Title

Essays on Influence of Information and Innovation in Digital Markets

Author

Li, Yilin

Publication Date

2024

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

ESSAYS ON INFLUENCE OF INFORMATION AND INNOVATION IN DIGITAL MARKETS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

Yilin Li

March 2024

The Dissertation of Yilin Li
is approved:

Professor Daniel Friedman, Co-Chair

Professor Kristian López Vargas, Co-Chair

Professor Natalia Lazzati

Peter Biehl
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Essays on Influence of Information and Innovation in Digital Markets

by

Yilin Li

This dissertation presents three experimental studies with an emphasis on the influence of information and innovation on digital markets including financial exchanges and online marketplaces. The first chapter focuses on the experimental evaluation of a new financial market design. The second and third chapters focus on how access to information affects people's behavior in the online marketplaces.

The first chapter provides a laboratory study of a newly proposed Flow Market format as a response to the design weaknesses of the continuous double auction used in most financial markets worldwide. We designed and deployed a laboratory experiment that compares the Flow Market and the CDA using several fundamental metrics. We find evidence that the flow market changes traders' behavior relative to CDA, allowing them to shred orders more effectively: compared to the CDA, the Flow market exhibits fewer and larger orders. We also find both formats perform similarly in terms of price and allocative efficiency. However, the Flow Market leads to lower price volatility. Interestingly, the total traded volume is lower under the Flow Market than under CDA. Still, this difference decreases with traders' experience, i.e., as they learn the mechanics

of the Flow format. Our findings provide initial insights regarding the feasibility of the Flow trade format and its potential to promote financial market stability and fairness.

The second chapter modifies the traditional sequential search models to consider the ex post uncertainty in which the uncertainty cannot be fully eliminated by the search. We derived players' optimal search strategy given their risk attitudes. We also test our theory in a laboratory experiment with a search game to track subjects' behavior and use a multiple price list and a bomb risk elicitation task to elicit subjects' risk preferences. We find that, in this scenario, risk-averse players tend to increase their reservation value and extend their search duration.

The last chapter investigates the informational barrier problem in the online marketplaces. We design a sequential game to study how informational barrier is formed due to the reputation system and propose a fractional searching mechanism to enable the entry by reliable firms. The model predicts the optimal behaviors of the firms and shows that the fractional searching enables cost-effective firms with no reputation to enter the market more easily. Additionally, we test our theory in a laboratory experiment using an interactive market game to track subjects' behaviors with a simplified reputation system. The experimental results indicate that fractional searching effectively alleviates the entry problem of new firms with superior quality.

To my families, friends, and advisers for their support and guidance,

through all times.

Acknowledgments

My journey at UCSC is full of unforgettable memories that will be cherished forever. I would like to express my sincere gratitude to those who have supported and encouraged me in the pursuit of my doctoral degree.

First of all, I want to thank my advisers, Professor Daniel Friedman and Professor Kristian López Vargas, without whose invaluable advice, continuous encouragement, and excellent guidance, the journey would have been much more difficult. Their expert knowledge in experimental design built my foundation in experimental research. They genuinely connected me to researchers working on related topics and shared their network in the industry. They are extremely supportive of my career choice and helpful in developing my new career after graduation.

I would also like to extend my thanks to professor Natalia Lazzati, an excellent member of my dissertation committee, for her constructive and timely feedback on my research work. Professor Lazzati has continuously offered her expertise in economic theory to give valuable input and guidance throughout my doctoral studies. I am especially grateful to professor Lazzati for encouraging and supporting me to explore other interesting fields from external campus.

Moreover, I would like to thank professor Todd Feldman for being a member on my qualification committee, offering valuable feedback on my early research work. I would also give a special thanks to my co-author, Shuchen Zhao, for refining our research topic and introducing me to conduct experiments in the lab.

Furthermore, I am grateful to the Economics Department at UCSC, and especially our department's Learning and Experimental Economics Projects Lab (LEEPS), for funding my experiments and research. I would not have completed my experiments on time without those funds and facilities. I would also like to thank professors Dong Wei, Michael Leung, Julian Martinez-Iriarte, and Gerelt Tserenjigmid for their insightful advice offered in our experimental workshops.

Additionally, I am incredibly thankful to my friends during my time at UCSC, including my cohorts, members of experimental workshop, and friends in the bay area. I will never forget the time we worked together on the preliminary exams, discussed interesting research ideas, and encouraged each other during the qualification exam and job market. I am blessed to become life-long friends with you all.

Last but not least, I would like to thank my parents for their unconditional financial and emotional support throughout the last six years. They always believe in me and support my decisions to the greatest extent. I am especially grateful to my husband for his understanding and encouragement through the good and the bad. He spent countless hours programming my complex experiments and has been the most reliable technical support for my first chapter. I would not be able to complete this journey without you all.

Chapter 1

Testing Flow Trading Format in the Laboratory [1]

1.1 Introduction

We present an experimental study of the Flow Market format, proposed by A. S. Kyle and Lee (2017) and extended by Budish, Cramton, A. Kyle, et al. (2023). Although the proposed flow format represents a significant departure from today's most widely used rules in stock exchanges, its study is relevant as economic theory predicts this new format may also correct several design flaws of the currently prevalent market rules.

Following or inspired by centuries-old rules of commodities and stock markets,

[1]The first chapter is a joint work with Daniel Friedman and Kristian López Vargas.

most modern stock exchanges operate under adaptations of the double auction. The prevailing adaptation, called continuous double auction (CDA) or the continuous limit order book (CLOB), is essentially the electronic implementation and streamlined version of the old centralized markets run manually. In the CDA, orders from traders are posted in a dynamic public order book and matched based on price-then-time priority. Orders are either market orders, executed immediately at the best available price, or limit orders, set at a specific price and waiting to be filled.

Although CDA/CLOB is called "continuous" because trade can occur at any instant, this market mechanism is not truly continuous. Prices, quantities, and time "continuity" are all limited. Prices can only move in fixed steps called tick sizes (say, $0.01), as do order sizes. Furthermore, orders are processed serially, one at a time, based on when they arrive. The discreteness in price creates demand and supply schedules that are step functions, making it typical to have unbalanced demanded and supplied quantities at an instantaneous clearing price. When not all the demanded (supplied) quantity can be instantly fulfilled, exchanges prioritize orders by when they arrive.

It follows that modern communication technology is remarkably consequential for the CDA. In particular, (1) High-frequency traders (HFTs) have strong incentives to exploit the CDA mechanism by using their speed to be first in line for orders when competing with slower traders at the same price. Since prices move in steps, it is hard to outcompete a trader with a speed advantage of even a fraction of a millisecond. HFTs can also react faster to new information and, before slower traders can respond, they can either cancel their own orders or pick off slower orders resting in the book. (2) To avoid price impact (obtaining a worse-than-necessary price for large orders), firms with better technology can send a larger number of smaller orders than firms with worse technology. Therefore, firms with worse technology also get worse prices on average, as they tend to have a larger price impact.

In sum, modern markets under CDA have built-in limitations that tilt the playing field in favor of those with cutting-edge technological capabilities, particularly high-frequency trading (HFT) firms. This emphasis on speed, especially in volatile market conditions where liquidity is crucial, raises concerns about market fairness and overall stability. The billions of investments in communication technology and infrastructure for trading, primarily driven by the race to exploit the time priority mechanism, are mostly unproductive from a broader economic perspective. HFT's strategies, heavily reliant on speed, might also exacerbate market volatility, making flash crashes more likely. This scenario underscores the need for a reassessment of market design.

There have been several proposals to correct the design limitations of the currently prevalent market institutions. Random priority (add reference), frequent batch auctions Budish, Cramton, and Shim (2015), delayed messaging with order pegging (Investor's Exchange, Aldrich and Friedman (2022)), micro-burst fees (Brolley and Zoican (2023)), priority rules (Degryse and Karagiannis (2022)). All of these proposals, however, target the first issue (HFTs exploiting time priority). One proposed format aims to correct both issues: the flow format proposed by A. S. Kyle and Lee (2017) and extended by Budish, Cramton, A. Kyle, et al. (2023).

Flow trading is based on a new type of order, the Continuous Scaled Limit Order (CSLO), which allows trading across a continuous price range. These orders are defined by five parameters: direction (buy or sell), maximum quantity ($Q_{\max}$), a price range between lower (P_L) and upper (P_H) limits, and a maximum trading speed ($U_{\max}$), in shares per unit of time.

CSLOs enable traders to express the desired immediacy in a richer fashion than standard limit orders in CDA, since CSLOs adapt their trading speed, $U(p(t))$, dynamically based on the market price, $p(t)$. Full trading speed is applied when the market price is optimal for the trader –below P_L for bids or above P_H for asks. Within the price range, the trading speed for buy orders decreases, and for sell orders increases linearly with the price, indicating a trader's varying willingness to execute orders depending on how favorable the price is. The agent's trading halts when prices move outside the trader's acceptable range, withdrawing from the market.

First, HFTs cannot exploit the mechanism significantly because their expected gains from sniping stale are orders of magnitude smaller under Flow trading than under CDA – they can pick off only tiny fractions of orders from slow traders. Second, the flow exchange allows investors to move order shredding into the exchange, leveling the field across many types of traders and improving the market outcomes related to the second issue.

Despite the relevance of the flow trading design proposal, there is no empirical study of its functioning and performance to our knowledge. Although studying this format in-depth and measuring to what degree it solves the above-mentioned issues will require extensive, more complex experiments, we postulate that a relatively simple lab experiment is the appropriate starting point. This is precisely the contribution of our paper; we study trader behavior and market performance within a relatively simple controlled laboratory setting, focusing only on two market formats: the continuous double auction (CDA) and the flow trading format (FLOW). In our setting, a single asset is traded on a single exchange, allowing traders to own or owe multiple units of this asset.

We adopt the induced value paradigm and adapt it slightly by providing contracts to traders at the start of each trading period. These contracts are of two types: buy contracts, committing the computer to buy a specified number of shares from the trader at a set price and time, and sell contracts, committing the computer to sell shares to the trader under similar conditions. This mechanism aims to generate trading incentives by creating situations where traders can profit by executing trades at prices more favorable than those specified in their contracts.

The experimental design utilized a between-group approach, with 80 participants divided into 40 per market format across five groups of eight traders. Each market session consisted of 20 trading periods, lasting two minutes each, with eight traders per market. The participants were evenly split between receiving buy and sell contracts. The 20 trading periods were split into five blocks, each containing four periods with constant contract settings for repeated underlying demand and supply conditions.

We find significant changes in traders' behavior emerging from the possibility of managing order sizes and frequencies more effectively under the Flow format. Specifically, the Flow Market format led to a significant reduction in the number of orders placed, with 21-24 fewer orders in each period in the Flow format compared to CDA (which, on average, displayed 47 orders). Moreover, the order size in the Flow format was larger by an average of 13 shares per order than in CDA (which, on average, receives 21 orders per period). Interestingly, despite these improvements in order management, the total traded volume in the Flow Market was notably lower than in CDA. However, this gap narrowed as trading periods progressed, suggesting a substantial learning pattern in the Flow Market. Notably, price volatility was also significantly lower in the Flow Market, with a difference of -0.7 in the absolute price change and about 7% lower standard deviation of the log price change compared to CDA. Our findings provide relevant insights into the viability of the Flow format as a design that promotes a more stable market.

Our paper is related and contributes to the literature on design alternatives to the

CDA/CLOB format in the presence of heterogeneous access to technology and high-frequency trading firms. In particular, our research contributes to understanding the basic properties and empirical feasibility of the Flow Market format, conceptualized by A. S. Kyle and Lee (2017) and extended by Budish, Cramton, A. Kyle, et al. (2023).

Along with the Flow format, there exist other market design proposals aimed at correcting the weaknesses of the CDA-based prevalent current market designs. These proposals include frequent batch auctions, delayed messaging and order pegging, burst fees, and changing the priority rules.

The frequent batch auctions (FBA, Budish, Cramton, and Shim (2015)) consist of aggregating orders over fixed, short intervals and conducting call auctions at the end of each batching period. By batching orders, the speed advantage of HFTs' is lost, and in this manner, this format mitigates the exploitative behavior of HFTs. Aldrich and López Vargas (2019) compare CDA and FBA formats using laboratory experiments. They examine the effects of switching CDA for FBA on predatory trading behaviors, technology investment, transaction costs, and volatility. The findings reveal that FBA reduces predatory behaviors, diminishes the investment in socially-inefficient speed-enhancing technologies, and lowers transaction costs and volatility compared to CDA. Haas, Khapko, and Zoican (2021) also study this matter theoretically and find that relative to continuous-time trading, periodic batch auctions reduce HFT informational rents.

Aldrich and Friedman (2022) presented an alternative approach through delayed messaging with order pegging, inspired explicitly by the Investor's Exchange (IEX) context. They explore theoretically and empirically the potential of using intentional messaging delays in financial markets to counteract the advantages of high-frequency traders (HFTs) over traditional traders. Their model and empirical analysis show that such delays can reduce the predatory behaviors of HFTs, safeguarding other traders' orders from being exploited. They also document that message delay introduces additional transactions and queuing costs.

Brolley and Zoican (2023) propose micro-burst fees that target the economic incentives fueling HFT. By imposing a fee on liquidity-taking orders during periods of high message traffic, the proposed format aims to deter latency arbitrage, improve market liquidity, and increase exchange revenues beyond what traditional colocation fees offer. Their theoretical model shows that such fees can reduce the incentive for HFTs to engage in predatory practices.

Degryse and Karagiannis (2022) explored the effects of priority rules. In particular, this paper explores the implications of different order execution rules, such as price-time (PT) priority versus price-broker-time (PBT) priority, which adds broker identity to the prioritization. Employing both theoretical and empirical analysis, the authors model investor behavior under these priority regimes. They find that the optimal design, between PT or PBT, depends on the relative size of the price tick with respect to the dispersion of investors' valuations.

In this literature, the current main contributions remain theoretical and the empirical assessment of the proposed designs is challenging due to the lack of observational data of non-existing market design. We argue that carefully designed lab experiments are helpful for test-bedding these designs in similar fashion we implement the study of the Flow format.

The rest of the paper is organized as follows: Section 1.2 presents the experiment and its implementation. Section 3.5 presents the results, and Section 3.6 presents the conclusions and discussion for future research.

1.2 Experiment

1.2.1 Environment

In our experiments, a single asset is traded on a single exchange, and traders can own or owe multiple units of this asset. All prices and monetary magnitudes are expressed in Experimental Currency Units ECUs.

