Collusive Market Making and Retail Investor Attention*†

Jin Li Ye Luo Yannan Sun Xiaowei Zhang

Abstract

Collusion in market making raises execution costs and reduces liquidity. Market makers increasingly exploit algorithms to facilitate trading, making collusion more difficult to detect. To provide potential evidence to detect collusion, this paper introduces a structural model of wholesaler-broker demand system combined with a multi-agent reinforcement learning (MARL) framework to quantify the level of under-competitiveness among market makers. We structurally estimate monthly order flow demand elasticities for six major wholesalers in the U.S. equity market w.r.t. effective over quotes (EFQs). Subsequently, we make use of the estimates to train two MARL systems: one allowing collusion and the other enforcing competition, to simulate the strategic EFQ-setting behavior of the market makers. Counterfactual markets under collusive and competitive scenarios are constructed to build a metric of the under-competitiveness level for each market maker. We validate the effectiveness of this metric using Payment for Order Flow (PFOF). Exploiting the exogenous surge in retail investor attention that was mainly driven by meme stocks during 2021, we demonstrate that heightened retail attention significantly but temporarily increases competition among market makers for the affected stocks. Finally, we quantify the welfare implications: while market makers collectively reduce investor execution costs relative to the NBBO benchmark, under-competitiveness increases aggregate execution costs by approximately 16% compared to the competitive counterfactual.

Keywords: Market Microstructure, Collusion Detection, Structural Demand Estimation, Reinforcement Learning, Retail Investor Attention, Payment for Order Flow (PFOF), Welfare Implications

JEL Codes: G12, G14, G50, G24, G28

*Contact information: Jin Li (jli1@hku.hk), Ye Luo (kurtluo@hku.hk), Yannan Sun (yannansun@connect.hku.hk) Faculty of Business and Economics, The University of Hong Kong, Pokfulam Road, Hong Kong SAR. Xiaowei Zhang (xiaoweiz@ust.hk), Department of Industrial Engineering and Decision Analytics, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong SAR.

†We are thankful to Jiantao Huang, Xing Hu, and Jiaheng Yu

1 Introduction

The U.S. equity market's retail trading landscape has undergone a dramatic transformation, with over 90% of marketable retail orders in National Market System (NMS) stocks now routed to a concentrated group of six off-exchange wholesalers, where the top two firms—Citadel Securities and Virtu Financial—account for approximately 66% of their share volume as of early 2022 (SEC, 2023). This extreme concentration, evidenced by a Herfindahl-Hirschman Index (HHI) exceeding 2,500 in most stocks and rising over time (Edwin and Dermot, 2022; Dyhrberg et al., 2025), signals a highly concentrated market-making sector that raises concerns about competition and potential welfare losses for retail investors. Compounding this, wholesalers operate under a bilateral model with retail brokers, internalizing over 80% of executed retail share volume by dollar value (SEC, 2023), shielding orders from broader market competition. Unlike exchanges, where visible liquidity pools foster dynamic price discovery, this isolation minimizes wholesalers' exposure to rival pricing pressures, particularly when brokers' routing decisions are less sensitive to execution quality, creating fertile ground for coordinated behavior among wholesalers.

These competitive dynamics are further complicated by demand shocks, which serve as critical catalysts by amplifying wholesalers' inventory and adverse selection risks. A sudden influx of orders may force wholesalers to accumulate unwanted positions, as classic inventory models suggest (Ho and Stoll, 1981). To mitigate the risk of adverse price movements, they may widen bid-ask spreads, reduce price improvements, or route excess orders externally, effects that are magnified by retail orders' aggressive liquidity consumption. From the other side, heightened demand intensifies adverse selection risk, as wholesalers struggle to differentiate uninformed retail sentiment from informed trading, prompting less aggressive quoting to avoid being "picked off" (Glosten and Milgrom, 1985). This dynamic adjustment to demand shocks shifts wholesalers' strategies beyond static equilibria, offering a window into how they respond to exogenous pressures and potentially coordinate to protect profits and mitigate risks.

Motivated by these insights, this paper aims to quantify the degree of competition among wholesalers in the market-making industry and investigate how sudden retail investor attention shapes their competitive behavior. By synthesizing structural demand estimation, reinforcement learning simulations, and causal analysis, we intend to answer such a research question: What is the level of competition among wholesalers, are they fully competitive, and how do their execution strategies dynamically respond to sudden demand shocks triggered by retail investor attention?

The market-making industry in which wholesalers participate operates as follows. When retail investors place orders through brokers like Robinhood, Charles Schwab, or E*TRADE, these orders are rarely sent to public exchanges such as NYSE or Nasdaq; instead, they are predominantly routed to wholesalers, who leverage advanced technology to internalize vast order volumes efficiently. Brokers adjust routing based on wholesalers' past execution quality, with heterogeneity in evaluation frequency (monthly or quarterly) and scope (stock-by-stock or baskets) (Huang et al., 2024). Wholesalers are drawn to retail order flow due to its low adverse selection costs—retail trades are less likely to reflect private information (Glosten and Milgrom, 1985; Easley et al., 1996)—prompting them to offer price improvement, executing orders slightly better than the best quoted prices, NBBO, as a key benefit to investors. Additionally, the payment for order flow (PFOF) practice, where wholesalers share market-making profits with brokers, introduces another layer of competition in order to obtain more order flow. And the extent to which PFOF influences routing decisions varies across brokers, depending on its weight in their revenue streams, potentially skewing incentives away from pure execution quality considerations.

The competitiveness of off-exchange retail order execution remains a subject of intense debate among academics and regulators. SEC has expressed concern that high concentration in the wholesaler market provides limited benefits to retail investors, and proposed a new rule in 2023 to improve order competiiton, see SEC (2023). The academic literature also presents different evidence and conclusions on this matter. One stream of research suggests the current structure is largely beneficial. Proponents argue that wholesalers deliver substantial value through low trading costs and consistently better execution prices compared to public exchanges (Kothari et al., 2021; Battalio and Jennings, 2023; Dyhrberg et al., 2025). Theoretically, wholesalers can leverage non-toxic retail order flow to offset inventory imbalances, thereby enhancing execution quality (Baldauf et al., 2024). Evidence from anonymous broker data further supports this view, showing that brokers route orders responsively based on wholesalers' past performance, suggesting active competition on price improvement (Ernst et al., 2024). Conversely, another body of work raises concerns about the exercise of market power. These studies argue that wholesalers may exploit their dominant positions to maintain wide spreads and enhance profits (Edwin and Dermot, 2022; Van Kervel and Yueshen, 2025). Furthermore, evidence from proprietary, self-generated order data indicates that wholesaler competition is far from perfect (Huang et al., 2024). The intricate nature of this broker-wholesaler marketplace, characterized by competing effects in the multiple dimensions and motivates the structural approach taken in this paper. Our results reveal that collusion is a continuous dynamic, which fluctuated with different market conditions, rather than a two-sided indicator of wholesalers.

These institutional practices, combined with the concentrated market structure, give rise to the intense discussions among regulators and academics. Regulators have expressed concern that high wholesaler concentration limits the benefits to retail investors, prompting proposals aimed at enhancing order-by-order competition (SEC, 2023). The academic literature offers conflicting conclusions. One stream of research suggests the current structure is largely efficient. Proponents argue that wholesalers provide substantial value through superior pricing and lower trading costs compared to public exchanges (Kothari et al., 2021; Battalio and Jennings, 2023; Dyhrberg et al., 2025). This view is supported by evidence that brokers route orders responsively to wholesalers' past performance, implying active competition on price improvement (Ernst et al., 2024), and is consistent with theories of internalization benefiting from non-toxic retail order flow (Baldauf et al., 2024). Conversely, another body of work raises concerns about the exercise of market power, arguing that dominant wholesalers may maintain wide spreads to enhance profits (Edwin and Dermot, 2022; Van Kervel and Yueshen, 2025). (Huang et al., 2024) uses self-generated experimental data and shows that competition is far from perfect. This literature thus presents a puzzle when we try to conclude the wholesalers' competition level. The root of this tension lies in the market's complexity, wherein wholesalers do not compete on price improvement alone but across a bundle of attributes—including execution speed and PFOF—that are difficult to measure in isolation. This multi-dimensional competition makes it challenging to recover the underlying demand elasticities that ultimately govern broker routing decisions.

Therefor, to properly measure the extent of competition, we specify and estimate a structural model of the demand system between brokers and wholesalers, in the spirit of **?** and Nevo (2001). The first step of our analysis is to recover the primitive demand elasticities of broker routing decisions with respect to price improvement (using a standard measure, Effective-Over-Quoted Spread, EFQ), execution speed, and other execution quality measures, and then understand how the brokers' heterogeneity varies on these execution qualities. We find that demand is elastic, with EFQ elasticity of -6.26, indicating that a 1% increase in EFQ reduces market share of retail order flow by 6.26%, implying a competitive oligopoly structure. This demand elasticity is heterogeneous across firm size, being substantially higher for S&P 500 stocks (-12.0) than for smaller, non-S&P 500 stocks (-5.7), indicating stronger price competition for larger firms. The heterogeneity of demand elasticity on execution quality is also stark across liquidity levels: the elasticity is most pronounced for the most liquid quintile of stocks (-9.6) and declines monotonically to its lowest level for the most illiquid quintile (-2.5). These results imply that brokers are more sensitive to the change of EFQ in the markets of larger and more liquid stocks.

With the structural model, we have a micro-founded brokers assignment function based on wholesalers' past performance and varies with brokers' heterogeneity. Then we employ model-based reinforcement learning (RL) agents to simulate credible counterfactuals by constructing AI wholesalers that act competitively or collusively under the same conditions. And then we can use the deviation of EFQ from the competitive counterfactual to construct a novel Under-Competitiveness Index (UCI), to quantify the degree of competition of wholesalers. Our analysis of the UCI reveals that while the largest wholesalers are more competitive on average, they exhibit significantly less competitive behavior when executing orders for smaller stocks.

Third, to causally identify that if the demand shock that may affect the competition dynamics of wholesalers, we use a difference-in-differences (DiD) design centered on the 2021 meme-stock event. We find that retail attention shocks lead to a temporary but significant decrease in competition, increasing our UCI by 0.109, implies that wholesalers act collusively to encounter the inventory risk and the adverse-selection risk.

Finally, we quantify the welfare implications of the current competition level of wholesalers. While wholesalers generate substantial cost savings for retail traders relative to the NBBO benchmark, we estimate that under-competitiveness imposes an additional 16% in aggregate execution costs compared to our simulated competitive counterfactual.

