory imply less penalty, we have that $LB \prec LS$. Moreover, the action value of LLB is

$$\upsilon_s(LLB) = p^b_\omega \left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}\right] + (1-p^b_\omega)\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right].$$

We now show that if condition (C1) holds, then LLB is never optimal in $s = (BH, q)$ where $q > 0$. We proceed by contradiction. Suppose LLB is optimal in $s = (BH, q)$ where $q > 0$. This implies the following claim is true:

Claim 2. *If LLB is optimal in $s = (BH, q)$ where $q > 0$, then $\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1} \geq -\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}$.*

Proof of Claim 2. We prove this claim by contradiction. Suppose that $-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q} > \vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}$, then we have

$$\begin{aligned}\upsilon_{BH,q}(LLB) &= p^b_{BH}\left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}\right] + (1-p^b_{BH})\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right]\\ &\leq p^b_{BH}\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right] + (1-p^b_{BH})\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right] = -\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\,,\end{aligned}$$

which follows from Lemma 1. However, $\upsilon^*_{BH,q} = \upsilon_{BH,q}(LLB)$ because LLB is optimal by assumption. Therefore, the inequality above becomes $\upsilon^*_{BH,q} \leq -\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}$, which implies

$$\upsilon^*_{BH,q} \leq -\frac{\alpha\,q^2}{1-\delta}\,.$$

Now, consider a suboptimal strategy σ that does nothing in all $\omega \in \Omega$ at inventory level q. The value of this strategy in $s = (BH, q)$ where $q > 0$ is given by

$$\upsilon_{BH,q}(\sigma) = -\frac{\alpha\,q^2}{1-\delta}\,.$$

Therefore, $\upsilon_{BH,q}(\sigma) \geq \upsilon^*_{BH,q}$ is a contradiction because the strategy σ is suboptimal as a consequence of Lemma 2. Hence, the claim follows. □

Next, the claim implies that $\upsilon_{SH,q}(LLB) \geq \upsilon_{BH,q}(LLB)$ because

$$\begin{aligned}\upsilon_{SH,q}(LLB) &= p^b_{SH}\left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}\right] + (1-p^b_{SH})\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right]\\ &\geq p^b_{BH}\left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}\right] + (1-p^b_{BH})\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right]\\ &= \upsilon_{BH,q}(LLB)\,,\end{aligned}$$

where the inequality follows as a result of $p^b_{SH} > p^b_{BH}$.
Use $1 - p^b_{BH} = (1-p^a_{BH}) + (p^a_{BH} - p^b_{BH})$ in the action value of LLB in $s = (BH, q)$ to obtain

$$\upsilon_{BH,q}(LLB) = p^b_{BH}\left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\,\upsilon^*_{BH,q+1}\right] + (1-p^b_{BH})\left[-\alpha\,q^2 + \delta\,\upsilon^*_{BH,q}\right]$$

$$
\begin{aligned}
&= p^b_{BH}\ \left[\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}\right] + (p^a_{BH} - p^b_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right] \\
&\qquad + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right] \\
&\le p^b_{BH} \max\left\{\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}, -\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right\} \\
&\quad + (p^a_{BH} - p^b_{BH}) \max\left\{\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}, -\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right\} \\
&\qquad + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right] \\
&= p^a_{BH} \max\left\{\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}, -\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right\} \\
&\qquad + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right]\ ,
\end{aligned}
$$

where the inequality above follows as a result of $p^b_{BH} < p^a_{BH}$ from condition (C1) and because $x \le \max\{x, y\}$ and $y \le \max\{x, y\}$ for all $x, y \in \mathbb{R}$.

Next, observe that both $\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}$ and $-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}$ are less than $\vartheta/2 - \alpha\,(q-1)^2 + \delta\, \upsilon^*_{SH,q-1}$ because $\upsilon^*_{SH,q-1} \ge \upsilon^*_{SH,q} \ge \upsilon_{SH,q}(LLB) \ge \upsilon_{BH,q}(LLB) = \upsilon^*_{BH,q} \ge \upsilon^*_{BH,q+1}$. Therefore,

$$
\begin{aligned}
\upsilon_{BH,q}(LLB) &\le p^a_{BH} \max\left\{\vartheta/2 - \alpha\,(q+1)^2 + \delta\, \upsilon^*_{BH,q+1}, -\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right\} \\
&\qquad + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right] \\
&\le p^a_{BH}\left[\vartheta/2 - \alpha\,(q-1)^2 + \delta\, \upsilon^*_{SH,q-1}\right] + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{BH,q}\right] \\
&\le p^a_{BH}\left[\vartheta/2 - \alpha\,(q-1)^2 + \delta\, \upsilon^*_{SH,q-1}\right] + (1 - p^a_{BH})\left[-\alpha\, q^2 + \delta\, \upsilon^*_{SH,q}\right] \\
&= \upsilon_{BH,q}(LLS)\,,
\end{aligned}
$$

where the last inequality follows from $\upsilon^*_{SH,q} \ge \upsilon_{SH,q}(LLB) \ge \upsilon_{BH,q}(LLB) = \upsilon^*_{BH,q}$. This implies that $\upsilon_{BH,q}(LLB) \le \upsilon_{BH,q}(LLS)$. Hence, we have a contradiction as LLB cannot be optimal in $s = (BH, q)$ where $q > 0$.

An analogous reasoning holds for $q < 0$. □ □

Lemma 7. *Let $\kappa = 0$, $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If condition* (C2) *holds, then $\alpha_1(\omega, q) > 0$ for all $q > 0$ and $\omega = SH$. Similarly, if condition* (C1) *holds, then $\alpha_1(\omega, q) > 0$ for all $q < 0$ and $\omega = BH$.*

Proof of Lemma 7. We focus on $q > 0$. We require the following claim.

Claim 3. *Let $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. If condition* (C2) *holds, then there exists $\alpha > 0$ such that $LS \prec LB$ for all $q > 0$ and $\omega = SH$. Therefore, for this $\alpha > 0$, LS is not optimal.*

Proof of Claim 3. The action values of LS and LB are given by

$$\begin{aligned} v_s(LS) &= p^a_\omega\,\vartheta/2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} + \delta\, p^a_\omega \sum_{\omega'} p_{\omega'|\omega}\, \left(\upsilon^*_{\omega,q-1} - \upsilon^*_{\omega,q}\right) \\ &\quad -\alpha\, p^a_\omega (q-1)^2 - \alpha(1-p^a_\omega) q^2 \\ v_s(LB) &= p^b_\omega\,\vartheta/2 + \delta \sum_{\omega'} p_{\omega'|\omega}\, \upsilon^*_{\omega,q} + \delta\, p^b_\omega \sum_{\omega'} p_{\omega'|\omega}\, \left(\upsilon^*_{\omega,q+1} - \upsilon^*_{\omega,q}\right) \\ &\quad -\alpha\, p^b_\omega (q+1)^2 - \alpha(1-p^b_\omega) q^2 \,. \end{aligned}$$

As $\alpha \to 0$, $v_s(LB) - v_s(LS) \to \left(p^b_\omega - p^a_\omega\right)\vartheta/2$ as a consequence of Lemma 4. Therefore, $LS \prec LB$ in $\omega = SH$ because of $p^a_{SH} < p^b_{SH}$ from condition (C2). Hence, LS cannot be optimal because $LS \prec LB$. □

From Claim 3, there exists $\alpha > 0$ such that $LS \prec LB$ so that LS is not optimal. We set the value of α so that MS is not optimal. For this value of α, either LB, LLB or LLS is the optimal action. If LB or LLB is the optimal action, then we have $\alpha_1(\omega, q) > 0$. To complete the proof we see that LLS cannot be the optimal action for the value of α. We proceed by contradiction. Assume LLS is optimal and take α close to zero so that the effect of inventory is negligible when considering the optimal continuation value. We denote $\upsilon_\omega = \upsilon_{\omega,q} = \upsilon_{\omega,q'}$ to obtain

$$\begin{aligned} \upsilon_{SH}(LLB) &= p^b_{SH}(\vartheta/2 + \delta\, \upsilon^*_{BH}) + (1-p^b_{SH})\,\delta\, \upsilon^*_{BH} = p^b_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{BH} \\ &\geq p^b_{SH}\,\vartheta/2 + \delta\, \upsilon_{BH}(LLS) = p^b_{SH}\,\vartheta/2 + \delta\,(p^a_{BH}\,\vartheta/2 + \delta\, \upsilon^*_{SH}) \\ &> p^b_{SH}\,\vartheta/2 + \delta\,(p^a_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{SH})\,, \end{aligned}$$

where the last inequality follows from $p^a_{BH} > p^a_{SH}$.

Observe that

$$p^b_{SH}\,\vartheta/2 + \delta\,(p^a_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{SH}) = p^b_{SH}\,\vartheta/2 + \delta\, \upsilon_{SH}(LLS) = p^b_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{SH}\,,$$

where the last equality follows from assuming that LLS is optimal at SH. Therefore, we have

$$\upsilon_{SH}(LLB) > p^b_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{SH} > p^a_{SH}\,\vartheta/2 + \delta\, \upsilon^*_{SH} = \upsilon^*_{SH}\,,$$

which is a contradiction. Therefore, LLS is not optimal. The same reasoning holds for $q < 0$. □ □

Lemma 8. *Let $\kappa = 0$, $p^a_{SH} < p^a_N < p^a_{BH}$ and $p^b_{SH} > p^b_N > p^b_{BH}$ hold. For $\omega = N$, we have $\alpha_1(\omega, q) > 0$ for either $q > 0$ or $q < 0$.*

Proof of Lemma 8. From Lemma 4, we have

$$\upsilon_N(LLS) = p_N^a\,\vartheta/2 + \delta\,\upsilon_{SH} \quad \text{and} \quad \upsilon_N(LLB) = p_N^b\,\vartheta/2 + \delta\,\upsilon_{BH}\,,$$

for a small enough value of α so that if

$$p_N^a\,\vartheta/2 + \delta\,\upsilon_{SH} > p_N^b\,\vartheta/2 + \delta\,\upsilon_{BH}\,,$$

then $\alpha_1(\omega, q) > 0$ for $q < 0$. Similarly, if

$$p_N^a\,\vartheta/2 + \delta\,\upsilon_{SH} < p_N^b\,\vartheta/2 + \delta\,\upsilon_{BH}\,,$$

then $\alpha_1(\omega, q) > 0$ for $q > 0$. □ □

Lemma 9. *For κ satisfying condition* (C3)*, Lemma 6 continues to hold.*

Proof of Lemma 9. Since κ does not play a role in the comparison between actions LS and LB, we still have that action LS dominates action LB at state (BH, q) for $q > 0$. Suppose that LLB is optimal at $(BH, q > 0)$, then

$$\frac{\vartheta}{2} - \alpha\,(q+1)^2 + \delta\,((1-\kappa)\,\upsilon^*_{BH,q+1} + \frac{\kappa}{2}\,\upsilon^*_{N,q+1} + \frac{\kappa}{2}\,\upsilon^*_{SH,q+1})$$
$$> -\alpha\,q^2 + \delta\,((1-\kappa)\,\upsilon^*_{BH,q} + \frac{\kappa}{2}\,\upsilon^*_{N,q} + \frac{\kappa}{2}\,\upsilon^*_{SH,q})\,.$$

Suppose the inequality above is not true. We know $\max\{\upsilon^*_{BH,q}, \upsilon^*_{N,q}, \upsilon^*_{SH,q}\} = \upsilon^*_{BH,q}$ because if the maximum was $\upsilon^*_{SH,q}$, then $\upsilon_{BH,q}(LLS) > \upsilon_{BH,q}(LLB)$, which cannot be true assuming LLB is optimal at $(BH, q > 0)$. Therefore, if both $\upsilon^*_{BH,q}$ is the maximum and LLB is optimal at (BH, q), then staying at $(BH, q > 0)$ forever would be optimal. However, this would give the same payoff as doing nothing forever, which is never optimal. Therefore the previous inequality holds. From the previous inequality, we have $\upsilon_{SH,q}(LLB) > \upsilon_{BH,q}(LLB)$ because $p^b_{SH} > p^b_{BH}$. Therefore, $\upsilon^*_{BH,q} < \upsilon^*_{SH,q}$. Now, observe that

$$\upsilon_{BH,q}(LLB) = p^b_{BH}\,\left(\frac{\vartheta}{2} - \alpha\,(q+1)^2 + \delta\,\left((1-\kappa)\,\upsilon^*_{BH,q+1} + \frac{\kappa}{2}\,\upsilon^*_{N,q+1} + \frac{\kappa}{2}\,\upsilon^*_{SH,q+1}\right)\right)$$
$$+(1-p^b_{BH})\,\left(-\alpha\,q^2 + \delta\,\left((1-\kappa)\,\upsilon^*_{BH,q} + \frac{\kappa}{2}\,\upsilon^*_{N,q} + \frac{\kappa}{2}\,\upsilon^*_{SH,q}\right)\right)$$
$$= p^b_{BH}\,\left(\frac{\vartheta}{2} - \alpha\,(q+1)^2 + \delta\,\left((1-\kappa)\,\upsilon^*_{BH,q+1} + \frac{\kappa}{2}\,\upsilon^*_{N,q+1} + \frac{\kappa}{2}\,\upsilon^*_{SH,q+1}\right)\right)$$
$$+(p^a_{BH} - p^b_{BH})\,\left(-\alpha\,q^2 + \delta\,\left((1-\kappa)\,\upsilon^*_{BH,q} + \frac{\kappa}{2}\,\upsilon^*_{N,q} + \frac{\kappa}{2}\,\upsilon^*_{SH,q}\right)\right)$$

$$+(1-p^b_{BH})\left(-\alpha\, q^2+\delta\,\left((1-\kappa)\,\upsilon^*_{BH,q}+\frac{\kappa}{2}\,\upsilon^*_{N,q}+\frac{\kappa}{2}\,\upsilon^*_{SH,q}\right)\right)$$
$$\le p^a_{BH}\left(\frac{\vartheta}{2}-\alpha\,(q+1)^2+\delta\,\left((1-\kappa)\,\upsilon^*_{BH,q+1}+\frac{\kappa}{2}\,\upsilon^*_{N,q+1}+\frac{\kappa}{2}\,\upsilon^*_{SH,q+1}\right)\right)$$
$$+(1-p^a_{BH})\left(-\alpha\, q^2+\delta\,\left((1-\kappa)\,\upsilon^*_{BH,q}+\frac{\kappa}{2}\,\upsilon^*_{N,q}+\frac{\kappa}{2}\,\upsilon^*_{SH,q}\right)\right)$$
$$\le p^a_{BH}\left(\frac{\vartheta}{2}-\alpha\,(q-1)^2+\delta\,\left((1-\kappa)\,\upsilon^*_{BH,q-1}+\frac{\kappa}{2}\,\upsilon^*_{N,q-1}+\frac{\kappa}{2}\,\upsilon^*_{SH,q-1}\right)\right)$$
$$+(1-p^a_{BH})\left(-\alpha\, q^2+\delta\,\left((1-\kappa)\,\upsilon^*_{BH,q}+\frac{\kappa}{2}\,\upsilon^*_{N,q}+\frac{\kappa}{2}\,\upsilon^*_{SH,q}\right)\right)=\upsilon_{BH,q}(LLS)\,.$$

Hence, a contradiction, so LLB is never optimal at $(BH, q > 0)$. The same reasoning holds for $q < 0$. □ □

Lemma 10. *For κ satisfying condition* (C3)*, Lemma 7 continues to hold.*

Proof of Lemma 10. It is straightforward to see that Claim 3 still holds because the transition probabilities of large limit orders do no play a role in its proof. To show that Lemma 7 continues to hold, we prove that LLS is not optimal in $(SH, q > 0)$ for a small enough value of α. We proceed by contradiction. Assuming that LLS is optimal in $(SH, q > 0)$, for a small enough value of α, we have

$$\upsilon_{SH}(LLB)=p^b_{SH}\,(\frac{\vartheta}{2}+\delta\,((1-\kappa)\,\upsilon^*_{BH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{SH})$$
$$+(1-p^b_{SH})\,\delta\,((1-\kappa)\,\upsilon^*_{BH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{SH})$$
$$=p^b_{SH}\,\frac{\vartheta}{2}+\delta\,((1-\kappa)\,\upsilon^*_{BH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{SH})\,,$$

and

$$\upsilon_{SH}(LLS)=p^a_{SH}\,(\frac{\vartheta}{2}+\delta\,((1-\kappa)\,\upsilon^*_{SH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{BH})$$
$$+(1-p^a_{SH})\,\delta\,((1-\kappa)\,\upsilon^*_{SH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{BH})$$
$$=p^a_{SH}\,\frac{\vartheta}{2}+\delta\,((1-\kappa)\,\upsilon^*_{SH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{BH})$$

We have that $\upsilon^*_{BH} > \upsilon^*_{SH}$ because $p^a_{BH} > p^a_{SH}$ and there is less penalty at $q-1$ than at q and because we assume that LLS is optimal at $(SH, q > 0)$. Therefore, use $p^b_{SH} > p^a_{SH}$, $\upsilon^*_{BH} > \upsilon^*_{SH}$, and $(1-\kappa\,) > \kappa/2$ to obtain

$$p^b_{SH}\,\frac{\vartheta}{2}+\delta\,\left((1-\kappa)\,\upsilon^*_{BH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{SH}\right)$$
$$>p^a_{SH}\,\frac{\vartheta}{2}+\delta\,\left((1-\kappa)\,\upsilon^*_{SH}+\frac{\kappa}{2}\,\upsilon^*_{N}+\frac{\kappa}{2}\,\upsilon^*_{BH}\right)$$

so that $\upsilon_{SH}(LLB) > \upsilon_{SH}(LLS) = \upsilon^*_{SH}$, and therefore a contradiction. The same reasoning holds for $q < 0$. □

Remark 1. *For κ satisfying condition* (C3)*, Lemma 8 continues to hold because the transition probabilities of the large limit orders do not play a role in the proof of Lemma 8.*

Lemma 11. *Let $p^a_{SH} < p^a_N < p^a_{BH}$, $p^b_{SH} > p^b_N > p^b_{BH}$, and* (C3) *and* (C4) *hold. If $(p^b_N - p^a_N) > \frac{\delta(1-\kappa-\frac{\kappa}{2})}{1+\delta(1-\kappa-\frac{\kappa}{2})}(p^b_{SH} - p^a_{BH})$, then $\alpha_1(\omega, q) > 0$ for $q > 0$ and $\omega = N$. Similarly, if $(p^a_N - p^b_N) > \frac{\delta(1-\kappa-\frac{\kappa}{2})}{1+\delta(1-\kappa-\frac{\kappa}{2})}(p^a_{BH} - p^b_{SH})$, then $\alpha_1(\omega, q) > 0$ for $q < 0$ and $\omega = N$.*

Proof of Lemma 11. To prove the lemma, we first establish the following claim.

Claim 4. *If $p^a_{SH} < p^a_N < p^a_{BH}$, $p^b_{SH} > p^b_N > p^b_{BH}$, and* (C3) *and* (C4) *hold, then for a small enough value of α, $LS \prec LLS$ for $\omega = BH$, $LB \prec LLB$ for $\omega = SH$, and $LB \prec LLB$ and $LS \prec LLS$ for $\omega = N$.*

Proof of Claim 4. If (C4) holds, then the following hold:

$$p^a_{BH} - p^b_{SH} < (p^b_{SH} - \max\{p^a_N, p^b_N\})\frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}},$$

$$p^b_{SH} - p^a_{BH} < (p^a_{BH} - \max\{p^a_N, p^b_N\})\frac{p_{N|SH} - \frac{\kappa}{2}}{p_{SH|SH} - \frac{\kappa}{2}},$$

$$p^a_{BH} - p^b_{SH} < (p^b_{SH} - \max\{p^a_N, p^b_N\})\frac{p_{N|N} - \frac{\kappa}{2}}{p_{BH|N} - \frac{\kappa}{2}},$$

$$p^b_{SH} - p^a_{BH} < (p^a_{BH} - \max\{p^a_N, p^b_N\})\frac{p_{N|N} - \frac{\kappa}{2}}{p_{SH|N} - \frac{\kappa}{2}}.$$

We first focus on $\omega = BH$. For a small enough value of α, the inventory penalty is negligible and the optimal action is LS or LLS. Moreover, LLS is preferred to LS if and only if

$$(1-\kappa)\, v_{SH} + \frac{\kappa}{2}\, v_N + \frac{\kappa}{2}\, v_{BH} > p_{BH|BH}\, v_{BH} + p_{N|BH}\, v_N + p_{SH|BH}\, v_{SH}$$

If $\max\{v_{SH}, v_N, v_{BH}\} = v_{SH}$, then the last inequality trivially holds because of (C3). On the other hand, if $\max\{v_{SH}, v_N, v_{BH}\} = v_{BH}$[21], we proceed as follows: Note the last inequality holds if and only if

$$(1-\kappa-p_{SH|BH})\, v_{SH} > (p_{BH|BH} - \frac{\kappa}{2})\, v_{BH} + (p_{N|BH} - \frac{\kappa}{2})\, v_N \iff .$$

$$(1-\kappa-(1-p_{BH|BH}-p_{N|BH}))\, v_{SH} > (p_{BH|BH} - \frac{\kappa}{2})\, v_{BH} + (p_{N|BH} - \frac{\kappa}{2})\, v_N \iff .$$

[21] It is not possible for both $v^*_N > v^*_{SH}$ and $v^*_N > v^*_{BH}$ to hold because $p^a_N < p^a_{BH}$ and $p^b_N < p^b_{SH}$. This is enough to exclude the case $\max\{v_{SH}, v_N, v_{BH}\} = v_N$.

$$(p_{BH|BH} + p_{N|BH} - \frac{\kappa}{2} - \frac{\kappa}{2})\, v_{SH} > (p_{BH|BH} - \frac{\kappa}{2})\, v_{BH} + (p_{N|BH} - \frac{\kappa}{2})\, v_N \iff .$$

$$(p_{N|BH} - \frac{\kappa}{2})\,(v_{SH} - v_N) > (p_{BH|BH} - \frac{\kappa}{2})\,(v_{BH} - v_{SH}) \iff .$$

$$v_{BH} - v_{SH} < (v_{SH} - v_N)\, \frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}}\,.$$

where the last inequality follows from (C3). If $\max\{v_{SH}, v_N, v_{BH}\} = v_{BH}$, then $v_{BH} < p^a_{BH}\,\vartheta/2 + \delta\, v_{BH}$. Moreover, the optimal action in $\omega = SH$ is LLB when $\max\{v_{SH}, v_N, v_{BH}\} = v_{BH}$. Hence,

$$v_{BH} - v_{SH} < p^a_{BH}\,\vartheta/2 + \delta\, v_{BH} - p^b_{SH}\,\vartheta/2 - \delta\, v_{BH} = (p^a_{BH} - p^b_{SH})\,\vartheta/2\,.$$

And also,

$$v^*_{SH} - v^*_N = v_{SH}(LLB) - v^*_N > p^b_{SH}\,\frac{\vartheta}{2} + \delta\, v^*_{BH} - \max\{p^a_N, p^b_N\}\,\frac{\vartheta}{2} - \delta\, v^*_{BH}$$

$$= (p^b_{SH} - \max\{p^a_N, p^b_N\})\,\frac{\vartheta}{2}$$

because $v_N < \max\{p^a_N, p^b_N\}\vartheta/2 + \delta v_{BH}$. Therefore, if

$$p^a_{BH} - p^b_{SH} < (p^b_{SH} - \max\{p^a_N, p^b_N\})\,\frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}}\,,$$

then LLS is preferred to LS because

$$(v^*_{SH} - v^*_N)\,\frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}} > (p^b_{SH} - \max\{p^a_N, p^b_N\})\frac{\vartheta}{2}\,\frac{p_{N|BH} - \frac{\kappa}{2}}{p_{BH|BH} - \frac{\kappa}{2}}$$

$$> (p^a_{BH} - p^b_{SH})\,\frac{\vartheta}{2} > v^*_{BH} - v^*_{SH}\,.$$

For $\omega = SH$, we follow a similar reasoning so that if

$$p^b_{SH} - p^a_{BH} < (p^a_{BH} - \max\{p^a_N, p^b_N\})\,\frac{p_{N|SH} - \frac{\kappa}{2}}{p_{SH|SH} - \frac{\kappa}{2}}\,,$$

then LLB is preferred to LB in $\omega = SH$.

For $\omega = N$, we first compare LLS with LS. As before, LLS is preferred to LS if and only if

$$(1-\kappa)\, v^*_{SH} + \frac{\kappa}{2}\, v^*_N + \frac{\kappa}{2}\, v^*_{BH} > p_{BH|N}\, v^*_{BH} + p_{N|N}\, v^*_N + p_{SH|N}\, v^*_{SH}\,.$$

If $\max\{v_{SH}, v_N, v_{BH}\} = v_{SH}$, then the former inequality holds because of (C3). On the other hand, if $\max\{v_{SH}, v_N, v_{BH}\} = v_{BH}$, proceeding as before we have that the last inequality holds if and only if

$$v_{BH} - v_{SH} < (v_{SH} - v_N)\,\frac{p_{N|N} - \frac{\kappa}{2}}{p_{BH|N} - \frac{\kappa}{2}}.$$

The remainder follows the same reasoning as that in the case with $\omega = BH$ so that if

$$p^a_{BH} - p^b_{SH} < (p^b_{SH} - \max\{p^a_N, p^b_N\})\,\frac{p_{N|N} - \frac{\kappa}{2}}{p_{BH|N} - \frac{\kappa}{2}}\,,$$

then LLS is preferred to LS in $\omega = N$.