We implement an induced value environment with some innovations with respect to the standard experiments. At the beginning of each trading period, we provide traders with *contracts*. There are two types of contracts: 1) *buy contracts*, and 2) *sell contracts*. A buy contract is a commitment from the computer to buy from the trader a specified number of shares at a certain price and at some point later in the trading period. For example, a buy contract might read "Buy 300 shares at a price of 14 ECUs in 80 seconds". This means that 80 seconds from now, the exchange will buy up to 300 shares from the traders inventory at a price of 14 ECUs per share. If the trader holds less than 300 units, the computer will purchase the whole inventory. Similarly, a sell contract that says "Sell 400 shares at a price of 13 ECUs in 90 seconds" means that, in 90 seconds, the exchange will sell to the trader up to 400 shares at the price of 14 ECUs per share. If the trader is short (owes) only 350 shares, then the computer will sell only 350 units to this trader.

Contracts generate induced-value incentives to trade in our environment similar to classic double auction experiments. For traders with a buy contract, it would be worthwhile to buy shares at a lower price than the contract price from other traders. Similarly, for traders with a sell contract, it would be profitable to sell shares at a higher price than the contract price. Buyers will earn the difference between the contract price and the average buying price per share for their inventory while sellers will earn the difference between their average selling price and the contract price per share instead. For simplicity, in this experiment, we provided each trader with a single contract (buy or sell) at the beginning of each trading period and with execution time right before the end of the period.

Traders can also profit from trading outside the contract. For example, a trader with a buy contract from the computer at 10 dollars per share, could sell part of her inventory to another trader at prices that are more convenient relative to the contract (i.e., if the market price is above 10 ECUs).

Two main differences with the standard induced-value designs of standard CDA experiments are worth noting. First, among other things, our study aims to measure how order fragmentation (shredding) differs between the two formats. This requires a design where traders are induced to value a much larger set of units than in the usual CDA markets where inventories consist of a few items. Second, unlike standard market experiments, we originally designed (value-inducing) contracts to have an issue and expiration times to accommodate the possibility of manipulating the urgency of trading during the same period by giving multiple consecutive contracts to the same trader within the period. For the sake of simplicity we ended up giving a single contract per trader-period in this paper.

Within the framework of these incentives to trade using contracts, our paper studies two different market formats, the continuous double auction (CDA), and the flow trading format (FLOW), and, in particular, the differences in traders' behavior and market performance between these two formats.

1.2.2 Market Formats

Continuous Double Auction (CDA)

The Continuous Double Auction mechanism, also known as the *consolidated limit order book*, is structured around the concept of standard limit orders, which are defined by three parameters: a buy-sell indicator, a specified quantity ($Q_{\max}$), and a limit price (P). A limit order expresses a trader's commitment to buy or sell up to the specified quantity of shares at the specified price or better.

At the core of an exchange operating under CDA is the *order book*, which is a dynamic repository that prioritizes all pending limit orders first by their price level and then by the time of order. That is, Buy orders or bids are arranged from the highest to the lowest price; sell orders, or asks, are ordered from the lowest to the highest price. When orders are at the same price level, the book prioritizes among them according to the time of order submission, with earlier orders receiving precedence over later ones. The difference between the highest bid and the lowest ask is called the *spread*.

When a new limit order is introduced with a price that matches or exceeds the best available opposite-side price (the best contra-side price), it triggers immediate execution at the contra-side price, and the executed quantity is determined by the lesser of the quantities specified in the matching buy and sell orders. In this case we say that the arriving order has crossed the market. This immediate transaction removes the executed quantity from the order book, reflecting a completed trade. If instead the new order does not cross the market, it is stored in the order book following the price-time priority described above.

CDA User Interface

In the experiment, we implement the CDA format with the interface shown in Panel (a) of Figure 1.1. We use a dynamic interactive graphical device to represent the limit order book (top-middle box). In the same device, the trader sets and submits their limit orders. The horizontal axis denotes the quantity of shares, and the vertical axis denotes prices in ECUs per share. To set a *buy* order, the trader drags the blue dot inside the grid to set the desired price and the quantity. Then, the trader submits this buy order by clicking on the “Send Buy” blue button. The blue-line segments graphically describe her buy order. Similarly, the trader sets a *sell* order by dragging the red square inside the grid to the desired price and the quantity and submits it by clicking the “Send Sell” red button. The active orders will appear in the same graph, and the Active Orders table on the left side of the screen. To submit a new buy (sell) order, the trader must cancel her active buy (sell) orders and then create and submit a new one.

In the top-right Market box, the trader can observe the current state of the market. The blue curve is the market demand, and the red is the market supply. The intersection of them determines the instantaneous trading price and the traded quantity. Trades happen immediately when orders come into the market, crossing the spread. The green dots indicate the most recent trades. More prominent dots represent more recent trades. In nearly all moments of the trading period, the demand and supply cross each other along the vertical axis ($q = 0$) representing the CDA's most common state where the best bid is below the best ask. Only when an order arrives at the best price on the contra side or beyond does the supply and demand cross for that instant to represent the trade.

The trader will see how her inventory and cash change within each trading period in the bottom right section. In the bottom middle section, the trader will see her projected profit, which indicates her profits if the active contract and the period were to expire at that very moment – i.e., it provides a summary of her current state of potential profits to the trader. The top left table shows an Active Contracts table, which describes the contract the trader received from the computer right after the start of the trading period. The table below describes the trader's active (unexpired) contract. Fully executed orders will appear in the third row of the same column. Finally, expired contracts will move to the Executed Contracts section at the bottom of that left column of tables.

Flow Trading (FLOW)

We now describe the Flow trading format proposed by **Kyle2017**. The flow trading departs from the standard limit order and creates the concept of *Continuous Scaled Limit Orders* (CSLO). This type of order is necessary for the interaction between demand and supply over a continuous range of prices and for the flow trading of shares.

CSLOs are characterized by a set of five parameters: a directional indicator specifying buy or sell orders, a maximum quantity ($Q_{\max}$) to be traded, a price range defined by lower (P_L) and upper (P_H) bounds with $P_L < P_H$, and a maximum trading speed ($U_{\max}$), measured in shares per hour. This framework allows traders to express their willingness to transact up to a specified quantity within a defined price range at a rate per unit of time that does not exceed $U_{\max}$.

The main feature of CSLOs lie in their adaptive trading speed or flow demand/supply, $U(p(t))$, which is linked to the market clearing price, $p(t)$, at every point in time during the trading period.

The trading speed adjusts as follows. Full speed ($U_{\max}$) is maintained when the market price is at or below the lower price threshold (P_L) for buy orders, or above the upper price threshold (P_H) for sell orders, signifying a trader's eagerness to execute orders outside their specified price range. Within the price range ($P_L \leq p(t) \leq P_H$), the trading speed decreases linearly with price for buy orders and increases linearly for sell orders. The linear interpolation ensures that the execution rate diminishes as prices approach the less favorable end of the specified range, reflecting a trader's diminishing willingness to trade as prices move away from the trader's best conditions. Trading ceases ($U(p(t)) = 0$) when prices fall outside the favorable range—i.e., above P_H for buy orders or below P_L for sell orders. This indicates a withdrawal from the market due to the clearing price exceeding the acceptable price range of the trader.

These conditions can be summarized in the following equations for a trader's flow demand and supply, respectively, at the clearing price at time t, $p(t)$. We removed the trader index for simplicity.

$$U_d(p(t)) = \begin{cases} U_{\max} & \text{if } p(t) < P_L \\ \frac{P_H - p(t)}{P_H - P_L} U_{\max} & \text{if } P_L \leq p(t) \leq P_H \\ 0 & \text{if } p(t) > P_H \end{cases} \tag{1.1}$$

$$U_s(p(t)) = \begin{cases} U_{\max} & \text{if } p(t) > P_H \\ \frac{p(t) - P_L}{P_H - P_L} U_{\max} & \text{if } P_L \leq p(t) \leq P_H \\ 0 & \text{if } p(t) < P_L \end{cases} \tag{1.2}$$

The market flow trading emerges from the intersection of aggregate individual flow demands and supplies. The resulting market flow demand and supply are piece-wise linear functions of the price at a certain point in time. This intersection determines the market's equilibrium (clearing) price and the overall transaction rate, balancing the willingness to trade of all agents. Individual trading speeds are, in turn, dictated by each trader's specific flow demand or supply curve in relation to the clearing price. This mechanism can accommodate a large set of trading strategies and preferences, enabling agents to tailor their market engagement to reflect not only their target prices but also their urgency of trading, making their trading adaptable to fluctuations in market conditions.

FLOW User Interface

The software and interface for the Flow market is analogous to that of the CDA market. Figure 1.1 panel (b) shows one trader's FLOW trading screen. The trader sets and submits buy and sell orders in the top-middle box (titled "Your Input"). The horizontal axis denotes *rate* in shares per second, and the vertical axis denotes prices in ECUs per share. The trader sets a *buy* order by moving the blue dot along the vertical axis to set the high price and dragging the blue square anywhere inside the grid to set the low price and the maximum rate. Then, she enters the total number of shares she wants to buy in the box next to the "Send Buy" blue button and submits it by clicking that button. The blue-line segments graphically describe her buy order, as shown in Figure 1.1 panel (b). Likewise, she sets a *sell* order by moving the red dot along the vertical axis to set the low price and dragging the red square anywhere inside the grid to set the high price and the maximum rate. Then, she enters the total number of shares she wants to sell in the box next to the "Send Sell" red button and submits it by clicking that button. She can submit a new order by canceling her active orders on the same side of the market and then creating a new order at any time.

The trader sees the market supply and demand in the top-right Market box. The blue curve is the market demand, and the red is the market supply. The intersection of them determines the clearing price and the aggregate clearing rate. The green horizontal line in the Market diagram extends into the *Your Input* box and indicates the current market price. The x-coordinate of the intersection of this line and the trader's order indicates the trader's clearing rate. The column of tables on the left side of the interface and the bottom boxes (projected profits, inventory, and cash) are analogous to those in the CDA format UI.

1.2.3 Implementation

We implemented a between-group design where each session consisted of 20 trading periods, each lasting two minutes. Each market consisted of eight traders. We have 80 participants, 40 trading in each market format and, therefore, five groups per format.

In each market of eight traders, four received a single buy contract and four a sell contract. As mentioned above, buy and sell contracts are analogous to induced values and costs in standard market experiments. In particular, a buy (sell) contract is a commitment from the computer to buy (sell) up to $Q_{contract}$ units at price $p_{contract}$ at the preset expiration time of the contract (all expiration times were set near the end of the trading period).

We divided the 20 trading periods into five blocks, each block with four periods. In each block, all four periods have the same contract configuration, which means that the underlying period-level demand and supply curves repeat, and so do the quantity and

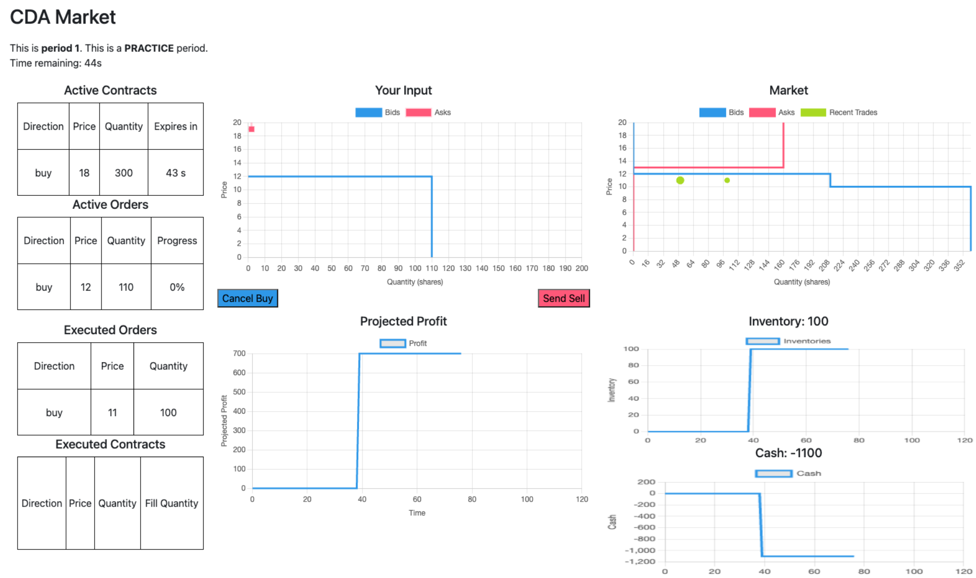

(i) CDA Format

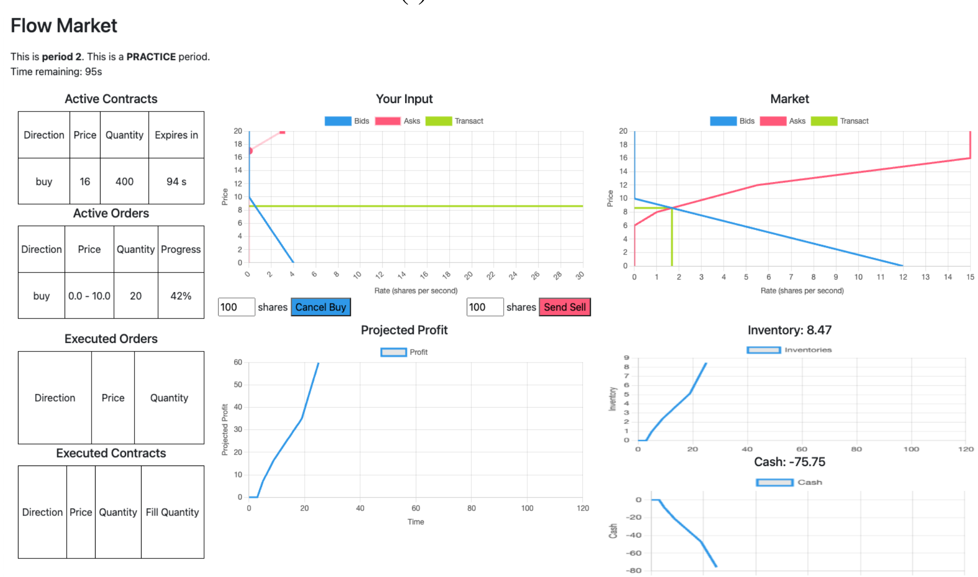

(ii) FLOW Format

Figure 1.1: Experiment Interfaces

Figure 1.2: Contract Design

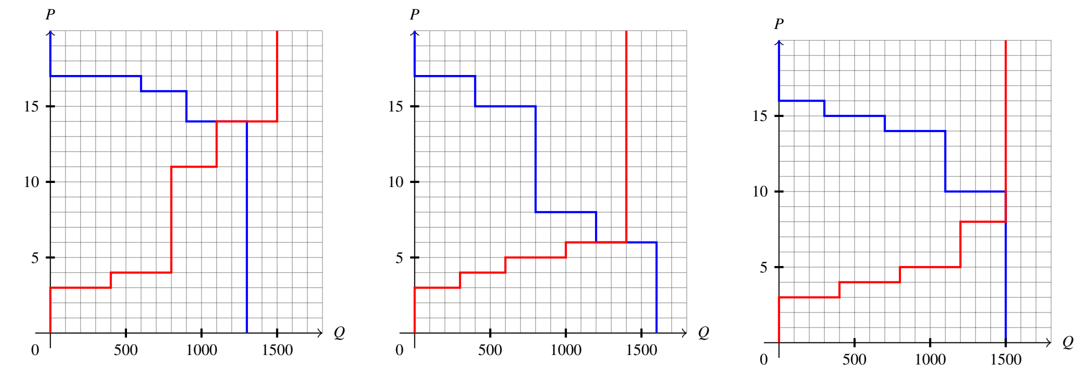

(i) Block 1 (T1-T4) and Block 5 (T17-T20)

(ii) Block 2 (T5-T8) and Block 4 (T13-T16)

(iii) Block 3 (T9 - T12)

price of equilibrium. After the four-period block, the configuration changes. We use three different configurations of contracts in the five blocks. The first three blocks have different configurations: 1, 2, and 3. Block four repeats the configuration 2, and the fifth block repeats the configuration 1.