The remainder of the paper proceeds as follows. Section 2 reviews the institution background to briefly introduce the current market-making industry, and also discusses related literature on wholesaler execution quality, U.S. equity retail trading competition, and algorithmic collusion. Section 3 describes the data sources and construction method for the major variables. Section 4 presents the structural model of the demand system between wholesalers and brokers and its theoretical foundations, as well as the estimation algorithms and the estimation results. Section 5 discusses the Multi-Agent Reinforcement Learning (MARL) algorithms that is used to construct the AI wholesalers that act collusively or competitively in the same historical market conditions, and thus build dynamic counterfactuals for the real wholesalers. Section 6 conducts the external validation of the under-competition level we constructed. Section 7 reports the causal inference results from the DiD analysis. Section 8 shows the counterfactual analysis on retail investors' welfare, proxied by the trading costs. Section 9 concludes with implications and avenues for future research.

2 Institutional Background and Related Literature

2.1 Institutional Background

In the U.S. equity market, retail investors submit orders through brokerage firms, which then route these orders to various execution venues, with the majority directed to off-exchange wholesalers (SEC, 2023). Wholesalers, such as Citadel Securities and Virtu Financial, favor retail order flow due to its lower adverse selection risk compared to institutional orders, as retail trades are typically less informed and more predictable (Glosten and Milgrom, 1985). To attract this flow, wholesalers often provide price improvements (PIs), offering executions better than the national best bid and offer (NBBO) by small increments, such as sub-penny amounts, which are not available on lit exchanges due to the limitation of minimum tick size.

Brokers evaluate wholesalers' past performance to determine routing decisions, with the effective-over-quoted spread (EFQ) emerging as a key metric of execution quality Ernst et al. (2024), alongside other indicators like execution likelihood, speed, and size of order fills (Financial Industry Regulatory Authority (FINRA), 2024). Routing adjustments occur periodically, often on a monthly or quarterly basis; some brokers route orders using a method close to proportional routing, by evaluating the past performance and route at the aggregate level of stocks, while others use selective routing on a stock-by-stock level (Huang et al., 2024). We find this assignment process can be modeled structurally using demand estimation methods, to capture brokers' discrete choices among wholesalers based on observed and unobserved characteristics (Berry et al., 1995; Nevo, 2000).

A critical institutional feature is the payment for order flow (PFOF) mechanism, whereby wholesalers share a portion of their market-making profits with brokers in exchange for directed order flow. While PFOF incentivizes brokers to route orders to high-paying wholesalers, it may distort incentives and conflict with brokers' best execution obligations, potentially prioritizing revenue over optimal execution quality. Upon receiving orders, the execution of retail orders faces no competition, as the wholesalers can decide whether to internalize the orders or send a portion of the orders to other venues. As wholesalers need to accept all the marketable orders that they receive, the wholesalers can choose to further route orders to other venues when they are facing severe inventory risk, like the case that the received orders are imbalanced.

This market-making industry for retail orders in the US equity market is concentrated, as in **?**, six wholesalers received over 90% of retail orders, and the top 2 (Citadel and Virtu) take over 66% in the first quarter of 2022. Dyhrberg et al. (2025) provides more evidence that the majority of the stocks' Herfindahl–Hirschman Index (HHIs) is greater than 1,500 and even 2,500, showing the highly concentrated environment.

2.2 Related Literature

There is a rising literature that investigates the execution quality of retail orders in the current market structure. Battalio and Jennings (2023) uses the data that includes all marketable orders routed to specific wholesalers to analyze the execution quality, and they find that compared to the SEC Rule 605 reports, the execution quality of wholesalers is substantially better in their data. Brown et al. (2024) use data from the broker Robinhood, and they find that though brokers receive a large amount of PFOF, retail trades still receive better execution quality in terms of lowered total costs. Ernst et al. (2024) uses proprietary data from three anonymous brokers, and they investigate the execution quality from the perspective of brokers. Their results show that, though have different routing strategies, the brokers are responsive to wholesalers' historical performance. Huang et al. (2024) conduct a field experiment where they use self-generated orders submitted to six brokers, and their results show that most brokers hardly change their routing. Due to this stickiness of routing strategies, they conclude that the competition among wholesalers is far from perfect. Edwin and Dermot (2022) uses the internalization data and finds that the internalization of wholesalers is associated with higher spreads and worse price improvement, and thus increases the cost of retail investors. Our work is closest to Dyhrberg et al. (2025), which investigates the execution quality from the wholesalers' perspective, and also uses the Rule 605 data. However, the find that wholesalers are competitive and that the better prices offered by wholesalers are due to economies of scale.

Our work is also related to the literature on algorithmic collusion, most of the papers use algorithmic experiments to simulate and uncover the black box of potential algorithmic collusion (Calvano et al., 2020; Dou et al., 2024; Asker et al., 2022; Li et al., 2025), these papers build the model-free reinforcement learning framework to let the algorithms learn a collusive policy or a competitive one without economic conditions to restrict the equilibrium. On the other hand, Assad et al. (2024) provides empirical evidence of algorithmic collusion in Germany's retail gasoline markets in a staggered DiD setting. We contribute to this literature by firstly obtaining a structural model to articulate the economic conditions, and then embedding the conditions in Multi-Agent Reinforcement Learning (MARL), in this way, the static equilibrium estimated from structural model can evolve to a dynamic that adapts to different market conditions, and the dynamics constructed by MARL can be bounded economically without falling to the black hole due to algorithmic facet.

3 Data

We obtain wholesaler-level monthly order execution quality data for all National Market System (NMS) stocks traded in the US from publicly available Rule 605 data, along with monthly routing and Payments for Order Flow (PFOF) data from Rule 606 reports. For stock-level characteristics, we utilize data from the Center for Research in Securities Prices (CRSP). Additionally, we extract quarterly-level broker characteristics from brokers' 10-K reports through the EDGAR API.

3.1 Order Execution Quality

We collect wholesaler-level monthly order execution quality data for all National Market System (NMS) stocks traded in the United States from publicly available Rule 605 reports, which are mandated by the Securities and Exchange Commission (SEC) under Regulation NMS. Rule 605 requires market centers to publicly disclose detailed monthly statistics on their order execution quality, as mentioned in Dyhrberg et al. (2025), Rule 605 data is widely recognized in the industry, which mainly covers the retail orders and therefore the best public data source to analyze retail execution quality. The data structure captures granular information at the stock-month-market center level, where each observation contains detailed execution metrics including the total number of orders, shares executed, and various measures of execution quality. Specifically, the dataset includes execution statistics disaggregated by order size categories, price improvement metrics (measured in both dollar terms and basis points), and speed of execution across different market centers. The data also distinguishes between market orders, limit orders, and other order types, with separate reporting for orders executed at the National Best Bid and Offer (NBBO), orders executed outside the NBBO, and orders that receive price improvement. This granular structure allows us to construct measures of execution quality at the wholesaler-stock-month level.

We collect the ten largest wholesalers' Rule 605 data from January 2019 to June 2024. According to SEC reports, industry practices, and Dyhrberg et al. (2025), the top six wholesalers' market shares are more than 90% in the first quarter of 2022 [1], and in Dyhrberg et al. (2025), the top two wholesalers, Citadel and Virtu, dominate for almost 70% of all retail flow. This concentration shows that SEC's concerns about the lack of competition in retail order execution are worthwhile for this research to understand wholesalers' under-competition level.

Our execution quality measures follow Dyhrberg et al. (2025) and Ernst et al. (2024) by

[1] See SEC proposed the order competition rule and analysis in 2023, Order Competition Rule

including quoted, effective, and realized spreads. Quoted spreads represent the publicly visible bid-ask spread at the time of order submission, reflecting the baseline cost of immediate execution in the market. Effective spreads measure the actual execution price relative to the midpoint of the prevailing quote at the time of order submission, capturing the realized cost of execution including any price improvement or deterioration that occurs during the execution process. Following Ernst et al. (2024), we construct an Effective-over-Quoted (EFQ) that provides a standardized measure of price improvement by normalizing the effective spread against the quoted spread, where lower EFQ values indicate larger savings for retail investors relative to the publicly available prices. Unlike effective spreads alone, which can vary due to market-wide liquidity conditions, EFQ isolates the wholesaler-specific contribution to execution quality by controlling for the prevailing market spread environment. Ernst et al. (2024) brings important insights by using the broker's proprietary data, showing that brokers are responsive to wholesalers' past performance in terms of offering better prices, and the most used measure is EFQ. Following their work, we use EFQ as the most important execution quality measure as the "price" in the demand estimation in the product market.

To capture the multi-dimensional nature of execution quality, we extend our analysis beyond EFQ by incorporating additional Rule 605 metrics that reflect distinct dimensions of execution performance. Execution speed, measured across order size categories ranging from 100 to 9,999 shares and disaggregated by order type, captures the temporal efficiency dimension of execution quality. This metric reflects wholesalers' operational capabilities and their ability to minimize market impact through rapid order processing. The ratio of shares executed at the quote (at-quote ratio) measures execution precision, capturing the wholesaler's ability to execute orders at the publicly displayed prices without experiencing price slippage. The realized spreads capture the execution cost relative to the quote midpoint at a specified time interval after execution (typically 5 minutes), and help to measure the sustained cost efficiency after market impact considerations, aligning with Dyhrberg et al. (2025) that realized spreads are toxicity-adjusted execution cost. These measures complement EFQ by addressing different aspects of execution quality: EFQ captures the price improvement dimension (cost efficiency), execution speed reflects the temporal dimension (operational efficiency), the at-quote ratio captures the precision dimension (execution reliability), and the realized spreads evaluate from the adverse selection dimension (information risk), which investors value because it measures the wholesaler's ability to absorb and manage the permanent price impacts from informed trading.

3.2 Wholesaler Liquidity Provision Costs

Providing liquidity to retail orders also induces costs to wholesalers. We construct the measures of costs to provide liquidity using Rule 605 data. We calculate the firm-month volatility of realized spreads as the standard deviation of market makers' average realized spreads, labeling this variable as "Firm-Month Spread Volatility." This measure serves as a proxy for the variability in costs associated with liquidity provision, particularly capturing elements of adverse selection risks as highlighted in seminal works (Copeland and Galai, 1983; Glosten and Milgrom, 1985; Easley and O'hara, 1987). Additionally, we derive an inventory imbalance proxy by computing the ratio of away executed shares to total shares (setting it to zero when total shares are zero to handle edge cases), labeled as "Inventory Imbalance Proxy." This variable reflects potential inventory holding costs that wholesalers incur when managing unbalanced positions, consistent with inventory management theories in market making (Stoll, 1978; Ho and Stoll, 1981).