Finally, using a similar reasoning, we have that if

$$p^b_{SH} - p^a_{BH} < (p^a_{BH} - \max\{p^a_N, p^b_N\})\,\frac{p_{N|N} - \frac{\kappa}{2}}{p_{SH|N} - \frac{\kappa}{2}}\,,$$

then LLB is preferred to LB in $\omega = N$. □

For $\omega = N$, if the value of α is sufficiently small, then we have the following action values

$$\begin{aligned}
\upsilon_N(LLB) = p^b_N\,(\frac{\vartheta}{2} + \delta(\,(1-\kappa)\upsilon_{BH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{SH})\\
+(1-p^b_N)\,\delta(\,(1-\kappa)\upsilon_{BH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{SH}) =\\
p^b_N\,\frac{\vartheta}{2} + \delta\,((1-\kappa)\upsilon_{BH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{SH})\,,\\
\upsilon_N(LLS) = p^a_N\,(\frac{\vartheta}{2} + \delta(\,(1-\kappa)\upsilon_{SH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{BH})\\
+(1-p^a_N)\,\delta(\,(1-\kappa)\upsilon_{BH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{SH})\\
= p^a_N\,\frac{\vartheta}{2} + \delta\,((1-\kappa)\upsilon_{SH} + \frac{\kappa}{2}\upsilon_N + \frac{\kappa}{2}\upsilon_{BH})\,,
\end{aligned}$$

For a sufficiently small value of α, the optimal action is either LLS or LS in $\omega = BH$, whereas the optimal action is either LLB or LB at $\omega = SH$. From Claim 4, we have $LS \prec LLS$ for $\omega = BH$ and $LB \prec LLB$ for $\omega = SH$. Moreover, we note that $\delta\frac{\kappa}{2}\upsilon_N$ appears both in the expressions for $\upsilon_N(LLS)$ and for $\upsilon_N(LLB)$, and also when computing both υ_{SH} and υ_{BH}. Hence, it does not play a role in the comparison between $\upsilon_N(LLS)$ and for $\upsilon_N(LLB)$. Therefore, $\upsilon_N(LLB) > \upsilon_N(LLS)$ happens if and only if (we get rid of terms involving υ_N, so that in particular we have to modify the terms υ_{BH} and υ_{SH} with new terms denoted x, y)

$$p^b_N\,\frac{\vartheta}{2} + \delta\,((1-\kappa)x + \frac{\kappa}{2}y) > p^a_N\,\frac{\vartheta}{2} + \delta\,((1-\kappa)y + \frac{\kappa}{2}x)\,,$$

with

$$x = p^a_{BH}\frac{\vartheta}{2} + \delta((1-\kappa)y + \frac{\kappa}{2}x)\,.$$

and

$$y = p^b_{SH}\frac{\vartheta}{2} + \delta((1-\kappa)x + \frac{\kappa}{2}y)\,.$$

The last 2 conditions allow to solve for x and y and therefore plugging them in the expressions for $\upsilon_N(LLB)$ and $\upsilon_N(LLS)$ we get that

$$\upsilon_N(LLB) > \upsilon_N(LLS) \iff (p_N^b - p_N^a) > \frac{\delta(1-\kappa-\frac{\kappa}{2})}{1+\delta(1-\kappa-\frac{\kappa}{2})}(p_{SH}^b - p_{BH}^a)$$

□

Proof of Theorem 1. We first consider the case $\kappa = 0$. Combining Lemma 5 and Lemma 6 ensures that there exist values of α small enough for which LS or LLS is optimal at (BH, q), for $q > 0$. Combining Lemma 7 and Claim 4 guarantees that there exist values of α small enough for which LLB is optimal at (SH, q), for $q > 0$. This proves the Theorem for states $(SH, q > 0)$. An analogous reasoning is used to prove the result for states $(BH, q < 0)$. Combining Lemma 8 and Lemma 11 proves the result for either states $(N, q > 0)$ and $(N, q < 0)$. For κ satisfying condition (C3), the result follows from Lemmas 9 and 10 and Remark 1. □ □

Proof of Lemma 3. We compare the value of two strategies. Specifically, we consider an optimal stationary pure Markov strategy σ^* and a strategy σ that is suboptimal.

We run both the suboptimal strategy and the optimal strategy until the states match, and then the suboptimal strategy plays according to the optimal strategy. Here, the target state continues to change as a consequence of running the optimal strategy; so the suboptimal strategy is defined to follow the inventory level of the optimal strategy. Specifically, if the optimal action leads to the inventory staying at the same level, then the suboptimal strategy does nothing; if the optimal action leads to the inventory increasing by one unit, then the suboptimal strategy submits a buy market order; finally, if the optimal action leads to the inventory decreasing by one unit, then the suboptimal strategy submits a sell market order. Therefore, the states of the two chains are always at the same level of inventory, so the difference in the payoff received (at each step) is less than ϑ.

Now, for the states to match, the volume imbalance regime ω also needs to be the same. Observe that at each time step, the probability that the two chains meet at the same volume imbalance regime (after one step) is greater than m, where m is the minimum element of the transition probability matrix corresponding to the baseline book regime transitions (when no large limit orders are submitted). Hence, the hitting time is dominated by a geometric random variable with success probability m; thus, the expectation of the hitting time is less than $1/m$. Therefore, we have

$$\upsilon^*_{\omega,q} - \upsilon^*_{\omega',q} \le \upsilon^*_{\omega,q} - \upsilon_{\omega',q}(\sigma) \le \vartheta/m$$

because the discount parameter $\delta < 1$. □ □

Proof of Proposition 2. From the action values, $LLB \prec LLS$ if and only if

$$\alpha > \frac{\left(p^b_\omega - p^a_\omega\right)\vartheta/2}{p^a_\omega(2q-1) + p^b_\omega(2q+1)}$$

$$+\delta\,\frac{p^b_\omega\left((1-\kappa)(\upsilon^*_{BH,q+1} - \upsilon^*_{BH,q}) + \frac{\kappa}{2}(\upsilon^*_{N,q+1} - \upsilon^*_{N,q}) + \frac{\kappa}{2}(\upsilon^*_{SH,q+1} - \upsilon^*_{SH,q})\right)}{p^a_\omega(2q-1) + p^b_\omega(2q+1)}$$

$$+\delta\,\frac{p^a_\omega\left((1-\kappa)(\upsilon^*_{SH,q} - \upsilon^*_{SH,q-1}) + \frac{\kappa}{2}(\upsilon^*_{N,q} - \upsilon^*_{N,q-1}) + \frac{\kappa}{2}(\upsilon^*_{BH,q} - \upsilon^*_{BH,q-1})\right)}{p^a_\omega(2q-1) + p^b_\omega(2q+1)}$$

$$+\delta\,\frac{\left((1-\frac{3\kappa}{2})(\upsilon^*_{BH,q} - \upsilon^*_{SH,q})\right)}{p^a_\omega(2q-1) + p^b_\omega(2q+1)}\,.$$

We use the upper bound from Lemma 3 and the upper bound $\upsilon^*_{\omega,q} - \upsilon^*_{\omega,q-1} \leq 0$ from Lemma 1, to obtain

$$\alpha_1(\omega,q) \leq \frac{\left(p^b_\omega - p^a_\omega\right)\vartheta/2 + \vartheta/m}{p^a_\omega(2q-1) + p^b_\omega(2q+1)} = \bar{\alpha}_1(\omega,q) \tag{2.15}$$

as an upper bound for $\alpha_1(\omega,q)$ that is strictly positive.

From Lemma 5, we have $\alpha_3(\omega,q) > 0$ for all $\omega \in \Omega$ and $q > 0$. Therefore, the result follows because $\alpha_1(\omega,q) \wedge \alpha_3(BH,q) \leq \bar{\alpha}_1(\omega,q)$ and $\alpha_1(\omega,q) \wedge \alpha_3(BH,q+1) \leq \bar{\alpha}_1(\omega,q)$. □ □

Proof of Proposition 3. For a fixed volume imbalance regime ω, the function $\bar{\alpha}_1(\omega,q)$ monotonically decreases as the absolute value of q increases. The choice of α ensures that $\alpha \notin I'(s)$ for all states $s = (\omega,q)$ where $q \neq 0$, so the result follows. □ □

Proof of Proposition 4. If $\vartheta \to 0$, then $\bar{\alpha}_1(\omega,q) \to 0$ for all $\omega \in \Omega$ and $q \neq 0$. Therefore, the result follows as a consequence of Proposition 3. □ □

Chapter 3

Empirical Aspects and Extensions of the Spoofing Model

3.1 Introduction

In this Chapter we provide the empirical evidence that supports the features of the model of Chapter 2 and the conditions in Theorem 1 to guarantee the presence of manipulation in many of the possible states. We also analyze some extensions of the model of Chapter 2 and see that manipulation can still hold on them.
In particular, in Section 3.2 we detail the estimation procedure for both the fill probabilities and the baseline transition probabilities required in our model of Chapter 2, and we see that the conditions of Theorem 1 of Chapter 2 are usually satisfied empirically.
In Section 3.3 we extend our results from the previous Chapter in five directions. First, when the trading horizon is finite, we use backward induction to solve numerically for an optimal strategy and find that algorithms can also learn to manipulate the order book. Second, we show that algorithms will learn to manipulate the book when there is a trend in the fundamental value of the asset. Third, we allow the learning algorithm to submit buy and sell limit orders (of either small or large volume) at different price levels, and find that manipulation also arises. Fourth, when the manipulative orders are filled for multiple units, the learning algorithm will also learn to manipulate the book. Fifth, we study the effect of introducing a competing market maker. We find that if both market makers train their algorithms offline, then their algorithms either coordinate or mis-coordinate depending on their initial inventory level. On the other hand, if both market makers train their algorithms online, then their algorithms learn to coordinate by either riding the manipulative sequences of each other or by allowing one market maker to ride the other market

maker's manipulative sequences to avoid mis-coordination.
Finally, in Section 3.4 we give some additional tables and figures regarding estimated transition probability matrices between book regimes for different assets, and concerning the extensions of Section 3.3.

3.2 Empirical Estimation

This section uses Nasdaq data to test the conditions derived in Section 2.4. We discuss the estimation procedure for the parameters of our model, and we use these estimates to determine if market conditions from Nasdaq are conducive for an algorithm to learn to manipulate the order book. We use the dataset from Section 2.2.

Table 3.1: Summary statistics for April 2023.

Ticker	Decision interval Δt	Ave. spread (ticks)	Ave. queue size best bid	Ave. queue size best ask	Ave. volume traded per Δt	Ave. volume imbalance
AAPL	5 seconds	1.168	583	600	1217	0.008
	1 second	1.167	583	600	243	0.006
	0.5 seconds	1.168	584	601	122	0.006
AMZN	5 second	1.205	532	572	1146	-0.027
	1 seconds	1.206	532	571	229	-0.028
	0.5 seconds	1.206	533	571	115	-0.027
CSCO	5 seconds	1.005	2088	2046	488	0.012
	1 second	1.006	2100	2060	98	0.012
	0.5 seconds	1.011	2120	2078	49	0.012
INTC	5 seconds	1.005	3378	3517	975	-0.014
	1 second	1.005	3385	3530	195	-0.016
	0.5 seconds	1.006	3405	3557	97	-0.017
MSFT	5 seconds	1.783	114	119	746	-0.021
	1 second	1.784	114	119	149	-0.022
	0.5 seconds	1.788	114	119	75	-0.022
TSLA	5 seconds	2.231	190	200	2089	-0.01
	1 second	2.235	195	198	418	-0.01
	0.5 seconds	2.232	194	198	209	-0.009

Table 3.1 provides summary statistics for six assets at three different decision intervals Δt: 5 seconds, 1 second, and 0.5 seconds. We estimate the statistics by sampling the relevant features at every decision interval.

3.2.1 Estimation Procedure

In our model, there are two sets of model parameters to estimate: the transition probabilities of the volume imbalance regime $p_{\omega'|\omega}$ for all $\omega, \omega' \in \Omega$, and the fill probabilities in each volume imbalance regime p^a_ω and p^b_ω for all $\omega \in \Omega$.

Transition probabilities. Let $n_{\omega,\omega'}$ denote the number of times that volume imbalance moved from state ω to ω'. Then, from standard results, we have that

$$\hat{p}_{\omega'|\omega} = \frac{n_{\omega,\omega'}}{\sum_{\omega''} n_{\omega,\omega''}}\,, \tag{3.1}$$

where transitions from the end of one trading day to the start of the next trading day are excluded from the count. Table 3.5 of Section 3.4 provides the estimates of the transition probability matrix for the assets at the three different decision intervals.[1]

Fill probabilities. We estimate the fill probabilities with counterfactual analysis. Following [10], we submit "hypothetical" limit orders (with unit volume) at the end of the best bid and best ask queues at time t, and we track if the hypothetical order was filled between t and $t+1$ following price-time priority. These hypothetical orders account for all the change in behavior of market participants described in Section 2.2.

3.2.2 Spoofing Conditions

Table 3.2 uses the estimates from Tables 2.7 and 3.5 to check if conditions in Theorem 1 are satisfied. The entry NA indicates that the conditions in Theorems 1 and 2 are not applicable because condition (C4) does not hold.

Table 3.2: Testable conditions from Theorem 1, with $\kappa = 0$ in (C4).

Ticker	5 seconds			1 second			0.5 seconds		
	(C1), (C2)	(C4)	Side	(C1), (C2)	(C4)	Side	(C1), (C2)	(C4)	Side
AAPL	✓	✓	$q > 0$	✓	✓	$q > 0$	✓	✓	$q > 0$
AMZN	✓	✓	$q < 0$	✓	✓	$q > 0$	✓	✓	$q > 0$
CSCO	✓	✓	$q < 0$	✓	✓	$q < 0$	✓	✓	$q > 0$
INTC	✓	✓	$q > 0$	✓	✓	$q > 0$	✓	✓	$q > 0$
MSFT	✓	✓	$q < 0$	✓	✓	$q > 0$	✓	✗	NA
TSLA	✓	✓	$q < 0$	✓	✓	$q < 0$	✓	✓	$q < 0$

For all assets and decision intervals considered, conditions (C1) and (C2) are satisfied. Therefore, there are values of the inventory aversion parameter α where a manipulative strategy is optimal for $s = (SH, q > 0)$, $s = (BH, q < 0)$, and $s = (N, q)$ for either $q > 0$ or $q < 0$. On the other hand, condition (C4) allows one to determine if manipulation occurs in the neutral regime when inventory is long or short.

[1] The rows of the estimates of the transition probability matrix do not always sum to unity because we round the estimates to the second decimal point.

Figures 3.1- 3.6 in Section 3.4 plots $\bar{\alpha}_1(\omega, q)$ as a function of inventory, so the area under the curves describes the intervals $I'(s)$, where $\alpha \in I'(s)$ is a necessary condition for the algorithm to learn to manipulate the order book. Therefore, for all assets and decision intervals considered, there exists a range of values of the inventory aversion parameter α below the curve where an algorithm will learn to manipulate the book.

Also, for most assets in Table 3.5 of Section 3.4, we see that condition (C3) holds for values of κ ranging from 0.2 to 0.5 depending on the stock. Therefore, in certain instruments, if a manipulative order allows one to transition to the appropriate heavy regime half of the time, then a manipulative strategy can become dynamically optimal. On the other hand, assets such as CSCO and INTC only hold for small values of κ ranging from 0.04 to 0.2 depending on the decision interval. In these cases, a manipulative strategy may no longer be dynamically optimal if a manipulative order does not allow one to transition to the appropriate heavy regime almost all of the time.

3.3 Extensions

This section analyzes five extensions to our model: finite trading horizon, trend in the value of the asset, limit orders posted at various price levels of the book, more than one order (or more than one lot) of the manipulative limit order is filled, and more than one market maker delegating trading to a learning algorithm. We show that in all these extensions the algorithms will learn to manipulate the limit order book, and show that when market makers compete, their algorithms can learn to coordinate their actions to manipulate the book.

3.3.1 Finite Trading Horizon

Above, we focused on the infinite-horizon case because most learning algorithms are designed for such a setting. However, a finite-horizon model best captures intraday trading because many market makers close inventories before the end of the trading day. Due to this, we assume that the optimal action at time T corresponds to DN. With a finite trading horizon, theoretical results guarantee that there exists an optimal non-stationary pure Markov strategy (see for example Proposition 4.4.3 in [68]), so the space of strategies to search over significantly increases. This is intuitive because the optimal action with a few minutes before the end of the trading horizon differs from the optimal action with a few hours before closing. Nonetheless, the problem can be solved through backward induction.

Intuitively, for a sufficiently long horizon T, the trading behavior at the start should resemble that from an infinite-horizon problem. Although we do not have testable conditions for a finite-trading horizon, we use dynamic programming to solve numerically for the optimal non-stationary pure Markov strategy. Figures 3.7-3.24 in Section 3.4 use the empirical estimates from Tables 2.7 and 3.5, and discount factor $\delta = 1$, to plot the optimal actions for each state and at each point in time $t = 0, 1, 2, ..., T = 30$. From the figures, one can string together the optimal action from one time step to another, and we show that quote-based manipulation can occur at every time point t, and that manipulation occurs for inventory levels closer to zero.

3.3.2 Trend in the Fundamental Value of the Asset

Here we assume that $2\,\beta - 1 \neq 0$. When $2\,\beta - 1 > 0$, there is an upward trend in the fundamental value of the asset, and when $2\,\beta - 1 < 0$ there is a downward trend. Optimality of action DN is mathematically difficult to study in this extension.[2] We cannot then apply the results obtained for the basic model and we have applied simulations. We consider the assets AMZN and CSCO, we set values of $\varphi \approx \vartheta/2$ and we see that both spoofing and more generally manipulation also hold in this setting, as we can see in Figure 3.25 in Section 3.4. Moreover, there is a gravitational pull towards a positive inventory level. The maximum of the optimal continuation values $v^*_{\omega,q}$ is for $q^* > 0$, as we can see in Figure 3.26 in Section 3.4.

3.3.3 Multiple Prices

Here, the action space includes the actions to submit buy and sell limit orders to K (finite) price levels. Following (2.3), the expected one-step utility from submitting buy or sell limit orders at price level $k \in \{1, 2, ..., K\}$ in state $s = (\omega, q) \in S$ is given by

$$\bar{u}(s,a) = \begin{cases} (\frac{\vartheta}{2} + (j-1)\,\varphi)\, p^{b,j}_{w} - \alpha\, p^{b,j}_{\omega}\,(q+1)^2 - \alpha\,(1 - p^{b,j}_{\omega})\, q^2 & \text{for } a = LB_j\,, \\ (\frac{\vartheta}{2} + (k-1)\,\varphi)\, p^{a,k}_{w} - \alpha\, p^{a,k}_{\omega}\,(q-1)^2 - \alpha\,(1 - p^{a,k}_{\omega})\, q^2 & \text{for } a = LS_k\,, \end{cases} \tag{3.2}$$

where $p^{a,k}_{\omega} \in (0,1)$ and $p^{b,j}_{\omega} \in (0,1)$ are the fill probabilities of a sell and buy limit order, respectively, at price level k in each regime $\omega \in \Omega$. The ordering $p^{a,1}_{\omega} \geq p^{a,2}_{\omega} \geq$

[2]Intuitively, there is at most one inventory level $q^* > 0$ in this setting for which DN could be optimal for a range of values of the inventory aversion parameter. This would generalize the original model, where it was easy to prove that $q^* = 0$ for any book regime ω because DN was never optimal at (ω, q) for $q \neq 0$. However, it is not straightforward to prove that there exists at most one level $q^* > 0$ for which DN is optimal in (ω, q^*).

... $\geq p_{\omega}^{a,K}$ and $p_{\omega}^{b,1} \geq p_{\omega}^{b,2} \geq ... \geq p_{\omega}^{b,K}$ holds for all $\omega \in \Omega$ because of price-time priority. The first price level $k = 1$ corresponds to the best bid and best ask, while price level $k > 1$ corresponds to $k - 1$ ticks below (above) the best bid (best ask).

The inventory and regime transitions for the buy (sell) limit orders at the different price levels are the same. The only difference between the different price levels is the expected one-step utility and the transition probabilities. We denote

$$k^a(s) = \underset{k\in\{1,...,K\}}{\arg\max} \left\{ \bar{u}(s, LS_k) + \delta \sum_{\omega'\in\Omega} p(\omega'|\omega, LS_k) \sum_{q'\in\mathcal{Q}} p(q'|q, LS_k)\, v^*_{\omega',q'} \right\}, \quad (3.3)$$

$$k^b() = \underset{j\in\{1,...,K\}}{\arg\max} \left\{ \bar{u}(, LB_j) + \delta \sum_{\omega'\in\Omega} p(\omega'|\omega, LB_j) \sum_{q'\in\mathcal{Q}} p(q'|q, LB_j)\, v^*_{\omega',q'} \right\}, \quad (3.4)$$

as the price levels on the ask side and bid side, respectively, that lead to the maximum expected stream of discounted payoffs in state $s = (\omega, q)$. If there are multiple price levels that maximize these equations, then all price levels will lead to the same expected stream of discounted payoffs in state $s = (\omega, q)$. With tiebreaks, there is one price level from the ask and one from the bid that maximize the equations above. Therefore, for each state s, price levels $k^a(s)$ and $k^b(s)$ strictly dominate the remaining price levels in state s. Thus, the remaining price levels in the ask and in the bid are not optimal actions in state s because they do not maximize (2.7). In the multiprice setting we consider the corresponding large limit orders: that is, for any possible LS_k, we also have the associated LLS_k, and for any possible LB_j, we also have the associated LLB_j. In this case, we ran simulations for $K = 3$ different levels in the ask and bid prices for AMZN and CSCO, we set $\varphi \approx \vartheta/2$ and find that the algorithm learns to manipulate and spoof the order book (recall that ϑ is the expected spread and φ is the size of the tick) as we see in Figure 3.27 in Appendix 3.4.

3.3.4 Multiple Filled Limit Orders

So far we assumed that when submitting a large number of limit orders or one limit order with a large volume, at most one unit or one lot is filled regardless of the volume of the limit order or the number of limit orders. Here we allow for two orders to be filled, and we run simulations to see what is its effect. As large limit orders are planned to be cancelled, we assume the probability of filling 1 limit order is higher than the probability of filling 2 limit orders in this scenario. In particular, the probability of filling one limit order is twice the probability of filling two orders. As above, we find that the algorithms learn to manipulate and spoof the order book, as shown in Figure 3.28 in Section 3.4.

3.3.5 Multiple Market Makers

Our analyses thus far focused on quote-based manipulation by a single market maker. An extension is to study multiple market makers who delegate their decision making processes to learning algorithms. How does the introduction of another market maker (who also uses a learning algorithm) affect a single algorithm's ability to manipulate the order book?

We ignore competition between limit orders from the different algorithms to simplify the analysis. Instead, we analyze the effect of multiple algorithms attempting to control the volume imbalance regime through manipulative orders. To formalize the new transition dynamics of the volume imbalance regime, we focus on two market makers where $a = (a^1, a^2)$ is the action profile. The transition dynamics now depend on the action profile $p(\omega' \mid \omega, a)$ for all $\omega, \omega' \in \Omega$, which we summarize below:

- If both market makers submit a small or no order, then the volume imbalance evolves according to its baseline dynamics.

- If one market maker submits a large limit order and the other market maker submits a small or no order, then the volume imbalance regime moves to buy-heavy or sell-heavy with probability one (depending on which side the large limit order is placed).

- If both market makers submit large limit orders on the same side of the book, then the volume imbalance regime moves to buy-heavy or sell-heavy with probability one (depending on which side the large limit orders are placed).

- If both market makers submit large limit orders on opposing sides of the book, then the volume imbalance evolves according its baseline dynamics.

The notion of optimality with multiple market makers is based on equilibrium solution concepts. In turn, the solution concepts depend on the game and the strategies used by the learning algorithms. Generic learning algorithms search for an optimal strategy in the space of stationary Markov strategies; hence, an algorithm conditions its behavior on the set of states encoded in the algorithm. In our setting, each algorithm conditions their behavior on the volume imbalance regime which is publicly observable, but they also condition their behavior on their own level of inventory which is private information.

With private information, the most appropriate equilibrium solution concept in this setting is a perfect Bayesian equilibrium, where the belief determines the opponent's level of inventory, and the optimal strategy should be optimal with respect to the belief. However, generic learning algorithms use stationary Markov strategies and do not account for an opponent's level of inventory, so a perfect Bayesian equilibrium is not appropriate for generic algorithms.[3] Nevertheless, we can analyze the effect of this misspecification on an algorithm's ability to manipulate the order book.

Table 3.3: Average number of manipulative sequences over 50 trading intervals.

Ticker	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 $q=-4$
AMZN	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
CSCO	Baseline	5 seconds	25.56	15.25	15.12	29.19	18.29	18.82
		1 second	32.98	13.88	13.70	36.48	11.59	10.40
		0.5 seconds	37.27	9.54	9.55	40.27	7.73	7.89
	Offline	5 seconds	29.72	29.65	20.50	18.91	21.67	19.99
		1 second	37.13	37.10	14.73	12.62	21.25	18.53
		0.5 seconds	40.90	41.04	9.40	8.66	21.07	19.69
	Online	5 seconds	22.16	29.33	12.75	17.94	11.69	18.78
		1 second	20.13	36.46	0.0	13.00	0.0	10.79
		0.5 seconds	32.59	39.78	0.0	14.91	0.0	14.49

To study the effect of introducing a second learning algorithm, we first establish a baseline with one market maker who uses an algorithm to learn the optimal trading strategy. The market makers solve for the optimal strategy with the policy iteration algorithm using the empirical estimates from Tables 2.7 and 3.5, discount factor $\delta = 0.95$, and inventory aversion $\alpha = 10^{-4}$ for market maker one and $\alpha = 10^{-5}$ for market maker two. Although we have two market makers for the baseline, we do not study their interaction for the baseline setting. For each market maker, we simulate their optimal strategy over 50 time steps 10,000 times when the market maker starts

[3] A Berk–Nash equilibrium (see [42]) is the most suitable equilibrium solution concept for this misspecified setting, but the analysis is beyond the scope of the paper, and also not applicable to generic learning algorithms.

with different values of the initial level of inventory. The initial volume imbalance regime is sampled with equal probability. Table 3.3 reports the average number of manipulative sequences for each market maker. We count a manipulative sequence as LLB at time t followed by LS or LLS at time $t+1$ for $q \geq 0$, or LLS at time t followed by LB or LLB at time $t+1$ for $q \leq 0$. The table reports the results of two representative assets AMZN and CSCO. The results of the other assets are reported in Tables 3.6–3.8 of Section 3.4.

3.3.5.1 Offline Learning

Here, both market makers train their algorithms offline with the misspecified model (the original model from Section 2.3) which ignores the strategic behavior of other algorithms, and we analyze the outcome of the interaction between the algorithms in the market. With the same setup as that of the baseline, we simulate the interaction of the market makers over 50 time steps 10,000 times when (i) both market makers start with zero inventory, (ii) both market makers start with the same level of inventory ($q = 4$), and (iii) the market makers start with opposing levels of inventory ($q = 4$ and $q = -4$).

When comparing the result with the baseline, the introduction of another market maker increases the number of times manipulation occurs in the market (i.e., market makers ride the manipulative sequences of each other) when the market makers start with zero inventory and with opposing levels of inventory. On the other hand, the introduction of another market maker decreases the number of times manipulation occurs in the market when the market makers start with the same level of inventory.

Table 3.4: Average manipulation statistics.

Ticker	Setup	Δt	Mismatching manipulative orders			Single manipulative order		
			Zero inv.	Same inv.	Opposing inv.	Zero inv.	Same inv.	Opposing inv.
AMZN	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
CSCO	Offline	5s	0.3008%	0.0591%	1.0232%	20.13	19.82	34.29
		1s	0.0115%	1.7566%	4.6456%	12.63	7.66	38.89
		0.5s	0.0109%	0%	4.5198%	8.82	2.98	41.07
	Online	5s	0.6417%	2.2775%	4.6541%	24.77	26.35	25.06
		1s	1.4149%	0%	0%	25.85	24.88	20.30
		0.5s	2.2263%	0%	0%	16.17	30.77	29.38

To analyze the impact of introducing another market maker, Table 3.4 reports the percentage of times when competing market makers submit large orders that cancel each other out divided by the number of times the market makers used a large order (mismatching manipulative orders), and the number of times where only one out of the two market makers submits a large order (single manipulative order). The results of the other assets are reported in Tables 3.9 and 3.10 of Section 3.4.