Figure 1.2 depicts the three contract configurations used in our experiment as period-level underlying demand and supply. The blue staircase represents the buy contracts, sorted to form the underlying demand curve, and the red underlying supply curve similarly represents the sell contracts. Each contract assigned to an individual trader is a step (a flat line segment) on either demand or supply curve. Within each block, each trader receives the same contract each period. Across blocks, buyers and sellers keep the same role, but we assign a new set of contacts as shown in Figure 1.2.

Experiment instructions were provided on the computer screen and paper.[2] After

[2]Instructions were discussed and piloted (with students and colleagues) to balance clarity and reasonable length.

15 minutes of reading instructions, a video summary was played with animations and annotations to clarify any confusion that subjects might have.[3] After the instructions, subjects participated in two 120-second trial periods where subjects became more familiar with the experiment interface and market rules. Subjects were then given the option to ask any questions.

In between trading periods, each subject received a summary screen displaying the profit history of each trading period. Each participant started each trading period with a zero endowment. Final payments were based on the final accumulated wealth of all twenty trading periods, using an exchange rate of 1000 ECUs = $1, plus a $7 participant fee. This information was provided to subjects in the written instructions.

Although ending a trading period with negative profits is possible (i.e., having negative ECUs in the period), this happened in less than 5% of cases and subjects never left the laboratory with less than $7. Payments were implemented following standard confidentiality procedures.

The experiment software was developed on oTree (**otree**). Sessions were conducted at the LEEPS Laboratory of the University of California, Santa Cruz. Recruitment was implemented through LEEPS' ORSEE instance (econlab.ucsc.edu; **subject**).

[3]Participants were not allowed to communicate with each other before, during, or after the experiment, nor did they learn the identity or characteristics of other participants.

1.3 Results

1.3.1 Descriptive statistics

Table 1.1 reports summary statistics for the experiment. Given that the transaction sizes vary quite largely in CDA and to better quantify the price impact in both formats, we weigh prices by the associated transaction quantity within each five-second interval across all periods and compute the average prices, price deviations, and price volatility. Panel (a) of the Table 1.1 shows measures of price and volatility. The next two lines of Table 1.1 show that on average price changes are roughly one and a half times larger in CDA as in FLOW and that price volatility is about twice as large.

Panel (b) shows measures of order sizes and transactions. Traded volume as a fraction of the CE quantity is notably larger in the CDA in periods 1-10, but the increase in periods 11-20 is larger in FLOW. The percentage of filled CE quantity also trends upward, and FLOW eliminates more than half of its shortfall from CDA in the last 10 periods. We consistently see more but smaller orders placed in CDA than in Flow, especially in periods 11-20. Figure 1.5 shows the underlying trends. Again there is considerable heterogeneity across groups, but there is a clear tendency towards the CE surplus (i.e., higher efficiency) and CE volume across blocks, perhaps especially in FLOW sessions. We see that the clearing rate in FLOW is 9.5 shares per second in periods 1-10 but decreases slightly by 0.3 share per second in periods 11-20. The mean transaction size in CDA is 82 shares in periods 1-10 but decreases slightly in periods 11-20 with a slight increase in the average number of transactions.

Panel (c) shows measures of efficiency. Recall that CE prices ranged from 6 to 14, depending on the block. The first line in this panel shows the mean absolute price deviation from CE were rather large but declined slightly in CDA markets, from 2.19 in periods 1-10 to 2.14 in periods 11-20. Price deviations in FLOW markets were larger at 2.32 in periods 1-10 but declined faster to 1.84 in periods 11-20. Figure 1.4 shows the evolution of prices over time in all sessions (in faint colors) and session averages (dark green). There is considerable variation across sessions, but one sees a general tendency to under-respond to shifts in CE price across blocks. Looking at the last block confirms the impression that FLOW markets eventually track CE a bit better than do CDA markets, although neither performance is especially impressive. The last row of the table gives a first glimpse of allocative efficiency. In the CDA, average efficiency rises from about 80% in the first 10 periods to almost 89% in the last 10. Efficiency also rises by about 9 percentage points in FLOW, but remains 2 or 3 percentage points behind. Figure 1.3 compares realized surplus across trader roles and market formats. Even though the realized surplus is always slightly higher in CDA than in FLOW, the difference between buyers and sellers is much smaller in FLOW. Buyers tend to earn more than competitive equilibrium profits in CDA, with around 40% earning at least one and a half times the CE profits, especially in the second half.

Table 1.1: Summary Statistics of Experimental Sessions

	T1 - T20		T1 - T10		T11 - T20	
	CDA	FLOW	CDA	FLOW	CDA	FLOW
(a) Price and volatility						
Average Price	8.62	9.37	9.04	9.57	8.21	9.18
(std.)	(2.77)	(2.54)	(2.95)	(2.75)	(2.51)	(2.29)
$\mid P_t^W - P_{t-1}^W \mid$	0.57	0.39	0.58	0.39	0.55	0.4
$\text{Std}(\ln P_t^W - \ln P_{t-1}^W)$	0.14	0.07	0.15	0.07	0.13	0.07
(b) Order size and transactions						
Traded/CE Quantity (%)	93.6	78.8	90.3	71.1	96.9	86.4
Filled CE Quantity (%)	85.2	77.1	81.0	69.8	89.3	84.5
#Orders	48	27	47	28	49	25
Order Size	21	36	22	32	20	40
Clearing Rate		9.37		9.54		9.2
(std.)		(2.6)		(2.8)		(2.4)
Mean Transaction Size	78		82		74	
(std.)	(61)		(61)		(60)	
#Transactions	14		13		15	
(c) Efficiency						
$\mid P_t - P_{CE} \mid$	2.16	2.08	2.19	2.32	2.14	1.84
Realized Surplus (%)	84.3	81.81	79.76	77.31	88.85	86.31
Realized Surplus (buy-sell)(%)	53.7	11.18	27.38	-5.16	80.03	27.51

Note: Panel (a) shows the measures of price and volatility. Average Price is the overall average quantity-weighted clearing price of each 5-second interval during each 120-second trading period over each 20 period group and over all 5 groups (Nobs = 2400) for each format. $\mid P_t - P_{t-1} \mid$ is the average absolute value of all quantity-weighted price changes within period, averaged again over 20 periods and 5 groups. $Std(\ln P_t - \ln P_{t-1})$ is the standard deviation of first differences of natural logs of quantity-weighted prices averaged over 20 periods and 5 groups. Panel (b) shows measures of order size and transactions. Traded/CE Quantity is the average of traded volume as a fraction of CE quantity averaged over 20 periods and 5 groups. Filled CE Quantity is the average of filled contract as a fraction of CE quantity averaged over 20 periods and 5 groups. #Orders, Order Size, and #Transactions are all averages over the 20 trading periods and 5 groups (Nobs = 100 each). Clearing Rate is the average non-zero clearing rate across over the 20 trading periods (excluding the first 10 seconds of each period) and 5 groups in the FLOW format. Mean Transaction Size is the average of number of shares per transaction and #Transactions is the mean number of trades per period in the 5 CDA groups. Panel (c) shows measures of efficiency. $\mid P_t - P_{CE} \mid$ is the average absolute deviation from CE price in each period of quantity-weighted price of each 5-second interval across 5 groups over 20 trading periods (Nobs = 2400). Realized Surplus (%) is the end-of-period realized surplus summed over all traders as a percentage of CE surplus, averaged over the 20 trading periods and 5 groups (Nobs = 100). Realized Surplus (buy-sell)(%) is the difference between realized surplus of buyers and sellers, averaged over the 20 trading periods and 5 groups (Nobs = 100).

Figure 1.3: Realized Surplus Distribution

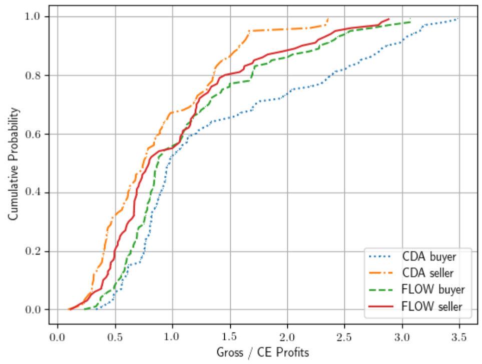

(i) T1-T20

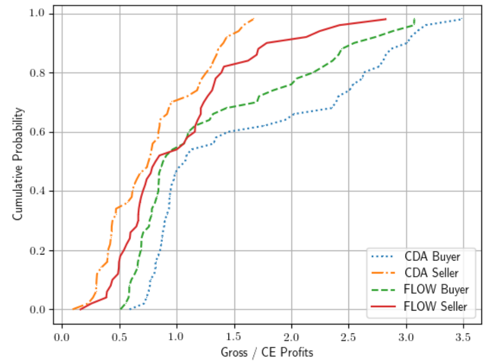

(ii) T11-T20

Figure 1.4: Clearing Price

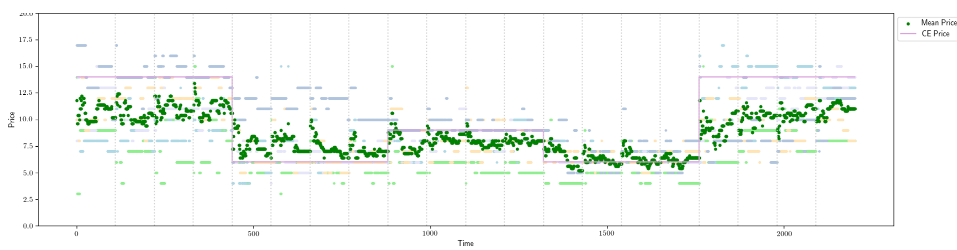

(i) CDA

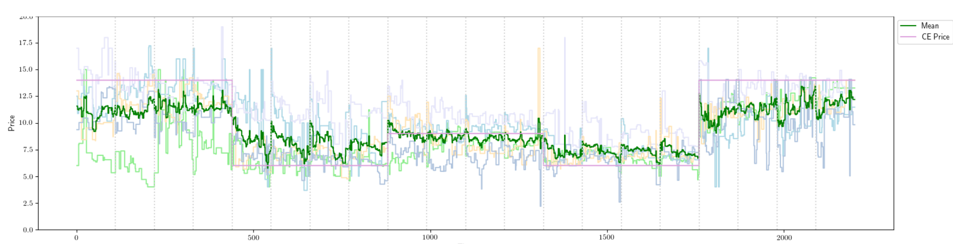

(ii) FLOW

Figure 1.5: Realized Surplus and Traded Volume

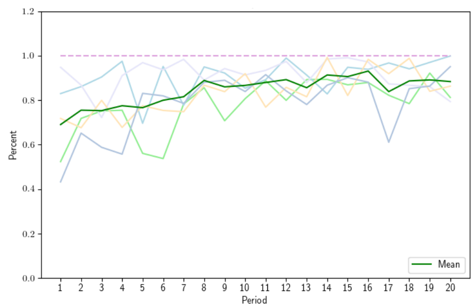

(i) CDA Surplus

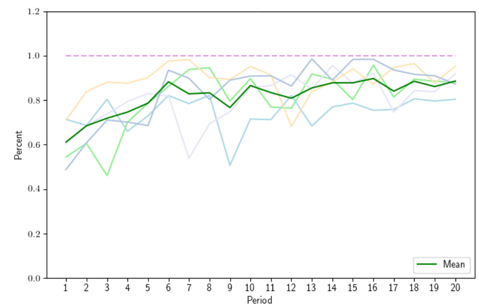

(ii) FLOW Surplus

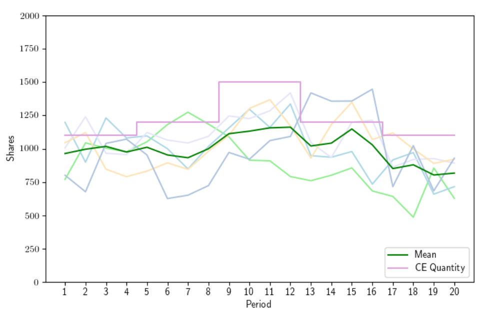

(iii) CDA Traded Volume

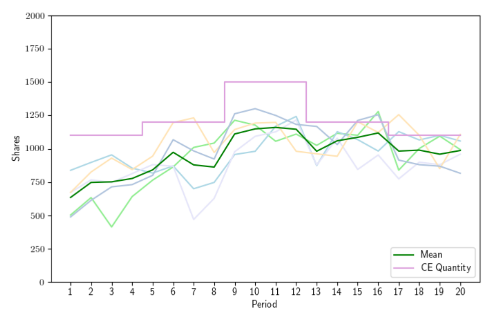

(iv) FLOW Traded Volume

1.3.2 Regressions

We compare the CDA and FLOW trading formats with the following performance metrics: (1) absolute deviations of clearing (or transaction) price from the competitive equilibrium price, $\mid P_t^W - P_{CE} \mid$; (2) volatility of prices, measured by $\mid P_t^W - P_{t-1}^W \mid$ and $Std\left(\ln P_t^W - \ln P_{t-1}^W\right)$; (3) order number; (4) order size; (5) traded volume Q; (6) market efficiency, measured by the fraction $\frac{\pi}{\pi_{CE}}$ of realized competitive equilibrium total surplus; (7) filled contract as a fraction $\frac{\text{Filled Contract}}{Q_{CE}}$ of the competitive equilibrium quantity; and (8) allocation efficiency, measured by the difference between buyer surplus and seller surplus.