3.3 Broker characteristics

To capture the economic incentives and strategic preferences that influence broker behavior in order routing decisions, we construct broker characteristics from two complementary data sources: Rule 606 reports and 10-K filings from the EDGAR database. Rule 606 reports, mandated by the SEC, provide monthly information on broker order routing practices and payments for order flow (PFOF) arrangements, which may distort the financial incentives that brokers receive from wholesalers for routing customer orders. These reports disclose the monetary payments that brokers receive from market makers in two subsamples, S&P 500 stocks and all others. From the 10-K filings, we obtain additional broker characteristics at the quarterly level, including assets under custody (AUC) that reflect brokers' scale and brokers' total revenue. Then we construct *PFOF Revenue Ratio*, dividing PFOF revenue by total revenue. This reflects the extent to which brokers rely on PFOF arrangements as a source of revenue, which is a concern that may influence their routing incentives and decision-making behavior and incur a fundamental principal-agent conflict.

We argue that PFOF represents an endogenous strategic choice by wholesalers operating at a different competitive dimension. Initial consideration of this issue might suggest that PFOF constitutes an important omitted variable that influences brokers' wholesaler routing decisions. However, deeper thinking reveals that PFOF is fundamentally an endogenous choice made by wholesalers at a strategic level distinct from their execution quality decisions. The strategic sequence operates as follows: after wholesalers determine their EFQ levels (execution quality), they strategically decide how much of their profits to allocate as PFOF payments to brokers. Furthermore, the extent to which PFOF can distort brokers' routing decisions is fundamentally determined by broker characteristics, specifically the degree to which PFOF contributes to their total revenue structure. This insight suggests that the effectiveness of PFOF as a strategic tool depends not on the absolute magnitude of PFOF payments, but on the relative importance of PFOF revenue within brokers' overall revenue streams. Therefore, we incorporate the PFOF revenue ratio as a critical broker characteristic in our structural estimation framework, help to estimate how this characteristic affects the brokers' taste of wholesalers' execution quality. Additionally, we utilize wholesaler-level PFOF data as external validation of wholesaler under-competition levels, and this can help to provide evidence that in two strategic dimensions, how wholesalers compete.

4 Structural Model of Wholesaler-Broker Demand System

In this paper, we develop a structural model to estimate the demand system for wholesaler services by retail brokers. We employ a random coefficients logit model, in the spirit of Berry et al. (1995), where brokers choose among wholesalers based on execution quality and cost characteristics. A central challenge in this setting is the principal-agent conflict: brokers receive Payment for Order Flow (PFOF) from wholesalers, an incentive not directly aligned with their clients' interests. Our framework does not assume this conflict away. Instead, we directly quantify its impact by explicitly modeling the trade-off brokers face. While competitive pressures (e.g., Ernst et al. (2024)) and the "duty of best execution" (SEC, 2023) compel brokers to serve client interests, our model allows for heterogeneity in this behavior. Specifically, we interact key execution quality attributes with broker-level characteristics, such as the ratio of PFOF to total revenue. This approach allows the data to reveal how a broker's sensitivity to execution quality changes with its reliance on PFOF.

4.1 Microfoundations

Suppose there are markets $t = 1, ..., T$, each with $i = 1, ..., I_t$ brokers. In the following estimation, a market will be defined as a stock-month combination. Let $u_{i,j,k,t}$ denote the indirect utility that broker i received from the execution service of wholesaler j in market t on stock k.

$$u_{i,j,k,t} = \alpha_{ik} EFQ_{jkt} + \mathbf{x}'_{jt}\beta_{ik} + \xi_{jt} + \epsilon_{ijkt}, \tag{1}$$

where $\mathbf{x}'_{jt}$ is a K-dimensional (row) vector of observable characteristics of product j (wholesaler j's execution service), ξ_{jt} is the unobserved (by researcher) product characteristics, and ϵ_{ijkt} is the mean-zero stochastic error term. So finally, α_{ik} is broker i's marginal utility from EFQ, and β_i is a K-dimensional (column) vector of broker-specific taste coefficients.

Examples of observed characteristics include execution Speed, the average time required to fill an order, realized spread, which captures the investor's implicit trading cost, and at-the-quote ratio, which reflects the percentage of shares executed at the prevailing National Best Bid and Offer (NBBO). The structural error term in our model, ξ_{jt}, captures all wholesaler characteristics that are unobserved in our data but are valued by brokers. A primary component of ξ_{jt} is the quality of the broker-wholesaler relationship, which encompasses the responsiveness of service teams, operational reliability, and long-term trust. Furthermore, it includes latent technological capabilities, such as the sophistication of a wholesaler's smart order router, its ability to handle illiquid or complex orders, and its capacity for risk absorption on large trades. Since these unobserved attributes are likely correlated with observed characteristics like realized spreads—for instance, a technologically superior wholesaler may offer both better service and better prices—they present a classic endogeneity problem, necessitating the use of instrumental variables for identification.

Beyond the unobserved characteristics of wholesalers ($\xi_j t$), brokers themselves exhibit heterogeneity in their preferences for these attributes. To capture this, we model each broker i's taste parameters as a function of both observable and unobservable components. We explicitly incorporate observable broker characteristics (which is referred to as demographics in demand estimation literature like Berry et al. (1995)) to explain systematic variation in preferences.

Crucially, we include a broker's ratio of PFOF to total revenue to quantify how the principal-agent conflict influences its routing decisions. By interacting this characteristic with the taste parameters for execution attributes, our model can empirically estimate the marginal rate of substitution between client price improvement and PFOF revenue, providing cross-sectional variation that helps to separately identify brokers' intrinsic preferences for quality from the influence of agency frictions. Furthermore, we include Assets Under Custody (AUC) to control for heterogeneity related to broker scale and sophistication. Larger brokers may possess superior routing technology or greater bargaining power with wholesalers, and controlling for

AUC prevents this variation from confounding our primary estimates.

Formally, we model the vector of a broker's individual taste parameters as a linear function of these broker characteristics, D_i, and a vector of unobserved taste shocks, v_i, which follows a multivariate normal distribution. This specification is as follows:

$$\begin{pmatrix} \beta_i \\ \alpha_i \end{pmatrix} = \begin{pmatrix} \beta \\ \alpha \end{pmatrix} + \Pi D_i + \Sigma v_i, \tag{2}$$

where β is a K-dimensional column vector and each of the entries represents the mean taste to each product characteristic, α is the mean utility to EFQ. Π is a $(K+1) \times d$ matrix of coefficients that measure how the taste coefficients vary with broker characteristics, and Σ is a $(K+1) \times (K+1)$ scaling matrix; v_i captures the additional characteristics that are unobserved.

The demand system consists of one outside good to account for the case that brokers may decide not to route orders to any of the top wholesalers. Without this, a simultaneous increase in EFQ of top wholesalers will not change the brokers' assignment function. The outside good may be provided by other small wholesalers and exchanges. And the indirect utility from this outside option is:

$$u_{i0kt} = \xi_{0t} + \pi_0 D_i + \sigma_0 v_{io} + \epsilon_{i0kt} \tag{3}$$

The mean utility from the outside good, ξ_{0t} is not identified, and the coefficients π_i and σ_0 are not identified separately from coefficients on an individual-specific constant term. The standard practice is to set $\xi_{0t}, \pi_0, \sigma_0$ to zero, as not making more assumptions, and this is equivalent to normalizing the utility from the outside good to zero.

4.2 Demand Derivation and Market Shares

Building on the utility specification, the choice probability for a given broker is derived. The full vector of parameters to be estimated is $\theta_1 = (\alpha, \beta)$, where β contains the mean taste parameters, and $\theta_2 = (\Pi, \Sigma)$ are the nonlinear parameters governing the influence of demographics and unobserved heterogeneity.

Define the set $A_{jkt} = (\epsilon_{kt} : u_{jkt} \geq u_{mkt}, \forall m \neq j)$, where $\epsilon_{kt} = (\epsilon_{ijk0}, ..., \epsilon_{iJkt})$, expands all the products and all the markets, the set A_{ijkt} includes all the error terms that when the utility is the largest among all products in each market.

Then the predicted market share of the jth product as a function of the mean utility levels of all the $J+1$ goods is

$$s_{jkt} = \int_{A_{jkt}} dP^*(D, v, \epsilon) = \int_{A_{jkt}} dP^*(\epsilon) dp^*(v) dp^*(D), \tag{4}$$

where P^* denotes the population distribution function among brokers.

4.3 Endogeneity and Instrumental Variables

A primary challenge in estimating the demand model is the endogeneity of wholesaler execution quality characteristics. A wholesaler's choice of execution quality is not random; it is strategically chosen and therefore likely to be correlated with unobserved, non-quantifiable attributes (ξ_{jt}), such as the quality of its service relationship with a broker or its technological reliability. This correlation would bias our parameter estimates. To address this, we employ a Generalized Method of Moments (GMM) framework using a comprehensive set of instruments designed to be correlated with the endogenous execution characteristics while remaining exogenous to the unobserved demand shocks.

The first and primary class of instruments consists of exogenous cost shifters that affect the wholesaler's marginal cost of providing liquidity. These are valid instruments because they are correlated with the endogenous quality variables but are driven by broad market and inventory dynamics rather than wholesaler-specific unobserved demand attributes. Specifically, we construct three shifters: *Idiosyncratic Execution Time Shocks* (ET_Resid), and we regress the deviation of a wholesaler's execution time from its monthly mean on a set of ticker-month fixed effects. The residual from this regression captures idiosyncratic operational delays, also aligns the order processing costs in Demsetz (1968) and Tinic (1972). These shocks increase inventory holding costs for the wholesaler, affecting its pricing, but are by construction orthogonal to systematic firm- and time-specific factors that might be part of ξ_{jt}. The second one is *Realized Spread Volatility* (Vol_Reliz_Spr), which is calculated as the monthly standard deviation of a wholesaler's realized spread for a given ticker. This variable proxies for the unpredictability of its profits and the risk of providing liquidity. Higher volatility reflects greater inventory risk, influencing the wholesaler's offered execution quality, but is driven by market dynamics, not unobserved service quality. *Order Imbalance* (Ord_Imb), which is defined as the proportion of shares a wholesaler routes away rather than internalizing. A high ratio indicates an inventory imbalance that the wholesaler cannot absorb, increasing its holding costs and risk. This reflects a cost shock driven by order flow composition, which is plausibly exogenous to the unobserved quality of the wholesaler itself. We also include squared terms of these cost shifters to capture potential non-linear effects of costs on quality.