When starting with zero inventory or opposing levels of inventory, the additional market maker does not lead to many instances where the large orders cancel each other out, but it does lead to more instances where the order book moves to a heavy regime. This allows the market makers to exploit the manipulative sequences of each other so that we have more manipulative sequences than would otherwise occur with only one market maker. When starting with the same level of inventory, the additional market maker does not lead to many instances where the large orders cancel each other out, but there are fewer instances where only one market maker submits a manipulative order. We see that market maker one disrupts market maker two because market maker one manipulates as often as the baseline setting but market maker two has significantly fewer manipulative sequences. In the offline learning setting, the algorithms either coordinate or mis-coordinate depending on their initial inventory.

3.3.5.2 Online Learning

Next, we assume that both market makers pre-train their algorithms offline with the misspecified model and use the results to initialize an online learning algorithm. With the same setup as that of the baseline, the market makers pre-train with the policy iteration algorithm and then use Q-learning to learn online (see [72], [16] for a basic explanation of Q-learning).[4]

For the online learning, we follow the experimental setup of [16]. The Q-learning algorithms learn have an ε-greedy choice rule with a time-declining exploration rate given by $\varepsilon_t = \exp(-\tau\, t)$, where the parameter $\tau > 0$ controls the rate of decay of exploration. The ε-greedy choice rule picks the (current) optimal action with probability $1 - \varepsilon$, and a random action is chosen with probability ε. The learning

[4] Training an algorithm online by interacting with the market is costly because the algorithm needs to experiment frequently to learn to behave optimally. In most situations with multiple algorithms, learning algorithms are longer guaranteed to behave optimally. Realistically, market makers train their algorithms offline or partially train their algorithms offline (with some online experimentation) to minimize the cost of experimentation. We use the policy iteration algorithm to solve for the optimal continuation value v in the misspecified model, which we use to compute the optimal action values to initialize the Q-values for Q-learning.

rate of the algorithms is 0.125 and the exploration parameter is $\tau = 10^{-5}$. Similar to [16], we say that the online learning converged if the optimal strategy for each player does not change for 100,000 consecutive periods.

To analyze the effect of online learning, we simulate the learning process until convergence 1,000 times. Once each learning process converges, we use the learned strategies to simulate the interaction of the market makers over 50 time steps 10 times when (i) both market makers start with zero inventory, (ii) both market makers start with the same level of inventory ($q = 4$), and (iii) the market makers start with opposing levels of inventory ($q = 4$ and $q = -4$). This produces a total of 10,000 interactions over 50 time steps as in the case of the baseline.

When comparing the number of manipulative sequences to the baseline, we see that online learning often leads to a reduction in manipulation by market maker one, but an increase in manipulation by market maker two when the market makers start with zero inventory and with opposing levels of inventory. When the market makers start with the same level of inventory, there is a reduction in manipulation from both market makers. When comparing the manipulation statistics to offline learning, we see that online learning leads to more instances where only one market maker sends a manipulative order when the market makers start with zero inventory and with the same level of inventory. When the market makers start with opposing levels of inventory, there are fewer instances where only one market maker sends a manipulative order.

In the online learning setting, we see that the market makers learn to coordinate. How they coordinate depends on their initial inventory. If the market makers start with zero inventory, then they coordinate by riding the sequences of each other to increase market manipulation. On the other hand, if the market makers start with the same level of inventory or with opposing levels of inventory, then they coordinate by allowing market maker one to ride market maker two's sequences to avoid their large limit orders cancelling each other out.

3.4 Additional Tables and Figures

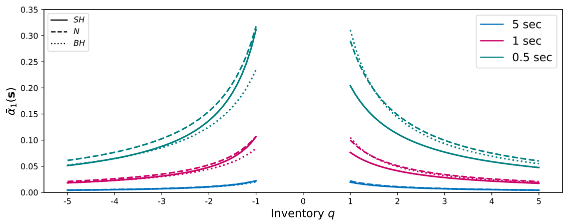

Figure 3.1: Intervals $I'(s)$ AAPL

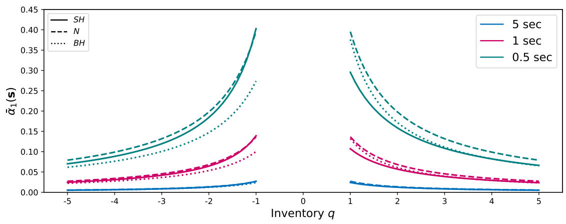

Figure 3.2: Intervals $I'(s)$ AMZN

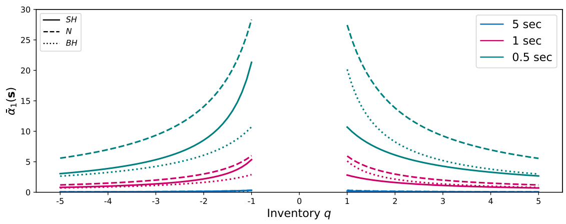

Figure 3.3: Intervals $I'(s)$ CSCO

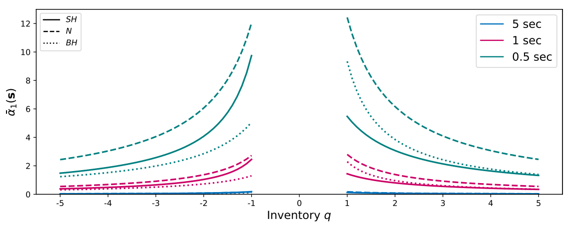

Figure 3.4: Intervals $I'(s)$ INTC

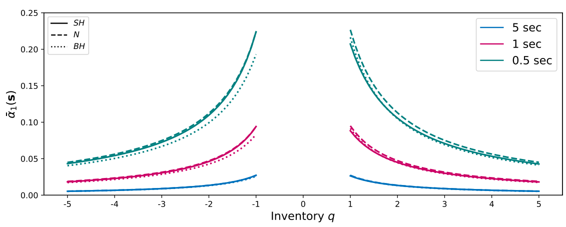

Figure 3.5: Intervals $I'(s)$ MSFT

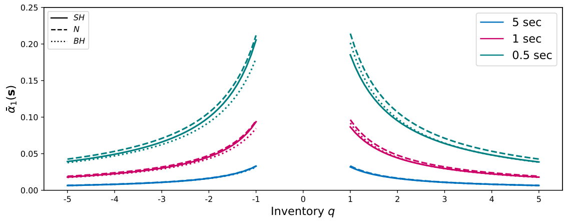

Figure 3.6: Intervals $I'(s)$ TSLA

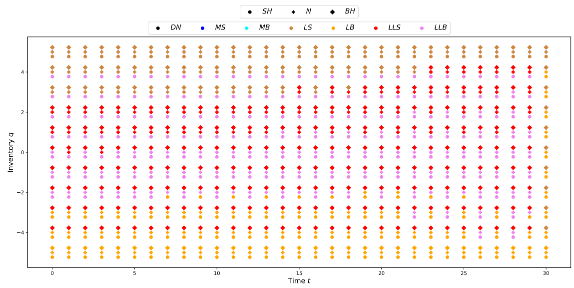

Figure 3.7: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AAPl: 5 sec

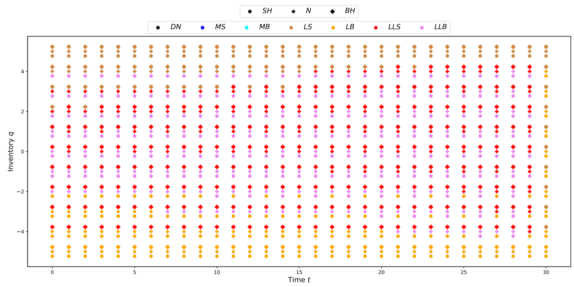

Figure 3.8: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AAPl: 1 sec

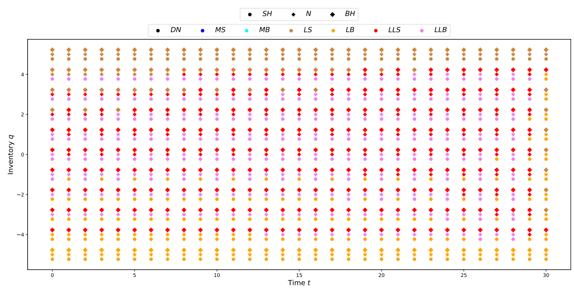

Figure 3.9: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AAPl: 0.5 sec

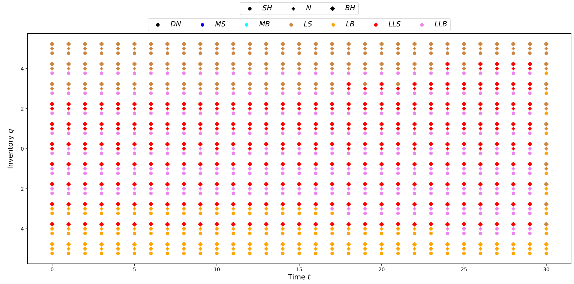

Figure 3.10: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AMZN: 5 sec

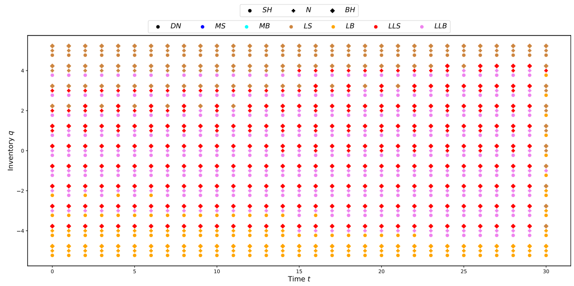

Figure 3.11: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AMZN: 1 sec

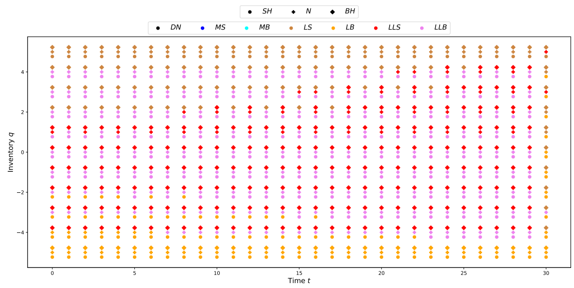

Figure 3.12: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for AMZN: 0.5 sec

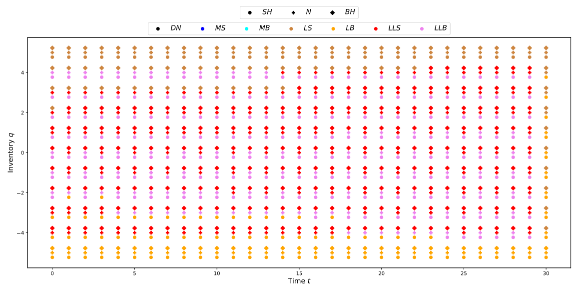

Figure 3.13: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for CSCO: 5 sec

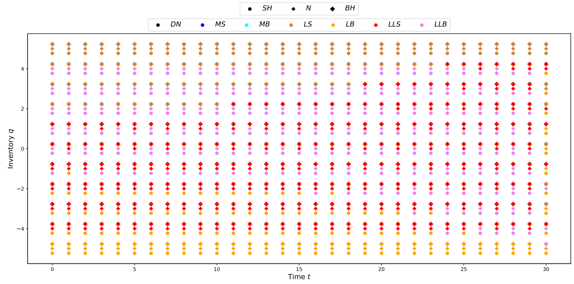

Figure 3.14: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for CSCO: 1 sec

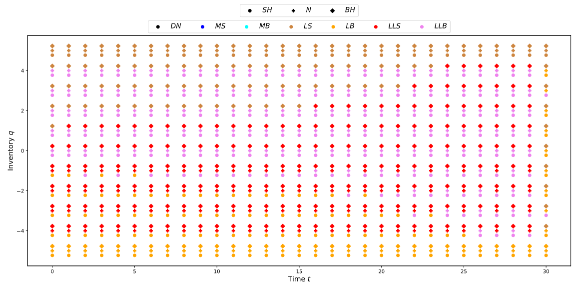

Figure 3.15: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for CSCO: 0.5 sec

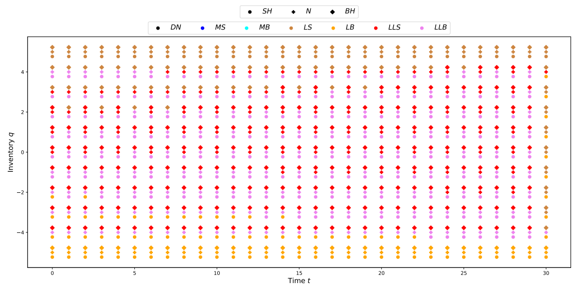

Figure 3.16: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for INTC: 5 sec

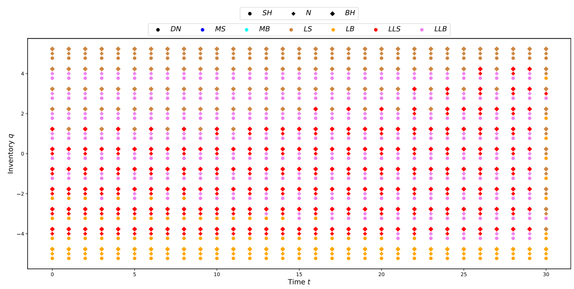

Figure 3.17: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for INTC: 1 sec

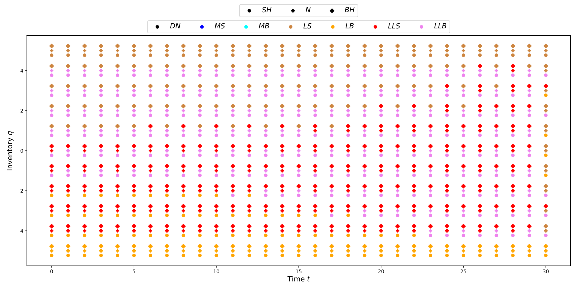

Figure 3.18: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for INTC: 0.5 sec

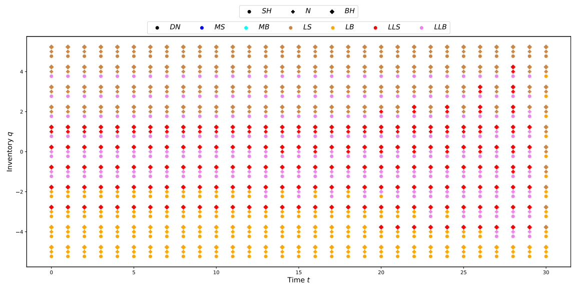

Figure 3.19: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for MSFT: 5 sec

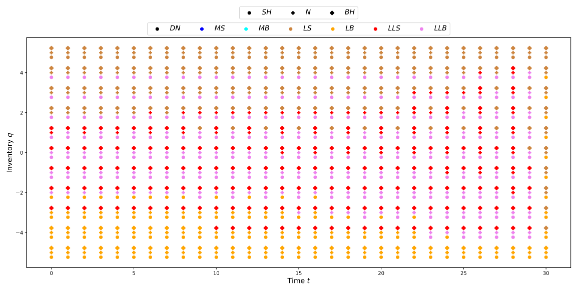

Figure 3.20: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for MSFT: 1 sec

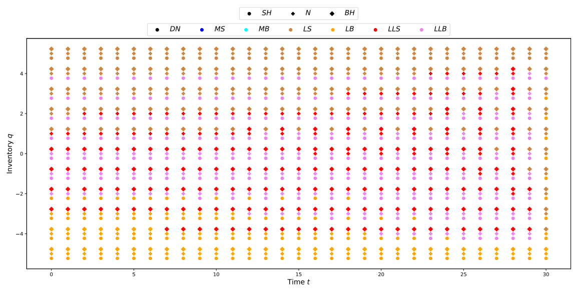

Figure 3.21: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for MSFT: 0.5 sec

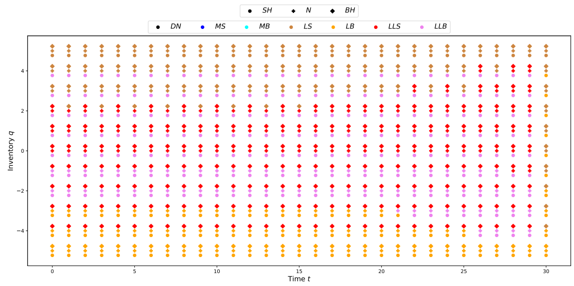

Figure 3.22: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for TSLA: 5 sec

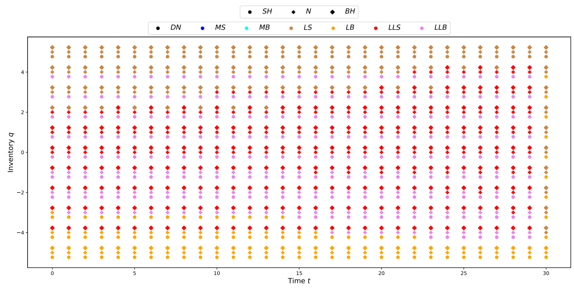

Figure 3.23: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for TSLA: 1 sec

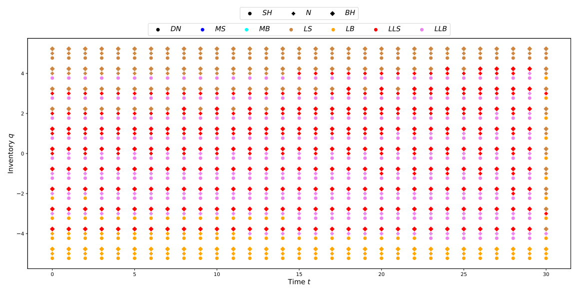

Figure 3.24: Optimal action for a finite trading horizon with $\alpha = 10^{-5}$ for TSLA: 0.5 sec

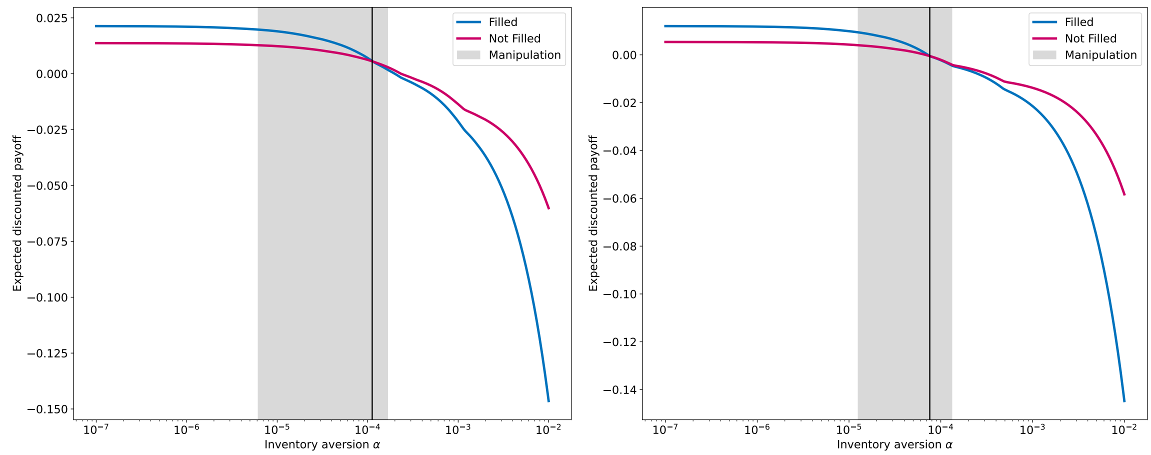

(a) AMZN (b) CSCO

Figure 3.25: Expected stream of discounted payoffs when the manipulative order is filled or not for $s = (SH, q = 2)$ with a price trend with $\beta = 0.51$.

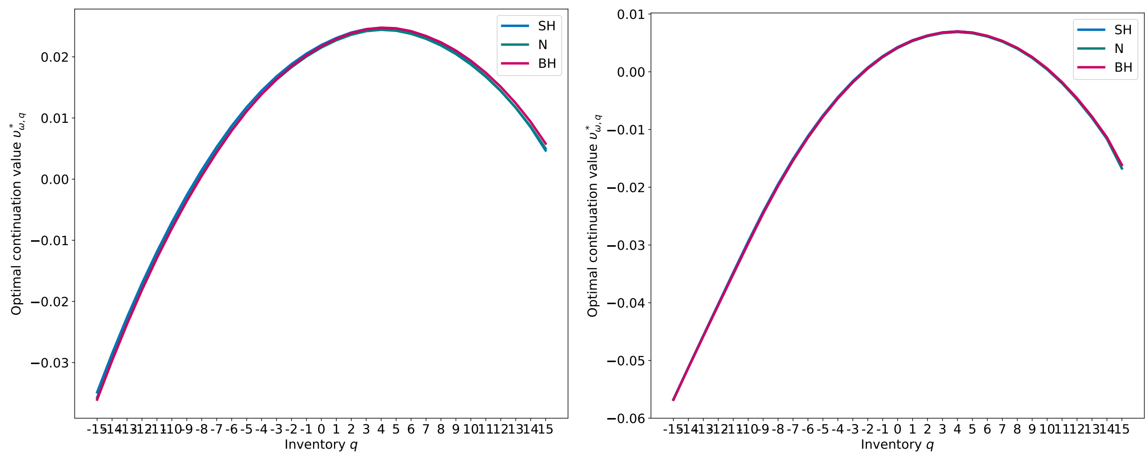

(a) AMZN (b) CSCO

Figure 3.26: Optimal continuation values $v_{\omega,q}$ with a price trend with $\beta = 0.51$.

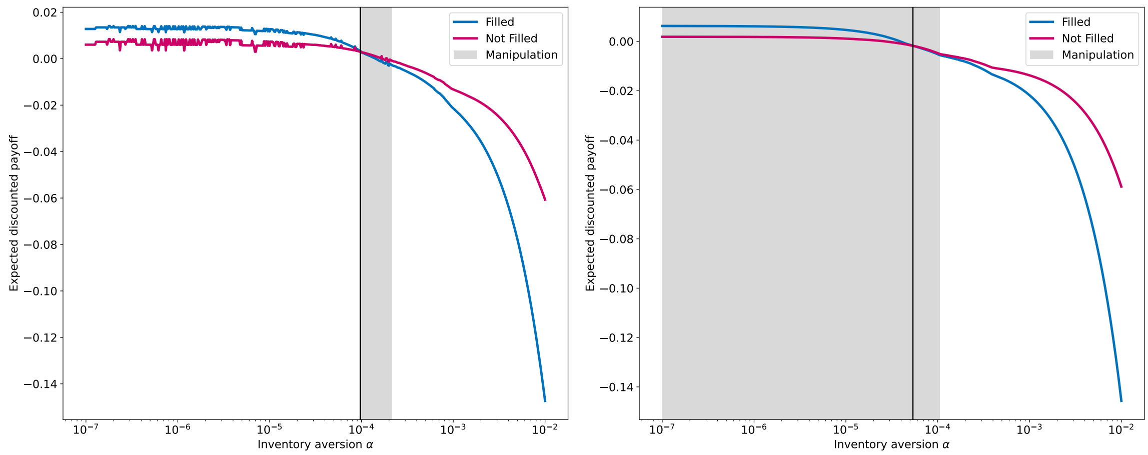

(a) AMZN (b) CSCO

Figure 3.27: Expected stream of discounted payoffs when the manipulative order is filled or not for $s = (SH, q = 2)$ for multiple prices with $K = 3$.

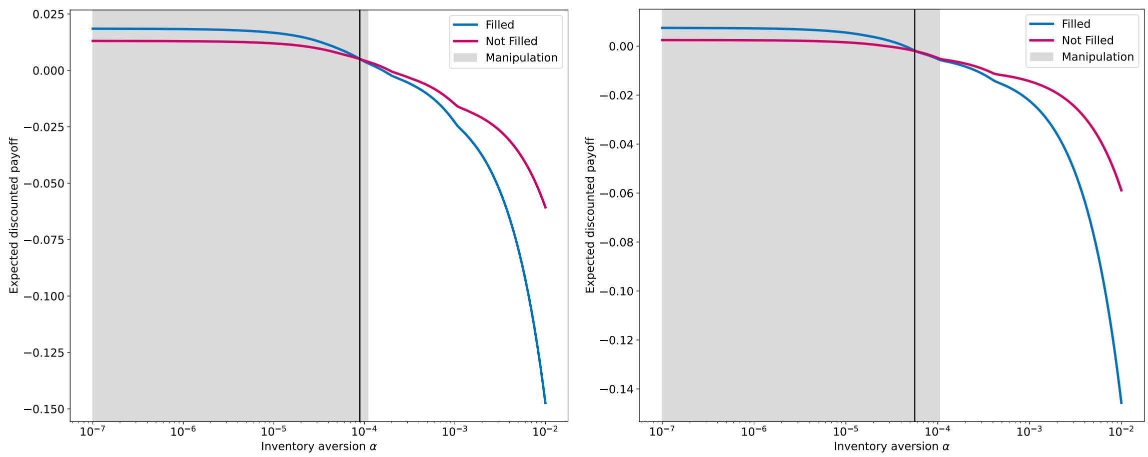

(a) AMZN (b) CSCO

Figure 3.28: Expected stream of discounted payoffs when we allow two manipulative orders to be filled for $s = (SH, q = 2)$.

Table 3.5: Transition probability matrix.

(a) AAPL: 5 seconds

	SH	N	BH
SH	0.32	0.41	0.28
N	0.28	0.43	0.29
BH	0.27	0.41	0.32

(b) AAPL: 1 second

	SH	N	BH
SH	0.42	0.37	0.21
N	0.26	0.48	0.26
BH	0.20	0.37	0.43

(c) AAPL: 0.5 seconds

	SH	N	BH
SH	0.52	0.32	0.16
N	0.22	0.55	0.23
BH	0.15	0.33	0.52

(d) AMZN: 5 seconds

	SH	N	BH
SH	0.35	0.41	0.24
N	0.3	0.44	0.26
BH	0.27	0.41	0.32

(e) AMZN: 1 second

	SH	N	BH
SH	0.47	0.35	0.18
N	0.26	0.51	0.23
BH	0.20	0.36	0.43

(f) AAPL: 0.5 seconds

	SH	N	BH
SH	0.56	0.31	0.13
N	0.23	0.57	0.21
BH	0.15	0.32	0.52

(g) CSCO: 5 seconds

	SH	N	BH
SH	0.52	0.39	0.09
N	0.15	0.7	0.15
BH	0.08	0.38	0.54

(h) CSCO: 1 second

	SH	N	BH
SH	0.76	0.21	0.03
N	0.08	0.84	0.08
BH	0.03	0.21	0.76

(i) CSCO: 0.5 seconds

	SH	N	BH
SH	0.82	0.16	0.02
N	0.06	0.87	0.07
BH	0.01	0.16	0.83

(j) INTC: 5 seconds

	SH	N	BH
SH	0.48	0.41	0.11
N	0.17	0.67	0.16
BH	0.11	0.43	0.45

(k) INTC: 1 second

	SH	N	BH
SH	0.71	0.24	0.04
N	0.10	0.81	0.09
BH	0.04	0.26	0.70

(l) INTC: 0.5 seconds

	SH	N	BH
SH	0.79	0.18	0.03
N	0.08	0.85	0.07
BH	0.02	0.20	0.78

(m) MSFT: 5 seconds

	SH	N	BH
SH	0.38	0.34	0.28
N	0.33	0.36	0.31
BH	0.31	0.34	0.35

(n) MSFT: 1 second

	SH	N	BH
SH	0.46	0.32	0.22
N	0.31	0.40	0.29
BH	0.24	0.33	0.43

(o) MSFT: 0.5 seconds

	SH	N	BH
SH	0.52	0.30	0.18
N	0.29	0.44	0.27
BH	0.20	0.30	0.49

(p) TSLA: 5 seconds

	SH	N	BH
SH	0.34	0.39	0.27
N	0.30	0.41	0.29
BH	0.28	0.39	0.32

(q) TSLA: 1 second

	SH	N	BH
SH	0.42	0.37	0.22
N	0.28	0.46	0.27
BH	0.23	0.37	0.41

(r) TSLA: 0.5 seconds

	SH	N	BH
SH	0.49	0.34	0.17
N	0.25	0.50	0.25
BH	0.18	0.33	0.48

Table 3.6: Baseline: Average number of manipulation sequences over 50 trading intervals.