To quantify treatment effects on choices made by subjects, we estimate the following model:

$$y_{g,t} = \alpha + \sum_{i=1}^{4} (\beta_i \, \text{Block}_i) + \gamma \, \text{FLOW}_{g,t} + \theta \, \text{Period}_{g,t} + \varepsilon_{g,t} \tag{1.3}$$

where $y_{g,t}$ is a performance metric indexed by group and time, Block_i is a dummy variables for contract configuration, $\text{FLOW}_{g,t}$ is the dummy variable for FLOW trading format, and $\text{Period}_{g,t}$ indicates the time period. The β_i's capture the contract block fixed effects, γ captures the market format treatment effect, and θ captures the period effect. For $y_{g,t} \in \{| P_t^W - P_{CE} |, | P_t^W - P_{t-1}^W |\}$, we estimated this model using five-second time intervals in order to have a balanced data between market formats. Within each time interval, we computed a quantity-weighted price to represent the clearing price for that time interval. The rest measures are at the group-period level. The estimation was done by combining data for both formats and all paid trading periods. The coefficients of models (1) - (4) in Table 1.2 are WLS estimates and the coefficients of the rest models in Table 1.2, Table 1.3, and Table 1.4 are OLS estimates.

Table 1.2: Regression Summary

	$\mid P_t^W - P_{CE} \mid$		$\mid P_t^W - P_{t-1}^W \mid$		$\mathrm{Std}\left(\ln P_t^W - \ln P_{t-1}^W\right)$	
	(1) T1-T20	(2) T11-T20	(3) T1-T20	(4) T11-T20	(5) T1-T20	(6) T11-T20
Intercept	7.813***	7.970***	2.456***	2.924***	0.345***	0.391***
	(0.840)	(0.997)	(0.462)	(0.905)	(0.069)	(0.124)
FLOW	-0.116	-0.191	-0.736***	-0.552***	-0.070***	-0.061***
	(0.218)	(0.328)	(0.119)	(0.146)	(0.016)	(0.023)
round	-0.228***	-0.234***	-0.062***	-0.094**	-0.010***	-0.013*
	(0.052)	(0.065)	(0.024)	(0.044)	(0.004)	(0.007)
Observations	4,800	2,400	4,600	2,300	200	100
R^2	0.318	0.447	0.140	0.120	0.328	0.311
Adjusted R^2	0.317	0.446	0.139	0.118	0.307	0.282

Note: *p<0.1; **p<0.05; ***p<0.01. The standard errors within parentheses are cluster-robust at the group level.

Table 1.3: Regression Summary

	order number		order size		traded volume	
	(1) T1-T20	(2) T11-T20	(3) T1-T20	(4) T11-T20	(5) T1-T20	(6) T11-T20
Intercept	53.4***	60.7***	2.0	-6.9	746.6***	728.0***
	(6.0)	(11.1)	(7.0)	(11.7)	(116.5)	(134.2)
FLOW	-21.1***	-23.7***	13.2***	17.5***	-178.8***	-132.3**
	(2.0)	(2.1)	(2.2)	(1.9)	(34.4)	(63.5)
round	-0.3	-0.7	1.2***	1.5**	18.9***	18.6**
	(0.3)	(0.6)	(0.4)	(0.6)	(6.0)	(7.9)
Observations	200	100	200	100	200	100
R^2	0.735	0.839	0.497	0.688	0.471	0.350
Adjusted R^2	0.727	0.833	0.481	0.674	0.454	0.322

Note: *p<0.1; **p<0.05; ***p<0.01. The standard errors within parentheses are cluster-robust at the group level.

Table 1.4: Regression Summary

	filled CE quantity		realized surplus		$\text{surplus}_{buy-sell}$	
	(1) T1-T20	(2) T11-T20	(3) T1-T20	(4) T11-T20	(5) T1-T20	(6) T11-T20
Intercept	0.634***	0.641***	0.521***	0.627***	2.390***	3.530***
	(0.117)	(0.122)	(0.109)	(0.138)	(0.554)	(0.940)
FLOW	-0.080*	-0.048	-0.025	-0.025	-0.425	-0.525*
	(0.044)	(0.034)	(0.035)	(0.030)	(0.276)	(0.275)
round	0.016***	0.014**	0.020***	0.014**	-0.019	-0.078
	(0.005)	(0.006)	(0.005)	(0.007)	(0.030)	(0.056)
Observations	200	100	200	100	200	100
R^2	0.286	0.184	0.312	0.090	0.738	0.796
Adjusted R^2	0.264	0.150	0.291	0.052	0.730	0.787

Note: *p<0.1; **p<0.05; ***p<0.01. The standard errors within parentheses are cluster-robust at the group level.

Result 1. The FLOW market format exhibits lower price volatility than CDA.

Table 1.2 presents models of with respect to pricing. Models (3) - (6) measure the price volatility. The absolute change of prices is around 0.7 smaller in FLOW than in CDA, and is significant at the $p = 0.0.1$ level. The difference still persists in that last ten periods. Additionally, model (5) suggests that the price is around 7% less volatile in FLOW than in CDA, and is significant at the $p = 0.0.1$ level. The difference is carried over to the second half of the periods. All models suggest there is some degree of learning among subjects.

Result 2. The FLOW market format effectively shreds orders as compared to CDA.

Table 1.3 presents order and volume related models. Models (1) and (2) show that FLOW market format has, on average, 21 - 24 fewer orders submitted in each period than CDA, and is significant at the $p = 0.01$ level. Models (3) and (4) further suggest that the order size in FLOW market format is 13 shares more per order on average than in CDA. There is also a significant learning pattern in terms of order size.

Result 3. The FLOW market format results in less total traded volume than CDA.

In models (5) and (6) of Table 1.3, the shortfall in traded volume is significant at the $p = 0.01$ level, but decreases in the last ten periods. Models (1) and (2) of Table 1.4 show that the shortfall in filled contract as a fraction of CE quantity is, on average, about 8% lower in FLOW than in CDA, but becomes smaller and insignificant in the last 10 periods. It also shows that subjects were learning across periods significantly.

Result 4. The FLOW market format has similar level of price and allocative efficiency as CDA.

In Table 1.2, models (1) and (2) reports estimates of price efficiency. Recall that CE prices ranged from 6 to 14. It shows that FLOW market format has a smaller absolute price deviation of 0.1 from the CE price than CDA on average. The difference increases to around 0.2 on average in the last ten periods, though still not significant.

Table 1.4 presents models of allocative efficiency. In our data, allocations may be less efficient in FLOW than in CDA. The average shortfall, of about 2.5%, is not significant. There is learning effect among subjects and is also shown in Figure 1.5. The surplus difference between buyers and sellers is around 42% lower in FLOW than in CDA. The allocative difference is over 50% lower in FLOW than in CDA, and is significant at the $p = 0.1$ level in the last ten periods.

1.4 Conclusions

Although the proposed Flow format is a significant departure from the standard and prevalent market designs, there are no empirical analyses of the behavior of traders in this format nor of its performance. To address this gap, our study examines using a laboratory environment the flow trading format and how it stands against the prevalent continuous double auction (CDA). Our exploration focuses on trader behavior and market performance metrics, in a framework where a single asset is transacted on a virtual exchange, and where traders are allowed to possess or owe multiple asset units.

We use the induced value paradigm with a modification: the introduction of contracts to traders at the onset of each trading interval. These contracts, are either buy and sell agreements and represent the commitment of the experimenter to respectively purchase or sell a predetermined quantity of shares from or to the trader at a stipulated price and deadline.

The experimental design is between-group, and the experiment used 80 participants distributed across two market formats. Each format was deployed in five distinct groups. Each session was organized into 20 trading periods of two minutes each, with an equitable allocation of buy and sell contracts among eight traders per market. The trading periods are further segmented into five blocks, maintaining uniform contract configurations to enable consistent analysis of underlying demand and supply dynamics.

Our findings show clear changes in traders' behavior within the Flow format, evidenced by a significant decrease in the frequency of orders and an increase in the average order size compared to the CDA format. Interestingly, the overall trading volume within the Flow Market initially lagged behind the CDA, a discrepancy that diminished progressively as participants seem to learn the Flow format's mechanics. Moreover, the Flow Market exhibited a reduced price volatility, (lower absolute price changes and a lower standard deviation in log price differences). That is, the Flow format seems to favor market price stability. These insights underline the Flow format's potential as a viable market design. This paper exhibits several limitations that are worth noting. First, we do not measure the differences in price impact between the two formats. Second, we do not test empirically the degree of protection against HFTs predatory behavior between the two formats. We plan to study these two important aspects of our research agenda in a subsequent paper.

Chapter 2

Sequential Search with Ex Post Uncertainty [1]

2.1 Introduction

In sequential search problems, players face ex ante uncertainty about product value. They may have information on the distribution of the value but only learn the exact utility they would receive after the search. In previous sequential search models, players choose the optimal stopping rule and purchase the searched product that provides the highest utility. This model has been widely discussed and applied in theoretical (e.g., Lippman and McCall (1976), Weitzman (1979)), empirical (e.g., Ghose, Ipeirotis, and B. Li (2012), Ursu (2018)), and experimental (e.g., Schotter and Braunstein (1981),

[1]The second chapter is a joint work with Shuchen Zhao.

Caplin, Dean, and Martin (2011)) studies.

Risk preference plays an important role in sequential search problems. Researchers have found that the probability of continuing a search decreases with the level of risk aversion (Nachman (1972), Lippman and McCall (1976)). Since players learn the exact utility of the product after each search, they make their search decision by comparing a current certain outcome with a future uncertain outcome. It is unsurprising that a risk-averse player tends to end the search earlier than a risk-neutral player.

However, whether people can learn the exact utility of a product after their search is questionable in experience goods markets. Imagine that you are looking for a lawyer to defend you in a lawsuit. Your utility depends on whether you win the lawsuit, and you can search for the quality and reputation of lawyers to maximize the probability of winning. Unlike the traditional commodity goods market, you cannot observe the result of the lawsuit when you hire the lawyer. Instead, you can only generate an expectation and gain the exact utility afterward. This is also a typical problem in financial markets, where most funds do not inform you of a guaranteed return, and there is always a possibility of default.

In those markets, the uncertainty of the search process exists not only ex ante but also ex post. Compared to traditional search models, a critical difference here is that the players face a lottery instead of a certain outcome after each search. In this scenario, both the current (a single lottery) and future outcomes (a complicated compound lottery) are uncertain. A risk-averse player may have an incentive to extend the search if the current outcome is risky, which contradicts the findings in the literature and thus motivates our study. How do players behave when such ex post uncertainty exists? How do risk preferences affect the search process in this scenario? We study these two research questions with a modified sequential search model and a laboratory experiment. This paper contributes to the literature by introducing ex post uncertainty into sequential search problems.

We first modify the sequential search model (Lippman and McCall (1976) and Schunk (2009)) and replace the exact searched value with lotteries that differ in their probability of success. If the player is risk-neutral, the two models will provide the same prediction. However, if the player is not risk-neutral, the two models offer different predictions; i.e., the expected search duration increases with the level of risk tolerance in the traditional model but decreases with the level of risk tolerance in our modified model.

We further examine the theory in a laboratory experiment with a between-subjects design. Laboratory subjects play either the traditional sequential search game with certain searched outcomes or the modified game with uncertain lotteries. Subjects' risk preference is elicited through both a multiple price list (MPL) (Holt and Laury (2002)) and a bomb risk elicitation task (BRET) (Crosetto and Filippin (2013)) to build the correlation between their search behavior and risk preferences. The results verify the theoretical predictions.

The remainder of this paper is organized as follows. Section 2 reviews the related literature concerning sequential search and risk elicitation methods. Section 3 introduces the modified sequential search model. Section 4 presents the experimental design and the hypotheses. Section 5 describes the results of the experiments. Finally, Section 6 concludes with the main findings and highlights unsolved questions to be addressed in future studies.

2.2 Literature

A large and growing literature has conclusively demonstrated consumers' optimal strategy in sequential search problems. Our paper adopts the consumer sequential search framework first developed by Lippman and McCall (1976). In the fundamental model, players actively seek an offer from a known distribution of values. Exactly one offer is presented after each search with a constant search cost. An individual chooses the optimal stopping rule and makes the purchase decision under risk neutrality, perfect recall, and an infinite time horizon. Lippman and McCall (1976) conclude that the optimal stopping rule to terminate the search is when the maximum sampled benefit exceeds the reservation prices of the unsampled products, which are determined by equating the marginal benefit of sampling one more product and the marginal cost of a search. Weitzman (1979) extends the fundamental model with heterogeneous product distributions and the optimal search order conditional on the distributions.

Although laboratory subjects do not perfectly follow the theoretical predictions, many laboratory experiments find that subjects' behaviors are close to those predicted by theoretical models. Papers that examine the effects of search costs, recall options, time horizon, searchers' knowledge, interest rates, search subsidy, wage distribution, and sunk costs have found that laboratory subjects behave similarly but not identically to theoretical predictions (e.g., Schotter and Braunstein (1981), Cox and Oaxaca (1989), Kogut (1990)).

To capture how subjects sequentially make search decisions in the laboratory, papers adopt methods such as tape recording (John D. Hey (1987)) and electronic information boards where subjects need use a mouse to click on the screen to obtain certain information (Sonnemans (1998)). Caplin, Dean, and Martin (2011)'s design also captures how subjects sequentially make decisions with choice process data and studies how complexity affects their behaviors. Other recent studies examine the effect of continuous time interactions (Brown, Flinn, and Schotter (2011)), the effect of the ambiguity of the offer distribution (Asano, Okudaira, and Sasaki (2015)) and the possibility of stock-out options (Kittaka and Mikami (2020)).

The previous literature mostly focuses on situations in which consumers are risk neutral. However, risk preference has proven to be a vital factor in many economic activities, including health (Barsky et al. (1997), Picone, Sloan, and Taylor (2004), Anderson and Mellor (2008)), investment (Barasinska, Schäfer, and Stephan (2008), Dohmen et al. (2011)), and consumption (Chetty and Szeidl (2007)). Our experiment is closely related to other search papers that focus on noisy signals and risk preference. Both Zwick et al. (2003) and Palley and M. Kremer (2014) propose a search environment in which consumers only observe the relative ranking of the searched results instead of their exact value. This limited information induces excessive searching in both studies. Nonetheless, the evidence on how risk preferences affect consumer search remains ambiguous. On the one hand, Schunk (2009) and Schunk and Winter (2009) find that individual risk attitude appears to be unrelated to decision heuristics, whereas loss aversion is related to search behavior. Individuals who avoid gambles tend to have a higher degree of loss aversion and tend to stop their search earlier. On the other hand, Miura, Inukai, and Sasaki (2017) also use laboratory experiments to test the effect of risk preference on search activities without recall and, consistent with their theoretical predictions, find a statistically significant effect of risk preference on the duration of search.