In addition to these cost shifters, we include two other classes of instruments to enhance identification. We use BLP-style instruments (Berry et al., 1995). This includes the sum of rivals' characteristics and the count of the number of rivals. These instruments are powerful predictors of a wholesaler's own quality choices due to strategic competitive responses. They satisfy the exclusion restriction under the standard assumption that the attributes of one wholesaler's rivals do not directly enter a broker's utility function for that wholesaler, except through their influence on the wholesaler's own chosen attributes.

And another group is interactions of cost shifters and product characteristics. In this group, we interact with the aforementioned cost shifters with observable product characteristics. The relevance of these instruments comes from the fact that the effect of a cost shock on a wholesaler's strategy likely depends on the product being traded. For example, a market-wide volatility shock will have a different impact on the quality offered for a liquid versus an illiquid stock. The exclusion restriction holds because the interaction of two exogenous variables (a valid cost shifter and an observable product characteristic) is itself plausibly uncorrelated with the unobserved shock ξ_{jt}.

4.4 Data and Estimation Algorithm

The data required to estimate the model consist of the following variables: market shares and EFQ in each market (in this paper we define it as a stock-month), product characteristics (other execution quality measures, including execution speed, average realized spread, and at-quote ratio), and broker characteristics.

Market shares and EFQ we obtained from Rule 605 reports, and these data are aggregated by wholesaler, stock, and month. The data covers up to 12,000 NMS stocks from January 2019 to June 2024. Execution speed (in the following analysis is the variable avg_time) is also obtained in Rule 605 data by doing a share-weight average of execution time across groups of 100-9,999 shares. Average realized spread (avg_reliz_spr) is obtained in Rule 605 data by doing share-volume weighted among different order-type groups to aggregate to wholesaler, stock, and month. And the at-quote ratio (at_quo_ratio) is the value of shares at the quote divided by the total shares, also from Rule 605 data.

Broker characteristics are obtained from Rule 606 data and 10-K reports on EDGAR. The PFOF revenue ratio ($PFOF\_rev\_ratio$) measures how much of the PFOF contributes to the broker's total revenue. And the asset under custody (AUC) proxies the broker size. The estimation is to solve

$$Min||s_{jkt} - \mathcal{S}||, \tag{5}$$

where s_{jkt} are the market shares given by equation 4. $\mathcal{S}$ are the observed market shares. Let $Z = [Z_1, ..., z_M]$ be a set of instruments such that

$$E[Z_m \omega(\theta^*)] = 0, \tag{6}$$

where ω, is a function of the model parameters θ, and θ^* denotes the true values of the model parameters. The GMM estimate is

$$\hat{\theta} = argmin_{\theta} \omega(\theta)' Z \phi^{-1} Z' \omega(\theta), \tag{7}$$

where ϕ is a consistent estimate of $E[Z' \omega \omega' Z]$. Following Berry et al. (1995), the error term is not defined as the difference between the observed and predicted market shares, rather, it is defined as the structural error, ξ_{jt}. In order to use equation 7, we need to express the error term as an explicit function of the parameters of the model and the data, the key insight is that as in equation 4, the error term ξ_{jt} only enters the mean utility level, δ_{jkt}. Furthermore, the mean utility level is a linear function of ξ_{jt}. Thus, we solve for each market the implicit system of equations

$$s(\delta_{jkt}; \theta_2) = \mathcal{S}_t \tag{8}$$

In solving the above equations, first we need to compute the left-hand side of equation 8, and which is defined by equation 4. And the integral defining the market shares in equation 4 has to be computed by simulation. Second, using the computation of the predicted market share, we invert the system of equations. Since the system of equations 8 is nonlinear, it is solved numerically using the method of contracting mapping suggested by BLP

$$\delta_{\cdot t}^{h+1} = \delta_{\cdot t}^{h} + ln\mathcal{S}_{\cdot t} - ln s_t(\cdot), h = 0, ..., H \tag{9}$$

where $s(\cdot)$ are the predicted market shares computed in the first step, H is the smallest integer such that $||\delta_{\cdot t}^{H} - \delta_{\cdot t}^{H-1}||$ is smaller than some tolerance level. Once the inversion has been computed, the error term is defined as

$$\omega_{jt} = \delta_{jt}(S_{\cdot t}; \theta_2) - (x_{jt}\beta + \alpha p_{jt}) \equiv \xi_{jt}, \tag{10}$$

and the observed market shares, $\mathcal{S}$, enter this equation.

The intuition behind the equation 10 is that, for given values of nonlinear parameters θ_2, we solve for the mean utility level $\delta_{.t}(\cdot)$ that set the predicted market share equal to the observed market shares, we define the residual as the difference between this valuation and the one predicted by the linear parameters α and β. The estimator in equation 7 is the one minimizes the distance between these different predictions.

The estimation proceeds via a nested fixed-point algorithm. The "outer loop" is a numerical search over the non-linear parameters (Π, Σ) to minimize the GMM objective function in equation 7. For each guess of (Π, Σ) in this outer loop, the "inner loop" solves the contraction mapping for each market to recover the mean utilities, δ_t. These mean utilities, in turn, allow us to compute the structural error vector ξ, form the sample moments $Z'\xi$, and evaluate the GMM objective. The final parameter estimates, $\hat{\theta}$, are those that set these sample moments as close to zero as possible, as determined by the weighting matrix.

4.5 Estimation Results

Figure 1 shows the kernel density estimate of the EFQ elasticity parameter α_{EFQ}. This figure reveals a distribution that is skewed toward more negative values, with a mean of -5.93 and a standard deviation of 3.70, showing heterogeneity in broker demand responsiveness to execution quality (EFQ) across ticker-month markets. The density peaks around -5 to -6, indicating that a 1% increase in EFQ typically reduces market share by approximately 5.93%, consistent with moderately elastic demand in a differentiated oligopoly where wholesalers face meaningful competitive constraints. Notably, the left tail extends beyond -10 (highlighted in red), encompassing about 15% of observations where elasticity is highly elastic, suggesting segments with intense competition and low market power, while the right tail (orange line between -1 and -15) captures the bulk of the distribution, and a small mass above -1 (yellow line) points to inelastic pockets potentially indicative of localized collusion or barriers exist in a small group of stocks.

Table 1 shows the summary statistics of estimation results across 8290 stocks. The price elasticity parameter, α_{EFQ}, which measures responsiveness to execution quality (EFQ), is statistically significant in 84.86% of the 8,290 valid stocks, with a mean value of -6.2639 among significant cases. This indicates moderately elastic demand overall, where a 1% increase in EFQ reduces market share by approximately 6.26%, consistent with a competitive oligopoly structure where wholesalers face substantial substitution pressures. In contrast, $\beta_{avg_t ime}$ (sensitivity to average execution time) is significant in 60.23% of cases (mean -0.0196),

Figure 1: Kernel Density of α_{EFQ}

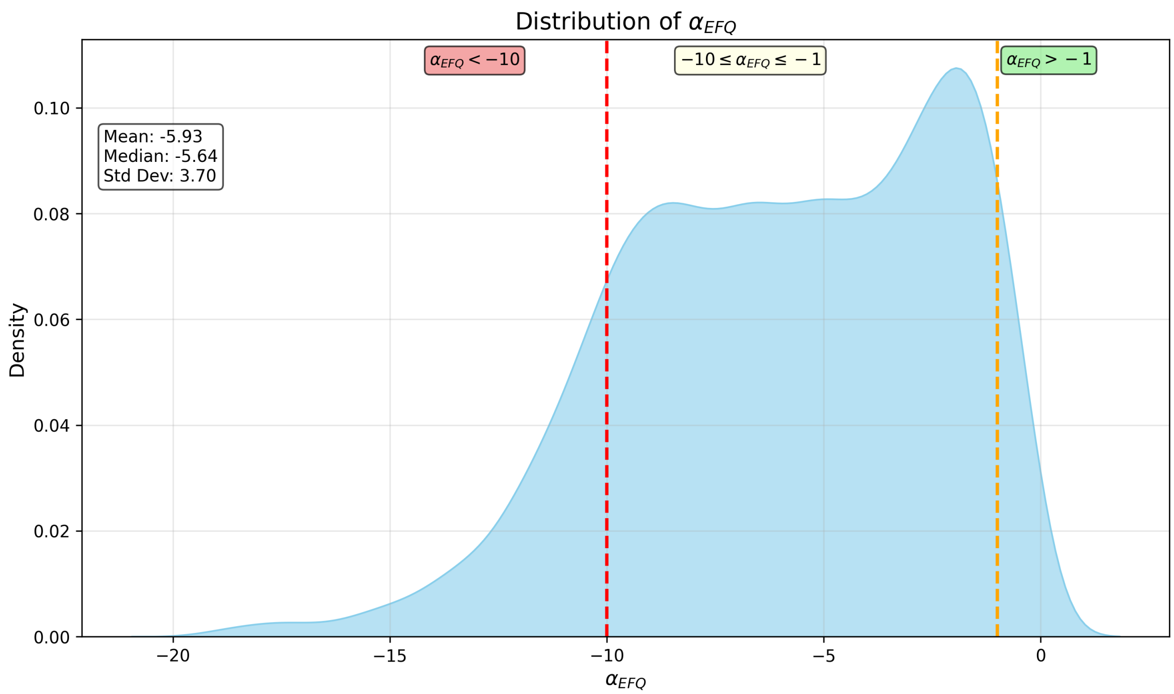

Note: This figure displays the kernel density estimate for the α_{EFQ} parameter, which represents the demand elasticity (i.e., brokers' order routing sensitivity to EFQ change) of EFQ. The distribution is constructed from the full cross-section of parameter estimates obtained from structural estimation among all available NMS stocks from January 2019 to June 2024. The sample distribution is left-skewed, with a mean of -5.93, a median of -5.64, and a standard deviation of 3.70. The dashed vertical lines at -10 (red) and -1 (orange) delineate three distinct regimes, categorizing the estimates into high sensitivity ($\alpha_{EFQ} < -10$), moderate sensitivity ($-10 \leq \alpha_{EFQ} \leq -1$), and low sensitivity ($\alpha_{EFQ} > -1$) groups.

suggesting brokers penalize delays, though less consistently than EFQ. The parameters $\beta_{avg_reliz_spr}$ (realized spread, 37.78% significant, mean -2.1643) and $\beta_{at_quo_ratio}$ (at-quote ratio, 55.83% significant, mean 0.4616) exhibit lower significance rates, implying that spread efficiency and quote reliability influence demand heterogeneously across markets, potentially moderated by broker demographics. Notably, the random coefficient variance σ_{EFQ} is significant in only 0.59% of cases (mean 4.1304), showing challenges in identifying heterogeneity in EFQ preferences; one of the possible reasons could be limited demographic variations, and another possible one is that the brokers truly have small taste heterogeneity. Demographic interactions, such as $\pi_{EFQ_PFOF_rev_ratio}$ (0.24% significant, mean 71.4266) and π_{EFQ_AUC} (2.26% significant, mean -0.8997), show even lower significance, suggesting that while PFOF reliance and assets under custody shape EFQ sensitivity for some brokers, these effects are not pervasive across the sample.