Ticker	Decision Interval Δt	Zero inventory		Same inventory		Opposing inventory	
		Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2
		$q = 0$	$q = 0$	$q = 4$	$q = 4$	$q = 4$	$q = -4$
AAPL	5 seconds	24.77	20.78	20.80	25.72	21.80	22.45
	1 second	25.04	13.88	14.09	28.72	18.56	19.10
	0.5 seconds	26.23	11.56	11.68	31.48	16.36	15.08
INTC	5 seconds	25.28	17.27	17.27	28.75	20.27	18.88
	1 second	32.00	17.00	17.05	34.83	12.59	13.14
	0.5 seconds	35.22	12.04	12.03	39.12	9.40	9.92
MSFT	5 seconds	25.07	21.68	21.70	25.43	22.20	21.97
	1 second	25.55	16.87	16.87	27.33	19.15	18.86
	0.5 seconds	26.31	10.32	10.30	29.82	15.30	15.95
TSLA	5 seconds	24.97	21.85	21.87	25.41	22.41	22.31
	1 second	25.32	18.50	18.55	27.38	20.65	21.01
	0.5 seconds	26.13	14.29	14.17	29.25	18.64	19.48

Table 3.7: Offline learning: Average number of manipulation sequences over 50 trading intervals.

Ticker	Decision Interval Δt	Zero inventory		Same inventory		Opposing inventory	
		Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2
		$q = 0$	$q = 0$	$q = 4$	$q = 4$	$q = 4$	$q = -4$
AAPL	5 seconds	24.86	26.09	20.96	22.20	20.79	23.00
	1 second	26.31	29.13	15.75	19.00	14.82	20.25
	0.5 seconds	28.92	31.90	14.34	16.79	13.54	17.38
INTC	5 seconds	27.79	28.89	21.76	20.62	22.16	19.44
	1 second	35.28	35.48	17.76	14.53	21.76	17.39
	0.5 seconds	39.53	39.49	12.18	10.66	21.17	18.93
MSFT	5 seconds	25.03	25.77	21.55	22.54	21.72	22.38
	1 second	25.94	28.10	17.56	20.08	17.42	19.88
	0.5 seconds	27.02	30.51	11.30	16.06	11.24	17.27
TSLA	5 seconds	25.07	25.91	21.98	22.97	21.90	22.86
	1 second	25.53	27.66	18.92	21.24	18.73	21.36
	0.5 seconds	26.32	29.22	14.83	19.35	14.91	20.09

Table 3.8: Online learning: Average number of manipulation sequences over 50 trading intervals.

Ticker	Decision Interval Δt	Zero inventory		Same inventory		Opposing inventory	
		Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2
		$q = 0$	$q = 0$	$q = 4$	$q = 4$	$q = 4$	$q = -4$
AAPL	5 seconds	24.68	25.76	20.80	21.81	20.63	22.40
	1 second	21.80	28.62	12.57	18.25	11.84	19.39
	0.5 seconds	20.42	31.51	1.49	15.05	1.41	15.10
INTC	5 seconds	22.89	28.53	14.26	19.96	13.75	18.8
	1 second	20.26	34.81	0.0	12.36	0.0	12.39
	0.5 seconds	27.04	38.39	0.0	14.77	0.0	16.13
MSFT	5 seconds	24.65	25.38	21.27	22.05	21.32	21.85
	1 second	23.84	27.33	15.48	18.88	15.41	18.67
	0.5 seconds	23.03	29.36	8.53	14.31	8.51	14.37
TSLA	5 seconds	24.56	25.44	21.44	22.31	21.47	22.30
	1 second	24.30	27.23	17.67	20.79	17.65	20.86
	0.5 seconds	24.18	29.19	12.78	18.91	12.67	19.02

Table 3.9: Offline learning: Average manipulation statistics.

(a) Percentage of large orders on opposite sides over 50 trading intervals.

Ticker	Δt	Zero inv.	Same inv.	Opposing inv.
AAPL	5s	0.1002%	0.1008%	0.4369%
	1s	0.1187%	1.2180%	0.0044%
	0.5s	0.0250%	1.0324%	0%
INTC	5s	0.1235%	1.1789%	0.6517%
	1s	0.0609%	0.0153%	4.9916%
	0.5s	0%	0.0005%	4.1832%
MSFT	5s	0.6766%	1.6972%	0.1863%
	1s	0.2604%	1.0269%	0.0401%
	0.5s	0.1993%	0.7633%	0.0095%
TSLA	5s	0.1195%	0.1283%	0.5394%
	1s	0.3506%	0.2716%	1.9706%
	0.5s	1.7434%	1.6606%	0.9519%

(b) Number of times where only one market maker submits a large order over 50 trading intervals.

Ticker	Δt	Zero inv.	Same inv.	Opposing inv.
AAPL	5s	13.28	18.26	18.47
	1s	23.10	25.81	27.52
	0.5s	20.75	24.17	28.98
INTC	5s	21.58	20.71	29.97
	1s	14.64	8.99	35.89
	0.5s	10.10	4.29	39.34
MSFT	5s	25.92	25.22	27.84
	1s	24.00	22.06	26.58
	0.5s	22.07	17.58	26.12
TSLA	5s	13.48	17.35	18.69
	1s	24.01	23.99	27.89
	0.5s	22.65	24.07	29.37

Table 3.10: Online learning: Average manipulation statistics.

(a) Percentage of large orders on opposite sides over 50 trading intervals.

Ticker	Δt	Zero inv.	Same inv.	Opposing inv.
	5s	0.2159%	0.4622%	0.5406%
AAPL	1s	0.3898%	1.3460%	3.5092%
	0.5s	0.6900%	0%	0%
	5s	0.3182%	1.6649%	1.7663%
INTC	1s	1.7169%	0%	0%
	0.5s	1.8299%	8.2368%	0%
	5s	0.0366%	0.2648%	0.5282%
MSFT	1s	0.1880%	2.0196%	2.4215%
	0.5s	0.3387%	0%	6.8309%
	5s	0.0415%	0.2238%	0.3103%
TSLA	1s	0.3340%	1.3380%	1.8276%
	0.5s	0.4559%	3.1344%	4.5598%

(b) Number of times where only one market maker submits a large order over 50 trading intervals.

Ticker	Δt	Zero inv.	Same inv.	Opposing inv.
	5s	19.24	23.47	19.97
AAPL	1s	25.29	26.97	24.09
	0.5s	26.78	27.36	24.69
	5s	24.01	26.20	25.41
INTC	1s	25.89	20.79	23.26
	0.5s	19.58	29.51	32.50
	5s	21.35	22.15	22.31
MSFT	1s	23.32	23.67	23.77
	0.5s	23.81	22.90	23.22
	5s	18.51	20.67	19.76
TSLA	1s	21.67	24.39	23.59
	0.5s	22.15	25.94	24.22

Chapter 4

Public and Private Order Flow

4.1 Introduction

Blockchains are distributed ledgers that record transactions sequentially in blocks created by specialised agents known as builders. At regular time intervals (called slots), a consensus protocol selects which block is appended to the chain. Pending transactions are first stored in *memory pools*, where they wait to be included in a block. In public memory pools, transaction information is visible to all network participants, while in private memory pools, transaction information visibility is restricted to specific builders. In this chapter we focus on the Ethereum blockchain, where block creation follows a *proposer–builder separation* (PBS) mechanism: blocks are built by competing builders, and validators act as proposers and select among the blocks. This institutional structure, combined with the coexistence of public and private memory pools, underpins the interaction between traders and builders studied in this chapter.

The order in which transactions are executed on the blockchain has direct consequences for the execution prices received by traders. Traders pay *priority fees* to builders to obtain better queue positions, and effectively compete in *priority gas auctions* (PGAs). Because all transactions in the public pool are visible before execution, some market participants can exploit this information to extract *maximal extractable value* (MEV) via techniques such as frontrunning, backrunning, or sandwich attacks. This generates additional costs and risks for traders. Recently, private pools emerged to mitigate these costs. By sending transactions directly to a builder, traders hide their orders from public view and avoid being targeted by MEV strategies. We note that submissions to private pools could be vulnerable to MEV strategies, since builders could extract MEV themselves. However, if a builder extracted MEV, the trust relation between her and traders that have submitted to her private pool would break. This reputational cost would make the builder to lose order flow in the future, due to the dynamic nature of the relation between traders and builders. It is then unlikely traders suffer from MEV strategies if they use private channels, and, in any case, the likelihood is much lower compared to when they submit orders to the public pool. Hence, we can assume for simplicity there is no MEV when using private pools. However, using private channels comes at a cost: orders sent privately can only be executed if the builder who received them is selected to create the next block. Thus, traders face a fundamental trade-off between execution certainty (public pools) and MEV protection (private pools). This trade-off is the fundamental problem we will tackle.

We analyse how the design of blockchain execution channels shapes the behaviour of traders and builders in Ethereum. In Ethereum, traders can route their transactions through two main types of channels: a public memory pool, where all pending transactions are visible to the network, or private memory pools, where transactions are hidden from public view and sent directly to specific builders. Builders, in turn, compete to create new blocks and decide which transactions to include and in what order. Execution of transactions in memory pools happens when the transactions are included in a block that is selected. Once this execution occurs, the transactions are removed from the memory pool. The coexistence of public and private order flow creates a fundamental trade-off: public pools offer faster and more reliable execution but expose traders to extraction through maximal extractable value (MEV), whereas private pools protect against MEV but involve the risk of non-execution if the corresponding builder is not selected for the next block. Therefore, it is fundamental for traders to decide whether to submit transactions to the public memory pool or to a private memory pool.

We develop two complementary models to analyse this trade-off. In both settings, traders decide where to submit their transactions and how much to pay in priority fees to improve their execution position, while builders compete for block inclusion and extract rents from transaction ordering. The models differ in their focus: the first examines competition between two private builders in the MEV-Boost auction; the second studies equilibrium submission behaviour when builders are randomly selected and only a specific fraction of them can process private order flow. Together, they characterise how liquidity conditions, builder composition, and market design determine whether order flow becomes concentrated or fragmented across execution channels.

Model I: competition between private builders. The first model studies how blockchain transaction flow fragments among competing execution channels—two private memory pools managed by builders and a public pool accessible to all. First, traders with random liquidity needs decide where to submit their transactions and compete in a priority gas auction to obtain better execution prices. Second, builders compete in an *MEV-Boost auction* to be selected by the proposer to create the next block. Key to order flow fragmentation is the expected cost of the priority gas auction. The magnitude of adverse price impact from competing traders in the DEX depends on the depth of liquidity. A low depth of liquidity increases the cost of trading late in the queue, thereby amplifying competition costs for priority in the block.

Our model uncovers a new channel that promotes centralisation in the blockchain beyond the avoidance of MEV attacks—namely, DEX liquidity. We show that DEX liquidity is a key factor in determining how order flow fragments among builders. When liquidity is abundant, traders concentrate in a single private pool to maximise their execution probability. Concentration in a private pool increases the likelihood that the associated builder wins the MEV-Boost auction. This strategic behaviour leads to market concentration in the hands of centralised entities. Conversely, when liquidity is scarce, competition for queue priority becomes the dominant force, and traders distribute their orders more evenly among private and public memory pools.

The model shows that consensus design and liquidity conditions have broader implications for market power in the blockchain. By internalising valuable order flow, builders with private pools provide MEV protection but gain a competitive advantage in the MEV-Boost auction. These effects reinforce market centralisation because the main builders are few (typically two) and privately operated. Hence, private pools create a fundamental tension between individual protection and system-wide concentration.

Model II: the role of builder composition. The second model abstracts from direct competition between builders and instead focuses on traders' strategic choice between public and private execution channels when the blockchain features an exogenous proportion α of public-only builders. Here, two traders interact, and builders are randomly selected to create the next block. A trader's expected payoff depends on whether her transaction is executed by a public builder (guaranteed inclusion but MEV exposure) or a private one (execution uncertainty but no MEV cost).

The analysis reveals how α governs equilibrium behaviour. When α is low, both traders optimally submit to the private pool. As α increases, one trader moves to the public pool to avoid execution risk, and for high enough α, both traders use the public pool exclusively. We also characterise the range of α values that can sustain builder equilibria. Equilibrium can occur either at $\alpha = 0$—where all builders accept private flow—or for sufficiently large α, where only public execution remains.

Related literature. This chapter contributes to the growing literature on decentralised finance, blockchain market design, and microstructure. It relates to three main strands of research: (i) the economics of blockchains, (ii) the microstructure of decentralised exchanges, and (iii) auction theory.

This work is closely related to the literature studying the allocation of order flow and the effects of MEV in blockchain systems. [20] analyse the trade-offs between public and private memory pools in the presence of frontrunning attacks, showing how attackers and heterogeneous traders affect welfare and execution outcomes. We complement their analysis by focusing instead on competition among builders as a driver of order flow fragmentation. The chapter also connects to the broader literature on the economics of blockchains (see [56]).

Several studies analyse the mechanics and welfare properties of AMMs, which form the backbone of trading in DEXs. The most popular AMM as of 2025 is Uniswap, which reached the v4 version in January 2025, and can operate on Ethereum. [11] provide an overview of the v4 protocol.[6], [62], [22], [29], and [30] show that existing AMM designs impose losses on liquidity suppliers due to adverse selection and price impact. [14] study the optimal exit time for liquidity providers in AMMs, and [12] examine optimal dynamic fees in AMMs. [7] focus on the equilibrium reward for liquidity providers in AMMs. [59] compare AMMs with traditional limit-order-book markets, while [66] analyse the various trading costs that arise in AMM environments. [23] study how gas fees affect trade informativeness in AMMs, and [32] propose design modifications to reduce liquidity providers' losses. Unlike these papers, which focus on liquidity provision mechanisms within DEXs, we examine the interaction between traders' submission strategies and the blockchain infrastructure itself—specifically, how public and private memory pools determine the allocation of order flow and the degree of market concentration. Our framework also connects to the literature on information revelation and competition among traders in traditional exchanges ([55], [43] and [8]). Finally, there is a broader literature that explores how blockchain technology transforms financial intermediation and market organisation ([52], [37]).

The remainder of the chapter is organised as follows. Section 4.2 provides background on blockchain protocols, DEXs, and private order flow. Section 4.3 introduces the first model with competing builders in the MEV-Boost auction and multiple traders. Section 4.4 develops the second model with exogenous builder composition. Section 4.5 concludes.

4.2 Institutional details

This section presents technical details and the characteristics of blockchain protocol components.

4.2.1 Blockchains

A blockchain is a distributed digital ledger stored by participants (referred to as nodes), within a public network. Any entity can be a node and maintain a copy of the ledger. A transaction is executed on the blockchain if it is in a block created by a builder. Blocks are sequentially added to the ledger at regular intervals called slots, whose duration is referred to as block time. For example, the block time in Ethereum, which is the main blockchain hosting the most liquid DEXs, is twelve seconds.

When agents submit transactions, they are not instantly included in a block. Instead, transactions are transmitted to network nodes and stored in the local memory, known as memory pool. The transactions propagate through the network until it is visible to all the nodes. Therefore, blockchain activity through memory pools is publicly visible before transactions are included in blocks. The block builder selected at each slot chooses transactions from their local memory pool to create the next block. Builders are paid gas fees by traders in order to have their transactions included in the block. Gas fees are paid in the blockchain's native currency. In the rest of this section, we focus on the particularities of the Ethereum blockchain.

In Ethereum's memory pools, gas fees consist of two components: a *base fee* and a *priority fee*. The base fee is paid to all nodes and is compulsory for inclusion in a block. Base fees represent a fixed cost for blockchain users. In contrast, the priority fee is optional and paid exclusively to the builder to incentivise the inclusion of a transaction in a block. Priority fees also determine the ordering of transactions: builders rank transactions according to priority fees. Higher priority fees imply better position in the queue.

Rent extraction in the form of MEV can take the form of builders including transactions submitted by "bots" or "attackers" who profit from the visibility of pending transactions in the memory pool. This *pre-trade transparency* also incentivises agents to (re)submit their transactions with higher priority fees, leading to *priority gas auctions* (PGAs) in which agents compete for execution priority within the block.

PGAs underpin the microstructure of DEXs on Ethereum. Traders compete by bidding higher priority fees to secure early execution and avoid adverse price impact from preceding trades. The equilibrium outcome is that traders submit their bids at the last possible moment so that private information is revealed only at block frequency (see [21]).

4.2.2 The Consensus Layer

Ethereum's Proof-of-Stake consensus protocol relies on a large set of validators, each of whom has staked at least 32 ETH[1] in the blockchain. For each block, validators are randomly selected to act as proposers, whose role is to propose a block for inclusion in the chain, and as attesters, who vote on the validity of proposed blocks. Every twelve seconds, a new slot begins, during which one validator is designated as proposer while others act as attesters. Attesters verify that the block proposed is valid and extend the chain. This provides finality through majority consensus; see [56] for more details.

The Ethereum blockchain adopted *proposer-builder separation* (PBS). In this consensus protocol, the proposer's responsibility is to select the most profitable block from among those provided by builders. Builders compete in an auction to have the block selected by the proposer. This auction, referred to as MEV-Boost auction, is key to determine the microstructural effects and welfare consequences of blockchains which we study below. In the MEV-Boost auction, builders continuously compete within each blockchain slot by submitting bids similar to an English auction, through which the proposer selects the winning builder for the next block. Importantly, the proposer does not know the full content of the block before signing. This prevents manipulation and ensures that transaction ordering is carried out by the builder.

Validators can also participate in block building, however, they face competition from dedicated and highly sophisticated block builders who can efficiently extract value from blocks and more easily win the MEV-Boost auctions. In practice, the builder market is highly concentrated: by 2024, the top three builders created more than 70% of blocks, while dozens of smaller builders competed with significantly less success.

PBS also introduced specialised agents such as searchers. This raises new strategic tradeoffs in transaction ordering which we study in the model below. Searchers, or MEV searchers, are automated bots that employ frontrunning or backrunning of transactions in the memory pool. In frontrunning, the searcher outbids the priority

[1]ETH, or Ether, is the native cryptocurrency of Ethereum.

fee of an agent to ensure that their transaction is processed first in the next block, and in backrunning, the searcher positions their transaction after the agent's transaction. Builders can decide which transactions to include in a block, including their own and bundles submitted by searchers who compete aggressively in priority fee auctions.

PBS is designed to decentralise MEV revenues from block creation, since they were concentrated in the hands of just a few sophisticated validators under the proof-of-work protocol. With PBS, Ethereum validators—namely, the thousands of nodes participating in the consensus layer—can enhance their profits by auctioning off block space to a competitive, yet relatively small, set of builders. In doing so, validators also capture a share of the MEV that was previously inaccessible to them.

4.2.3 Private flow

Agents operating on the Ethereum blockchain face costs due to builders, searchers, and competition for block priority. Private channels recently emerged in the blockchain ecosystem as a protection for blockchain users. Agents that use these private channels send their transactions directly to participants in the consensus layer such as builders, validators, or third-party service providers. The objective of private channels is to hide transaction details from being publicly accessible and being targeted in predatory MEV extraction like sandwich attacks. Private order flow refers to transactions that bypass the public memory pool and are transmitted directly through these private communication channels.

At present, Ethereum supports multiple mechanisms for submitting private transactions, including Flashbots Protect, MEV-Blocker, BackRunMe, Wallet-to-Builder, and Order Flow Auctions. In what follows, we focus on the most common form of private order flow, whereby traders submit their transactions directly to a reputed builder. The builder can include and execute the transaction if it is selected to produce the next block.

4.3 Order flow fragmentation and PBS

In this section, we study order flow fragmentation in the blockchain between decentralised channels (the public memory pool) and centralised channels (private memory pools of builders). At present, the two largest builders in the Ethereum ecosystem—Titan Builder and BuilderNet—produce more than 86% of blocks and both offer private order flow channels. Here, we model the strategic interaction between blockchain traders and the two builders as a two-stage game.

In the first stage, $M > 0$ traders, each with random liquidity needs, choose the channel to submit their orders: one of the two private memory pools or the public memory pool. In the second stage, the two builders compete in an MEV-Boost auction. Each builder maintains exclusivity over the private pool it operates. Their valuations of the block consist of a random component, representing uncertainty in the estimation of MEV within the block, and a deterministic component that scales with the volume of private order flow each builder receives.

4.3.1 Decentralised Exchanges

In our model, the orders of traders are executed on a DEX. DEXs are among the most significant financial applications of blockchain technology. By 2023, their trading volumes reached roughly $40 billion per month, with automated market makers (AMMs) constituting the dominant design ([28, 31]).

We consider a DEX operating on a blockchain that provides liquidity for a risky asset whose initial value in stage one is normalised to $V_0 = 0$. DEXs employ algorithmic trading functions that define iso-liquidity curves based on the reserves of two assets (see [32]). Unlike traditional order books, DEXs provide continuous liquidity, but at the cost of *impermanent losses* when prices move (see [29, 30]).

In DEXs, a set of deterministic rules governs how limited liquidity supply, represented by the aggregate reserves in the liquidity pool, determines execution prices and price impact (see [41]). In particular, price impact increases with the traded quantity of the risky asset and decreases with the available liquidity. In our model, the price impact of a buy (resp. sell) order of size δ is $2\,\eta\,\delta$ (resp. $-2\,\eta\,\delta$), where $\eta > 0$ measures the liquidity depth of the DEX. A higher value of η corresponds to lower liquidity depth and thus higher price impact (see [21]) for further details.

4.3.2 Only two private memory pools

To better highlight the main economic trade-offs and analytical results, in this section we solve a simplified version of the model in which only the two private pools are available as channels for submitting transactions, while the public memory pool is not accessible.

In stage one, $0 \leq N \leq M$ traders choose private pool 1, while the remaining $M - N$ traders choose private pool 2, where N is determined in equilibrium. This stage is crucial for traders: if they select a private memory pool whose builder is not chosen to create the next block, their transactions remain unexecuted, yielding a payoff of zero. Conversely, if they select a private pool whose builder is chosen, they compete in a priority gas auction (PGA) to fulfil their liquidity needs at the best expected price. We assume traders know M (the total number of them) and play simultaneously.

We show that there exists a pure-strategy Nash equilibrium describing the fragmentation of order flow across the blockchain, in which there is no incentive for traders to deviate from their chosen memory pool. This equilibrium reflects the trade-off between the losses arising from competition in PGAs—when too many traders concentrate in the same pool—and the gains from a higher probability that the corresponding builder is selected. We model the PGA as a first-price sealed-bid (FPSB) auction, where traders compete to avoid the adverse price impact of competing orders that would otherwise erode their profits.

In stage two, builders take the equilibrium fragmentation of order flow as given and compete to build the next block through the MEV-Boost auction. We model this competition as a first-price sealed-bid auction, where each builder's valuation of the block consists of two components: a deterministic term that scales linearly with the volume of private transactions submitted to its pool, and a random term capturing uncertainty in expected MEV.

Next, the model is solved by backward induction.

4.3.2.1 Stage two: MEV-Boost auction

In stage two, builder 1 takes as given the number N of transactions submitted to its private pool, while builder 2 takes as given the number $M - N$ of transactions submitted to its private pool. We assume that both builders know the total number of traders M, so that each builder is aware of the volume of private order flow received by both. Hence, both builders know the values M, N and $M - N$. However, builder 2 cannot observe the random component ϵ_1 in the valuation of builder 1, and similarly builder 1 cannot observe the random component ϵ_2 in the valuation of builder 2. We assume builders do not use MEV strategies, to avoid the reputational cost they would likely face, which would imply losing order flow in the future. Therefore, their value of the block does not account for the possibility of rent extraction in the form of MEV.

Builder 1 values the block as

$$V_1 = \alpha N + \epsilon_1, \tag{4.1}$$

where ϵ_1 represents an estimation error uniformly distributed over $[-\alpha N, \, \overline{\epsilon}\, \alpha N]$. The estimation error depends on N in order to guarantee that builders never have negative valuations. We assume $\alpha > 0$ is just a positive multiplicative coefficient. The term $\alpha N > 0$ captures the value extracted from the privately observed transactions in the builder's memory pool, which increases with the number of submitted transactions. The term $\overline{\epsilon}$ is related to the upper bound of the estimation error and encapsulates the maximum percentage increase over αN, with $\overline{\epsilon} = 1$ implying a maximum increase of 100% in the valuation. Similarly, builder 2 values the block as

$$V_2 = \alpha(M - N) + \epsilon_2, \tag{4.2}$$

where ϵ_2 represents an estimation error uniformly distributed over $[-\alpha\,(M-N), \overline{\epsilon}\,\alpha\,(M-N)]$. Thus, V_1 is uniformly distributed over $[0, \alpha\, N\,(1+\overline{\epsilon})]$, and V_2 is uniformly distributed over $[0, \alpha(M-N)(1+\overline{\epsilon})]$.

For simplicity, we do not include the explicit formulae for the priority fees paid by traders in stage one. Instead, we assume that these fees are not observed by the builder, who applies the deterministic pricing scheme with coefficient α described above. The key economic mechanism we aim to capture is that a builder's expected revenue scales with the number of privately received transactions.

We model the MEV-Boost auction as a first-price sealed-bid auction with heterogeneous bidders. In particular, one builder is *strong*, in the sense that it values the block more highly than the *weak* builder (the case with homogeneous bidders only happens when $N = M - N$, and hence the probability builder 1 wins the auction would be the same as the probability builder 2 wins the auction, which is $\frac{1}{2}$ by symmetry). For the harder case $N \neq M - N$ we proceed as follows: Since V_1 is uniformly distributed over $[0, \alpha\, N\,(1+\overline{\epsilon})]$ and V_2 is uniformly distributed over $[0, \alpha(M-N)(1+\overline{\epsilon})]$, comparing the values of N and $M - N$ allows to determine the *strong* and the *weak* builder. For example, if $N < M - N$, the first builder is *weak* and the second one is *strong*. We define $\overline{B}$ as the maximum possible valuation of the *weak* builder, which is $\alpha\, N\,(1+\overline{\epsilon})$ when $N < M - N$. We also define δ as the difference between the maximum possible valuation of the strong builder and the maximum possible valuation of the weak builder. For example, if $N < M - N$, $\delta = \alpha\,(M-N)\,(1+\overline{\epsilon}) - \alpha\, N\,(1+\overline{\epsilon}) = \alpha\,(M-2N)\,(1+\overline{\epsilon})$. Therefore, when $N < M-N$, V_1 is uniformly distributed over $[0, \overline{B}]$ and V_2 is uniformly distributed over $[0, \overline{B} + \delta]$. Then, we let $b_1(v, \delta)$ and $b_2(v, \delta)$ be the equilibrium bid functions of builders 1 and 2 respectively (recalling that $0 \leq v \leq \overline{B}$ for b_1 and $0 \leq v \leq \overline{B} + \delta$ for b_2). Moreover, the equilibrium inverse bid functions are denoted as $v_1(b,\delta)$ and $v_2(b,\delta)$. We characterise the equilibrium bidding strategies in this auction in the next result.