Similar to the three papers above, we remove the assumption of risk neutrality to study the effect of risk preferences in consumers' sequential search problems. Furthermore, we introduce ex post uncertainty into our model; the searched results are lotteries instead of exact values. The paper contributes to the existing literature by considering the effect of individual risk preferences when there is ex post uncertainty, which is closely related to the real-life search circumstances in experience goods markets. To our best knowledge, this is the first paper to focus on this specific setup and include laboratory experiments.

There are several methods to elicit laboratory subjects' risk preferences, including MPL (Holt and Laury (2002)), random lottery pairs (John D Hey and Orme (1994)), ordered lottery selection (Eckel and Grossman (2002)), Becker-DeGroot-Marschak (Becker, DeGroot, and Marschak (1964)), trade-off (Abdellaoui (2000)), portfolio choice and investment (Gneezy and Potters (1997), Choi et al. (2007)), and a recently developed method termed BRET (Crosetto and Filippin (2013)). Although the consistency of the risk elicitation methods and the validity of the expected utility theorem are under debate (see Friedman, Isaac, et al. (2014) and Friedman, Habib, et al. (2022) for an overview), these methods are widely accepted in the literature for estimating the relative relationship of risk preferences among individuals. Specifically, we apply two methods, namely, the MPL of Holt and Laury (2002) and the BRET of Crosetto and Filippin (2013). The former method is a classic approach to risk preference elicitation and has been frequently applied in the sequential search literature. In the latter method, subjects observe the information updated in a sequential order, which is similar to how they search and purchase in sequential search games. The subjects' constant relative risk tolerance preference is directly measured by the number of risky choices in MPL and the number of boxes collected in BRET.

2.3 Model

The model is based on the sequential search models of Lippman and McCall (1976) and Schunk (2009). Homogeneous consumers search from an infinite number of differentiable products and purchase one of the products. The consumers start with no product on the searched list. They can either choose to purchase a product from the current searched list or pay a constant cost c to search for the next product. Only one product is revealed in each search. The model assumes an infinite time horizon and perfect recall, so consumers can search an unlimited number of times and choose any searched product on the list.

To introduce ex post uncertainty, we allow the value of the products to be either high or low. Consumers receive v_H when the product has a high value and v_L when the product has a low value, with $v_H > v_L$ and $v_H > c$. The products differ in the probability of receiving the high value p_i. To simplify the computation, we define that the probability follows a uniform distribution, which is $p_i \sim U[0,1], \forall i$ with cumulative distribution function $F(p_i) = p_i$. Consumers observe the probability p_i after each search and earn either v_H or v_L when they purchase the product.

2.3.1 Reservation probability and comparative statics

In infinite horizon sequential search models (Lippman and McCall (1976), Schunk (2009)), the optimization problem is solved by comparing the expected benefit of continuing the next search with the current best product, and consumers treat the previous search costs as sunk costs. Consumers then seek to maximize their expected utility with an optimal stopping rule. We continue to apply the same method in our model.

Suppose that the next product to be searched for is product i, and the maximal probability of a product being of high value in the current searched list is z_i. If consumers choose not to search for this product i, they obtain the product with z_i, and their expected utility is $u(v_H)z_i + u(v_L)(1-z_i)$. If they choose to search for product i, they can expect the following marginal benefit.

$$\begin{aligned} G(z_i) =& [u(v_H-c)z_i + u(v_L-c)(1-z_i)] \int_0^{z_i} dF(p_i) \\ &+ \int_{z_i}^1 u(v_H-c)p_i + u(v_L-c)(1-p_i)dF(p_i) \end{aligned} \tag{2.1}$$

The first term in the above expression represents the case in which a probability smaller than z_i is found, and the second term represents the case in which a probability greater than z_i is found. Consumers continue their search if $G(z_i) > u(v_H)z_i + u(v_L)(1-z_i)$, as the expected gain from the next search is higher than accepting the current best product.

As proven in the literature, consumers apply a *reservation value* strategy, where they stop their search when the current best offer exceeds their reservation value. Similarly, we define *reservation probability* z^* where $G(z^*) = u(v_H)z^* + u(v_L)(1-z^*)$. Consumers continue their search if $z_i < z^*$, whereas they stop searching and choose the product with z_i if $z_i \geq z^*$. Let $\psi(z_i) = G(z_i) - u(v_H)z_i - u(v_L)(1-z_i)$. We have

$$\begin{aligned}
\psi(z_i) =& [u(v_H - c)z_i + u(v_L - c)(1 - z_i)] \int_0^{z_i} dF(p_i) \\
& + \int_{z_i}^1 u(v_H - c)p_i + u(v_L - c)(1 - p_i)dF(p_i) - u(v_H)z_i - u(v_L)(1 - z_i) \\
=& [u(v_H - c) - u(v_L - c)]z_i^2 + 2[u(v_L) - u(v_H)]z_i \\
& + u(v_H - c) + u(v_L - c) - 2u(v_L) \qquad (2.2)
\end{aligned}$$

The reservation probability z^* is computed by $\psi(z^*) = 0$. We have

$$[u(v_H - c) - u(v_L - c)]z^{*2} + 2[u(v_L) - u(v_H)]z^* + u(v_H - c) + u(v_L - c) - 2u(v_L) = 0 \qquad (2.3)$$

Using the quadratic formula, we can solve for z^*.

When $u(v_H - c) + u(v_L - c) - 2u(v_L) \leq 0$, $z^* = 0$, consumers immediately stop searching after the initial search, or choose not to search at all. The consumer's decision depends on whether the expected return of the initial search generates a positive payoff.

When $u(v_H - c) + u(v_L - c) - 2u(v_L) > 0$, there exists a unique $z^* \in (0, 1)$ with

$$z^* = \frac{u(v_H) - u(v_L) - \sqrt{[u(v_L) - u(v_H)]^2 - [u(v_H - c) - u(v_L - c)][u(v_H - c) + u(v_L - c) - 2u(v_L)]}}{u(v_H - c) - u(v_L - c)} \qquad (2.4)$$

The intuition for the solutions is straightforward. Given the marginal cost c, when the difference between v_H and v_L is small, consumers are less willing to chase a high-probability product; thus, they simply play the initial search or do not search at all. When the difference between v_H and v_L is sufficiently large, consumers are motivated to search for a product with a high probability of receiving v_H.

Based on eq (2.3), we can compute the comparative statistics. First, we replace each utility function with its second-order Taylor approximation at $x = x_0$. Without loss of generality, we assume that $u'(x_0) = 1$, $x_0 = v_L - c$ and use $a = -u''(x_0)$ as a proxy for the level of risk aversion. Through the implicit function theorem, we have

$$\begin{aligned}
\frac{\partial z^*}{\partial a} &= \frac{\frac{1}{2}(v_H + v_L - 2x_0)(z^* - 1)^2 - c(z^{*2} + \frac{1}{2}) - \frac{(v_H + 3v_L - 4x_0 - 2c)c}{2(v_H - v_L)}}{2(z^* - 1) + a(1 - z^*)(v_H + v_L - 2x_0) + 2acz^*} \\
&= \frac{\frac{1}{2}(v_H - v_L + 2c)(z^* - 1)^2 - c(z^{*2} + \frac{1}{2}) - \frac{(v_H - v_L + 2c)c}{2(v_H - v_L)}}{2(z^* - 1) + a(1 - z^*)(v_H - v_L + 2c) + 2acz^*}
\end{aligned} \tag{2.5}$$

Our conjecture is that the probability of continuing the search decreases with the level of risk tolerance; thus, $\frac{\partial z^*}{\partial a} > 0$. Although our conjecture does not hold true globally, there are interesting sufficient conditions in which the conjecture is true. See Appendix A for an example.

2.3.2 Model without ex post uncertainty

Following Schunk (2009), here we demonstrate the model without ex post uncertainty for comparison. Define $v_i = p_i v_H + (1-p_i)v_L$ as the expected value of the search outcomes for product i in the previous section; we have $v_i \sim U[v_L, v_H]$. In the model with certain outcomes, the consumers observe the expected value of the products and receive v_i if they choose product i, thus eliminating the ex post uncertainty from the model. To make this model comparable with that in the previous section, we continue to represent v_i with p_i and solve for z^*, which is the reservation probability and can be transferred directly to the reservation value.

The model modifies the outcomes of the search but does not affect consumers' decision-making process. Apparently, risk-neutral consumers are indifferent between the lottery and its expected value. For consumers, $\psi(z_i)$ appears as follows.

$$\begin{aligned}\psi(z_i) =& u(v_H z_i + v_L(1-z_i) - c)\int_0^{z_i} dF(p_i) \\ &+ \int_{z_i}^1 u(v_H p_i + v_L(1-p_i) - c)dF(p_i) - u(v_H z_i + v_L(1-z_i))\end{aligned} \tag{2.6}$$

By solving for z^* with $\psi(z^*) = 0$, we can calculate the reservation probability in the model without ex post uncertainty. Without ex post uncertainty, the common prediction in the literature is that $\frac{\partial z^*}{\partial a} < 0$ (e.g., Nachman (1972)). That is, the reservation value decreases with the level of risk aversion. Additional graphical comparison with simulations can be found in Appendix B and in the discussion of the experimental hypotheses.

2.4 Experimental design

Our experiment applies a two-group between-subjects design. Each experimental session is divided into three parts. In part one, the subjects randomly complete one of the two risk elicitation tasks, i.e., either the MPL or the BRET. The tasks are discussed in detail later in this section. In part two, the subjects play 25 rounds of the sequential search game, with 5 practice rounds and 20 paid rounds. In each round, the subjects start with no item on the search list. They can choose to purchase one of the items from the current searched list or to pay a constant search cost c to continue searching for the next item. If they continue searching, another item (with an uncertain or a certain outcome) will be added to the search list. The game is played in an infinite horizon in which subjects can search an unlimited number of times, and the round ends when the subjects eventually make a purchase. The subjects' payoff for the current round when choosing item i is given by (p_i, v_H, v_L, n, c), where p_i is the probability of item i being a high-value item, (v_H, v_L) are the payoffs for high- and low-value items, respectively, n is the total number of searches, and c is the constant marginal search cost. In all sessions, we set $(v_H, v_L, c) = (500, 100, 5)$. In part three, the subjects play the other risk elicitation task that was not played in part one.

The order of the three tasks is designed to mitigate the order effect between the search task and the risk elicitation tasks. One risk elicitation task is played before and the other after the search. To control for the order effect between the two risk elicitation tasks, their order is chosen randomly.

2.4.1 Treatment variables

The major treatment variable is the comparison between searching for certain values or uncertain lotteries.

In the *Uncertainty (UC)* treatment, it is publicly known that the probability of the items being of high value follows a uniform distribution $p_i \sim U[0,1]$. The subjects observe p_i after each search. If they choose to purchase item i, the system will randomly draw the final value based on p_i. In the UC treatment, the subjects receive $v_H - nc$ with probability p_i and $v_L - nc$ with probability $1 - p_i$. Figure 2.1 provides a sample user interface from the UC treatment.

In the *Certainty (C)* treatment, the subjects are paid directly based on the expected value $v_i = p_i v_H + (1 - p_i) v_L$ after each search, which is the same as in the standard sequential search model in the literature. We still have $p_i \sim U[0,1]$, but the public information is about v_i, that is, $v_i \sim U[v_L, v_H]$. This treatment serves as the baseline for this set of comparisons. The subjects' end-of-round payoff is $v_i - nc$ in the C treatment.

Figure 2.1: Uncertainty treatment sample user interface.

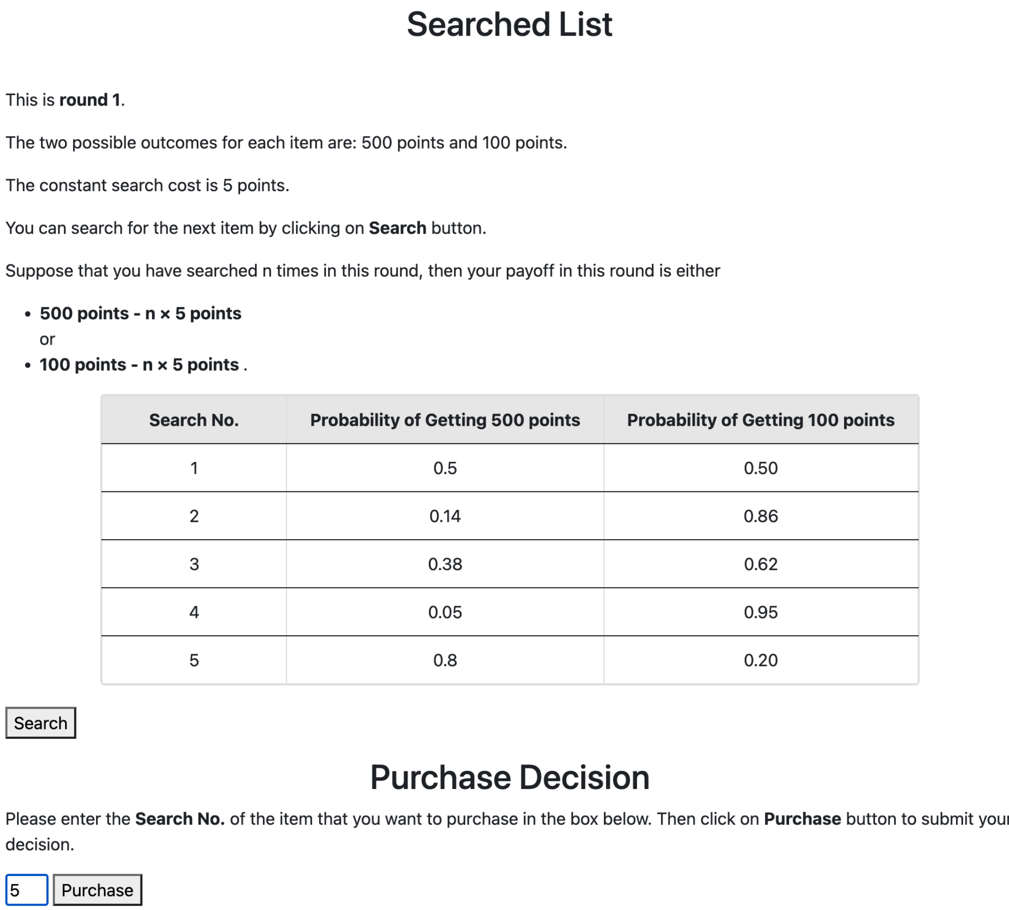

Searched List

This is **round 1**.