Table 1: Summary Statistics of Estimation Results

Parameter	**Num. Significant**	**Total Valid**	**% Significant**	**Mean if Significant**
α_{EFQ}	7035	8290	84.86%	-6.2639
β_{avg_time}	4993	8290	60.23%	-0.0196
$\beta_{avg\_reliz\_spr}$	3132	8290	37.78%	-2.1643
$\beta_{at\_quo\_ratio}$	4628	8290	55.83%	0.4616
σ_{EFQ}	49	8290	0.59%	4.1304
$\pi_{EFQ\_PFOF\_rev\_ratio}$	20	8290	0.24%	71.4266
π_{EFQ_AUC}	187	8290	2.26%	-0.8997

Note: This table reports summary statistics for the estimated parameters across the full sample of 8,290 valid estimations. Statistical significance is determined at the 5% level, requiring a t-statistic with an absolute value greater than 1.96. The column "Num. Significant" reports the number of parameter estimates meeting this criterion. "% Significant" expresses this count as a percentage of the "Total Valid" estimations. The final column, "Mean if Significant", presents the cross-sectional average of the parameter calculated using only the subset of statistically significant estimates.

The estimation results of the structural model, summarized in Table 2, reveal significant heterogeneity in key parameters across size-based subgroups, providing insights into broker demand responsiveness and wholesaler market power. The price elasticity parameter, α_{EFQ}, which captures the sensitivity of market share to changes in execution quality, especially the price improvement dimension (EFQ), is highly negative across all groups, ranging from -11.9547 in the S&P 500 subgroup to -3.2723 in the smallest tercile (Tercile 3), with a full-sample mean of -6.2574. This indicates elastic demand overall, where a 1% increase in EFQ leads to a 6.26% decrease in quantity demanded, consistent with a competitive market structure characterized by low wholesaler pricing power. Notably, elasticity is more pronounced in larger subgroups (S&P 500 and Tercile 1), suggesting brokers are more sensitive to quality in high-volume markets, potentially due to greater substitution options. The β parameters for execution characteristics, such as β_{avg_time} (-0.0195 full-sample) and $\beta_{avg\_reliz\_spr}$ (-2.2678 full-sample), exhibit negative signs in most cases, implying that brokers prefer faster execution and lower realized spreads, though the positive $\beta_{at\_quo\_ratio}$ (0.4568 full-sample) highlights a premium on quote reliability. Heterogeneity in EFQ preferences, captured by σ_{EFQ} (4.1304 full-sample), is higher in larger subgroups, showing broker taste variations on unobserved characteristics. Demographic interactions, such as $\pi_{EFQ\_PFOF\_rev\_ratio}$ (71.4266 full-sample) and π_{EFQ_AUC} (-0.8997 full-sample), further indicate that brokers with higher PFOF reliance are less sensitive to EFQ changes, while those with larger assets under custody exhibit greater sensitivity.

Table 2: Mean of Significant Parameters Across Size Subgroups

	SP500	**Tercile 1**	**Tercile 2**	**Tercile 3**	**Full Sample**
α_{EFQ}	-11.9547	-8.7700	-6.0324	-3.2723	-6.2574
β_{avg_time}	-0.0223	-0.0342	-0.0161	-0.0134	-0.0195
$\beta_{avg\_reliz\_spr}$	3.3246	-2.4174	-0.5606	-4.5125	-2.2678
$\beta_{at\_quo\_ratio}$	3.4555	2.0935	-0.1656	-0.6902	0.4568
σ_{EFQ}	6.3200	4.4282	4.7680	3.0313	4.1304
$\pi_{EFQ\_PFOF\_rev\_ratio}$	77.7865	79.0772	81.7440	48.9011	71.4266
π_{EFQ_AUC}	-0.7391	-1.8564	-0.6152	-0.1393	-0.8997

Note: This table presents the mean of statistically significant parameter estimates for different subgroups based on firm size. The reported value for each parameter is the cross-sectional average calculated using only the subset of estimates that are significant at the 5% level (t-statistic absolute value > 1.96). The "SP500" group includes all firms in the S&P 500 index. The remaining firms are sorted by market capitalization into three terciles. "Tercile 1" comprises the largest firms (mega, large, and mid-caps). "Tercile 2" contains medium-to-small cap firms, and "Tercile 3" consists of the smallest micro-cap firms. "Full Sample" reports the significant mean across all firms.

Table 3 shows how the estimation results differ across stocks with different liquidity. The stocks are clustered into different liquidity groups based on quoted spread (in percentage), and the quoted spread row corroborates the liquidity grouping, increasing from 0.0214% in Q1 to 0.4385% in Q5. The price elasticity α_{EFQ}, which measures responsiveness to execution quality (EFQ), exhibits a clear gradient, ranging from -9.5915 in the most liquid quintile (Q1) to -2.4826 in the most illiquid (Q5), with a full-sample mean of -6.2574. This pattern indicates that brokers in liquid markets are highly sensitive to EFQ changes, consistent with greater substitution opportunities and competitive pressures, while illiquid markets display inelastic demand, potentially facilitating collusion among wholesalers. Similarly, β_{avg_time} (sensitivity to average execution time) is more negative in liquid quintiles (-0.0397 in Q1) than illiquid ones (-0.010993 in Q5), suggesting brokers penalize delays more severely where alternatives abound. The coefficient $\beta_{avg\_reliz\_spr}$ (realized spread) flips from negative in liquid markets (-1.2958 in Q1) to strongly negative in illiquid (-5.9684 in Q5), implying that spread efficiency becomes a critical differentiator in low-liquidity environments. That is, brokers are more sensitive to realized spread in illiquid stocks. In contrast, $\beta_{at\_quo\_ratio}$ (at-quote ratio) is positive in liquid quintiles (2.3410 in Q1), reflecting a premium on reliability, but turns negative in illiquid ones (-0.7553 in Q5), possibly due to constrained execution options. The variance of random coefficients on EFQ, σ_{EFQ}, decreases with illiquidity (4.8822 in Q1 to 2.9195 in Q5), indicating reduced broker heterogeneity in less liquid markets. Demographic interactions further highlight nuances: $\pi_{EFQ\_PFOF\_rev\_ratio}$ (EFQ sensitivity modulated by PFOF revenue ratio) peaks in moderately liquid quintiles (88.8573 in Q2) but drops sharply in illiquid ones (44.6727 in Q5), suggesting PFOF mitigates EFQ penalties more in competitive settings, while π_{EFQ_AUC} (assets under custody) is most negative in Q1 (-1.9909), implying larger brokers are more EFQ-sensitive in liquid markets.

Table 3: Mean of Significant Parameters by Liquidity Quintiles

	Q1 (Most Liquid)	**Q2**	**Q3**	**Q4**	**Q5** (Most Illiquid)	**Full Sample**
α_{EFQ}	-9.5915	-7.6404	-5.7844	-4.1434	-2.4826	-6.2574
β_{avg_time}	-0.0397	-0.0180	-0.0141	-0.0136	-0.0110	-0.0195
$\beta_{avg\_reliz\_spr}$	-1.2958	-0.4464	-2.7587	-1.0180	-5.9684	-2.2678
$\beta_{at\_quo\_ratio}$	2.3410	1.5529	-0.3780	-0.8059	-0.7553	0.4568
σ_{EFQ}	4.8822	4.8273	4.3669	3.1579	2.9195	4.1304
$\pi_{EFQ\_PFOF\_rev\_ratio}$	67.9705	88.8573	72.4838	74.4243	44.6727	71.4266
π_{EFQ_AUC}	-1.9909	-0.3992	-0.1610	0.0037	-0.9481	-0.8997
Quoted Spread (%)	0.0214	0.0683	0.0672	0.2360	0.4385	0.3257

Note: This table reports the mean of statistically significant parameter estimates, conditional on stocks being sorted into quintiles based on liquidity. All stocks are ranked by their quoted spread measure and grouped into five quintiles, where Q1 represents the most liquid stocks (lowest quoted spread) and Q5 represents the most illiquid stocks (highest quoted spread). The parameter values shown are the cross-sectional averages for each quintile, calculated using only estimates that are significant at the 5% level (t-statistic absolute value > 1.96). The final row, "Quoted Spread (%)", reports the simple average of the quoted spread for all stocks within each quintile to provide context on the liquidity characteristics of the groups.

5 Simulating Competition with MARL

Understanding the competition level of wholesalers is hard empirically. Traditional methods like in Nevo (2001) compare the observed price-cost margins [2] to the price-cost margins calculated in the static Nash-bertrand equilibrium. We find it hard to use the same method in the financial market analysis, as the wholesalers interact with other wholesalers and brokers in so many markets repeatedly. And wholesalers are known as technically advantageous firms while actively using the most advanced algorithms and hardware to adjust prices. So in this paper, we want to construct a dynamic counterfactual by utilizing AI wholesalers that act collusively or competitively in the same market conditions, and compare the actions of real wholesalers to the AI counterparts, and thus to understand wholesalers' market power in a more dynamic perspective. This methodology draws from recent advances in algorithmic collusion literature, where Multi-Agent Reinforcement Learning (MARL) has been used to demonstrate how independent agents can tacitly coordinate on supra-competitive prices without explicit communication (Calvano et al., 2020; Asker et al., 2022; Li et al., 2025).

[2]Price-cost margins is the gap between price and marginal cost as a fraction of price, meaning to what percentage of the price is profit; higher price-cost margins reflect larger market power.

5.1 MARL Systems: Competitive and Collusive AI Wholesalers

We design two distinct MARL systems to simulate the AI wholesaler's behavior when they are designed to learn to act collusively or competitively :

- **Enforcing Competition**: In this setup, each AI wholesaler acts independently to maximize its immediate reward by providing the best possible execution quality. Specifically, the agent reduces its EFQ ($EFQ_{j,k,t}$) to attract the largest expected market share ($w_{j,k,t+1}$). The state space S_t is limited to the agent's own market share from the previous period, $w_{j,k,t-1}$, emphasizing short-term optimization without regard for rivals' actions. This enforces a competitive intention, where wholesalers focus solely on outperforming competitors in the current period.