Proposition 5. *The strong builder's bid distribution stochastically dominates that of the weak builder. Moreover, assuming $N < M - N$ (the reverse case is analogous), the equilibrium inverse bid functions in the MEV-Boost auction are*

$$\begin{cases} v_1(b,\delta) = \dfrac{2b\overline{B}^2(\delta+\overline{B})^2}{b^2\delta(\delta+2\overline{B}) + \overline{B}^2(\delta+\overline{B})^2}, \\ v_2(b,\delta) = \dfrac{2b\overline{B}^2(\delta+\overline{B})^2}{b^2\delta(\delta+2\overline{B}) - \overline{B}^2(\delta+\overline{B})^2}, \end{cases} \tag{4.3}$$

where

$$\overline{B} = \alpha N(1+\bar{\epsilon}) \qquad \textit{and} \qquad \delta = \alpha(1+\bar{\epsilon})(M-2N).$$

The previous result shows that the builder with the higher range of valuations for the block—that is, the builder with the greater number of transactions in its private memory pool—tends to submit a higher bid in the MEV-Boost auction and, consequently, to win the block. This finding is consistent with results from related auction-theoretic settings (see, for example, Lemma 1 in [18]).

While the previous result characterises the block-building winning probability for a fixed order-flow fragmentation, the next result examines how the probability of winning the MEV-Boost auction varies with the difference in the number of transactions submitted to each builder's private memory pool. In particular, we show that the greater this difference, the higher the probability that the strong builder wins the auction.

Proposition 6. *Let the builders' valuations be independently distributed as*

$$V_1 \sim U\left[0, \overline{B}\right], \qquad V_2 \sim U\left[0, \overline{B}+\delta\right].$$

Then the probability that the stronger builder wins the auction increases when δ increases if $\overline{B} > 0$ and $0 < \delta < \overline{B}/2$.

The previous result shows that the stochastic dominance of the bids submitted by the builder with the greater number of private orders increases with its excess number of trades. In particular, if all traders choose the same private pool, that pool is selected with probability 1.

4.3.2.2 Stage one

In stage one, M traders decide whether to submit their orders to one private pool or the other. Traders recognise that concentrating orders in a single pool increases the probability this pool is selected, but also intensifies competition for priority in the next block. Let $0 \leq N \leq M$ denote the number of traders who choose private pool 1, while the remaining $M - N$ traders choose private pool 2. We denote by $g(N)$ the probability that private pool 1 is selected, as characterised in Proposition 6. Recall that g is increasing in N, satisfies $g(0) = 0$, $g(M) = 1$, and, if M is even, $g(M/2) = 1/2$.

Next, we describe the PGA among traders when private builder 1 is selected. For simplicity, we assume that all traders face random liquidity needs that are independent and identically distributed according to a uniform distribution on $[0, \bar{\delta}]$, where $\bar{\delta} > 0$. Accordingly, all traders wish to buy the risky asset in the DEX; see [21] for a similar modelling approach. The case in which all traders wish to sell the asset is symmetric and leads to equivalent results.

Let $\{\varphi_1, \ldots, \varphi_N\}$ denote the priority fees submitted by the N traders who chose private pool 1, and let $\{\delta_1, \ldots, \delta_N\}$ denote their individual liquidity needs (if private builder 2 is selected instead, the analysis proceeds analogously by replacing N with $M - N$).

The PGA is modelled as a FPSB auction, where only the trader submitting the highest priority fee wins. The remaining $N - 1$ traders are placed randomly within the block. Although an all-pay auction would more accurately reflect the actual mechanism, our simplified setting focuses on the key economic effect: the expected priority fee paid by each trader increases with the number of competitors in the PGA. We formally establish this result below.

If trader i wins the auction, that is, if $\varphi_i > \varphi_j$ for all $j \neq i$, her buy order is executed first, thereby avoiding the costs associated with a worse position in the queue described in Section 4.3.1. When a trader loses the PGA, the price impact on her orders depends on the number of competitors. More precisely, on average, $(N - 1)/2$ orders are executed before trader i's order. Consequently, her execution price is given by

$$V_0 + 2\,\eta\,\frac{N-1}{2}\,\delta_i = \eta\,(N-1)\,\delta_i, \tag{4.4}$$

where V_0 denotes the initial value of the risky asset, and, as previously defined, $\eta > 0$ measures the liquidity depth of the DEX.

We further assume that the final value V of the risky asset is independent of the liquidity needs $\{\delta_1, \ldots, \delta_N\}$ and the mean of V is μ. Under these assumptions, trader i obtains the following payoffs. If she wins,

$$\Pi_i^{\text{win}} = -\varphi_i + \delta_i V, \tag{4.5}$$

and if she loses,

$$\Pi_i^{\text{lose}} = -\eta(N-1)\delta_i^2 + \delta_i V. \tag{4.6}$$

Hence, her expected payoff is

$$\Pi_i = \mathbb{P}[\varphi_i > \varphi_z] \left[-\varphi_i + \eta(N-1)\delta_i^2\right] + \delta_i V - \eta(N-1)\delta_i^2, \tag{4.7}$$

where φ_z is the second largest bid of all submitted transactions. To derive the expected payoff of each trader in private memory pool 1, we first establish the following auxiliary result.

Proposition 7. *Let $\delta_1, \ldots, \delta_N$ be i.i.d. random variables uniformly distributed on $[0, \bar{\delta}]$. Then, the expected maximum of their squared values is given by*

$$\mathbb{E}\left[\max\{\delta_1^2, \ldots, \delta_N^2\}\right] = \bar{\delta}^2 - \bar{\delta}^2 \frac{2}{N+2} = \bar{\delta}^2 \frac{N}{N+2}. \tag{4.8}$$

The following result characterises the equilibrium priority fees and expected payoffs of traders.

Proposition 8. *Consider $N \geq 1$ traders competing in a private pool whose corresponding private builder has been selected to create the next block. The optimal priority fee for trader i is*

$$\varphi_i^* = \eta(N-1)\delta_i^2 \frac{N-1}{N} = \frac{\eta(N-1)^2}{N} \delta_i^2. \tag{4.9}$$

The expected payoff for any trader i, conditional on submitting the optimal priority fee, is

$$f(N) = \frac{\eta(N-1)\bar{\delta}^2}{N(N+2)} + \frac{\mu\bar{\delta}}{2} - \frac{\eta\bar{\delta}^2(N-1)}{3}. \tag{4.10}$$

The optimal priority fee (4.9) *is increasing in the number of competitors N and the payoff* (4.10) *is decreasing in N.*

Proposition 8 characterises the payoff of each trader in a pool with N traders, managed by a builder who is selected by the blockchain protocol. In particular, this payoff is decreasing in the number of competitors within the same pool. However, this payoff is only realised when the builder is selected, which occurs with a probability that increases in the number of traders that have chosen the private pool associated to this builder. Hence, when deciding which pool to submit to, traders face a fundamental trade-off between a higher probability of builder selection and lower competition that improves their expected payoff.

The two private pools are symmetric. Hence, if traders split their orders such that N submit to private pool 1 and $M - N$ submit to private pool 2, this equilibrium is equivalent to the one in which $M - N$ traders submit to private pool 1 and N to private pool 2. Our main result below shows that, given this trade-off, there always exists a pure-strategy Nash equilibrium for the number of traders N.

Proposition 9. *Given M traders who allocate their order flow between two private memory pools, there exists a Nash equilibrium in pure strategies.*

Proof. We begin by considering the case where M is even. Let (A, B) denote the state in which A traders submit their orders to private pool 1 and B traders submit their orders to private pool 2, with $A + B = M$, $A \geq 0$, and $B \geq 0$. If either $(M/2,\ M/2)$ or $(M, 0)$ constitutes a Nash equilibrium, the result follows immediately. Suppose, however, that neither $(M/2,\ M/2)$ nor $(M, 0)$ is a Nash equilibrium. We show that in this case there must exist a Nash equilibrium $(N, M - N)$ with $\frac{M}{2} < N < M$.

We proceed by contradiction. Assume that $(N, M - N)$ is not a Nash equilibrium for any $M/2 < N < M$. Some trader in private pool 1 must have a profitable deviation to private pool 2, moving the system to state $(M - 1, 1)$, since $(M, 0)$ is not a Nash equilibrium. Because $(M - 1, 1)$ is not a Nash equilibrium, and the trader in private pool 2 does not have a profitable deviation back to pool 1, this implies that at least one trader in private pool 1 has a profitable deviation to private pool 2, leading to state $(M - 2, 2)$. Continuing in this manner, traders sequentially deviate, generating the sequence of states

$$(M - 1, 1), (M - 2, 2), (M - 3, 3), \ldots, \left(\frac{M}{2} + 1, \frac{M}{2} - 1\right).$$

At state $\left(\frac{M}{2} + 1, \frac{M}{2} - 1\right)$, since this state is not a Nash equilibrium, by the same reasoning, some trader in private pool 1 must profitably deviate to private pool 2, reaching state $(M/2,\ M/2)$. However, in this state no trader would have a profitable deviation, which contradicts the assumption that $(M/2,\ M/2)$ is not a Nash equilibrium.

When M is odd, the reasoning is analogous. Assume that no Nash equilibrium exists for any $(N, M - N)$ with $0 \leq N \leq M$. Starting from state $(M, 0)$ and applying the same logic as before, we obtain the sequence of profitable deviations

$$(M-1,1),(M-2,2),\ldots,\left(\frac{M+1}{2},\frac{M-1}{2}\right).$$

At state $\left(\frac{M+1}{2}, \frac{M-1}{2}\right)$, some trader in private pool 1 must have a profitable deviation to private pool 2, reaching state $\left(\frac{M-1}{2}, \frac{M+1}{2}\right)$, because state $\left(\frac{M+1}{2}, \frac{M-1}{2}\right)$ is not a Nash equlibrium and no trader in private pool 2 has a profitable deviation (as any move would reduce her payoff). But this leads to a contradiction because the deviating trader would then be in a pool with the same number $(\frac{M+1}{2})$ of traders as before and thus would obtain the same payoff. Therefore, a Nash equilibrium in pure strategies must exist. □

The previous result establishes the existence of an equilibrium. The following proposition characterises which Nash equilibrium order-flow fragmentation arises when the depth of liquidity is either significantly large or significantly small.

Proposition 10. *For $\eta \to 0$, the unique Nash equilibrium is $(M, 0)$. For $\eta \to \infty$, the unique Nash equilibrium is $(M/2,\ M/2)$ if M is even, and $(\frac{M+1}{2}, \frac{M-1}{2})$ if M is odd.*

Proof. First, consider the case $\eta \to 0$. In this limit, for all N, we have $f(N) \to \frac{\mu\bar{\delta}}{2}$, where $\frac{\mu\bar{\delta}}{2} > 0$. Therefore, each trader prefers to join the pool with the higher probability of being selected. Hence, $(M, 0)$ is a Nash equilibrium, because traders do not have an incentive to deviate to the pool that has a lower probability of being selected.

Regarding uniqueness, note that $(N, M-N)$ cannot be a Nash equilibrium for any $1 \leq N \leq M-1$, because any trader in the smaller pool—the one with $\min\{N, M-N\}$ traders—has a profitable deviation by moving to the other pool, which offers a higher selection probability.

Now, consider the case $\eta \to \infty$. In this limit, for all $N > 1$, we have $f(N) \to -\infty$, implying that traders now prefer to join the pool with a lower probability of being selected to avoid excessive competition. Consequently, $(M/2,\ M/2)$ (for M even) and $(\frac{M+1}{2}, \frac{M-1}{2})$ (for M odd) are Nash equilibria, because traders do not have a profitable deviation: if a trader deviates, that would place her in a pool with an equal or greater number of competitors.

Regarding uniqueness, all other configurations $(N, M-N)$ cannot be Nash equilibria, because any trader in the larger pool—the one with $\max\{N, M-N\}$ traders—has a profitable deviation by moving to the smaller pool, which has a lower probability of being selected. □

The previous result is key, it illustrates how order flow fragments across private pools depending on liquidity conditions. When liquidity is cheap, the adverse price impact from trades by competitors is negligible, and the dominant economic force is the pool's probability of being selected. In this case, traders concentrate their orders in a single private pool. Conversely, when liquidity is expensive, the dominant effect arises from the PGA. To minimise competition and reduce expected priority fees that increase with η, traders split their orders evenly between the two available execution channels.

The next result compares the Nash equilibria that arise under different levels of liquidity depth in the DEX. In particular, it shows that as the value of η increases, equilibrium outcomes tend to exhibit greater fragmentation of order flow, with both pools receiving a more balanced number of submissions.

Proposition 11. *Let $\eta > 0$. Suppose that state (A, B) with $A \geq B$ is a Nash equilibrium such that A is maximal—that is, the states $(A+1, B-1)$, $(A+2, B-2)$, ..., $(M, 0)$ are not Nash equilibria. Then, for any $\eta' > \eta$, the states $(A+1, B-1)$, $(A+2, B-2)$, ..., $(M, 0)$ are not Nash equilibria.*

Proof. We begin by analysing the auxiliary function

$$\tilde{f}(N) = \frac{N-1}{3} - \frac{N-1}{N(N+2)}, \quad N \geq 1, \tag{4.11}$$

and deriving some of its properties. Its derivative is given by

$$\tilde{f}'(N) = \frac{1}{3} + \frac{N^2 - 2N - 2}{N^2(N+2)^2}. \tag{4.12}$$

Since $N^2 - 2N - 2 = 0$ if and only if $N = 1 \pm \sqrt{3}$, the function $h(N) = -\frac{N-1}{N(N+2)}$ decreases on $(1, 1+\sqrt{3})$ and increases thereafter. Next, $h(N)$ (and hence $\tilde{f}(N)$) is increasing for $N \geq 3$ because $1 + \sqrt{3} < 3$. Moreover, for integer values of N, we have

$$\tilde{f}(1) = 0,\ \tilde{f}(2) = \frac{1}{3} - \frac{1}{8},\ \tilde{f}(3) = \frac{2}{3} - \frac{2}{15}.$$

Therefore, $\tilde{f}(N) < \tilde{f}(N+1)$ for all integer $N \geq 1$. We have $\tilde{f}(N) > 0$ for all $N > 1$ and $\tilde{f}(1) = 0$.

Next, we use these properties to show that, for $N > 1$, $f(N)$ decreases as η increases. Indeed, for $\eta' > \eta$,

$$f(N;\eta) - f(N;\eta') = \bar{\delta}^2(\eta' - \eta)\left(\frac{N-1}{3} - \frac{N-1}{N(N+2)}\right) > 0 \tag{4.13}$$

because $\tilde{f}(N) > 0$ for all integer $N > 1$. Furthermore, for $N_1 > N_2$, we have

$$f(N_1;\eta) - f(N_1;\eta') > f(N_2;\eta) - f(N_2;\eta') \iff \tilde{f}(N_1) > \tilde{f}(N_2),$$

which happens because $\tilde{f}(N)$ is increasing in N.

Given η, and using the same reasoning as in the proof of Proposition 9, any trader in private pool 1 has a profitable deviation to private pool 2 in states $(A+1, B-1)$, $(A+2, B-2)$, ..., $(M, 0)$. Thus, for a given state $(A+k, B-k)$ with $k \geq 1$,

$$f(B-k+1;\eta)\, g(B-k+1) > f(A+k;\eta)\, g(A+k).$$

If $\eta' > \eta$, then from the previous results we have

$$\begin{aligned} f(B-k+1;\eta') &= f(B-k+1;\eta) - \varepsilon, \\ f(A+k;\eta') &= f(A+k;\eta) - \varepsilon - \varepsilon', \end{aligned}$$

for some $\varepsilon \geq 0$, and $\varepsilon' > 0$. Hence, state $(A+k, B-k)$ cannot be a Nash equilibrium for $\eta' > \eta$, because any trader in private pool 1 has a profitable deviation to private pool 2. In particular,

$$\begin{aligned} & f(B-k+1;\eta')\, g(B-k+1) > f(A+k;\eta')\, g(A+k) \\ \iff & \left(f(B-k+1;\eta) - \varepsilon\right)\, g(B-k+1) > \left(f(A+k;\eta) - \varepsilon - \varepsilon'\right)\, g(A+k). \end{aligned}$$

Rearranging terms, this implies

$$\begin{aligned} f(B-k+1;&\eta)g(B-k+1) - \varepsilon g(B-k+1) \\ & > f(A+k;\eta)g(A+k) - \varepsilon g(A+k) - \varepsilon' g(A+k), \end{aligned}$$

which holds because

$$f(B-k+1;\eta)g(B-k+1) \geq f(A+k;\eta)g(A+k),$$

and

$$g(A+k) > \frac{1}{2} \geq g(B-k+1).$$

Thus,

$$-\varepsilon g(B-k+1) \geq -\varepsilon g(A+k) > -\varepsilon g(A+k) - \varepsilon' g(A+k),$$

confirming that the inequality is preserved and that the state $(A+k, B-k)$ cannot be a Nash equilibrium for $\eta' > \eta$. □

We have established in Proposition 9 the existence of a Nash equilibrium in pure strategies, and in Proposition 10 that for $\eta \to 0$ the unique equilibrium is $(M, 0)$. Combining them with Proposition 11, we can see the following: Define

$$\eta_1^* = \inf \{\eta > 0 : (M, 0) \text{ is not a Nash equilibrium}\}.$$

At $\eta = \eta_1^*$, there exists a Nash equilibrium (A, B) with maximal $A < M$. Proceeding analogously, let

$$\eta_2^* = \inf \{\eta > \eta_1^* : (A, B) \text{ is not a Nash equilibrium}\}.$$

Then, at $\eta = \eta_2^*$, there exists a Nash equilibrium (A', B') with $A' < A$. Repeating this reasoning iteratively, we find that, as the value of η increases (for M even), there exists a sufficiently large η for which the unique Nash equilibrium is symmetric, namely

$$\left(\tfrac{M}{2}, \tfrac{M}{2}\right).$$

At this threshold liquidity depth level, the dominant force in the trade-off faced by traders becomes minimising losses from PGA competition.

4.3.3 Centralised versus Decentralised order flow

In this section, we characterise order-flow fragmentation in the presence of a public memory pool competing with private channels. Let N_1 denote the number of traders who submit their transactions to private pool 1, N_2 the number of traders who submit to private pool 2, and $M - N_1 - N_2$ the number of traders who submit to the public pool, where

$$0 \le N_1 \le M, \quad 0 \le N_2 \le M, \quad \text{and} \quad 0 \le N_1 + N_2 \le M.$$

In practice, both private builders have access to the transactions sent to the public memory pool. For transactions submitted to the public pool, traders incur a cost $\pi > 0$, representing the per-unit MEV extracted by bots and searchers.

Similar to the reasoning in the previous section, we now characterise the payoff obtained by a trader who submits to private pool 1, conditional on builder 1 being selected. If trader i wins the PGA, her payoff is

$$\Pi_i^{\text{win}} = -\varphi_i + \delta_i V. \tag{4.14}$$

If she loses, her order competes not only with traders in the same private pool but also with those in the public memory pool whose transactions are visible to the builder. In this case, she competes with $M - N_2$ traders, and her payoff is

$$\Pi_i^{\text{lose}} = -\eta\, \delta_i^2 (M - N_2) + \delta_i V. \tag{4.15}$$

Hence, the trader's expected payoff is given by

$$\mathbb{E}[\Pi_i] = g(N_1, N_1 + N_2)\, f(M - N_2), \tag{4.16}$$

where $g(\,\cdot\,)$ denotes the probability that builder 1 wins the MEV-Boost auction, and $f(\,\cdot\,)$ denotes the payoff function from the PGA.

Similarly, a trader who submits to private pool 2 obtains the following payoff if private builder 2 is selected. If she wins the PGA, her payoff is

$$-\varphi_i + \delta_i V. \tag{4.17}$$

If she loses, her order competes with other traders in private pool 2 and with traders in the public memory pool whose transactions are visible to the builder. In this case, she competes with $M - N_1$ traders, and her payoff is

$$-\eta\, \delta_i^2 (M - N_1) + \delta_i V. \tag{4.18}$$

Hence, her expected payoff is

$$g(N_2, N_1 + N_2)\, f(M - N_1), \tag{4.19}$$

where, as before, $g(\,\cdot\,)$ denotes the probability that the corresponding builder is selected, and $f(\,\cdot\,)$ denotes the expected payoff from the PGA.

For a trader who submits to the public memory pool, the payoff depends on which private builder wins the MEV-Boost auction. With probability $g(N_1, N_1 + N_2)$, she participates in Game 1, and with probability $g(N_2, N_1 + N_2)$, she participates in Game 2. The payoffs in each game are as follows:

In **Game 1** (builder 1 wins), if the trader wins the PGA, her payoff is

$$-\pi\delta_i - \varphi_i + \delta_i V, \tag{4.20}$$

and if she loses, her payoff is

$$-\pi\delta_i - \eta\, \delta_i^2 (M - N_2) + \delta_i V. \tag{4.21}$$

Thus, her expected payoff in Game 1 is $f(M - N_2) - \pi\delta_i$.

In **Game 2** (builder 2 wins), if she wins the PGA, her payoff is

$$-\pi\delta_i - \varphi_i + \delta_i V, \tag{4.22}$$

and if she loses, her payoff is

$$-\pi\delta_i - \eta\,\delta_i^2(M - N_1) + \delta_i V. \tag{4.23}$$

Hence, her expected payoff in Game 2 is $f(M - N_1) - \pi\delta_i$.

Finally, the overall expected payoff for a trader submitting to the public memory pool is

$$\mathbb{E}[\Pi_i^{\text{public}}] = g(N_1, N_1 + N_2)\left(f(M - N_2) - \pi\delta_i\right) + g(N_2, N_1 + N_2)\left(f(M - N_1) - \pi\delta_i\right). \tag{4.24}$$

In this setting, we obtain an analogous result to that in Proposition 10. We denote by $(N_1, N_2, M - N_1 - N_2)$ the state in which N_1 traders submit to private pool 1, N_2 traders submit to private pool 2, and $M - N_1 - N_2$ traders submit to the public pool. The states $(N_1, N_2, M - N_1 - N_2)$ and $(N_2, N_1, M - N_1 - N_2)$ are equivalent because the two private pools are symmetric, and we can restrict attention to one representative of the pair.

We characterise the Nash equilibrium for sufficiently low and sufficiently high values of η. In particular, we obtain the following result.

Proposition 12. *For $\eta \to 0$, the unique Nash equilibrium is $(M, 0, 0)$.*

For $\eta \to \infty$, the unique Nash equilibrium is $(M/2,\ M/2,\ 0)$ if M is even, and $(\frac{M+1}{2}, \frac{M-1}{2}, 0)$ if M is odd.

Analogously to Proposition 11, we obtain the following result for this setting.

Proposition 13. *Let $\eta > 0$. Suppose that the state $(A, B, 0)$ with $A \geq B$ is a Nash equilibrium such that A is maximal—that is, the states $(A+1, B-1, 0)$, $(A+2, B-2, 0)$, ..., $(M, 0, 0)$ are not Nash equilibria. Then, for any $\eta' > \eta$, the states $(A+1, B-1, 0)$, $(A+2, B-2, 0)$, ..., $(M, 0, 0)$ cannot be Nash equilibria.*

4.4 The role of builder composition in order flow fragmentation

The model of Section 4.3 studies the fragmentation of order flow among builders and focuses on the interaction between private memory pools operated by competing builders. In that framework, traders decide which private pool to use, anticipating that each builder's probability of being selected to produce the next block depends on the total order flow it attracts. The analysis captures how competition between builders in the MEV-Boost auction and competition between traders in priority gas auctions jointly determine market fragmentation of order flow across builders.

By contrast, here, we abstract from competition between builders in the MEV-Boost auction and instead examines traders' strategic choice between public and private execution channels given an exogenous proportion of standard (public-only) builders and builders who offer a private channel for execution. In this section, traders face a coordination problem: choosing the public pool ensures execution but exposes to MEV extraction, while the private pool offers better terms but entails a risk of non-execution. The model characterises equilibrium thresholds for the proportion of private channels and the overall degree of decentralisation.

We consider a DEX that runs on a blockchain where traders can submit transactions either to the public memory pool or to a private one. We present a two-stage game to study how trading separates between public and private pools. In stage one, each builder decides whether to join the network of private builders or to remain a public builder. In this section, private builders can see transactions from both pools in case they are selected to build the block, while public builders can only see transactions from the public pool. Moreover, MEV extraction is limited to transactions from the public memory pool. In equilibrium, builders decide how to fragment between private and public builders, using rational expectations about the equilibrium behaviour of traders competing for queue priority in stage two based on how builders split. The two-stage game is solved by backward induction.

4.4.1 Stage two: Competition between traders.

In stage two, two representative traders, i and j, decide whether to submit their transactions to the public pool or the private pool, and they also compete through priority fees. In particular, they take as given the proportion of private and public builders, and thus know the probability that a builder of either type will be selected to create the next block. We study the traders' payoffs given their submission choices. Each trader has two possible strategies, denoted as *Public* and *Private*, and both make their decisions simultaneously. Hence, their interaction can be modelled as a 2×2 normal-form game. Specifically, there are three distinct subgames: one in which both traders submit their transactions to the public memory pool, one in which both submit to the private pool, and one in which each trader chooses a different builder type. We analyse the equilibrium priority fees and the expected wealth of traders in each subgame, and then characterise the Nash equilibrium of the overall normal-form game.

At the beginning of stage two, each trader observes a private and independent random liquidity need, $\{\delta_i, \delta_j\}$, drawn from a common and known uniform distribution with support $[0, \bar{\delta}]$. As before, we assume that both traders are buyers.

4.4.1.1 Both traders in the public memory pool

If both traders choose to submit their orders to the public pool, their transactions are certainly included in the next block and executed. In this case, traders engage in a sealed-bid first-price auction. Let $\{\varphi_i, \varphi_j\}$ denote the priority fees submitted by trader i and trader j, respectively, and let $\{\delta_i, \delta_j\}$ denote their private liquidity needs.

If trader i wins the auction, that is, if $\varphi_i > \varphi_j$, her buy order is executed first, so the execution price for her trade is V_0. The trader's terminal wealth in this case is

$$W_i(\text{win}) = -\varphi_i - \delta_i V_0 + \delta_i \mu - \pi \delta_i = -\varphi_i + \delta_i \mu - \pi \delta_i,$$

where μ denotes the expected value of the executed unit, and π represents the MEV extracted by the builder per executed unit.

For simplicity, we assume $\mu > \pi$, since otherwise traders would always face negative expected payoffs. In what follows, we normalise $V_0 = 0$ without loss of generality.