The two possible outcomes for each item are: 500 points and 100 points.

The constant search cost is 5 points.

You can search for the next item by clicking on **Search** button.

Suppose that you have searched n times in this round, then your payoff in this round is either

- **500 points - n × 5 points**
  or
- **100 points - n × 5 points** .

Search No.	Probability of Getting 500 points	Probability of Getting 100 points
1	0.5	0.50
2	0.14	0.86
3	0.38	0.62
4	0.05	0.95
5	0.8	0.20

Search

Purchase Decision

Please enter the **Search No.** of the item that you want to purchase in the box below. Then click on **Purchase** button to submit your decision.

5 Purchase

(Note: The instructions are displayed prior to this page, so the information on the screen just serves as a reminder.)

2.4.2 Eliciting risk preferences

To elicit their risk preference, subjects are asked to complete an MPL task (Holt and Laury (2002)) and a BRET task (Crosetto and Filippin (2013)) in a random order.

In MPL, each subject will first be given a multiple price list in text format, where each row in the list has the same two allocations, but different rows have different state probabilities. Subjects begin by indicating a preference, Option A or Option B, for each of the ten paired lottery choices in Table 2.1, with the understanding that only one of these choices will be selected at random ex post to determine the earnings for the option selected. Subjects' risk preference can be estimated by using the proportion of safe choices (Option A) among the ten decisions as an indicator of risk aversion. At the end of the risk elicitation task, one of the rows will be selected, and the subjects will be paid based on the realization of their choice for that row. In MPL, the subjects observe all ten paired lottery choices at the beginning of the task.

Table 2.1: Multiple price list.

Option A	Option B	Expected Payoff Difference	Open α interval if subject switches to Option B (assume EUT)
280pts	1/10 of 500pts, 9/10 of 100pts	140pts	(3.944, ∞)
280pts	2/10 of 500pts, 8/10 of 100pts	100pts	(2.687, 3.944)
280pts	3/10 of 500pts, 7/10 of 100pts	60pts	(1.896, 2.687)
280pts	4/10 of 500pts, 6/10 of 100pts	20pts	(1.277, 1.896)
280pts	5/10 of 500pts, 5/10 of 100pts	-20pts	(0.734, 1.277)
280pts	6/10 of 500pts, 4/10 of 100pts	-60pts	(0.211, 0.734)
280pts	7/10 of 500pts, 3/10 of 100pts	-100pts	(-0.335, 0.211)
280pts	8/10 of 500pts, 2/10 of 100pts	-140pts	(-0.974, -0.335)
280pts	9/10 of 500pts, 1/10 of 100pts	-180pts	(-1.888, -0.974)
280pts	10/10 of 500pts, 0/10 of 100pts	-220pts	($-\infty$, -1.888)

"pts" refers to the points subjects earn in the experiment. α refers to the risk-aversion parameter in $u(x) = x^{\alpha}$.

The BRET is a risk elicitation method recently developed by Crosetto and Filippin (2013). In its dynamic version, subjects decide sequentially how many boxes to collect among 100 in continuous time, one of which contains a bomb (hidden during the game). One box is automatically collected each second, and the subjects can stop the collection process at any time. The subjects' earnings increase linearly with the number of boxes collected but decrease to zero if the bomb is also collected. In short, the decision can be represented as a sequence of binary choices. At each second, the subjects compare two lotteries, namely, collecting k or $k+1$ boxes.

The BRET has several advantages. It requires minimal numeracy skills, avoids truncation of the data, allows the precise estimation of both risk aversion and risk seeking, and is not affected by the degree of loss aversion or by violations of the reduction axiom. Another reason we select BRET is that the information is updated in sequential order, which is similar to the sequential search tasks. In our sequential search models, the players' choices can also be regarded as comparing two risky outcomes, namely, choosing the current best offer or continuing the search. We believe that the risk preference that the method elicits best describes an individual's risk preference in the search tasks. Table B.1 in Appendix C shows the estimate of α for the BRET from the original research.

2.4.3 Session information

Ten sessions were held between December 2021 and June 2022. In total, 110 subjects participated in our sessions. The subjects were paid randomly based on one round's payoff of all 20 paid rounds, plus their $2 show-up fee and their payment in the two risk preference elicitation tasks. On average, the subjects earned $10 during a 40-minute session. The experiments were developed on oTree (Chen, Schonger, and Wickens (2016)), and the subjects were recruited on Orsee (Greiner (2015)), which is operated by the Learning and Experimental Economics Projects of Santa Cruz (Leeps Lab). Given the two between-subjects groups, the session information can be summarized as follows in Table 2.2.

Table 2.2: Session table.

Group	Search Outcomes	#Session	#Subjects
1	Certainty	5	54
2	Uncertainty	5	56

2.4.4 Hypotheses

In our experiment, we set $(v_H, v_L, c) = (500, 100, 5)$. Consistent with Crosetto and Filippin (2013) and our simulation in Appendix B, we use the CRRA utility function $u(x) = x^{\alpha}$ to generate the hypotheses. Figure 2.2 shows how z^* varies with α. The figure shows a clear distinction between the two models, where the reservation probability decreases with the level of risk tolerance when the search outcomes are uncertain but increases with the level of risk tolerance when the search outcomes are certain.

Figure 2.2: Reservation probability vs CRRA parameter.

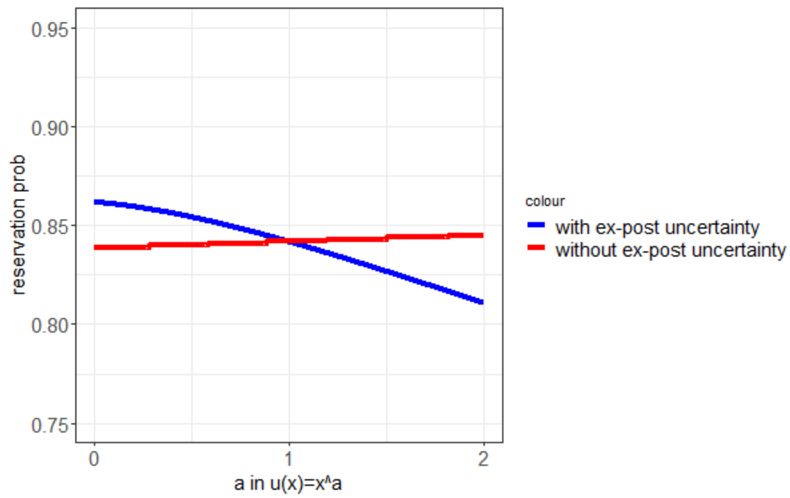

(Note: The plot shows how the reservation probability z^* varies with the CRRA parameter when $(v_H, v_L, c) = (500, 100, 5)$. The blue curve refers to the model predictions with ex post uncertainty. The red curve refers to the model predictions without ex post uncertainty.)

Although we cannot directly observe the reservation probability/value of the subjects, the search duration can be calculated given the chosen probability/value and observed in the experiment. From the model, we learned that the reservation value increases with the level of risk tolerance in the C treatment but that the reservation probability decreases with the level of risk tolerance in the UC treatment. Accordingly, in the C treatment, risk-averse consumers would shorten their search duration. In contrast, in the UC treatment, risk-averse consumers would be unlikely to accept a risky lottery with ex post uncertainty and thus would extend their search.

Hypothesis 1. *The average search duration increases with the level of risk tolerance in the C treatment.*

Hypothesis 2. *The average search duration decreases with the level of risk tolerance in the UC treatment.*

In addition, Figure 2.2 shows that risk neutral players choose the same reservation probability/value under both the C and the UC treatments. The statement can be easily verified in the model if we apply a linear utility function. We will also examine this hypothesis in our experiment.

Hypothesis 3. *Risk-neutral players have the same average search duration in the C and the UC treatments.*

2.5 Results

2.5.1 Data overview

We dropped one subject from the UC treatment because she selected 100 boxes in the BRET task, which guarantees a 0 payoff. Clearly, the subject did not pay attention to the tasks. After the deletion, we have 54 subjects in the C treatment and 55 subjects in the UC treatment. There are also 10 subjects with multiple crossing points in the MPL task, 6 from the UC treatment and 4 from the C treatment. In the main text of the paper, we use their last crossing point to determine their level of risk tolerance. Appendix D also provides another approach that uses their first crossing point (see Friedman, Habib, et al. (2022) for an overview). The main results are robust across the two specifications.

Table 2.3 presents the summary statistics of the experiment. Each row refers to an average value of a specific variable by treatments in the columns. The third column shows the p-value of the t-test of the comparison between the two treatments. “Chosen probability” is the subjects’ final decision, where their chosen value in the C treatment is converted to the equivalent value of probability in the UC treatment. “Recall rate” refers to the fraction of decisions where the subjects did not choose their last searched result, thus indicating that the subjects adjust their reservation probability/value during the game. “Optimal choice rate” refers to the fraction of time the subjects choosing the best searched outcome. The high optimal choice rate shows that the subjects rarely make mistakes in the experiment.

As indicated in Table 2.3, most differences of average data between treatments are insignificant. On average, the subjects search for slightly longer in the C treatment than in the UC treatment, but this difference becomes insignificant when we focus on the second half of the experiments. Subjects also have a higher recall rate in the C treatment than in the UC treatment. Nonetheless, both rates are close to the average experimental findings in the literature. In addition, the average search duration and chosen probability vary little between the data for all rounds and that for the last 10 rounds, which indicates weak learning between rounds.

Table 2.3: Descriptive Statistics of average data.

	C	UC	C-UC p-value
search duration (all rounds)	3.81	3.27	0.013**
search duration (last 10 rounds)	3.75	3.55	0.502
chosen probability (all rounds)	0.83	0.8	0.111
chosen probability (last 10 rounds)	0.84	0.83	0.422
# of risky choices in MPL	5.07	4.29	0.014**
# of boxes collected in BRET	39.11	43.20	0.243
Recall rate in search task	23.06%	13.64%	0.000***
Optimal choice rate in search task	99.07%	97.00%	0.001***
Number of subjects	54	55	–
Number of decisions	1080	1100	–

Note: *p<0.1; **p<0.05; ***p<0.01.

2.5.2 Search duration

The summary statistics reveal similarity between treatments, but the same average data may result from distinct relations between the subjects' behavior and their risk preference. For each subject, we calculate her average search duration in all 20 paid rounds and plot the relation between the average search duration and the subject's choices in the risk elicitation tasks. Figures 2.3 and 2.4 show the relations; both support hypotheses 1 and 2. The dots represent the relation for each subject, and the lines are estimated by OLS regression given these subject-level data (the shaded band refers to the 95% confidence interval of the estimated parameters). As shown in both figures, the search duration increases with the level of risk tolerance in the C treatment but decreases with the level of risk tolerance in the UC treatment, thus supporting hypotheses 1 and 2. Figures 2.5 and 2.6 confirm the hypotheses with the average of the last 10 rounds instead of all 20 rounds. Figures 2.5 and 2.6 also confirm that learning between rounds weakly exists in our experiments. However, hypothesis 3 is not supported, as the intersection of the two lines is located on the left side of the figures. Both figures show that a particular type of risk-averse subject, not risk-neutral subjects, plays the two treatments in the same manner.

Figure 2.3: Search Duration vs BRET (all 20 rounds).

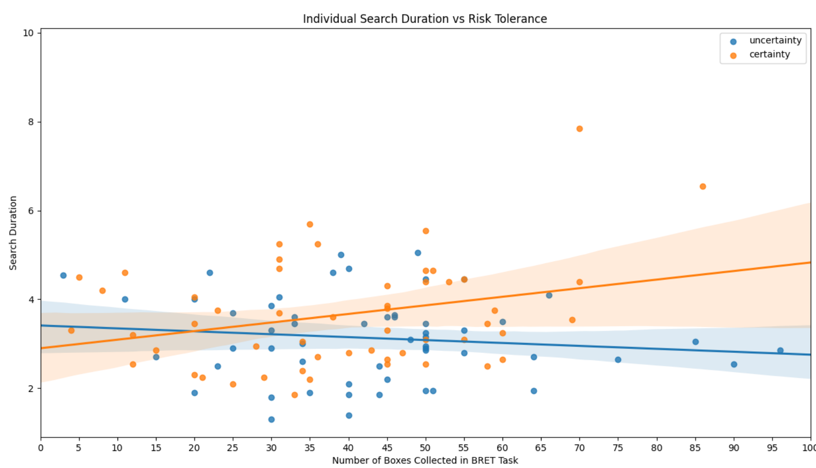

(Note: The scatter plot shows how search duration changes with the level of risk tolerance when using the average behavior of 20 paid rounds and number of boxes collected in BRET. The dots represent the relation for each subject, and the lines are estimated by OLS regression given these subject-level data (the shaded band refers to the 95% confidence interval of the estimated parameters).)

Tables 2.4 and 2.5 report the results of the estimation of the following regression (2.7). "Risk Tolerance" refers to the number of boxes collected in the BRET and the number of risky choices in the MPL. To make the two tasks comparable, we normalize

Figure 2.4: Search Duration vs MPL (all 20 rounds).

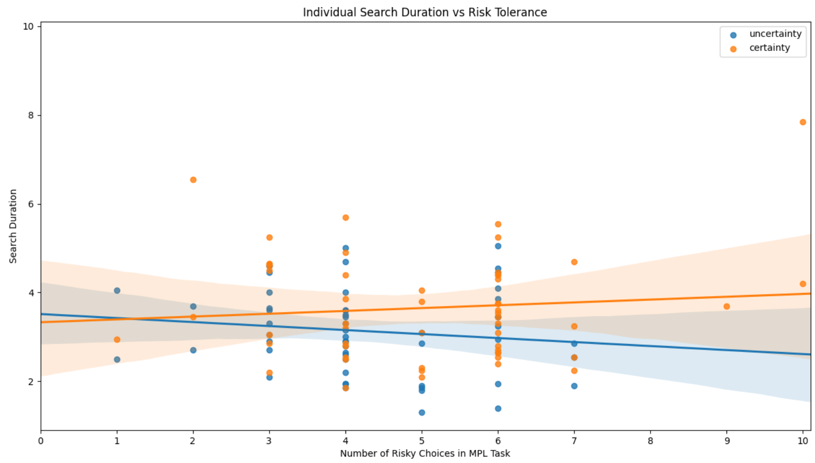

(Note: The scatter plot shows how search duration changes with the level of risk tolerance, when using the average behavior of 20 paid rounds and number of risky choices in MPL. The dots represent the relation for each subject, and the lines are estimated by OLS regression given these subject-level data (the shaded band refers to the 95% confidence interval of the estimated parameters).)