- **Allowing Collusion**: Here, the AI wholesalers are permitted to incorporate historical information, enabling the potential emergence of tacit collusion through memory-based strategies. The state space for agent j at time t is expanded to:

  $$S_{j,t} = \left(\{w_{j,t-k}\}_{k=1}^{K}, \{a_{j,t-k}, \overline{a}_{-j,t-k}\}_{k=1}^{K}\right),$$

  where $a_{j,t-k}$ denotes agent j's action (EFQ adjustment) in period $t-k$, and $\overline{a}_{-j,t-k}$ is the average action of all other agents in that period. The parameter K represents the memory horizon, allowing agents to learn from past interactions. As demonstrated in Li et al. (2025), this memory mechanism facilitates collusion by enabling agents to signal and punish deviations, leading to supra-competitive EFQs without explicit coordination.

For both systems, we utilize an Actor-Critic algorithm to train the agents. The Critic updates its parameters θ_C to estimate the value function:

$$\theta_{C,t+1} = \theta_{C,t} + \alpha_{C,t}[R_t + \gamma\phi(S_{t+1}, A_{t+1})^T\theta_{C,t} - \phi(S_t, A_t)^T\theta_{C,t}]\phi(S_t, A_t),$$

where R_t is the reward (e.g., profit from market making), γ is the discount factor, ϕ is the feature map, and $\alpha_{C,t}$ is the learning rate. The Actor updates its policy parameters θ_A:

$$\theta_{A,t+1} = \theta_{A,t} + \alpha_{A,t}\phi(S_t, A_t)^T\theta_{C,t}\nabla_{\theta_A}\ln\pi_{\theta_{A,t}}(A_t|S_t),$$

with the action space A_t corresponding to possible EFQ adjustments. The state space S_t varies by system, as defined above.

5.2 Integrating Simulated Brokers and Dynamic Equilibrium

To simulate realistic market dynamics, we introduce a representative "simulated broker" that operates based on the estimated broker assignment function from the demand model (Section 4). This simulated broker assigns order flow shares $w_{j,k,t+1}$ to each wholesaler j for stock k in time t using the estimated elasticity of price and product characteristics α_{EFQ} and $\beta_{avg\_time}, \beta_{avg\_reliz\_spr}, \beta_{at\_quo\_ratio}$.

We first train the RL to converge to learn the policy in the collusive system and the competitive system. Then, with the learned policy, in each RL episode, the system is initialized with historical data on demand for stock j, trading volume (V_{kt}), and liquidity of stock j, bid-ask spread ($A_{kt} - B_{kt}$). The simulated broker allocates shares based on maximizing utility from prior-period EFQs and other execution quality, while the six AI wholesalers respond with updated $\widehat{EFQ}_{j,k,t}^{Collusive}$ or $\widehat{EFQ}_{j,k,t}^{Competitive}$, depending on the environment. This iterative process generates dynamic counterfactual trajectories, enabling us to compute the under-competitiveness metric as the deviation between observed EFQs and simulated counterfactuals.

5.3 Comparing Simulated and Real EFQs: Insights from Dynamic Equilibrium Analysis

The MARL simulations yield counterfactual EFQs under both collusive and competitive regimes, enabling a direct comparison with observed real EFQs and the NBBO benchmark. Figure 2 illustrates these dynamics across six major wholesalers: Citadel (CDRG), G1 Execution (ETMM), Virtu (VIRT), Jane Street (JNST), Two Sigma (SOHO), and UBS (UBSS) over relative months, highlighting deviations that inform the extent of under-competitiveness in the market-making landscape.

In the competitive scenario (labeled as Oligopolistic Competitive Eq.), AI wholesalers aggressively minimize EFQs to maximize immediate market share, implying that retail investors could reduce trading costs when wholesalers competitively improve price improvements. Conversely, the collusive regime (labeled as Oligopolistic Collusive Eq.), exhibits elevated EFQs, as agents incorporate historical actions and market shares into their state space, fostering tacit coordination that sustains supra-competitive pricing. The real EFQ series (labeled as Real EFQ) consistently lies between the competitive counterfactual and the collusive counterfactual, but below the NBBO benchmark all the time.

These patterns help us understand that real wholesalers are persistently better than NBBO (trading on the exchange) by providing better prices in all cases. Though some wholesalers are deviating from most collusive EFQs, they still have a distance to provide the most competitive pricing to retail investors.

The advantage of comparing the real wholesalers' actions to their collusive/competitive counterpacts is that the dynamic counterfactuals generated by AI wholesalers has already incorporated the influence of market structure and demand fluctuations, thus the deviation of real EFQ to the oligopolistic competitive EFQ merely indicates the changes in wholesalers' competition level.

5.4 Under-Competition Level of Wholesalers

We further want to quantify the under-competition level of wholesalers and explore its heterogeneity on different stocks. We develop the Under-Competitiveness Index (UCI) based on counterfactual EFQs generated from the multi-agent reinforcement learning (MARL) simulations and real EFQs. The UCI is defined as:

$$UCI_{j,k,t} = \frac{EFQ_{j,k,t} - \widehat{EFQ}_{j,k,t}^{Competitive}}{\widehat{EFQ}_{j,k,t}^{Collusive} - \widehat{EFQ}_{j,k,t}^{Competitive}},$$

where $EFQ_{j,k,t}$ represents the observed effective quoted spread for wholesaler j on stock k at time t, $\widehat{EFQ}_{j,k,t}^{Competitive}$ is the simulated competitive EFQ, and $\widehat{EFQ}_{j,k,t}^{Collusive}$ is the simulated collusive EFQ. This normalization ensures that $UCI_{j,k,t} = 0$ when real wholesalers' behavior aligns with the competitive counterfactual, and $UCI_{j,k,t} = 1$ when it mirrors the collusive counterfactual. The index effectively captures deviations from a competitive equilibrium, with intermediate values indicating partial under-competitiveness. Across the sample, the UCI exhibits a mean of 0.291 and a standard deviation of 0.561, suggesting that, on average, wholesalers operate approximately 29.1% toward the collusive end of the spectrum, with significant heterogeneity driven by market and firm-specific factors.

To explore this heterogeneity, we regress the UCI on a dummy variable for the top two wholesalers (Citadel and Virtu, D_i^{Top2}) and stock characteristics—log trade volume, log stock price, and volatility—while controlling for stock and year-month fixed effects, as specified in the following model:

$$\begin{aligned} UCI_{j,k,t} =& \alpha_0 + \alpha_1 D_i^{Top2} + \alpha_2 \ln\text{ (Trade Volume)}_{j,t} + \alpha_3 \ln\text{ (Stock Price)}_{j,t} + \\ & \alpha_4 \text{Volatility}_{j,t} + \text{Stock FE} + \text{Year-month FE} + \varepsilon_{j,k,t}. \end{aligned}$$

The results, presented in Table 4, reveal distinct patterns across stock size segments. In

Figure 2: Wholesalers' Execution Quality (EFQ) Compared to AI Counterparts and NBBO

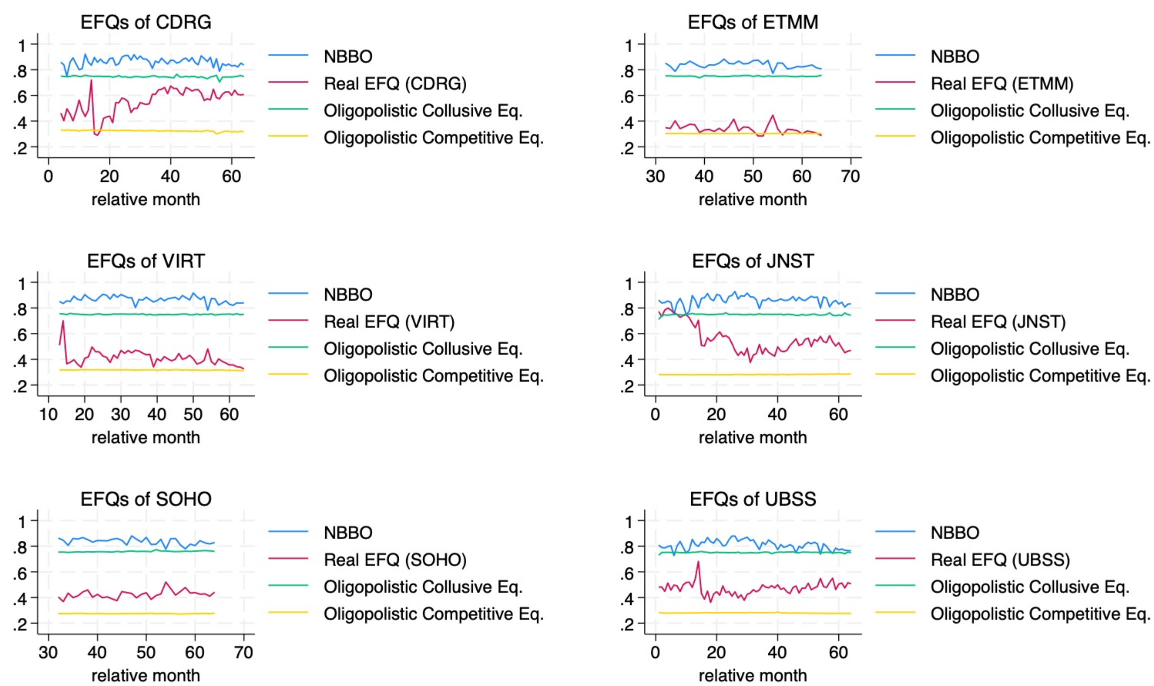

Note: This figure compares the empirically observed Execution Quality (EFQ) for the six largest off-exchange wholesalers: Citadel (CDRG), G1 Execution (ETMM), Virtu (VIRT), Jane Street (JNST), Two Sigma (SOHO), and UBS (UBSS), against dynamic counterfactuals and simulated AI agent outcomes. The x-axis represents the relative month in reality. The y-axis measures EFQ. For each wholesaler, four series are plotted: (1) Real EFQ: The actual, realized monthly EFQ for the named wholesaler. (2) NBBO: The actual benchmark representing the best available price on public exchanges. (3) Oligopolistic Collusive Eq.: The simulated EFQ achieved by AI agents operating under a joint profit-maximizing (collusive) equilibrium. (4) Oligopolistic Competitive Eq.: The simulated EFQ achieved by AI agents operating under a competitive (Nash) equilibrium. The plots show that the wholesalers' real-world performance consistently lies between the collusive and competitive equilibria predicted by the AI counterparts.

the full sample, the coefficient on D_i^{Top2} is negative and highly significant (-0.108 in Panel B, $p < 0.01$), indicating that Citadel and Virtu exhibit greater competitiveness overall, likely due to their scale and technological advantages, this finding align with Dyhrberg et al. (2025). However, this competitiveness varies by stock size. For S&P 500 stocks, the coefficient is -0.348, suggesting intense rivalry among large wholesalers in liquid markets. In contrast, for small non-S&P 500 stocks, the coefficient turns positive (0.050, $p < 0.01$), implying that these top wholesalers behave less competitively on smaller, less liquid stocks, potentially exploiting limited substitution options. Middle and large non-S&P 500 stocks show intermediate effects (-0.083 and -0.230, respectively), showing the trend that competitiveness diminishes as liquidity and market depth decrease.