Because the game is symmetric, we describe it from the perspective of trader i. From her viewpoint, δ_j is a random variable. If trader j's buy order is executed first, i.e., if $\varphi_i < \varphi_j$, then trader i's execution price becomes $V_0 + \eta\delta_j$. Hence, the competitor's order of size δ_j moves the price by $\eta\delta_j$, where $\eta > 0$ measures the liquidity depth in the DEX. In particular, we assume that each traded unit impacts the price by η. The wealth of trader i in case of losing the auction is then

$$W_i(\text{lose}) = -\varphi_i - \eta\, \delta_i\, \delta_j + \delta_i \mu - \pi \delta_i.$$

Note that $W_i(\text{lose})$ is a random variable, as δ_j is unknown to trader i. Consequently, the expected wealth of trader i is

$$\mathbb{E}_i[W_i] = -\varphi_i + \delta_i(\mu - \pi) - \eta\, \delta_i \int_0^{\bar{\delta}} 1_{\{\varphi_j > \varphi_i\}}\, x\, f(x)\, dx.$$

The optimisation problem that trader i solves to determine her optimal priority fee is therefore

$$\sup_{\varphi_i} \mathbb{E}_i[W_i] = \sup_{\varphi_i} \left\{ -\varphi_i + \eta\, \delta_i \int_0^{\bar{\delta}} 1_{\{\varphi_j < \varphi_i\}}\, x\, f(x)\, dx \right\}, \tag{4.25}$$

which is equivalent to a first-price sealed-bid auction where the *item* corresponds to the adverse price impact avoided by winning transaction priority.

In problem (4.25), the trader chooses her priority fee to maximise the product of the probability of winning the auction and the surplus obtained from queue priority. For each possible value of δ_j, she updates her estimate of the item's value, $\eta\,\delta_i\,\delta_j$, and the corresponding probability of winning the auction. We obtain the following result.

Proposition 14. *The optimal priority fee when both traders compete in the public memory pool is*

$$\varphi_i = \frac{\eta\,\delta_i^3}{3\,\bar{\delta}}\,, \tag{4.26}$$

for trader i, and the expected payoff of trader i is

$$\mathbb{E}[W_i] = \delta_i\left(\mu - \pi - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}} = \delta_i(\mu - \pi) + \left(\frac{\eta\,\delta_i^3}{6\,\bar{\delta}} - \frac{\eta\,\delta_i\,\bar{\delta}}{2}\right). \tag{4.27}$$

To ensure that the expected payoff remains positive for any realisation of the random liquidity need δ_i, we assume

$$\mu > \pi + \frac{\eta\,\bar{\delta}}{2}, \tag{4.28}$$

that is, transactions are profitable after accounting for both MEV extraction and adverse price impact.

4.4.1.2 Both traders submit transactions to the private memory pool

We assume that a proportion $1-\alpha$ of builders participate in the private channel. If both traders submit their orders to the private channel, then with probability α, their transactions are not executed in the next block, in which case their payoff is zero. With probability $1-\alpha$, both traders participate in an auction, and the builder does not extract MEV in the form of π.

When trader j follows the strictly increasing strategy φ, the expected wealth of trader i is given by

$$\begin{aligned}\mathbb{E}[W_i] &= (1-\alpha)\left(-\varphi_i + \mu\,\delta_i - \eta\,\delta_i \int_{\varphi^{-1}(\varphi_i)}^{\bar{\delta}} x\,f(x)\,dx\right)\\ &= (1-\alpha)\left(-\varphi_i + \mu\,\delta_i - \frac{\eta\,\delta_i}{2\,\bar{\delta}}\left(\bar{\delta}^2 - \varphi^{-1}(\varphi_i)^2\right)\right).\end{aligned}$$

Following similar steps as in the previous section, we obtain the next result, which characterises the equilibrium priority fees and payoffs in this case.

Proposition 15. *The optimal priority fee when both traders compete in the private memory pool is*

$$\varphi_i = \frac{\eta\,\delta_i^3}{3\,\bar{\delta}}\,, \tag{4.29}$$

and the expected payoff of each trader is

$$\mathbb{E}[W_i] = (1-\alpha)\left(\delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right). \tag{4.30}$$

The optimal priority fee is the same as in the previous section, where both traders submitted their orders to the public pool. This result is intuitive: both environments are symmetric, and the competitive structure of the auction for queue priority is identical. The differences in expected payoffs stem from MEV extraction and from the distinct probabilities of execution (1 in the public pool versus $1-\alpha$ in the private pool), but these factors do not affect the traders' strategic incentives in the auction. This symmetry no longer holds in the remaining case, where traders are asymmetric—one submits her order to the private pool and the other to the public pool.

4.4.1.3 One trader submits privately, the other publicly

In this case, traders play asymmetrically. Assume without loss of generality that trader i submits her order to the public pool, while trader j submits to the private pool. With probability $\alpha \in [0,1]$, a standard builder is selected by the blockchain protocol; in this case, only the transaction from trader i is included in the next block. Therefore, trader i wins with probability 1, and trader j obtains a zero payoff. With probability $1-\alpha$, a builder from the private channel is selected and creates the next block using transactions from both memory pools. In this event, the traders compete for queue priority. However, trader i incurs MEV losses equal to $-\pi\delta_i$, while trader j does not.

Assuming trader i follows the strictly increasing strategy $\varphi_i(\delta_i)$ and trader j follows the strictly increasing strategy $\varphi_j(\delta_j)$, and that both random liquidity needs are independently drawn from a uniform distribution on $[0,\bar{\delta}]$, we obtain the following result.

Proposition 16. *The optimal priority fees when trader i submits her order to the public memory pool and trader j submits to the private memory pool are*

$$\varphi_i = \frac{\eta(1-\alpha)\delta_i^3}{3\bar{\delta}}\,, \qquad \varphi_j = \frac{\eta(1-\alpha)\delta_j^3}{3\bar{\delta}}\,. \tag{4.31}$$

The expected payoff of trader i is

$$\mathbb{E}[W_i] = \frac{\eta(1-\alpha)\delta_i^3}{6\bar{\delta}} + \delta_i\left(\mu - \pi - \frac{(1-\alpha)\eta\bar{\delta}}{2}\right) \tag{4.32}$$

$$= (1-\alpha)\frac{\eta\delta_i}{6\bar{\delta}}\left(\delta_i^2 - 3\bar{\delta}^2\right) + \delta_i(\mu - \pi). \tag{4.33}$$

The expected payoff of trader j is

$$\mathbb{E}[W_j] = (1-\alpha)\left[(1-\alpha)\eta\delta_j\frac{\delta_j^2 - 3\bar{\delta}^2}{6\bar{\delta}} + \mu\delta_j\right]. \tag{4.34}$$

We note that the payoff of trader i is increasing in α, since $\delta_i^2 - 3\bar{\delta}^2 < 0$. This is intuitive: higher values of α increase the probability that trader i wins the auction directly without facing competition, as the other trader submits her order to the private channel.

However, the payoff of trader j is quadratic in α and can be expressed as the product of a decreasing function of α and an increasing function of α. This is also intuitive: as α increases, trader j benefits more when she participates in the auction (since she bids a lower price), but the probability that she actually takes part in the auction decreases.

The quadratic function describing $\mathbb{E}[W_j]$ has roots

$$\alpha = 1 \qquad \text{and} \qquad \alpha = 1 - \frac{6\bar{\delta}\mu}{\eta(3\bar{\delta}^2 - \delta_j^2)} < 1. \tag{4.35}$$

Moreover, the maximum is achieved when α satisfies

$$2(1-\alpha)\left(\eta\delta_j\frac{\delta_j^2 - 3\bar{\delta}^2}{6\bar{\delta}}\right) = -\mu\delta_j, \tag{4.36}$$

which is necessarily positive, since both $-\mu\delta_j$ and $\delta_j^2 - 3\bar{\delta}^2$ are negative. Denoting this value by $\tilde{\alpha}$, we necessarily have $\tilde{\alpha} < 1$, since $\alpha = 1$ is the larger root of the quadratic function. Therefore, $\mathbb{E}[W_j]$ is increasing on the interval $(0, \tilde{\alpha})$ and decreasing on $(\tilde{\alpha}, 1)$.

4.4.1.4 Nash equilibrium

We now characterise the Nash equilibria that arise in the 2×2 game, using the payoffs derived in the preceding sections. This is a normal-form game where trader i chooses her strategy along the rows and trader j along the columns. The first matrix below shows the payoffs for trader i, and the second matrix shows those for trader j. In both cases, the payoffs depend on the parameter α.

Trader i's payoffs

	Public	Private
Public	$\delta_i\left(\mu-\pi-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_i^3}{6\,\bar{\delta}}$	$(1-\alpha)\left(\frac{\eta\,\delta_i}{6\,\bar{\delta}}\right)(\delta_i^2-3\bar{\delta}^2)+\delta_i(\mu-\pi)$
Private	$(1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2-3\bar{\delta}^2}{6\bar{\delta}}+\mu\,\delta_i\right]$	$(1-\alpha)\left(\delta_i\left(\mu-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right)$

Trader j's payoffs

	Public	Private
Public	$\delta_j\left(\mu-\pi-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_j^3}{6\,\bar{\delta}}$	$(1-\alpha)\left[(1-\alpha)\eta\,\delta_j\frac{\delta_j^2-3\bar{\delta}^2}{6\bar{\delta}}+\mu\,\delta_j\right]$
Private	$(1-\alpha)\left(\frac{\eta\,\delta_j}{6\,\bar{\delta}}\right)(\delta_j^2-3\bar{\delta}^2)+\delta_j(\mu-\pi)$	$(1-\alpha)\left(\delta_j\left(\mu-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_j^3}{6\,\bar{\delta}}\right)$

We now derive the optimal strategy for trader i when trader j plays *Public*. If trader j plays *Public*, then trader i obtains a payoff of

$$\delta_i\left(\mu-\pi-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_i^3}{6\,\bar{\delta}} \tag{4.37}$$

when she also plays *Public*, and a payoff of

$$(1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2-3\bar{\delta}^2}{6\bar{\delta}}+\mu\,\delta_i\right] \tag{4.38}$$

when she instead plays *Private*. We seek the values of α for which these two payoffs are equal, that is, for which

$$\delta_i\left(\mu-\pi-\frac{\eta\,\bar{\delta}}{2}\right)+\frac{\eta\,\delta_i^3}{6\,\bar{\delta}}=(1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2-3\bar{\delta}^2}{6\bar{\delta}}+\mu\,\delta_i\right]. \tag{4.39}$$

Note that the left-hand side does not depend on α, whereas the right-hand side does. We obtain the following result.

Proposition 17 (**Best response to competitor choosing Public**). *Assume trader j submits to the public pool. If $\alpha<\alpha^*$, the best response of trader i is to submit her order to the private pool; otherwise, her best response is to submit to the public pool. The threshold value α^* satisfies*

$$\alpha^*(\delta_i)\in\left(\frac{\pi}{\mu},1\right). \tag{4.40}$$

We now derive the optimal strategy for trader i when trader j plays *Private*. If trader j plays *Private*, then trader i obtains a payoff of

$$(1-\alpha)\left(\frac{\eta\,\delta_i}{6\bar{\delta}}\right)(\delta_i^2 - 3\bar{\delta}^2) + \delta_i(\mu - \pi) \tag{4.41}$$

if she plays *Public*, and a payoff of

$$(1-\alpha)\left(\delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right) \tag{4.42}$$

if she plays *Private*. We then determine the values of α for which these two payoffs are equal. The result is summarised as follows.

Proposition 18 (Best response to competitor choosing Private). *Assume trader j submits to the private pool. If $\alpha < \frac{\pi}{\mu}$, it is optimal for trader i to submit her order to the private pool; otherwise, it is optimal for her to submit to the public pool.*

The previous two propositions allow us to characterise the optimal strategy for any trader as a function of her random liquidity need δ_i.

Corollary 1. *Assume trader i has a random liquidity need δ_i. Then the following hold:*

1. *If $0 < \alpha < \frac{\pi}{\mu}$, it is a dominant strategy to play Private.*
2. *If $\frac{\pi}{\mu} < \alpha < \alpha^*(\delta_i)$, playing Public is a best response when the other trader plays Private, and playing Private is a best response when the other trader plays Public.*
3. *If $\alpha^*(\delta_i) < \alpha < 1$, it is a dominant strategy to play Public.*

It is intuitive that playing *Private* is a dominant strategy for small values of α, since in that case there is a high probability that the randomly selected builder will be able to execute orders from the private memory pool, where traders avoid paying π. Similarly, for high values of α, there is a high probability that the builder will only execute orders submitted to the public memory pool, and thus playing *Public* becomes a dominant strategy.

It is also intuitive that $\alpha^*(\delta_i) > \frac{\pi}{\mu}$, since when the opponent plays *Private*, there is an incentive to deviate and play *Public*. The reason is that there exists a positive probability that the builder selected by the protocol will only execute orders from the public pool. In that case, the trader who sent her order to the public pool faces no competition in priority fees, as she knows she will win the auction with certainty.

To characterise the Nash equilibrium when both traders interact—so that we have two thresholds, $\alpha_1^*(\delta_i)$ and $\alpha_2^*(\delta_j)$, each corresponding to one trader and both belonging to the interval $\left(\frac{\pi}{\mu}, 1\right)$ by virtue of Proposition 17—the following result is relevant.

Proposition 19.

$$\partial_{\delta_i} \alpha^*(\delta_i) < 0. \tag{4.43}$$

We note that it was not necessary to obtain an explicit expression for $\alpha^*(\delta_i)$ to establish Propositions 17 and 19. However, we can also derive that

$$\alpha^*(\delta_i) = \frac{\frac{\eta\,\delta_i^3}{3\bar{\delta}} - \eta\,\delta_i\bar{\delta} + \mu\,\delta_i - \sqrt{\left(\frac{\eta\,\delta_i^3}{3\bar{\delta}} - \eta\,\delta_i\bar{\delta} + \mu\,\delta_i\right)^2 + 4\,\delta_i\pi\left(\frac{\eta\,\delta_i\bar{\delta}}{2} - \frac{\eta\,\delta_i^3}{6\bar{\delta}}\right)}}{\frac{\eta\,\delta_i^3}{3\bar{\delta}} - \eta\,\delta_i\bar{\delta}}. \tag{4.44}$$

Defining

$$\tilde{\delta}_i = -\frac{\eta\,\delta_i^3}{3\bar{\delta}} + \eta\,\delta_i\bar{\delta}, \tag{4.45}$$

we have $\tilde{\delta}_i > 0$ since $\bar{\delta} \geq \delta_i$. Under this notation, the expression above can be rewritten as

$$\alpha^*(\delta_i) = \frac{\tilde{\delta}_i - \mu\,\delta_i + \sqrt{(\mu\,\delta_i - \tilde{\delta}_i)^2 + 2\,\tilde{\delta}_i\,\delta_i\pi}}{\tilde{\delta}_i}. \tag{4.46}$$

Since $\alpha^*(\delta_i)$ is decreasing, and the minimum possible value of δ_i is 0, we can study the behaviour of $\alpha^*(\delta_i)$ as $\delta_i \to 0$. We have the following result.

Proposition 20. *When $\delta_i \to 0$,*

$$\alpha^*(\delta_i) \longrightarrow \frac{\eta\bar{\delta} - \mu + \sqrt{\mu^2 + \eta^2\bar{\delta}^2 - 2(\mu - \pi)\eta\bar{\delta}}}{\eta\bar{\delta}}. \tag{4.47}$$

We denote this limit by γ_1.

Moreover, if $\mu > \pi + \frac{\eta\bar{\delta}}{2}$, then the previous expression belongs to the interval $(0, 1)$.

We can also study the behaviour of $\alpha^*(\delta_i)$ when $\delta_i = \bar{\delta}$.

Proposition 21. *We have*

$$\alpha^*(\bar{\delta}) = \frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}}. \tag{4.48}$$

We denote this expression by γ_2.

If $\frac{\pi}{\mu} \in (0, 1)$ and $\mu > \pi + \frac{\eta\bar{\delta}}{2}$, then

$$1 > \frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} > \frac{\pi}{\mu}. \tag{4.49}$$

Combining Propositions 19, 20, and 21, we obtain the following result.

Corollary 2. *If* $\mu - \pi - \frac{\eta\bar{\delta}}{2} > 0$*, then*

$$\frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} \leq \alpha^*(\delta_i) \leq \frac{\eta\bar{\delta} - \mu + \sqrt{\mu^2 + \eta^2\bar{\delta}^2 - 2(\mu - \pi)\eta\bar{\delta}}}{\eta\bar{\delta}}. \tag{4.50}$$

Combining Corollary 1 and Proposition 19, we obtain a full characterisation of the Nash equilibria that arise when two traders decide whether to submit their orders to the public or the private memory pool.

Proposition 22. *Let* δ_i *and* δ_j *denote the random liquidity needs of traders* i *and* j*, and assume without loss of generality that* $\delta_i < \delta_j$*. Then:*

1. *If* $\alpha < \frac{\pi}{\mu}$*, there is a unique Nash equilibrium: (Private, Private).*

2. *If* $\frac{\pi}{\mu} < \alpha < \alpha^*(\delta_j)$*, there are two pure Nash equilibria, given by (Public, Private) and (Private, Public), that is, one trader plays Public and the other plays Private. In this region, there also exists a mixed Nash equilibrium.*

3. *If* $\alpha^*(\delta_j) < \alpha < \alpha^*(\delta_i)$*, there is a unique Nash equilibrium: (Public, Private), where trader* j *(the one with the higher liquidity need) plays Public and trader* i *plays Private.*

4. *If* $\alpha > \alpha^*(\delta_i)$*, there is a unique Nash equilibrium: (Public, Public).*

Here, both thresholds $\alpha^*(\delta_i)$ *and* $\alpha^*(\delta_j)$ *satisfy the inequalities in Corollary 2.*

We also note that, in particular, for

$$\alpha > \frac{\eta\bar{\delta} - \mu + \sqrt{\mu^2 + \eta^2\bar{\delta}^2 - 2(\mu - \pi)\eta\bar{\delta}}}{\eta\bar{\delta}}, \tag{4.51}$$

the unique Nash equilibrium is *(Public, Public).*

The intuition behind the four possible cases is as follows. For sufficiently small values of α (Case 1), it is optimal for both traders to submit their orders to the private pool, since with high probability the randomly chosen builder will be able to take orders from the private pool, where traders avoid paying MEV. For large values of α (Case 4), it is optimal for both traders to submit their orders to the public pool, as the probability that the selected builder can execute private orders becomes very low. In Case 3, the trader with the higher liquidity need faces greater pressure to execute her order, and it is therefore optimal for her to submit to the public pool, avoiding the risk of non-execution. Finally, in Case 2, there exist two asymmetric pure Nash equilibria and a mixed equilibrium: because both traders have similar (intermediate) liquidity needs, either trader may benefit from submitting her order to the public pool when the other submits to the private pool.

The mixed Nash equilibrium is characterised by probabilities $(p_i, 1-p_i)$ and $(p_j, 1-p_j)$, where p_i denotes the probability that trader i plays *Public*, and p_j denotes the probability that trader j plays *Public*. We obtain the following result.

Proposition 23. *For the mixed Nash equilibrium, which can only arise when* $\alpha \in \left(\frac{\pi}{\mu}, \alpha^*(\delta_j)\right)$*, both the probability* p_i *that trader* i *plays Public and the probability* p_j *that trader* j *plays Public are increasing in* α*. Their expressions are given by*

$$p_j(\alpha) = \frac{\pi - \alpha\mu}{\alpha(2-\alpha)\left(\frac{\eta\delta_i^2}{6\bar{\delta}} - \frac{\eta\bar{\delta}}{2}\right)}, \qquad p_i(\alpha) = \frac{\pi - \alpha\mu}{\alpha(2-\alpha)\left(\frac{\eta\delta_j^2}{6\bar{\delta}} - \frac{\eta\bar{\delta}}{2}\right)}. \tag{4.52}$$

Proposition 23 shows that for higher values of α, there is a lower probability that the randomly chosen builder can execute orders submitted to the private memory pool. Therefore, as α increases, both traders become more inclined to play *Public* rather than *Private* in the mixed Nash equilibrium. We also note that $p_i\left(\frac{\pi}{\mu}\right) = 0$ and $p_j\left(\frac{\pi}{\mu}\right) = 0$, since for $\alpha \leq \frac{\pi}{\mu}$ it is optimal for both traders to play the pure strategy *Private*, assigning zero probability to *Public*.

Moreover, we observe that $p_i(\alpha) > p_j(\alpha)$ whenever $\delta_i < \delta_j$: p_j is the probability that makes trader i indifferent between playing *Public* or *Private*, while p_i is the probability that makes trader j indifferent between the two. Because trader j has a higher liquidity need ($\delta_j > \delta_i$), she requires trader i to play *Public* more frequently than herself (that is, $p_i(\alpha) > p_j(\alpha)$ in the mixed Nash equilibrium). Otherwise, trader j would prefer to play the pure strategy *Public* rather than mixing, given her stronger urgency to execute.

4.4.2 Stage one: The builders' problem

We assume that builders decide whether to operate as *standard builders*—who only execute transactions from the public memory pool—or as *general builders*—who can execute transactions from both the public and the private pools. Each builder is assumed to have an equal probability of being selected by the blockchain protocol. By executing transactions from the public pool, a builder earns a fee π per executed unit. Choosing to also execute transactions from both pools (in particular, from the private pool) is a weakly dominant strategy for builders, since it expands the set of orders they can execute. This weak dominance gives rise to an equilibrium in which all traders submit their orders to the private pool, and all builders choose to execute from both pools.

We now study which values of α can arise as equilibrium outcomes, and which ones are optimal from the perspective of the builders—that is, those that maximise total rent extraction. We refer to the former as *equilibrium values* of α, and to the latter as *optimal values*, where optimality is defined in terms of builders' aggregate revenue.

We first consider the equilibrium values of α, distinguishing four possible cases:

1. $\alpha \in [0, \frac{\pi}{\mu})$
2. $\alpha \in [\frac{\pi}{\mu}, \gamma_2)$
3. $\alpha \in [\gamma_2, \gamma_1)$
4. $\alpha \in [\gamma_1, 1]$

We obtain the following results.

Proposition 24. *For $\alpha \in [0, \frac{\pi}{\mu})$, the only equilibrium value is $\alpha = 0$.*

Proposition 25. *There exists no equilibrium for any value of $\alpha \in (\frac{\pi}{\mu}, \gamma_2)$.*

Proposition 26. *There exists no equilibrium for any value of $\alpha \in (\gamma_2, \gamma_1)$.*

Proposition 27. *For any $\alpha \in (\gamma_1, 1]$, an equilibrium can exist.*

The candidate values of α that are optimal in terms of rent extraction are those that maximise the total expected revenue earned by the population of builders. It follows that the values of α maximising rent extraction are those in the interval $(\gamma_2, 1)$. For such values, both traders optimally submit their orders to the public pool, where builders extract MEV from each trader. Moreover, priority fees are highest when both traders play the same strategy—either both play *Public* or both play *Private*—but builders prefer the former case, since in that case they also collect the MEV fee, which is not available when traders play *Private*.

4.5 Conclusion

Private pools emerged as a mechanism to protect users from MEV-extractive strategies—such as frontrunning—that exploited the transparency of public memory pools. We studied how the presence of private memory pools influenced traders' decisions on where to submit their transactions in blockchains. In our framework, traders not only competed for queue priority but also decided whether to submit their orders to the public pool, where they faced an additional fixed cost from MEV extraction, or to a private pool, where execution was uncertain because the corresponding builder might not be selected to produce the next block.

Our results showed that liquidity conditions played a central role in determining the equilibrium fragmentation of order flow. When liquidity was deep, traders concentrated their orders in a single private pool because competition for block priority is less fierce, decreasing decentralisation. As liquidity becomes shallower, competition for priority fees intensifies and traders spread their orders more evenly across pools. For sufficiently low depth of liquidity, the unique Nash equilibrium featured a balanced allocation across all available pools.

In a second model, we introduced a setting with one private pool and a proportion of builders who had access to it. We showed how changes in the proportion of builder who run the private pool affects traders' optimal transaction submission strategies and the resulting mix between public and private execution. Together, the two models illustrated how blockchain design parameters shaped the distribution of order flow between public and private channels, and ultimately the degree of centralisation in decentralised markets.

4.6 Proofs

Proof of Proposition 5. The results follow from standard findings in asymmetric first-price auctions with independent private values. Regarding the stochastic dominance, since $V_1 \sim U[0, \overline{B}]$ and $V_2 \sim U[0, \overline{B} + \delta]$ with $\delta > 0$, V_2 first-order stochastically dominates V_1, and hence the equilibrium bid distribution of bidder 2 stochastically dominates that of bidder 1 (see Lemma 1 in [18]). Regarding the equilibrium bids, when

$$V_1 \sim U[0, \overline{B}], \qquad V_2 \sim U[0, \overline{B} + \delta],$$

the equilibrium bid functions are strictly increasing in the valuation, and the bids lie in the interval (see [61])

$$b \in \left[0, \tfrac{\overline{B}(\delta+\overline{B})}{\delta+2\overline{B}}\right].$$

Letting $b_1(v,\delta)$ and $b_2(v,\delta)$ be the equilibrium bid functions and $v_1(b,\delta)$ and $v_2(b,\delta)$ their inverses with respect to the valuation, we use the general expressions from [57] for supports $[\underline{v}_1, \bar{v}_1]$ and $[\underline{v}_2, \bar{v}_2]$. Setting $\underline{v}_1 = 0$ and $\underline{v}_2 = \epsilon$ and letting $\epsilon \to 0$, we obtain

$$\begin{cases} v_1(b,\delta) = \dfrac{2b\overline{B}^2(\delta+\overline{B})^2}{b^2\delta(\delta+2\overline{B})+\overline{B}^2(\delta+\overline{B})^2}, \\ v_2(b,\delta) = \dfrac{2b\overline{B}^2(\delta+\overline{B})^2}{b^2\delta(\delta+2\overline{B})-\overline{B}^2(\delta+\overline{B})^2}. \end{cases} \tag{4.53}$$

□

Proof of Proposition 6. Denote the CDFs of the equilibrium bids by

$$G_1(b,\delta) = \mathbb{P}[b_1(v,\delta) \le b] = \mathbb{P}[v \le v_1(b,\delta)] = \frac{v_1(b,\delta)}{\overline{B}},$$
$$G_2(b,\delta) = \frac{v_2(b,\delta)}{\overline{B}+\delta}.$$

By Lemma 1 in [18], bidder 2's bids stochastically dominate bidder 1's, which can be readily verified with the equations above in our special case.