Figure 2.5: Search Duration vs BRET (last 10 rounds).

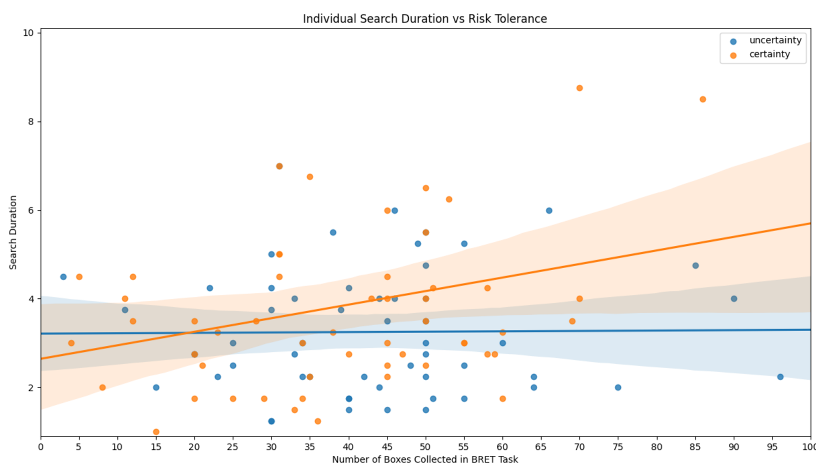

(Note: The scatter plot shows how search duration changes with the level of risk tolerance when using the average behavior in the last 10 paid rounds and number of boxes collected in BRET.)

Figure 2.6: Search Duration vs MPL (last 10 rounds).

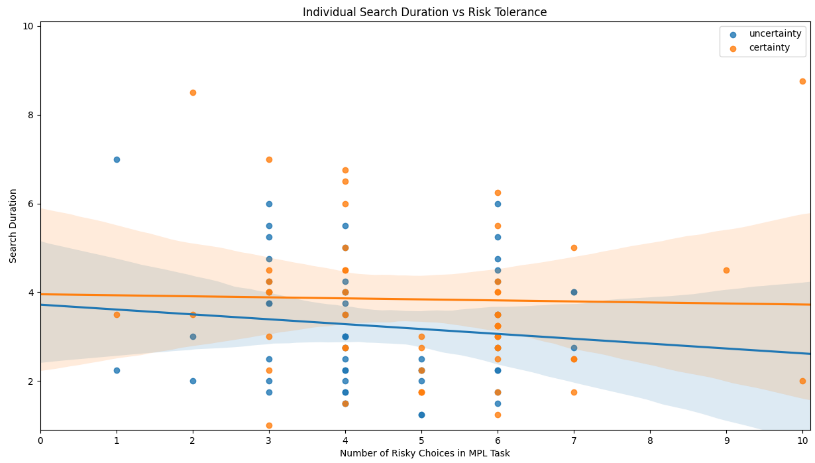

(Note: The scatter plot shows how search duration changes with the level of risk tolerance when using the average behavior in the last 10 paid rounds and number of risky choices in MPL.)

the choices to $[0,1]$ and set the risk-neutral point (50 boxes in BRET and 6 risky choices in MPL) at 0.5. The higher the risk tolerance is, the higher the level of risk seeking. “UC” is an indicator dummy for the subjects in the UC treatment. The dependent variable “Duration” is defined as the average search duration of subjects. In Tables 2.4 and 2.5, columns (1) and (2) apply BRET, columns (3) and (4) use MPL, columns (1) and (3) cover the all 20 paid rounds, and columns (2) and (4) only count the last 10 paid rounds. Based on our hypotheses 1 - 3, we should have $\beta_1 > 0$, $\beta_1 + \beta_2 < 0$, and $0.5\beta_2 + \beta_3 = 0$, respectively.

Overall, the signs of the regression outcomes satisfy hypotheses 1 and 2: the data reveal a positive β_1 and a negative $\beta_1 + \beta_2$ in all conditions. However, the significance is not always strong. Hypothesis 1 is rejected when using the BRET data, and hypothesis

2 is rejected when using the MPL data. The main reason for the insignificant results is that the slope of the lines in Figure 2.2 is relatively flat, and it is therefore difficult to detect significance in the data. Nonetheless, the experimental data already show a larger difference between the two treatments than the difference in the theoretical predictions.

$$\text{Duration} = \alpha + \beta_1 \text{ Risk Tolerance} + \beta_2 \text{ Risk Tolerance} \times \text{UC} + \beta_3 \text{ UC} + \varepsilon \tag{2.7}$$

Table 2.4: Regression summary.

	Dependent Variable: Average Search Duration			
	BRET	BRET	MPL	MPL
	1-20 rounds	11-20 rounds	1-20 rounds	11-20 rounds
Risk Tolerance	2.173**	2.241*	0.917	1.225
	(0.835)	(1.189)	(0.858)	(1.196)
Risk Tolerance x UC	-2.815**	-3.043*	-2.707*	-4.439**
	(1.172)	(1.669)	(1.507)	(2.101)
UC	0.583	1.018	0.498	1.484*
	(0.527)	(0.750)	(0.619)	(0.864)
const	2.965***	2.878***	3.417***	3.223***
	(0.359)	(0.512)	(0.403)	(0.561)
Observations	109	109	109	109
R^2	0.118	0.041	0.084	0.045
Note:		*p<0.1; **p<0.05; ***p<0.01.		

Table 2.5: Tests of regression coefficients.

	BRET		BRET		MPL		MPL	
	1-20 rounds		11-20 rounds		1-20 rounds		11-20 rounds	
H_0	t-stat	p-value	t-stat	p-value	t-stat	p-value	t-stat	p-value
$\beta_1 > 0$	2.603	0.005***	1.885	0.031**	1.069	0.144	1.024	0.154
$\beta_1 + \beta_2 < 0$	-0.781	0.218	-0.684	0.248	-1.445	0.076*	-1.860	0.033**
$0.5\beta_2 + \beta_3 = 0$	-3.484	0.001***	-1.826	0.071*	-3.031	0.003***	-2.131	0.035**

Result 1. *Search duration increases with the level of risk tolerance in the C treatment in all conditions, but the hypothesis is weak when using the BRET data.*

Result 2. *Search duration decreases with the level of risk tolerance in the UC treatment in all conditions, but the hypothesis is weak when using the MPL data.*

Similarly to what we observe from the figures, hypothesis 3 is rejected in all conditions. The current intersection deviates from the theoretical prediction toward the risk-averse side, either because players systematically search more in the C treatment or search less in the UC treatment.

Result 3. *Hypothesis 3 is rejected in all conditions. Risk-neutral subjects play differently in the two treatments.*

Although the experiment is not designed to explain the rejection of hypothesis 3, we list three possible explanations for further exploration. First, the downward-shifting blue lines show that people are more risk averse in the search tasks than in the risk elicitation tasks (Friedman, Habib, et al. (2022)). There is a spatial difference between BRET and the search task, and the subjects become more risk averse when seeing lotteries as numbers instead of images. Comparing MPL and the search task, we find that there could be a cardinality issue because the items appear sequentially in random order in the search task but are monotonic in MPL. Second, following the first explanation, the risk aversion in the UC treatment could also result from aversion to compound lotteries. In the UC treatment, the aggregated lotteries for the next search can be considered compound lotteries, and thus, subjects may behave more risk aversely in the search tasks in the UC treatment (see Abdellaoui, Klibanoff, and Placido (2011) for an overview of aversion to compound lotteries) than under the C treatment. Third, the subjects might subjectively weight the probabilities of the lotteries in the UC treatment, following either prospect theory or salience theory. The subjective weighting function could also bias the theoretical predictions.

2.5.3 Reservation strategies

Do the subjects follow the constant reservation probability strategy? Figure 2.7 shows the quantiles and the mean of the chosen outcome by search duration. Note that we transfer the chosen value to the equivalent probability in the C treatment. In this section, we take each subject's decision as a single observation instead of averaging them by subject. Overall, the subjects' decisions are highly consistent over search durations.

We do not observe any increasing or decreasing trend between the chosen probabilities and the search duration, which supports the notion that subjects have consistent reservation probability strategies. The fluctuation on the right side of the figure is due mainly to the small sample size with long search duration. The result is also strengthened by the subjects' recall rate, which is the fraction of time in which the subjects did not choose their last searched result. When the subjects apply the reservation strategy, they will not recall a previously searched outcome but will immediately stop searching once the searched outcome exceeds the reservation probability. The recall is 13.64% in the UC treatment and 23.06% in the C treatment, which supports the notion that a large fraction of decisions are motivated by reservation strategies.

Figure 2.7: Individual chosen probabilities at each search duration.

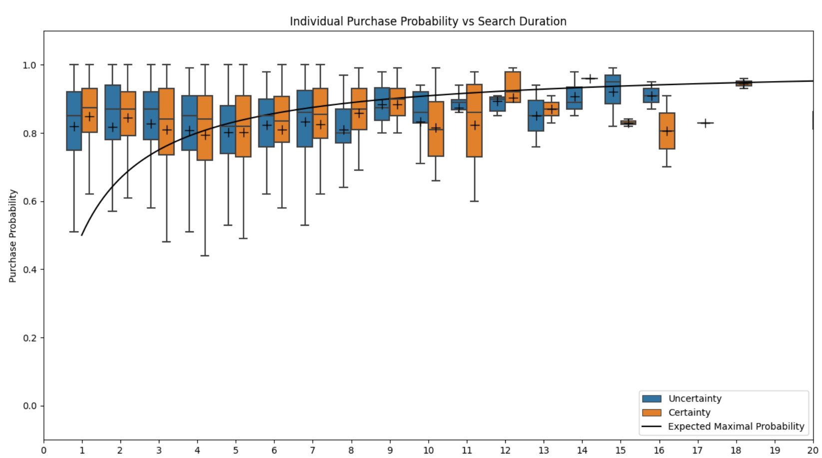

(Note: The boundaries of the boxes refer to the quantiles of the data given the search duration q_1 and q_3. Define $IQR = q_3 - q_1$. The boundaries of the extended lines refer to $q_1 - 1.5IQR$ and $q_3 + 1.5IQR$. The $+$ sign refers to the mean, and the horizontal line inside the box refers to the median.)

Conjecture 1. *In general, players follow the constant reservation probability strategies described in Section 3.*

2.6 Conclusion

Do risk-averse people search more or less? In traditional sequential search models, risk-averse players consistently search less than risk-neutral and risk-seeking players. However, consumers could also extend their search in experience goods and financial markets, where they face ex post uncertainty and receive uncertain outcomes after the search. We have shown and proven the results with a modified sequential search model and a laboratory experiment.

Intuitively, both the present best outcome and potential future outcomes are uncertain in our new scenarios. The risk from the current best outcome may exceed the aggregated risk from all the possible future outcomes and motivate risk-averse people to extend their search and reduce their risk. This paper adds new insights into both theoretical models and laboratory experiments of sequential search problems with non-risk-neutral players. Furthermore, our findings provide insights for empirical problems, such as advertising and recommendation mechanism design, in related markets. If people search longer in these markets, firms could adopt different strategies, and the market equilibrium, including the salience and prominence of the top advertising positions, might be substantially different between the experience goods markets and the commodity goods market. Our results can be applied to various types of mechanism design problems.

Chapter 3

Informational Entry Barrier and Fractional Searching

3.1 Introduction

The fast developing online marketplaces is becoming a powerful tool to help traders exchange goods across a variety of industries. Platforms such as eBay and Amazon lower the traders' search cost and have already become worldwide popular. Among the mechanisms that improve the market efficiency, reputation system is a key component. Reputation is generated from participants' transaction history and thus makes participants' previous actions observable. On the one hand, reputation system can greatly mitigate problems caused by asymmetric information and increase market efficiency.

Over the past years almost every online marketplaces have built up their own feedback and rating system. On the other hand, current reputation systems suffer from problems. We are interested in the question whether the history of incumbent firms create entry barriers for new entrants and propose a fractional searching method to solve the problem.

Consider a recently-developed product that is listed on Amazon. Initially, its quality is unknown. What is worse is that a similar product with hundreds of reasonable reviews on Amazon might create an entry barrier for this new product to enter the market even when the new product is more cost efficient or with higher quality. The intuition is straightforward as consumers will more likely choose a product from a supplier who has already established a credible reputation. The new entrant, who has not yet had a chance to build it, is assumed to be of average quality only. Meanwhile, there is another argument that the incumbent has a good review and thus is likely to provide the high quality good, leading to the market efficiency. However, this neglects the possibly better quality the new entrants can provide for consumers, which would not have a chance to be proven with the entry deterrence. In order to mitigate such problems, we propose the fractional searching mechanism to promote the product from new entrants so that they would be able to establish good reputation for their superior quality. By fractional searching, we mean to list only a fraction of sellers to consumers' search so as to increase the relative exposure of new entrants.

This project studies the phenomenon of entry barrier under reputation system in the laboratory environment. First, based on the previous studies on optimal recommendation policy, this project adds to the possible solutions to "cold start" problem. Promoting the early exploration for users to discover potentially valuable products affects the market efficiency. Second, the reputation barrier affects firms when they enter the market sequentially and the experiment simulates this by letting subjects enter the market at two separate timings. Furthermore, the paper aims at studying the mechanisms that help firms with no or low reputation profiles to overcome the entry and reputation barrier. This paper has two major contributions to the literature. To my best knowledge, it is the first paper that discusses in general firms' entry and exit decisions and entry barrier in a competitive market in the laboratory environment. It is also the first paper that discusses potential policy impact on such markets.

We first established a game theoretical model to consider the entrant's entry problem when there are incumbents in a market with reputation system. We choose oligopoly model to serve as the baseline where the incumbent and the entrant simultaneously set their quality and then price. To account for the reputation system, we consider a sequential move game where entrants join the market occupied by incumbents and their quality information is unknown to consumers initially. In the first stage, entrants choose the quality after observing the incumbent's quality decision. Then in the next stage, they simultaneously set their prices to compete against each other. In this setting, we study two conditions in which incumbents and entrants interact with each other. In the first condition with complete searching, entrants and incumbents simply compete as described above. In the second condition with fractional searching, entrants are provided an opportunity of not competing against the incumbent in the first place. Instead, they are separated into a fraction of the entire market and get an opportunity to produce better quality. Under both conditions, the entrant's quality is revealed as public information to consumers if they can get a positive demand in the second stage. We derived the Nash equilibrium under the above conditions and further examine the theory in a laboratory experiment with a between-subjects design. Laboratory subjects were in either oligopoly, complete searching or fractional searching environment and took a role of incumbent or entrant randomly assigned at the beginning of each round. The results show that most subjects follow theoretical predictions and fractional searching effectively alleviates the informational barrier to entry resulted from the reputation system.