Additional controls for trade volume, stock price, and volatility further refine these insights. Higher trade volume and lower volatility are associated with reduced UCI (e.g., -0.040 and -0.097 for volume and volatility in the full sample), consistent with competitive pressures in active markets. The positive UCI shift in small stocks aligns with findings in Dyhrberg et al. (2025), which suggest that large wholesalers may leverage market power in fragmented segments, supporting the hypothesis that under-competitiveness is more pronounced where demand elasticity is lower (as observed in the structural estimation results).

Table 4: Under-Competitiveness across Stocks

	(1) Full Sample UCI	(2) S&P500 UCI	(3) Non-S&P500 (Small) UCI	(4) Non-S&P500 (Middle) UCI	(5) Non-S&P500 (Large) UCI
Panel A: without controls					
Dummy(Top 2)	-0.098**	-0.318***	0.115**	-0.090***	-0.239***
	(0.0008)	(0.0020)	(0.0016)	(0.0015)	(0.0014)
Observations	2,217,251	148,041	656,880	691,379	720,951
Adjusted R-squared	0.165	0.279	0.101	0.123	0.172
Panel B: with controls					
Dummy(Top 2)	-0.108***	-0.348***	0.050***	-0.083***	-0.230***
	(0.0007)	(0.0021)	(0.0016)	(0.0013)	(0.0011)
Trade Volume(log)	-0.040***	-0.070***	-0.036***	-0.032***	-0.041***
	(0.0006)	(0.0041)	(0.0009)	(0.0012)	(0.0014)
Price(log)	-0.116***	-0.110***	-0.120***	-0.120***	-0.109***
	(0.0009)	(0.0049)	(0.0018)	(0.0020)	(0.0019)
Volatility	-0.097***	-0.372***	-0.027**	-0.079***	-0.253***
	(0.0059)	(0.0273)	(0.0106)	(0.0104)	(0.0108)
Observations	2,100,400	126,682	606,672	677,441	689,605
Adjusted R-squared	0.168	0.309	0.107	0.127	0.177
Year-Month FE	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES

Note: Robust standard errors in parentheses.* $p < .10$, ** $p < .05$, *** $p < .01$.

6 Validation with Payment For Order Flow

We externally validate the effectiveness of the Under-Competitiveness Index (UCI) by examining its relation to wholesalers' Payment for Order Flow (PFOF). The central idea is that wholesalers sequentially compete along two economically linked dimensions: (i) execution quality delivered to retail investors–especially in the dimension of price improvement (EFQ)–and (ii) monetary transfers to retail brokers via PFOF that reflect profits from market making shared with brokers. If a higher UCI signals weaker competition over execution quality (i.e., wholesalers deliver less price improvement relative to the competitive counterfactual), we should observe that wholesalers also remit less PFOF, holding other factors constant.[3] This logic yields a cross-check on whether the UCI captures meaningful variation in wholesale market competitiveness.

[3] In this context, while we assume both EFQ and PFOF are influenced by competitive pressure, we do not assume that PFOF directly affects EFQ.

Conceptually, except for execution-focused competition (on EFQ), wholesalers also compete for retail order flow by sharing a portion of their expected market-making profits with brokers via PFOF. In a more competitive environment, wholesalers have stronger incentives to bid for order flow either by improving execution quality to end-investors or by paying higher PFOF to brokers. A key testable implication follows: if the UCI correctly captures cross-sectional and time-series variation in competitive pressure, then higher UCI (weaker competition on execution quality) should be associated with lower PFOF paid to brokers, all else equal. A negative co-movement would indicate that wholesalers under weaker competitive pressure economize on both dimensions—price improvement to investors and cash rebates to brokers. We implement this implication using the following panel specification at the wholesaler-month level:

$$\text{PFOF}_{j,t} = \gamma_0 + \gamma_1 \overline{\text{UCI}}_{j,t} + \gamma_2 \,\text{ExecutionTime}_{j,t} + \text{Firm FE} + \text{Year–Month FE} + \varepsilon_{jt}. \quad (11)$$

The dependent variable, $\text{PFOF}_{j,t}$, measures payments for order flow from wholesaler j at time t to its broker counterparties. The key regressor, $\overline{\text{UCI}}_{j,t}$, is wholesaler j's average under-competitiveness index at time t; higher values indicate weaker competition on execution quality (lower price improvement). We include the wholesaler's average execution time, $\text{ExecutionTime}_{j,t}$, to capture an additional dimension of execution quality orthogonal to price improvement. Slower execution may reduce wholesalers' ability to earn profits (e.g., via adverse selection or missed price opportunities), thereby affecting both EFQ and PFOF. Firm fixed effects absorb time-invariant wholesaler characteristics such as technology, clientele, and business model, while year–month fixed effects absorb common shocks to market-making profitability and trading conditions.

Table 5 reports panel regressions of PFOF on $\overline{\text{UCI}}$ and controls. The coefficient on $\overline{\text{UCI}}$ is negative and statistically significant across specifications. This finding supports the interpretation that higher UCI–indicative of weaker competition on execution quality–corresponds to less aggressive bidding for order flow via PFOF. The documented negative relation between $\overline{\text{UCI}}$ and PFOF provides external validation that the UCI captures meaningful variation in wholesalers' competitive posture. When competition over execution quality diminishes (higher $\overline{\text{UCI}}$), wholesalers also pay less to brokers, suggesting a coordinated contraction along both investor-facing (price improvement) and broker-facing (PFOF) margins. This pattern helps to understand that the UCI is not merely an artifact of measurement noise in EFQ but tracks underlying competitive intensity among wholesalers.

Table 5: Competition on PFOF

	(1) PFOF(log)	(2) PFOF(log)
$\overline{\text{UCI}}$	-0.271*	-0.270*
	(0.1623)	(0.1566)
Execusion Time		0.396*
		(0.2164)
Firm FE	YES	YES
Year-Month FE	YES	YES
Observations	244	244
Adjusted R-squared	0.846	0.848

Note: Roust standard errors in parentheses.

* $p < .10$, ** $p < .05$, *** $p < .01$

7 Empirical Results

7.1 Retail Investor Attention and Wholesalers' Competition

We exploit the sudden spike in social media attention in mid-January 2021 as a quasi-experimental shock to retail investor attention. Multiple contemporaneous media sources document extraordinary increases in social media activity—largely driven by the WallStreetBets forum on Reddit—concentrated in a set of stocks that include GameStop (GME), AMC Entertainment Holdings (AMC), Koss (KOSS), Pfizer (PFE), Moody's (MCO), and Disney (DIS). We treat wholesalers' interaction with these affected stocks as exposed to an exogenous surge in retail attention and order submission, and use a difference-in-differences (DiD) design to test the causal impact on wholesalers' competition as captured by the Under-Competitiveness Index (UCI).

Formally, we estimate at the wholesaler–stock–time level:

$$\begin{aligned}\mathrm{UCI}_{i,j,t} = \beta_0 + \beta_1 D_{ij}^{\mathrm{Treated}} \times \mathrm{Post}_t + \beta_2 \ln(\text{Trade Volume})_{jt} + \beta_3 \ln(\text{Stock Price})_{jt} \\ + \beta_4 \text{Volatility}_{jt} + \text{Firm FE} + \text{Stock FE} + \text{Year–Month FE} + \varepsilon_{ijt}, \end{aligned} \tag{12}$$

where $D_{ij}^{\mathrm{Treated}}$ indicates wholesaler–stock pairs in the set of affected stocks, while Post_t indicates the post-shock period beginning in mid-January 2021. The controls absorb contemporaneous changes in trading intensity and risk conditions. Firm fixed effects account for time-invariant wholesaler heterogeneity (e.g., technology, broker relationships), stock fixed effects capture security-specific characteristics (e.g., tick size, baseline liquidity, clientele), and year–month fixed effects purge aggregate shocks to market-making conditions.

Column (1) in Table 6 reports the DiD estimates. The coefficient on $D_{ij}^{\text{Treated}} \times \text{Post}_t$ is positive and statistically significant, indicating that wholesalers are less competitive–i.e., UCI rises–in the affected stocks once social media attention spikes. In coluimn (2), the effects remain stable across specifications with progressively richer controls and fixed effects.

Table 6: Treated Stock Effect on MMs' Competition

	(1) UCI	(2) UCI
Dummy(Treatment)*Post	0.045*	0.109***
	(0.0264)	(0.0315)
Trade Volume(log)		-0.050***
		(0.0037)
Price(log)		-0.121***
		(0.0051)
Volatility		0.013
		(0.0332)
Stock FE	YES	YES
Firm FE	YES	YES
Year-Month FE	YES	YES
Observations	38,840	35,576
Adjusted R-squared	0.254	0.260

Note: Robust standard errors in parentheses.
* $p < .10$, ** $p < .05$, *** $p < .01$

To assess the temporal profile, Figure 3 plots event-time coefficients from a specification that replaces the single DiD interaction with a full set of leads and lags. The estimates show a sharp but temporary increase in UCI that peaks about two months after the shock. This dynamic supports a transient mechanism tied to short-lived order-flow imbalances rather than persistent structural changes in competition. The validity of the DiD design is further supported by the absence of pre-trends in the event-time plot.

Figure 3: Dynamic Impacts of Receiving High Search in Twitter

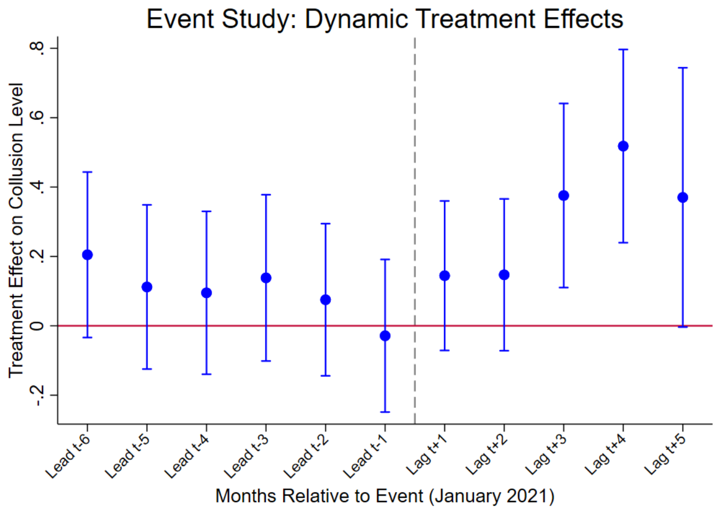

7.2 Mechanism

The documented decline in competition aligns with classic microstructure predictions under extreme, one-sided order-flow pressure. Meme-stock episodes are characterized by sudden, massive inflows of marketable retail orders that seek immediate execution. Theory predicts that when dealers face heightened inventory and adverse-selection risks, they reduce depth (e.g. Stoll, 1978) and widen effective spreads (e.g. Glosten and Milgrom, 1985) to compensate.