Using the expressions for $v_1(b,\delta)$ and $v_2(b,\delta)$ from Proposition 5, the derivative of the pointwise difference $G_1(b,\delta) - G_2(b,\delta)$ with respect to δ (for fixed b) is

$$\partial_\delta(G_1(b,\delta) - G_2(b,\delta)) \tag{4.54}$$
$$= \frac{2\left(-b^7\delta^2\overline{B}^2(\delta+2\overline{B})^4 - b^5\delta^2\overline{B}^4(\delta+\overline{B})^2\left(\delta^2-4\overline{B}^2\right)\right)}{\left(\overline{B}^4(\delta+\overline{B})^4 - b^4\delta^2(\delta+2\overline{B})^2\right)^2} \tag{4.55}$$
$$+ \frac{2\left(b^3\overline{B}^6(\delta+\overline{B})^4\left(\delta^2-4\overline{B}^2\right) + b\overline{B}^8(\delta+\overline{B})^6\right)}{\left(\overline{B}^4(\delta+\overline{B})^4 - b^4\delta^2(\delta+2\overline{B})^2\right)^2} \ge 0 \tag{4.56}$$

where the inequality holds for $\overline{B} > 0$, $0 < \delta < \overline{B}/2$, and $b \in \left[0, \frac{\overline{B}(\delta+\overline{B})}{\delta+2\overline{B}}\right]$. Hence, the (pointwise) dominance of bidder 2's bid distribution strengthens as δ increases.

Finally, note that the upper bound of the bid support is increasing in δ:

$$\partial_\delta\left(\tfrac{\overline{B}(\delta+\overline{B})}{\delta+2\overline{B}}\right) = \frac{\overline{B}(\delta+2\overline{B}) - \delta\overline{B} - (\overline{B})^2}{(\delta+2\overline{B})^2} = \frac{\overline{B}^2}{(\delta+2\overline{B})^2} > 0,$$

which further amplifies the separation between G_1 and G_2 as δ grows.

□

Proof of Proposition 7. Since $\max\{\delta_1^2, \ldots, \delta_N^2\}$ is a positive random variable, we can write

$$\mathbb{E}\big[\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\}\big] = \int_0^{\bar{\delta}^2} \mathbb{P}\big(\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\} > x\big)\, dx. \tag{4.57}$$

This can be expressed as

$$\begin{aligned}\mathbb{E}\big[\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\}\big] &= \int_0^{\bar{\delta}^2} \big(1 - \mathbb{P}\big(\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\} < x\big)\big)\, dx \\ &= \int_0^{\bar{\delta}^2} \big(1 - \mathbb{P}(\delta_1^2 < x)^N\big)\, dx. \end{aligned} \tag{4.58}$$

Because δ_1 is uniformly distributed on $[0, \bar{\delta}]$, we have $\mathbb{P}(\delta_1 < \sqrt{x}) = \frac{\sqrt{x}}{\bar{\delta}}$. Substituting this into the integral gives

$$\mathbb{E}\big[\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\}\big] = \int_0^{\bar{\delta}^2} \left(1 - \left(\frac{\sqrt{x}}{\bar{\delta}}\right)^N\right) dx. \tag{4.59}$$

Evaluating the integral yields

$$\begin{aligned}\mathbb{E}\big[\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\}\big] &= \left[x - \frac{x^{\frac{N}{2}+1}}{\bar{\delta}^N \left(\frac{N}{2}+1\right)}\right]_0^{\bar{\delta}^2} \\ &= \bar{\delta}^2 - \bar{\delta}^2 \frac{2}{N+2} = \bar{\delta}^2 \frac{N}{N+2}. \end{aligned} \tag{4.60}$$

□

Proof of Proposition 8. The payoff for trader i can be written as

$$P(\text{win})\,(-\varphi_i + \eta(N-1)\delta_i^2) + \delta_i V - \eta(N-1)\delta_i^2. \tag{4.61}$$

In the event of winning, we have $\delta_i = \max\{\delta_1, \delta_2, \ldots, \delta_N\}$. Moreover, the optimal priority fee is

$$\varphi^* = \eta(N-1)\delta_i^2 \frac{N-1}{N} = \frac{\eta(N-1)^2}{N}\,\delta_i^2. \tag{4.62}$$

Substituting this into the payoff expression gives

$$\begin{aligned}-\varphi^* + \eta(N-1)\delta_i^2 &= -\frac{\eta(N-1)^2}{N}\delta_i^2 + \eta(N-1)\delta_i^2 \\ &= \eta(N-1)\delta_i^2\left(1 - \frac{N-1}{N}\right) = \frac{\eta(N-1)}{N}\,\delta_i^2. \end{aligned} \tag{4.63}$$

Since $P(\text{win}) = \frac{1}{N}$, the expected payoff can be expressed as

$$f(N) = \frac{\eta(N-1)}{N^2}\,\mathbb{E}\big[\max\{\delta_1^2, \delta_2^2, \ldots, \delta_N^2\}\big] + \mathbb{E}\big[\delta_i V - \eta(N-1)\delta_i^2\big]. \tag{4.64}$$

Using the result from Proposition 7, $\mathbb{E}[\max\{\delta_1^2, \ldots, \delta_N^2\}] = \bar{\delta}^2(1 - \frac{2}{N+2})$, and the fact that $\mathbb{E}[\delta_i] = \frac{\bar{\delta}}{2}$, we obtain

$$\begin{aligned} f(N) &= \frac{\eta(N-1)}{N^2}\bar{\delta}^2\left(1 - \frac{2}{N+2}\right) + \frac{\bar{\delta}}{2}\mu - \frac{\eta\bar{\delta}^2(N-1)}{3} \\ &= \frac{\eta(N-1)\bar{\delta}^2 N}{N^2(N+2)} + \frac{\bar{\delta}}{2}\mu - \frac{\eta\bar{\delta}^2(N-1)}{3} \\ &= \frac{\eta(N-1)\bar{\delta}^2}{N(N+2)} + \frac{\bar{\delta}}{2}\mu - \frac{\eta\bar{\delta}^2(N-1)}{3}. \end{aligned} \tag{4.65}$$

□

Proof of Proposition 12. If $\eta \to 0$, then for all N we have $f(N) \to \frac{\bar{\delta}}{2}\mu$, with $\frac{\bar{\delta}}{2}\mu > 0$. Consider the state $(M, 0, 0)$. Since $f > 0$ for all N, no trader has a profitable deviation by switching to the other private pool, as it does not have a higher probability of being selected. Similarly, no trader has an incentive to deviate to the public pool, because in state $(M-1, 0, 1)$ all traders have the same probability of being selected, while the trader in the public pool incurs the additional MEV cost π.

Uniqueness for $\eta \to 0$. No state $(N, M-N, 0)$ with $1 \leq N \leq M-1$ can be a Nash equilibrium, since any trader in the private pool with $\min\{N, M-N\}$ traders has a profitable deviation by switching to the other private pool, which has a higher probability of being selected. Moreover, no state $(N, 0, M-N)$ with $M-N \geq 1$ can be a Nash equilibrium, because any trader in the public pool would deviate to private pool 1 to avoid paying π. Finally, states $(N_1, N_2, M-N_1-N_2)$ with $N_1 \geq 1$, $N_2 \geq 1$, and $M-N_1-N_2 \geq 1$ cannot be Nash equilibria either, since any trader in the private pool with $\min\{N_1, N_2\}$ traders would profitably deviate by switching to the other private pool.

If $\eta \to \infty$, then $f \to -\infty$. In this case, in the state $(\frac{M}{2}, \frac{M}{2}, 0)$ (if M is even) or $(\frac{M+1}{2}, \frac{M-1}{2}, 0)$ (if M is odd), no trader has a profitable deviation to the other private pool (by the same reasoning as in the proof of Proposition 10). Likewise, no trader has a profitable deviation to the public pool as the probability of being selected increases (since traders in the public pool compete regardless of which private pool is selected), and moreover by deviating to the public pool there is the additional cost π.

Uniqueness for $\eta \to \infty$. No other state $(N_1, N_2, 0)$ can be a Nash equilibrium, because any trader in the private pool with $\max\{N_1, N_2\}$ traders would have a profitable deviation to the other private pool (again, as in Proposition 10). Furthermore, no state $(N_1, N_2, M-N_1-N_2)$ with $M-N_1-N_2 \geq 1$ can be a Nash equilibrium, since any trader in the public pool has a profitable deviation to either private pool—doing so decreases her probability of being selected and also allows her to avoid paying π. ☐

Proof of Proposition 13. The proof is analogous to that of Proposition 11. ☐

Proof of Proposition 14. Assume trader j follows a continuously increasing strategy $\varphi_j = \varphi(\delta_j)$. The expected wealth of trader i is then

$$\begin{aligned}\mathbb{E}[W_i] &= -\varphi_i + \delta_i(\mu - \pi) - \eta\,\delta_i \int_{\varphi^{-1}(\varphi_i)}^{\bar{\delta}} x\,f(x)\,dx \\ &= -\varphi_i + \delta_i(\mu - \pi) - \frac{\eta\,\delta_i}{2\bar{\delta}}\left(\bar{\delta}^2 - \varphi^{-1}(\varphi_i)^2\right). \end{aligned} \tag{4.66}$$

The first-order condition with respect to φ_i gives

$$0 = -1 + \frac{\eta\,\delta_i\,\varphi^{-1}(\varphi_i)}{\bar{\delta}\,\varphi'(\varphi^{-1}(\varphi_i))}. \tag{4.67}$$

In equilibrium, symmetry implies that trader i employs the same strategy as trader j, so $\varphi_i = \varphi(\delta_i)$. Thus, for all $\delta_i \in [0, \bar{\delta}]$, we have

$$0 = -1 + \frac{\eta\,\delta_i^2}{\bar{\delta}\,\varphi'(\delta_i)}. \tag{4.68}$$

Hence, the bidding function φ satisfies the differential equation

$$\varphi'(\delta_i) = \frac{\eta\,\delta_i^2}{\bar{\delta}}, \tag{4.69}$$

subject to the boundary condition $\varphi(0) = 0$. Solving this differential equation yields

$$\varphi_i = \frac{\eta\,\delta_i^3}{3\,\bar{\delta}}. \tag{4.70}$$

Since $\varphi^{-1}(\varphi_i(\delta_i))^2 = \delta_i^2$, the expected payoff of trader i becomes

$$\begin{aligned}\mathbb{E}[W_i] &= -\frac{\eta\,\delta_i^3}{3\bar{\delta}} + \delta_i(\mu - \pi) - \frac{\eta\,\delta_i}{2\bar{\delta}}\left(\bar{\delta}^2 - \delta_i^2\right) \\ &= -\frac{\eta\,\delta_i^3}{3\bar{\delta}} + \delta_i(\mu - \pi) - \frac{\eta\,\delta_i\bar{\delta}}{2} + \frac{\eta\,\delta_i^3}{2\bar{\delta}} \\ &= \frac{\eta\,\delta_i^3}{6\bar{\delta}} + \delta_i\left(\mu - \pi - \frac{\eta\,\bar{\delta}}{2}\right). \end{aligned} \tag{4.71}$$

☐

Proof of Proposition 15. Deriving the optimal priority fees and the traders' payoffs follows the same steps as in the proof of Proposition 14. ☐

Proof of Proposition 16. In this setting, trader i holds the belief $f_j \sim U(0, \bar{\delta})$ over trader j's trading volume δ_j. Similarly, trader j holds the belief $f_i \sim U(0, (1-\alpha)\bar{\delta})$ over trader i's adjusted trading volume $\tilde{\delta}_i$. Trader i plays

$$\varphi_i : \tilde{\delta}_i \mapsto b, \qquad [0, (1-\alpha)\bar{\delta}] \to [0, \bar{b}],$$

and trader j plays

$$\varphi_j : \delta_j \mapsto b, \qquad [0, \bar{\delta}] \to [0, \bar{b}].$$

The expected wealth of trader i is

$$\begin{aligned}
\mathbb{E}[W_i] &= \alpha\left(-\varphi_i + \delta_i(\mu - \pi)\right) + (1-\alpha)\left(-\varphi_i + \delta_i(\mu-\pi) - \frac{\eta\,\delta_i}{2\bar{\delta}}\left(\bar{\delta}^2 - \varphi_j^{-1}(\varphi_i)^2\right)\right) \\
&= -\varphi_i + \delta_i(\mu-\pi) - (1-\alpha)\frac{\eta\,\delta_i}{2\bar{\delta}}\left(\bar{\delta}^2 - \phi_j(\varphi_i)^2\right) \\
&= -\varphi_i + (1-\alpha)\frac{\eta\,\delta_i}{2\bar{\delta}}\phi_j(\varphi_i)^2 + \delta_i(\mu-\pi) - (1-\alpha)\frac{\eta\,\delta_i}{2\bar{\delta}}\bar{\delta}^2. \qquad (4.72)
\end{aligned}$$

Similarly, the expected wealth of trader j is

$$\begin{aligned}
\mathbb{E}[W_j] &= (1-\alpha)\left(-\varphi_j + \mu\delta_j - \frac{\eta\,\delta_j}{2(1-\alpha)\bar{\delta}}\left((1-\alpha)^2\bar{\delta}^2 - \phi_i(\varphi_j)^2\right)\right) \\
&= (1-\alpha)\left(-\varphi_j + \mu\delta_j - \frac{\eta\,\delta_j}{2}(1-\alpha)\bar{\delta} + \frac{\eta\,\delta_j}{2\bar{\delta}(1-\alpha)}\phi_i(\varphi_j)^2\right), \qquad (4.73)
\end{aligned}$$

where we denote $\phi_i(\varphi) = \varphi_i^{-1}(\varphi)$ and $\phi_j(\varphi) = \varphi_j^{-1}(\varphi)$.

The optimisation problem for trader i is therefore

$$\sup_{\varphi_i}\left\{-\varphi_i + (1-\alpha)\delta_i\frac{\eta}{2\bar{\delta}}\phi_j(\varphi_i)^2\right\}. \qquad (4.74)$$

Defining $\tilde{\delta}_i = (1-\alpha)\delta_i$, the problem becomes

$$\sup_{\varphi_i}\left\{-\varphi_i + \eta\tilde{\delta}_i\int_0^{\phi_j(\varphi_i)} x f_j(x)\,dx\right\}, \qquad f_j \sim U(0, \bar{\delta}). \qquad (4.75)$$

Similarly, trader j solves

$$\sup_{\varphi_j}\left\{-\varphi_j + \eta\delta_j\int_0^{\phi_i(\varphi_j)} x f_i(x)\,dx\right\}, \qquad f_i \sim U(0, (1-\alpha)\bar{\delta}). \qquad (4.76)$$

These problems simplify to

$$\begin{cases}
\sup_{\varphi_i}\left\{-\varphi_i + \dfrac{\eta\tilde{\delta}_i}{2\bar{\delta}}\phi_j(\varphi_i)^2\right\}, \\
\sup_{\varphi_j}\left\{-\varphi_j + \dfrac{\eta\delta_j}{2(1-\alpha)\bar{\delta}}\phi_i(\varphi_j)^2\right\}.
\end{cases}$$

The first-order conditions for both traders are

$$\begin{cases} \dfrac{\bar{\delta}}{\eta} = \tilde{\delta}_i \, \phi_j'(\varphi_i) \, \phi_j(\varphi_i), \\ \dfrac{\bar{\delta}(1-\alpha)}{\eta} = \delta_j \, \phi_i'(\varphi_j) \, \phi_i(\varphi_j). \end{cases}$$

Since

$$\begin{cases} \tilde{\delta}_i = \phi_i(\varphi_i), \\ \delta_j = \phi_j(\varphi_j), \end{cases}$$

the system becomes

$$\begin{cases} \dfrac{\bar{\delta}}{\eta} = \phi_i(b) \, \phi_j'(b) \, \phi_j(b), \\ \dfrac{\bar{\delta}(1-\alpha)}{\eta} = \phi_j(b) \, \phi_i'(b) \, \phi_i(b). \end{cases}$$

We propose the functional forms (ansatz)

$$\begin{cases} \phi_i(b) = l_i b^{1/3}, \\ \phi_j(b) = l_j b^{1/3}. \end{cases}$$

Substituting into the system gives

$$\begin{cases} \dfrac{\bar{\delta}}{\eta} = \dfrac{l_i l_j^2}{3}, \\ \dfrac{\bar{\delta}(1-\alpha)}{\eta} = \dfrac{l_j l_i^2}{3}, \end{cases} \qquad \Longrightarrow \qquad \begin{cases} \dfrac{l_i}{l_j} = 1-\alpha, \\ l_j^3(1-\alpha) = \frac{3\bar{\delta}}{\eta}. \end{cases}$$

Thus,

$$\begin{cases} l_i = \left(\dfrac{3\bar{\delta}(1-\alpha)^2}{\eta} \right)^{1/3}, \\ l_j = \left(\dfrac{3\bar{\delta}}{\eta(1-\alpha)} \right)^{1/3}. \end{cases}$$

Consequently,

$$\begin{cases} \phi_i(b) = \left(\dfrac{3\bar{\delta}(1-\alpha)^2 b}{\eta} \right)^{1/3}, \\ \phi_j(b) = \left(\dfrac{3\bar{\delta} b}{\eta(1-\alpha)} \right)^{1/3}. \end{cases}$$

Inverting yields the equilibrium bidding strategies:

$$\begin{cases} \varphi_i(\tilde{\delta}_i) = \dfrac{\eta \tilde{\delta}_i^{\,3}}{3\bar{\delta}(1-\alpha)^2}, \\ \varphi_j(\delta_j) = \dfrac{\eta(1-\alpha)\delta_j^3}{3\bar{\delta}}. \end{cases}$$

Since $\tilde{\delta}_i = (1-\alpha)\delta_i$, we obtain

$$\begin{cases} \varphi_i(\tilde{\delta}_i) = \dfrac{\eta(1-\alpha)\delta_i^3}{3\bar{\delta}}, \\ \varphi_j(\delta_j) = \dfrac{\eta(1-\alpha)\delta_j^3}{3\bar{\delta}}. \end{cases}$$

Using these bidding strategies, we compute the expected payoffs. Noting that

$$\phi_i(\varphi_j(\delta_j)) = \phi_i\left(\frac{\eta(1-\alpha)\delta_j^3}{3\bar{\delta}}\right) = (1-\alpha)\delta_j, \qquad \phi_j(\varphi_i(\tilde{\delta}_i)) = \phi_j\left(\frac{\eta(1-\alpha)\delta_i^3}{3\bar{\delta}}\right) = \delta_i,$$

we find that the expected payoffs are

$$\begin{aligned} \mathbb{E}[W_i] &= \frac{\eta(1-\alpha)\delta_i^3}{6\bar{\delta}} + \delta_i\left(\mu - \pi - \frac{(1-\alpha)\eta\bar{\delta}}{2}\right) \\ &= (1-\alpha)\frac{\eta\delta_i}{6\bar{\delta}}\left(\delta_i^2 - 3\bar{\delta}^2\right) + \delta_i(\mu-\pi), \end{aligned} \tag{4.77}$$

and

$$\mathbb{E}[W_j] = (1-\alpha)\left[(1-\alpha)\eta\delta_j\frac{\delta_j^2 - 3\bar{\delta}^2}{6\bar{\delta}} + \mu\delta_j\right]. \tag{4.78}$$

□

Proof of Proposition 17. Assume that trader j plays *Public*. If trader i also plays *Public*, her expected payoff is

$$\delta_i\left(\mu - \pi - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}. \tag{4.79}$$

If instead trader i plays *Private*, her expected payoff is

$$(1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2 - 3\bar{\delta}^2}{6\bar{\delta}} + \mu\,\delta_i\right]. \tag{4.80}$$

Trader i is indifferent between playing *Public* or *Private* when α satisfies

$$\delta_i\left(\mu - \pi - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}} = (1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2 - 3\bar{\delta}^2}{6\bar{\delta}} + \mu\,\delta_i\right]. \tag{4.81}$$

Simplifying, we obtain

$$\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2} - \delta_i\pi = (1-\alpha)^2\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) - \alpha\mu\delta_i. \tag{4.82}$$

The equation in (4.82) has no solution in $\alpha \in [0,1]$ if $\pi > \alpha\mu$, since in this case the left-hand side is more negative than the right-hand side, using that $\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2} < 0$. In

this situation, *Private* is a dominant strategy, which is intuitive for small α. Hence, any possible solution $\alpha \in [0, 1]$ must satisfy $\alpha > \frac{\pi}{\mu}$.
Furthermore, the right-hand side of the original equality is zero when $\alpha = 1$, and thus smaller than the left-hand side, which is assumed positive. When $\alpha = 0$, the right-hand side is larger than the left-hand side, since $\pi > \alpha\mu = 0$. Therefore, there exists a value $\alpha^* < 1$ that satisfies the equality.
Additionally, the right-hand side of (4.82) becomes more negative as α increases beyond 1, implying that no root can exist for $\alpha \geq 1$. Using again that $\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2} < 0$, we find

$$\lim_{\alpha \to -\infty} \left[(1-\alpha)^2 \left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2} \right) - \alpha\mu\delta_i \right] = -\infty. \tag{4.83}$$

Combining this with the fact that the right-hand side of (4.82) is greater than the left-hand side when $\alpha = 0$ (since $-\delta_i\pi < 0$) ensures the existence of a negative root. Consequently, there exists a unique positive root of the equation, and it lies in the interval $(\frac{\pi}{\mu}, 1)$. We denote this value by α^*. □

Proof of Proposition 18. We look for values of α such that trader i is indifferent between submitting her order to the public memory pool and to the private pool—that is, when the payoffs in (4.41) and (4.42) are equal:

$$(1-\alpha)\left(\frac{\eta\,\delta_i}{6\bar{\delta}}\right)(\delta_i^2 - 3\bar{\delta}^2) + \delta_i(\mu - \pi) = (1-\alpha)\left(\delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right). \tag{4.84}$$

We show that there exists exactly one value of $\alpha \in [0, 1)$ for which this equality holds.
Define the functions

$$g(\alpha) = (1-\alpha)\left(\frac{\eta\,\delta_i}{6\bar{\delta}}\right)(\delta_i^2 - 3\bar{\delta}^2) + \delta_i(\mu - \pi), \tag{4.85}$$

$$h(\alpha) = (1-\alpha)\left(\delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right). \tag{4.86}$$

It is clear that $h(\alpha)$ is strictly decreasing in α, while $g(\alpha)$ is strictly increasing in α. Hence, there can be at most one value of $\alpha \in [0, 1)$ such that playing *Private* or *Public* yields the same payoff.
To show that such a value exists, it suffices to verify that $g(0) < h(0)$ and $g(1) > h(1)$. Indeed, we have

$$h(1) = 0 < \delta_i(\mu - \pi) = g(1),$$

and

$$h(0) = \delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}} > \left(\frac{\eta\,\delta_i}{6\bar{\delta}}\right)(\delta_i^2 - 3\bar{\delta}^2) + \delta_i(\mu - \pi) = g(0), \tag{4.87}$$

where the inequality holds because

$$\delta_i\mu - \frac{\delta_i\eta\bar{\delta}}{2} + \frac{\eta\,\delta_i^3}{6\bar{\delta}} > \frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\delta_i\eta\bar{\delta}}{2} + \delta_i(\mu-\pi) \iff 0 > -\delta_i\pi, \tag{4.88}$$

which is clearly true.
Thus, using the intermediate value theorem, there exists a unique $\alpha \in [0,1)$ satisfying $g(\alpha) = h(\alpha)$. Solving the equation yields the value

$$\alpha = \frac{\pi}{\mu}. \tag{4.89}$$

□

Proof of Corollary 1. The result follows directly by combining Propositions 17 and 18, and by noting that $\alpha^* > \frac{\pi}{\mu}$. □

Proof of Proposition 19. From the proof of Proposition 17, $\alpha^*(\delta_i)$ satisfies

$$\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2} - \delta_i\pi = (1-\alpha)^2\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) - \alpha\mu\delta_i. \tag{4.90}$$

Rearranging terms yields

$$\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right)\left[1-(1-\alpha)^2\right] = \delta_i(\pi - \alpha\mu). \tag{4.91}$$

We restrict attention to values $\alpha \in (\frac{\pi}{\mu}, 1)$, since we know that $\alpha^*(\delta_i) \in (\frac{\pi}{\mu}, 1)$. We see that $\alpha^*(\delta_i)$ is an injective function. Suppose, to the contrary, that there exists a common α^* satisfying the equation for more than one $\delta_i \in (0, \bar{\delta})$. Then, for such values of δ_i, we would have

$$\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right)k = \delta_i b, \tag{4.92}$$

where $k = 2\alpha^* - (\alpha^*)^2 > 0$ and $b = \pi - \alpha^*\mu < 0$.
This simplifies to

$$\left(\frac{\eta\,\delta_i^2}{6\bar{\delta}} - \frac{\eta\,\bar{\delta}}{2}\right)k = b, \tag{4.93}$$

so that

$$\delta_i = \sqrt{(\frac{b}{k} + \frac{\eta\bar{\delta}}{2})(\frac{6\bar{\delta}}{\eta})}$$

which cannot have two positive solutions for δ_i. Hence $\alpha^*(\delta_i)$ in injective, and due to continuity, it is monotone. To see that $\alpha^*(\delta_i)$ is decreasing, it is enough to see it is decreasing for a given fixed set of parameters η, $\bar{\delta}$, μ and π, due to continuity (otherwise, if for a set of parameters it is decreasing and for another set it is increasing, we would get a contradiction using Bolzano's Theorem). One can check, taking a given set of parameters that it is decreasing (checking for example that $\alpha(\bar{\delta}) < \alpha(\bar{\delta}/2)$).

□

Proof of Proposition 20. We know that $\alpha^*(\delta_i)$ satisfies

$$\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right)\left[1-(1-\alpha)^2\right] = \delta_i(\pi - \alpha\mu). \tag{4.94}$$

When $\delta_i \to 0$ (and $\delta_i \neq 0$), this expression simplifies to

$$\left(\frac{\eta\,\delta_i^2}{6\bar{\delta}} - \frac{\eta\,\bar{\delta}}{2}\right)\left[1-(1-\alpha)^2\right] = \pi - \alpha\mu, \tag{4.95}$$

which can be rewritten as

$$\alpha^2\left(\frac{\eta\bar{\delta}}{2}\right) + \alpha(\mu - \eta\bar{\delta}) - \pi = 0. \tag{4.96}$$

This quadratic equation in α has solutions

$$\alpha = \frac{\eta\bar{\delta} - \mu \pm \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}}}{\eta\bar{\delta}}. \tag{4.97}$$

We now examine the two possible roots. The negative root is

$$\begin{aligned}\alpha &= \frac{\eta\bar{\delta} - \mu - \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}}}{\eta\bar{\delta}} \leq \frac{\eta\bar{\delta} - \mu - \sqrt{(\mu-\eta\bar{\delta})^2}}{\eta\bar{\delta}} \\ &= \frac{\eta\bar{\delta} - \mu - |\eta\bar{\delta} - \mu|}{\eta\bar{\delta}} \leq 0,\end{aligned} \tag{4.98}$$

and thus is not relevant, since it is negative. The positive root is

$$\alpha = \frac{\eta\bar{\delta} - \mu + \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}}}{\eta\bar{\delta}} = \frac{\eta\bar{\delta} - \mu + \sqrt{\mu^2 + \eta^2\bar{\delta}^2 - 2(\mu-\pi)\eta\bar{\delta}}}{\eta\bar{\delta}}. \tag{4.99}$$

We can show that this root satisfies $0 \leq \alpha \leq 1$. Indeed,

$$\alpha = \frac{\eta\bar{\delta} - \mu + \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}}}{\eta\bar{\delta}} \geq \frac{\eta\bar{\delta} - \mu + \sqrt{(\mu-\eta\bar{\delta})^2}}{\eta\bar{\delta}} = \frac{\eta\bar{\delta} - \mu + |\mu - \eta\bar{\delta}|}{\eta\bar{\delta}} \geq 0. \tag{4.100}$$

Moreover, to see $\alpha \leq 1$, note that

$$\begin{aligned}\frac{\eta\bar{\delta} - \mu + \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}}}{\eta\bar{\delta}} \leq 1 &\iff \eta\bar{\delta} - \mu + \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}} \leq \eta\bar{\delta} \\ &\iff \sqrt{(\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta}} \leq \mu \\ &\iff (\mu-\eta\bar{\delta})^2 + 2\pi\eta\bar{\delta} \leq \mu^2 \\ &\iff \eta\bar{\delta} \leq 2(\mu - \pi),\end{aligned} \tag{4.101}$$

which holds by the assumption $\mu - \pi - \frac{\eta\bar{\delta}}{2} > 0$. □

Proof of Proposition 21. We first show that

$$\frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} > \frac{\pi}{\mu} \tag{4.102}$$

holds if and only if

$$2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}} > \frac{2\eta\bar{\delta}\pi}{\mu}. \tag{4.103}$$

This is equivalent to

$$\sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}} > \frac{2\eta\bar{\delta}\pi}{\mu} + 3\mu - 2\eta\bar{\delta}, \tag{4.104}$$

$$\iff (3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta} > \left(\frac{2\eta\bar{\delta}\pi}{\mu} + 3\mu - 2\eta\bar{\delta}\right)^2. \tag{4.105}$$

Expanding the right-hand side, we obtain

$$(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta} > (3\mu - 2\eta\bar{\delta})^2 + \left(\frac{2\eta\bar{\delta}\pi}{\mu}\right)^2 + 4\frac{\eta\bar{\delta}\pi(3\mu - 2\eta\bar{\delta})}{\mu}. \tag{4.106}$$

Simplifying terms gives

$$0 > \left(\frac{2\eta\bar{\delta}\pi}{\mu}\right)^2 - \frac{8\eta^2\bar{\delta}^2\pi}{\mu} \iff \frac{2\pi}{\mu} > \frac{\pi}{\mu}\cdot\frac{\pi}{\mu} \iff 2 > \frac{\pi}{\mu}, \tag{4.107}$$

which holds since $\frac{\pi}{\mu} \in (0,1)$.