We formalize a two-stage mechanism test. First, we verify that the attention shock raises the intensity of marketable orders on the affected stocks. Specifically, we estimate:

$$\begin{aligned} \text{Marketable Orders}_{i,j,t} = \beta_0 + \beta_1 D_{ij}^{\text{Treated}} \times \text{Post}_t + \beta_2 \ln(\text{Trade Volume})_{jt} + \beta_3 \ln(\text{Stock Price})_{jt} \\ + \beta_4 \text{Volatility}_{jt} + \text{Firm FE} + \text{Stock FE} + \text{Year–Month FE} + \varepsilon_{ijt}. \end{aligned} \tag{13}$$

Table 7 shows that the interaction term is positive and significant: retail attention increases the share or level of marketable orders in treated stocks relative to controls. This confirms that the shock translates into immediate-execution demand that dries up displayed liquidity and heightens dealers' inventory exposure.

Table 7: Social Attention's Influence on Marketable Orders

	(1)	(2)
	Marketable Orders (log)	Marketable Orders (log)
Dummy(Treatment)*Post	0.074***	0.114***
	(0.0268)	(0.0343)
Total Orders(log)	0.924***	0.804***
	(0.0031)	(0.0053)
Trade Volume(log)		0.202***
		(0.0057)
Price(log)		-0.032***
		(0.0045)
Volatility		0.380***
		(0.0309)
Stock FE	YES	YES
Firm FE	YES	YES
Year-Month FE	YES	YES
Observations	38,762	35,879
Adjusted R-squared	0.954	0.958

Note: Roust standard errors in parentheses.

* $p < .10$, ** $p < .05$, *** $p < .01$

Second, we examine whether higher marketable-order pressure correlates with weaker competition among wholesalers. We estimate specifications that relate UCI to measures of marketable order intensity, conditioning on the same set of controls and fixed effects as above. Table 8 reports that greater marketability is associated with higher UCI, consistent with wholesalers behaving less competitively when liquidity pressure intensifies.

Table 8: Marketable Orders' Influence on Competition

	(1)	(2)
	UCI	UCI
Marketable Orders(log)	0.061***	0.076***
	(0.0068)	(0.0074)
Total Orders(log)	-0.076***	-0.049***
	(0.0069)	(0.0071)
Trade Volume(log)		-0.077***
		(0.0053)
Price(log)		-0.117***
		(0.0052)
Volatility		-0.045
		(0.0343)
Stock FE	YES	YES
Firm FE	YES	YES
Year-Month FE	YES	YES
Observations	37,785	35,025
Adjusted R-squared	0.255	0.262

Note: Roust standard errors in parentheses.

* $p < .10$, ** $p < .05$, *** $p < .01$

Taken together, these results help us understand the potential mechanism. The social media attention shock induces a surge in marketable retail flow, which raises inventory and adverse-selection risks for wholesalers. In response, wholesalers relax competition on price improvement–captured by an increase in UCI–leading to temporarily worse execution outcomes relative to pre- and post-shock norms. What's more, the dynamic analysis, which combines the MARL and a structural model, brings us innovative insights: the collusion can be motivated by risk mitigation instead of the traditional view of profit maximization; and the transience of the effect in the event study takes us new implication that is hard to obtain from static analysis: collusion could be relaxed from a persistent action pattern but a strategic response to temporal market conditions.

8 Counterfactual Analysis

We quantify the welfare and efficiency implications of wholesalers' competitive conduct by conducting counterfactual exercises on execution costs, using the simulated results from

Section 5.3. The counterfactuals are constructed under three benchmark regimes. First, a full-collusion regime represents the least competitive scenario consistent with wholesalers jointly minimizing price improvement. Second, a full-competition regime represents an upper bound on competitive pressure in which wholesalers aggressively bid for order flow through price improvement. Third, an NBBO regime removes wholesalers from the trading stack and routes retail marketable orders directly to the public markets for execution at the National Best Bid and Offer (NBBO), absent internalization-based price improvement.

Table 9 summarizes the results. Relative to the realized baseline, execution costs would increase by 58.3% in the full-collusion regime, decrease by 15.9% in the full-competition regime, and increase by 765.6% in the NBBO-only regime. These magnitudes imply two substantive conclusions. First, wholesalers as currently operating deliver meaningful price improvement relative to a public-market benchmark without internalization; removing wholesalers and routing directly to NBBO would substantially worsen retail execution outcomes in our sample. Second, there remains a sizable gap between current outcomes and the outcome achievable under full competition: moving from the status quo to the full-competition frontier lowers execution costs by roughly one-sixth.

From a policy perspective, the counterfactual analysis implies that wholesalers' market making can be investor-beneficial on average relative to NBBO execution, but that equilibrium competition among wholesalers is important in determining the magnitude of those benefits. Interventions that strengthen competitive pressure–through improved transparency, enhanced routing incentives, or order-by-order auctions for retail order flow—are likely to move outcomes toward the full-competition counterfactual and reduce execution costs for retail investors.

Table 9: Counter-factual Analysis for the Execution Cost

	(1)	(2)	(3)
	Counter Factual Scenarios		
	Full Collusion	Full Competition	NBBO
Execution Cost	↑58.3%	↓15.9%	↑765.6%

9 Conclusion

This paper provides novel evidence on under-competitiveness in the U.S. equity market making sector. Our methods firstly construct a structural model with theoretical foundations to estimate the demand system between brokers and wholesalers, and then use reinforcement learning to simulate the dynamic equilibrium counterfactual with AI wholesalers make decisions under same market conditions, and finally, we use causal inference in a DiD setting to examine the impact of sudden retail investors' attention, following increased demand, make the wholesalers act less competitive to encounter inventory risk and adverse selection risk. The primary findings include that, the mean EFQ elasticities is -6.26, indicating elastic demand on EFQ, and this elasticity is more pronounced in S&P 500 tickers (-12) than non-S&P 500 (-5.7), meaning that brokers are more sensitivity to execution quality change on larger stocks; and based on the index of under-competitiveness we construct, the wholesalers are providing better execution than NBBO but still worse than AI wholesalers' competitive counterfactual. This under-competitiveness index is also validated against PFOF data. By conducting the counterfactual analysis, we obtain the welfare implications that, in general, wholesalers reduce retail investors' trading costs by 80% relative to NBBO, but still worse than the fully competitive counterfactual by 16%.

This paper brings a new methodology into view: using structural estimation to get the economic conditions, and embedding the conditions in Reinforcement Learning to obtain dynamic implications. This method has its potential to be used in multiple scenarios, for example, to extend to analyze the market power of financial intermediaries in the lending market. This hybrid method allows for the simulation of credible counterfactuals—such as perfect competition—against which real-world outcomes can be measured. The potential applications of this methodology are broad, extending beyond market microstructure. For instance, it could be used to analyze the market power of financial intermediaries in corporate and consumer lending markets, to model competition in insurance underwriting, or to understand the dynamic interactions between different trading venues.

References

John Asker, Chaim Fershtman, Tel Aviv, and Ariel Pakes. Artificial Intelligence and Pricing: The Impact of Algorithm Design. *AEA Papers and Proceedings*, 2022.

Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu. Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy*, 132(3):723–771, 2024.

Markus Baldauf, Joshua Mollner, and Yueshen Bart Zhou. Siphoned apart: A portfolio perspective on order flow segmentation. *Journal of Financial Economics*, 2024.

R. Battalio and R. Jennings. Wholesaler execution quality. 2023.

Steven T Berry, James A Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 1995.

Thomas Brown, Travis Johnson, S.P. Kothari, and Etic C. So. Anatomy of trading costs for retail investors: Savings from off-exchange execution. 2024.

Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110:3267–3297, 2020.

Thomas E Copeland and Dan Galai. Information effects on the bid-ask spread. *the Journal of Finance*, 38(5):1457–1469, 1983.

Harold Demsetz. The cost of transacting. *The quarterly journal of economics*, 82(1):33–53, 1968.

Winston Dou, Itay Goldstein, and Yan Ji. AI-Powered Trading, Algorithmic Collusion, and Price Efficiency. Working paper, 2024.

Anne Haubo Dyhrberg, Andriy Shkilko, and Ingrid M Werner. The retail execution quality landscape. *Journal of Financial Economics*, 168:103863, 2025.

David Easley and Maureen O'hara. Price, trade size, and information in securities markets. *Journal of Financial economics*, 19(1):69–90, 1987.

David Easley, Nicholas Kiefer, and Maureen O'Hara. Cream-skimming or profit-sharing? the curious role of purchased order flow. *Journal of Finance*, 1996.

Hu Edwin and Murphy Dermot. Competition for retail order flow and market quality. 2022.

Thomas Ernst, Andrey Malenko, Chester Spatt, and Jian Sun. What does best execution look like? 2024.

Financial Industry Regulatory Authority (FINRA). Rule 5310. best execution and interpositioning. https://www.finra.org/rules-guidance/rulebooks/finra-rules/5310, 2024. Accessed: [Insert Date You Accessed It].

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.

Thomas Ho and Hans R Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial economics*, 9(1):47–73, 1981.

Xing Huang, Philippe Jorion, Jeongmin Lee, and Christopher Schwarz. Who is minding the store? order routing and competition in retail trade execution. 2024.

S. Kothari, E. So, and T. Johnson. Commission savings and execution quality for retail trades. 2021.

Jin Li, Ye Luo, Yannan Sun, and Xiaowei Zhang. Learning to compete or learning to collude? emergent behavior of ai market makers. *Working Paper*, 2025.

Aviv Nevo. A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of economics & management strategy*, 9(4):513–548, 2000.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69 (2):307–342, 2001.

SEC. Order competition rule, release no. 34–96495. 2023.

Hans R Stoll. The supply of dealer services in securities markets. *The Journal of Finance*, 33 (4):1133–1151, 1978.

Seha M Tinic. The economics of liquidity services. *The Quarterly Journal of Economics*, 86 (1):79–93, 1972.

Vincent Van Kervel and Bart Zhou Yueshen. Anticompetitive price referencing. *Available at SSRN 4545730*, 2025.