Next, we verify that

$$\frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} < 1. \tag{4.108}$$

Indeed,

$$\begin{aligned}\frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} &= \frac{2\eta\bar{\delta} - 3\mu + \sqrt{9\mu^2 + 4\eta^2\bar{\delta}^2 - 12\mu\eta\bar{\delta} + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} \\ &< \frac{2\eta\bar{\delta} - 3\mu + \sqrt{9\mu^2 + 4\eta^2\bar{\delta}^2 - 6\eta^2\bar{\delta}^2}}{2\eta\bar{\delta}} \\ &= \frac{2\eta\bar{\delta} - 3\mu + \sqrt{9\mu^2 - 2\eta^2\bar{\delta}^2}}{2\eta\bar{\delta}}.\end{aligned} \tag{4.109}$$

Since

$$\sqrt{9\mu^2 - 2\eta^2\bar{\delta}^2} \leq 3\mu,$$

we conclude that

$$\frac{2\eta\bar{\delta} - 3\mu + \sqrt{(3\mu - 2\eta\bar{\delta})^2 + 12\pi\eta\bar{\delta}}}{2\eta\bar{\delta}} \leq \frac{2\eta\bar{\delta} - 3\mu + 3\mu}{2\eta\bar{\delta}} = 1. \tag{4.110}$$

The first inequality above follows from the assumption $\mu > \pi + \frac{\eta\bar{\delta}}{2}$ which implies $\pi - \mu < -\frac{\eta\bar{\delta}}{2}$. □

Proof of Proposition 22. The result follows directly from combining Corollary 1 and Proposition 19. □

Proof of Proposition 23. Let p_i denote the probability that trader i plays *Public*, and p_j the probability that trader j plays *Public*. For a mixed Nash equilibrium to exist, the expected payoff of trader i from playing *Public* or *Private* must be equal, given that trader j plays the mixed strategy $(p_j, 1 - p_j)$. By symmetry, an analogous condition holds for trader j. We focus on the first condition, as the second follows identically.

Trader i is indifferent when

$$\begin{aligned} &p_j \left[\delta_i\left(\mu - \pi - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right] + (1 - p_j)\left[(1-\alpha)\left(\frac{\eta\,\delta_i}{6\bar{\delta}}(\delta_i^2 - 3\bar{\delta}^2)\right) + \delta_i(\mu - \pi)\right] \\ &\quad = p_j(1-\alpha)\left[(1-\alpha)\eta\,\delta_i\frac{\delta_i^2 - 3\bar{\delta}^2}{6\bar{\delta}} + \mu\,\delta_i\right] + (1-p_j)(1-\alpha)\left[\delta_i\left(\mu - \frac{\eta\,\bar{\delta}}{2}\right) + \frac{\eta\,\delta_i^3}{6\,\bar{\delta}}\right]. \end{aligned} \tag{4.111}$$

Simplifying this equality, we obtain

$$\begin{aligned} &p_j\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + p_j\delta_i(\mu - \pi) + (1-\alpha)\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + \delta_i(\mu - \pi) \\ &-p_j(1-\alpha)\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + p_j\delta_i(\pi - \mu) = p_j(1-\alpha)^2\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + p_j(1-\alpha)\delta_i\mu \\ &\qquad + (1-\alpha)\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + (1-\alpha)\delta_i\mu \\ &\qquad - p_j(1-\alpha)\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) - p_j(1-\alpha)\delta_i\mu. \end{aligned} \tag{4.112}$$

After cancelling terms, we are left with

$$p_j\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + \delta_i(\mu - \pi) = p_j(1-\alpha)^2\left(\frac{\eta\,\delta_i^3}{6\bar{\delta}} - \frac{\eta\,\delta_i\bar{\delta}}{2}\right) + (1-\alpha)\delta_i\mu. \tag{4.113}$$

Solving for p_j, we find

$$p_j = \frac{\pi - \alpha\mu}{\alpha(2-\alpha)\left(\frac{\eta\,\delta_i^2}{6\bar{\delta}} - \frac{\eta\,\bar{\delta}}{2}\right)}. \tag{4.114}$$

By symmetry, we also have

$$p_i = \frac{\pi - \alpha\mu}{\alpha(2-\alpha)\left(\frac{\eta\,\delta_j^2}{6\bar{\delta}} - \frac{\eta\,\bar{\delta}}{2}\right)}. \tag{4.115}$$

We note that

$$\frac{1}{\frac{\eta\,\delta_i^2}{6\bar{\delta}} - \frac{\eta\,\bar{\delta}}{2}} < 0, \tag{4.116}$$

since $\delta_i \leq \bar{\delta}$.
Next, consider the function

$$\frac{\pi - \alpha\mu}{\alpha(2-\alpha)}. \tag{4.117}$$

Its derivative with respect to α is

$$\frac{d}{d\alpha}\left(\frac{\pi - \alpha\mu}{\alpha(2-\alpha)}\right) = \frac{-\alpha^2\mu - 2\pi(1-\alpha)}{(2\alpha - \alpha^2)^2}, \tag{4.118}$$

which is negative for all $\alpha \in (0,1)$, the range of interest. Since the denominator of p_j is negative and the numerator is decreasing in α, it follows that

$$\frac{\partial p_j(\alpha)}{\partial \alpha} > 0, \qquad \frac{\partial p_i(\alpha)}{\partial \alpha} > 0.$$

Hence, both p_i and p_j are increasing in α. □

Proof of Proposition 24. We cannot have an equilibrium with $\alpha \in (0, \frac{\pi}{\mu})$, since by Proposition 22 both traders play *Private* in this setting. Therefore, it is optimal for builders to be able to take orders from the private memory pool, and thus builders find it optimal to set $\alpha = 0$. □

Proof of Proposition 25. For $\alpha \in (\frac{\pi}{\mu}, \gamma_2)$, consider two types of builders: (i) those who only execute orders from the public memory pool, and (ii) those who can execute orders from both the public and private pools.

If a fraction α of builders execute only public orders, their expected revenue, conditional on being chosen, is

$$\begin{aligned} R_1(\alpha) &= MEV + \int_0^{\bar{\delta}} \frac{\eta\,\delta_j^3(1-\alpha)}{3\bar{\delta}}\frac{1}{\bar{\delta}}\,d\delta_j \\ &= MEV + \frac{\eta(1-\alpha)\bar{\delta}^4}{12\bar{\delta}^2}. \end{aligned} \tag{4.119}$$

For builders that can execute orders from both pools, the expected revenue is

$$
\begin{aligned}
R_2(\alpha) &= MEV + \int_0^{\bar{\delta}} \frac{\eta\, \delta_j^3 (1-\alpha)}{3\bar{\delta}} \frac{1}{\bar{\bar{\delta}}}\, d\delta_j + \int_0^{\bar{\delta}} \frac{\eta\, \delta_i^3 (1-\alpha)}{3\bar{\delta}} \frac{1}{\bar{\bar{\delta}}}\, d\delta_i \\
&= MEV + \frac{\eta (1-\alpha) \bar{\delta}^4}{6 \bar{\delta}^2}.
\end{aligned} \tag{4.120}
$$

We seek values of $\alpha \in (\frac{\pi}{\mu}, \gamma_2)$ for which $R_1(\alpha) = R_2(\alpha)$. However, it is immediate that

$$
R_1(\alpha) < R_2(\alpha), \qquad \forall\, \alpha \in \left(\tfrac{\pi}{\mu}, \gamma_2\right), \tag{4.121}
$$

so there are no possible equilibrium values of α in this interval. □

Proof of Proposition 26. For $\alpha \in [\gamma_2, \gamma_1]$, the procedure to obtain candidate equilibrium values of α proceeds as follows:

1. Find the value δ^* such that $\alpha = \alpha^*(\delta^*)$.

2. Since α^* is decreasing in δ, we then have:

$$
\begin{cases}
\text{With probability } \left(\frac{\delta^*}{\bar{\delta}}\right)^2, & \text{we are in Case 2 of Proposition 22;} \\
\text{With probability } \frac{2\delta^*(\bar{\delta}-\delta^*)}{\bar{\delta}^2}, & \text{we are in Case 3 of Proposition 22;} \\
\text{With probability } \left(\frac{\bar{\delta}-\delta^*}{\bar{\delta}}\right)^2, & \text{we are in Case 4 of Proposition 22.}
\end{cases}
$$

Candidate equilibrium values of α are those for which the expected revenue of builders who only consider orders from the public pool equals the expected revenue of builders who consider orders submitted to both pools. These expected revenues are derived from the cases described in Proposition 22, weighted by the corresponding probabilities above.

Cases 3 and 4 each have a single Nash equilibrium, while Case 2 has two pure Nash equilibria and one mixed one. For simplicity, we focus on the pure equilibria. Since both pure equilibria in Case 2 yield identical expected revenues for builders (as both traders have liquidity needs below δ^* and thus integration limits are 0 and δ^*), the choice of which equilibrium to use does not matter.

Given $\alpha \in (\gamma_2, \gamma_1)$, the revenue for a builder that can execute orders only from the public pool is

$$
\begin{aligned}
R_1(\alpha) = \Bigg[&\left(\frac{\delta^*}{\bar{\delta}}\right)^2 \left(\pi + \int_0^{\delta^*} \frac{\eta\, \delta_j^3 (1-\alpha)}{3\bar{\delta}} \frac{1}{\bar{\bar{\delta}}}\, d\delta_j\right) + \frac{2\delta^*(\bar{\delta}-\delta^*)}{\bar{\delta}^2} \left(\pi + \int_{\delta^*}^{\bar{\delta}} \frac{\eta\, \delta_j^3 (1-\alpha)}{3\bar{\delta}} \frac{1}{\bar{\bar{\delta}}}\, d\delta_j\right) \\
&+ \left(\frac{\bar{\delta}-\delta^*}{\bar{\delta}}\right)^2 \left(2\pi + 2\int_{\delta^*}^{\bar{\delta}} \frac{\eta\, \delta_i^3}{3\bar{\delta}} \frac{1}{\bar{\bar{\delta}}}\, d\delta_i\right) \Bigg].
\end{aligned} \tag{4.122}
$$

Evaluating the integrals gives

$$R_1(\alpha) = \Bigg[\left(\frac{\delta^*}{\bar{\delta}}\right)^2 \left(\pi + \frac{\eta(1-\alpha)(\delta^*)^4}{12\bar{\delta}^2}\right) + \frac{2\delta^*(\bar{\delta}-\delta^*)}{\bar{\delta}^2} \left(\pi + \frac{\eta(1-\alpha)}{12\bar{\delta}^2}\left(\bar{\delta}^4 - (\delta^*)^4\right)\right)$$

$$+ 2\left(\frac{\bar{\delta}-\delta^*}{\bar{\delta}}\right)^2 \left(\pi + \frac{\eta}{12\bar{\delta}^2}\left(\bar{\delta}^4 - (\delta^*)^4\right)\right) \Bigg]. \tag{4.123}$$

The integrals above represent averages over traders' random liquidity needs. The term π appears when only one trader submits to the public pool, while 2π arises when both do. The limits of integration depend on the case of Proposition 22, as described earlier.

Similarly, the revenue for a builder who can execute orders from both pools is

$$R_2(\alpha) = \Bigg[\left(\frac{\delta^*}{\bar{\delta}}\right)^2 \left(\pi + 2\int_0^{\delta^*} \frac{\eta\, \delta_i^3(1-\alpha)}{3\bar{\delta}}\frac{1}{\bar{\delta}}\, d\delta_i\right) + \frac{2\delta^*(\bar{\delta}-\delta^*)}{\bar{\delta}^2} \left(\pi + \int_0^{\bar{\delta}} \frac{\eta\, \delta_i^3(1-\alpha)}{3\bar{\delta}}\frac{1}{\bar{\delta}}\, d\delta_i\right)$$

$$+ \left(\frac{\bar{\delta}-\delta^*}{\bar{\delta}}\right)^2 \left(2\pi + 2\int_{\delta^*}^{\bar{\delta}} \frac{\eta\, \delta_i^3}{3\bar{\delta}}\frac{1}{\bar{\delta}}\, d\delta_i\right)\Bigg]$$

$$= \Bigg[\left(\frac{\delta^*}{\bar{\delta}}\right)^2 \left(\pi + \frac{\eta(1-\alpha)(\delta^*)^4}{6\bar{\delta}^2}\right)$$

$$+ \frac{2\delta^*(\bar{\delta}-\delta^*)}{\bar{\delta}^2} \left(\pi + \frac{\eta(1-\alpha)\bar{\delta}^2}{12}\right) + 2\left(\frac{\bar{\delta}-\delta^*}{\bar{\delta}}\right)^2 \left(\pi + \frac{\eta}{12\bar{\delta}^2}\left(\bar{\delta}^4 - (\delta^*)^4\right)\right)\Bigg]. \tag{4.124}$$

In this scenario, both traders' orders are executed by the same builder, so each term includes two integrals. Equilibrium candidate values $\alpha \in (\gamma_2, \gamma_1)$ must satisfy

$$R_1(\alpha) = R_2(\alpha). \tag{4.125}$$

However, it is immediate that

$$R_1(\alpha) < R_2(\alpha), \qquad \forall\, \alpha \in (\gamma_2, \gamma_1), \tag{4.126}$$

so no equilibrium values of α exist in this range. □

Proof of Proposition 27. For any $\alpha \in (\gamma_1, 1]$, it is optimal for both traders to submit their orders to the public pool. In this case, builders are indifferent between executing only public orders or orders from both pools, since the outcomes are identical. Hence, any $\alpha \in (\gamma_1, 1]$ can serve as an equilibrium. For $\alpha < \gamma_1$, however, it is no longer optimal for both traders to submit to the public pool. □

Bibliography

[1] Ibrahim Abada and Xavier Lambin. Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided? *Management Science*, 69(9):5042–5065, 2023.

[2] AFM. Machine Learning in Trading Algorithms: Application by Dutch Proprietary Trading Firms and Possible Risks, 2023. Avaialble at: `www.afm.nl/en/sector/actueel/2023/maart/her-machine-learning`.

[3] Franklin Allen and Douglas Gale. Stock-Price Manipulation. *The Review of Financial Studies*, 5(3):503–529, 1992.

[4] Franklin Allen and Gary Gorton. Stock Price Manipulation, Market Microstructure and Asymmetric Information. *European Economic Review*, 36(2):624–630, 1992.

[5] Yakov Amihud and Haim Mendelson. Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2):223–249, 1986.

[6] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of uniswap markets. *Available at arXiv 1911.03380v7*, 2021.

[7] Alif Aqsha, Philippe Bergault, and Leandro Sánchez-Betancourt. Equilibrium reward for liquidity providers in automated market makers. *arXiv preprint arXiv:2503.22502*, 2025.

[8] Alif Aqsha, Fayçal Drissi, and Leandro Sánchez-Betancourt. Strategic learning and trading in broker-mediated markets. *arXiv preprint arXiv:2412.20847*, 2024.

[9] Matteo Aquilina, Eric Budish, and Peter O'Neill. Quantifying the High-Frequency Trading "Arms Race". *The Quarterly Journal of Economics*, 137(1):493–564, 09 2021.

[10] Álvaro Arroyo, Álvaro Cartea, Fernando Moreno-Pino, and Stefan Zohren. Deep Attentive Survival Analysis in Limit Order Books: Estimating Fill Probabilities with Convolutional-Transformers. *Available at SSRN 4432087*, 2023.

[11] Brad Bachu, Joel Hasbrouck, Fahad Saleh, and Xin Wan. An overview of uniswap v4 for researchers. *Available at SSRN*, 2025.

[12] Leonardo Baggiani, Martin Herdegen, and Leandro Sánchez-Betancourt. Optimal dynamic fees in automated market makers. *arXiv preprint arXiv:2506.02869*, 2025.

[13] Mark Bagnoli and Barton L. Lipman. Stock Price Manipulation Through Takeover Bids. *The RAND Journal of Economics*, 27(1):124–147, 1996.

[14] Philippe Bergault, Sébastien Bieber, and Leandro Sánchez-Betancourt. Optimal exit time for liquidity providers in automated market makers. *arXiv preprint arXiv:2509.06510*, 2025.

[15] Jonathan Brogaard, Dan Li, and Jeffrey Yang. Does High Frequency Market Manipulation Harm Market Quality? *Available at SSRN 4280120*, 2022.

[16] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicoló, and Sergio Pastorello. Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10):3267–97, 2020.

[17] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicoló, and Sergio Pastorello. Algorithmic Collusion with Imperfect Monitoring. *International Journal of Industrial Organization*, 79:102712, 2021.

[18] Estelle Cantillon. The effect of bidders' asymmetries on expected revenue in auctions. *Games and Economic Behavior*, 62:1–25, 2008.

[19] Charles Cao, Oliver Hansch, and Xiaoxin Wang. The Information Content of an Open Limit-Order Book. *Journal of Futures Markets*, 29(1):16–41, 2009.

[20] Agostino Capponi. Maximal extractable value and allocative inefficiencies in public blockchains. *Available at SSRN 4931619*, 2024.

[21] Agostino Capponi, Álvaro Cartea, and Fayçal Drissi. Do longer block times impair market efficiency in decentralized markets? *Available at SSRN 5290232*, 2025.

[22] Agostino Capponi, Garud Iyengar, Jay Sethuraman, et al. Decentralized finance: Protocols, risks, and governance. *Foundations and Trends® in Privacy and Security*, 5(3):144–188, 2023.

[23] Agostino Capponi, Ruizhe Jia, and Shihao Yu. Price discovery on decentralized exchanges. *Available at SSRN 4236993*, 2024.

[24] Álvaro Cartea, Patrick Chang, Mateusz Mroczka, and Roel Oomen. AI-Driven Liquidity Provision in OTC Financial Markets. *Quantitative Finance*, 22(12):2171–2204, 2022.

[25] Álvaro Cartea, Patrick Chang, and José Penalva. Algorithmic Collusion in Electronic Markets: The Impact of Tick Size. *Available at SSRN 4105954*, 2022.

[26] Álvaro Cartea, Patrick Chang, José Penalva, and Harrison Waldon. Algorithms can Learn to Collude: A Folk Theorem from Learning with Bounded Rationality. *Available at SSRN 4293831*, 2023.

[27] Álvaro Cartea, Ryan Donnelly, and Sebastian Jaimungal. Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics*, 8(1):635–671, 2017.

[28] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Execution and statistical arbitrage with signals in multiple automated market makers. In *2023 IEEE 43rd International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pages 37–42. IEEE, 2023.

[29] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Predictable losses of liquidity provision in constant function markets and concentrated liquidity markets. *Applied Mathematical Finance*, 30(2):69–93, 2023.

[30] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralized finance and automated market making: Predictable loss and optimal liquidity provision. *SIAM Journal on Financial Mathematics*, 15(3):931–959, 2024.

[31] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralised finance and automated market making: Execution and speculation. *Journal of Economic Dynamics and Control*, pages 105–134, 2025.

[32] Álvaro Cartea, Fayçal Drissi, Leandro Sánchez-Betancourt, David Siska, and Lukasz Szpruch. Strategic bonding curves in automated market makers. *Available at SSRN 5018420*, 2024.

[33] Álvaro Cartea, Sebastian Jaimungal, and Yixuan Wang. Spoofing and Price Manipulation in Order-Driven Markets. *Applied Mathematical Finance*, 27(1-2):67–98, 2020.

[34] Archishman Chakraborty and Bilge Yılmaz. Informed Manipulation. *Journal of Economic Theory*, 114(1):132–152, 2004.

[35] Archishman Chakraborty and Bilge Yılmaz. Manipulation in Market Order Models. *Journal of Financial Markets*, 7(2):187–206, 2004.

[36] Jean-Edouard Colliard, Thierry Foucault, and Stefano Lovo. Algorithmic pricing and liquidity in securities markets. *HEC Paris Research Paper*, 2022.

[37] Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797, 2019.

[38] Thomas E. Copeland and Dan Galai. Information Effects on the Bid-Ask Spread. *The Journal of Finance*, 38(5):1457–1469, 1983.

[39] Dodd-Frank. Dodd–Frank Wall Street Reform and Consumer Protection Act, 2010. Public Law 111–203. Avaialble at: https://www.govinfo.gov/app/details/PLAW-111publ203.

[40] Winston Wei Dou, Itay Goldstein, and Yan Ji. AI-Powered Trading, Algorithmic Collusion, and Price Efficiency. *Available at SSRN 4452704*, 2023.

[41] Fayçal Drissi. Models of market liquidity: Applications to traditional markets and automated market makers. *Available at SSRN 4424010*, 2023.

[42] Ignacio Esponda and Demian Pouzo. Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models. *Econometrica*, 84(3):1093–1130, 2016.

[43] F Douglas Foster and S Viswanathan. Strategic trading when agents forecast the forecasts of others. *The Journal of Finance*, 51(4):1437–1478, 1996.

[44] Merritt B. Fox, Lawrence R. Glosten, and Sue S. Guan. Spoofing and its Regulation. *Colum. Bus. L. Rev.*, page 1244, 2021.

[45] Tatiana Franus, Malvina Marchese, and Richard Payne. A machine learning ensemble approach to predict financial markets manipulation. *Working Paper*, 2024.

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

[47] Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

[48] Ben Hambly, Renyuan Xu, and Huining Yang. Recent advances in reinforcement learning in finance. *Mathematical Finance*, 33(3):437–503, 2023.

[49] Lars Peter Hansen and Thomas J. Sargent. *Robustness*. Princeton University Press, 2007.

[50] Joseph E. Harrington. Developing Competition Law for Collusion by Autonomous Artificial Agents. *Journal of Competition Law & Economics*, 14(3):331 – 363, 2018.

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

[52] Campbell R Harvey. Cryptofinance. *Available at SSRN 2438299*, 2016.

[53] Joel Hasbrouck and George Sofianos. The trades of market makers: An empirical analysis of nyse specialists. *The Journal of Finance*, 48(5):1565–1593, 1993.

[54] Thomas Ho and Hans R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.

[55] Craig W Holden and Avanidhar Subrahmanyam. Long-lived private information and imperfect competition. *The Journal of Finance*, 47(1):247–270, 1992.

[56] Kose John, Barnabé Monnot, Peter Mueller, Fahad Saleh, and Caspar Schwarz-Schilling. Economics of ethereum. *Journal of Corporate Finance*, 91:102718, 2025.

[57] Todd R. Kaplan and Shmuel Zamir. Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

[58] 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.

[59] Alfred Lehar and Christine A Parlour. Decentralized exchange: The uniswap automated market maker. *Available at SSRN 3905316*, 2021.

[60] Ananth Madhavan and Seymour Smidt. An analysis of changes in specialist inventories and quotations. *The Journal of Finance*, 48(5):1595–1628, 1993.

[61] Eric Maskin and John Riley. Asymmetric auctions. *The review of economic studies*, 67(3):413–438, 2000.

[62] Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046*, 2022.

[63] Brian Ning, Franco Ho Ting Lin, and Sebastian Jaimungal. Double Deep Q-Learning for Optimal Execution. *Applied Mathematical Finance*, 28(4):361–380, 2021.

[64] OECD. Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers. Available at: `https://www.oecd.org/finance/artificial-intelligence-machine-learning-big-data-in-finance.htm`, 2021.

[65] Maureen O'Hara and George S. Oldfield. The microeconomics of market making. *The Journal of Financial and Quantitative Analysis*, 21(4):361–376, 1986.

[66] Andreas Park. The conceptual flaws of decentralized automated market making. *Management Science*, 69(11):6731–6751, 2023.

[67] Christine A. Parlour. Price Dynamics in Limit Order Markets. *The Review of Financial Studies*, 11(4):789–816, 1998.

[68] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994.

[69] Satinder Singh, Tommi Jaakkola, Michael L. Littman, and Csaba Szepesvári. Convergence Results for Single-Step On-Policy Reinforcement-Learning Algorithms. *Machine Learning*, 38(3):287–308, 2000.

[70] Thomas Spooner, John Fearnley, Rahul Savani, and Andreas Koukorinis. Market Making via Reinforcement Learning. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, page 434–442, Richland, SC, 2018. International Foundation for Autonomous Agents and Multiagent Systems.

[71] Hans R. Stoll. The Supply of Dealer Services in Securities Markets. *The Journal of Finance*, 33(4):1133–1151, 1978.

[72] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

[73] Csaba Szepesvári. *Algorithms for Reinforcement Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Springer Cham, 2010.

[74] Jos Van Bommel. Rumors. *The Journal of Finance*, 58(4):1499–1519, 2003.

[75] Jean-Luc Vila. Simple Games of Market Manipulation. *Economics Letters*, 29(1):21–26, 1989.

[76] Yun-Yi Wang. Strategic Spoofing Order Trading by Different Types of Investors in Taiwan Index Futures Market. *Journal of Financial Studies*, 27(1):65–103, 2019.

[77] Christopher J. C. H. Watkins and Peter Dayan. Q-Learning. *Machine Learning*, 8(3):279–292, 1992.

[78] Basil Williams and Andrzej Skrzypacz. Spoofing in Equilibrium. *Available at SSRN 3742327*, 2021.