The rest of the paper is organized as follows. Section 3.2 provides a discussion of the related literature on reputation system and entry barrier problems. Section 3.3 introduces an adapted model to examine the informational entry barrier. Section 3.4 introduces the experimental design and treatment definition. Section 3.5 presents the experimental results. Finally, section 3.6 concludes with the main findings and presents a brief discussion about future research.

3.2 Literature

A large and growing literature has conclusively demonstrated the importance of reputation within online marketplaces and empirically study the mechanism (e.g. Resnick et al. (2000); Einav, Farronato, and Levin (2016)). Though designed differently, reputation systems have the same aims at reducing adverse selection and moral hazard problems and empirical evidence seems to demonstrate it to have worked well (e.g. Houser and Wooders (2006)). However, the systems suffer from problems such as manipulation (Mayzlin, Dover, and Chevalier (2014)), inflation (Horton and Golden 2015), and under-provision (Fradkin et al. (2015)). Moreover, well-constructed reputation could become an entry barrier of new merchants, which prevents the marketplaces from growing (e.g. Bagwell (1990)).

This project focuses on the entry barrier caused by the reputation system and informational uncertainty in the market entry problems and propose a fractional searching method to solve the problem. The project is related to laboratory and theoretical studies on reputation system as it is the backgroud of the project. Lab studies have found that both direct reputation (fixed matching) and indirect reputation (observable past) improve the efficiency of the trust game, which is the most frequently used game for studying reputation system (Bohnet and Huck (2004); G. E. Bolton, Katok, and Ockenfels (2005)). While the majority of paper studying related mechanisms of reputation (e.g. Charness, Du, and Yang (2011); G. Bolton, Greiner, and Ockenfels (2013); L. Li and Xiao (2014)) and various signalling mechanisms (Mailath and Samuelson (2001); Board and Meyer-ter-Vehn (2013)), few papers put the reputation system in the market and early market-related studies focused more on adverse selection and signalling (Miller and Plott (1985); Lynch et al. (1986)). Hörner (2002) first introduces the reputation system to a competitive market with firms choosing effort and timing of exit. The study shows how competition mitigates the inefficiencies of moral hazard with the existence of high-effort equilibria. Dana and Fong (2011) also shows that equilibrium exhibiting high quality may exist in oligopoly markets even when low-quality is the unique equilibrium outcome in monopoly and competitive markets. Huck, Lünser, and Tyran (2012) and Huck, Lünser, and Tyran (2016) run lab experiments on the efficiency of reputation system. They find that both reputation and competition improve market efficiency but adding the price competition leads the market to the opposite inefficient direction.

However, literature in reputation system focus on long-run interactions but our project studies the short-run entry problem, which is related to the literature of market entry. Since Heflebower (1957), a major direction of researches on market entry studies various ways of entry barrier and entry deterrence strategies. Informational barrier caused by consumers' uncertainty of entrants' product quality is one of the main types of barriers. Schmalensee (1982) highlights how superior information regarding product quality can endow incumbent firms with a first-mover advantage, leading to inefficient entry choices, even when price competition is added. His finding is further confirmed and developed by Conrad (1983) with a dominant firm price leadership model and by Bagwell (1990) with inefficient incumbents. Farrell (1986) models moral hazard as entry barriers from another prospective when producers have incentive to provide low quality and quit the market. Consumers foresee this possibility and do not purchase from new entrants. Recent papers from Rouviere and Soubeyran (2008), Jeon and Lovo (2012), and Atkeson, Hellwig, and Ordoñez (2014) study reputation barrier respectively with collective reputation, with credit rating agencies, and with entry taxes. In the marketing literature, similar idea has also been discussed as brand names and several papers explained the phenomenon from consumers' prospective by risk aversion (Krouse (1984)) and brand loyalty (Villas-Boas (2004)). Lab studies on market entry problems mainly focus on price deterrence (Cooper, Garvin, and Kagel (1997); Müller and Götz (2017)) and capacity deterrence (Mason and Nowell (1992); Brandts, Pezanis-Christou, and Schram (2008)). To our best knowledge, there is no lab experiment that studies the informational entry barrier.

The entry barrier we aim to solve can be considered as a form of "cold start" problem in computer science context, which has also been highlighted in recent years in mechanism design studies. There is a great presence of entry barriers in web search engines which can be thought of as recommendation systems as search engines repeatedly return currently popular pages and those newly-created but high-quality pages are effectively shut out since users usually focus on the top few results (Cho and Roy (2004); Cho, Roy, and Adams (2005)). They propose a new ranking metric, named page quality, and it effectively alleviates the information imbalance problem and identifies high-quality pages much easier. Pandey et al. (2005) also identifies the similar entrenchment problems suffered by search engines and proposes a randomized rank promotion scheme to offer new pages a chance to prove their worth, in which the promoted pages are assigned randomly-chosen rank positions. They show that a modest amount of randomness leads to improved search results, as is measured by the aggregate results quality amortized over time. While the above studies focus on the web search engines, Einav, Kuchler, et al. (2015) and Fradkin (2015) explore the inefficiencies of online marketplaces caused by that consumers cannot consider all options available, and demonstrate better ranking algorithms can improve transaction probabilities and generate large gains in volume, revenue and consumer surplus. Che and Hörner (2017) and I. Kremer, Mansour, and Perry (2014) examine the intertemporal informational externality that consumer choices generate, and thus also identify policies that are consistent with Pandey et al. (2005). They treat the range of products as exogenous, and thus abstract from firms' incentives. Dinerstein et al. (2018) specifically examines the platform's role in guiding search and provides evidence for a search redesign by prioritizing product quality to achieve higher efficiency. Vellodi (2018) studies the role of information design in shaping industry dynamics through endogenous participation and incorporates the social learning feature of information diffusion. The optimal rating design he proposes involves the upper censorship to stimulate participation.

While our project is related to previous studies, we also make contribution to the literature. First, instead of studying the long-run reputation effect, we take it as the backgroud of our project and focus on the short-run entry problem when new entrants enter the market with no history. Second, we build a theoretical framework where both incumbents and entrants in our model compete by price but their type of quality is also endogenously determined. Third, we first introduces fractional searching mechanism to solve reputation problems. Fourth, none of the above topics has been studied in the laboratory environment and we believe we are the first project to study reputation barrier and introduce related market mechanisms in the lab.

3.3 Model

3.3.1 Market Environment

To show the existence of reputation barrier and the efficiency of fractional searching, we construct a model of infinitely repeated market with two long-lived firms: incumbent I and entrant E, and a continuum of short-lived consumers in each period. Firms compete with price and quality of their products and consumers pick the firm that provides them with the highest expected utility.

The game is played as follows. Before the repeated game starts, incumbent and entrant sequentially decide the quality of their products. Then firms start a Bertrand competition for an infinite number of periods while keeping their quality consistent during the game. Incumbent moves first so her quality is observed by entrant when entrant chooses her quality. At the beginning of the repeated interaction, consumers only observe incumbents' quality. The payoff for both firms is their inter-temporal discounted total payoff in the repeated interaction.

Firms choose to produce their products with quality $q_i \in \{L, H\}$ and sell it on the market with price p_{it} in each period, where $i = I, E$ and $t = 1, 2, ..., \infty$. We normalize the cost of producing L for both firms to 0. Incumbent produces H with cost c_I and entrant produces H with c_E. Firms share the same inter-temporal discount factor δ and maximize the inter-temporal payoff $\sum_{t=1}^{\infty}(p_{it} - c_i)d_{it}$, where $d_{it} \in [0, 1]$ is the demand for firm i at time t. The continuum of consumers live uniformly on $[0, 1]$ in each period. Each consumer chooses to purchase a unit of product from the firm which gives them the highest non-negative expected utility in each period. Consumers receive u_H for H and u_L for L. Consumers' expected utility for firm i at time t can be represented as $Eu_{it} - p_{it}$. For firms and consumers, we make the following assumptions.

Assumption 1. *Producing H is more beneficial to both consumers and the social welfare but more costly to the firm than producing L. We have*

(1) $c_I > 0, c_E > 0$*;*

(2) $u_H > u_L$*;*

(3) $u_H - c_I \geq u_L$*, and* $u_H - c_E \geq u_L$*.*

Assumption 2. *Predatory pricing is not allowed:* $\forall t$*,* $p_{it} \geq 0$ *if firm i chooses L and* $p_{it} \geq c_i$ *if firm i chooses H.*

Assumption 3. *Consumers* Eu_{it} *changes when firms' quality is revealed. If firms produced H,* $Eu = u_H$*. If firms produced L,* $Eu = u_L$*. For firms with no history,* $Eu = \pi u_H + (1-\pi)u_L$*, where parameter* $\pi \in [0,1]$ *and* π *is exogenously determined.*

Assumption 4. *Incumbent's quality can be directly observed by the consumers from* $t = 1$*. At period t, entrant's quality is revealed to consumers if entrant received positive demand in any period before* t*.*

Assumption 5. *Both firms and consumers prefer H to L at the decision boundaries:*

(1) Consumers will choose the firm with a higher probability of producing H when firms provide the same expected utility;

(2) Firms will choose to produce H when producing H and L earn the same inter-temporal total payoff.

Assumption 6. *Firms know all exogenous parameters* $\{u_H, u_L, \pi, c_I, c_E, \delta\}$ *while consumers only know their own parameters* $\{u_H, u_L, \pi\}$*. Both firms and consumers have access to the full market history.*

3.3.2 Baseline Model with Simultaneous Entry

The market environments reveals some important principles. First, as both firms know all the market information, consumers' expected utility can be perfectly predicted before firms choose the price in each period. Second, the quality of both firms is determined before repeated game starts, so both firms know the other firm's quality when they are playing the Bertrand competition. Third, given the consumers' expectation and the quality of the firms, both firms will try to set the price such that the expected utility they provide is slightly higher than their opponents', which is exactly what Bertrand competition predicts. With the principles above, the infinitely repeated game can be seen as a 2-stage extensive-form super-game where incumbent and entrant sequentially choose their quality. The inter-temporal discounted payoff they get from their quality decisions can be directly calculated based on the principles. We follow the base model of Bertrand competition and rule out the possibility of tacit collusion in the infinitely repeated interaction.

When either incumbent or entrant plays as the monopolist of the market, they produce H as it is more profitable than producing L when they can extract all market surplus. They will also set the price of their products equal to consumers' expected utility. If both firms enter the market simultaneously, they both serves as incumbents and consumers have access to both of their quality in period 1. Let ε be defined as the minimum unit of price. Similar to the general prediction of Bertrand competition, the

Nash equilibrium (q_I^*, q_E^*) can be display as follows.

Proposition 1. *When both firms enter the market simultaneously,* $(q_I^*, q_E^*) = (H, H)$. *Whoever with a lower cost sets the price slightly lower than the other firm's cost and takes the entire market.*

3.3.3 Sequential Entry and Reputation Barrier

What if incumbent and entrant enter the market sequentially? The question focuses on the short-term interaction when entrants join the market occupied by incumbents in the infinitely repeated game. As described in the previous subsection, the infinitely repeated game can be seen as a 2-stage extensive-form super-game where incumbent and entrant sequentially choose their quality and price and payoff are automatically optimized by Bertrand competition. The game can be solved by backward induction and the equilibrium is conditional on the exogenous parameters.

Proposition 2. *When both firms enter the market sequentially,*

1) *If* $c_I \leq (1-\pi)(u_H - u_L)$, *incumbent sets up a strong entry barrier by producing H and take the market.* $(q_I^*, q_E^*) = (H, H)$.

2) *Entrant will choose the fly-by-night strategy by producing L when*

   $c_I > (1-\pi)(u_H - u_L)$ *and* $c_I - c_E \leq (1-\pi)(u_H - u_L)$, *or*

   $c_I - c_E > (1-\pi)(u_H - u_L)$ *and* $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - c_E < 0$.

$(q_I^*, q_E^*) = (H, L)$. *Consumers choose entrant at* $t = 1$ *and go back to incumbent when* $t > 1$.

3) If $c_I - c_E > (1-\pi)(u_H - u_L)$ *and* $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - c_E \geq 0$, $(q_I^*, q_E^*) = (H, H)$. *Entrant successfully enters and takes the market.*

Proof. See Appendix C.1. □

3.3.4 Sequential Entry with Fractional Searching

Fractional searching provides an opportunity for entrants to producing H without competing with incumbents. Assume that τ of consumers can only search for and trade with entrant while the other $1-\tau$ of consumers can only trade with incumbent at $t = 1$. The restriction is removed from $t = 2$ and consumers share the product information at $t = 1$. To set up the environment for fractional searching, some additional assumptions are made below.

Assumption 7. *(1) Market can be perfectly separated at* $t = 1$*;*

(2) Information can be perfectly shared from $t = 2$.

Releasing both assumptions affect the initial market power of both firms and the spread of information, which in fact affect τ when both assumptions are true. As a result, we can follow these two assumptions and simplify the environment with only one parameter and no uncertainty.

Under fractional searching, entrant has the chance to build up her reputation at $t = 1$ and earn a positive payoff. Meanwhile, entrant earns less fly-by-night payoff if they produce L when $\tau < 1$ because she faces a smaller market at $t = 1$. As a result, it is more likely that entrant will produce H compared to the market with no fractional searching. As consumers share the same reputation information, entrant's quality is revealed and consumers' Eu becomes u_H after $t > 1$. Cost-efficient entrant is able to compete with incumbent and take the market. However, it is also possible that entrant chooses fly-by-night strategy when the incentive of producing H is weak or when they are not allowed to play H.

Proposition 3. *When both firms enter the market sequentially, let $\tau \in (0,1)$ be the fractional searching parameter.*

1) *Entrant will choose the fly-by-night strategy by producing L and consumers choose incumbent after $t > 1$ if either one of the conditions is satisfied:*

   $c_I \leq c_E$, $\pi u_H + (1-\pi)u_L < c_E$, *or* $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - \tau c_E < 0$.

   In these conditions, $(q_I^*, q_E^*) = (H, L)$.

2) *If* $c_I > c_E$, $\pi u_H + (1-\pi)u_L \geq c_E$, *and* $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - \tau c_E \geq 0$, $(q_I^*, q_E^*) = (H, H)$. *Entrant will produce H and successfully enter and take the market.*

3) *It is easier for entrant to enter the market with H when the market allows fractional searching.*

Proof. See Appendix C.2. □

3.4 Experimental Design

Based on the model, we design a game to test the information barrier and a simplified fractional searching algorithm. The game is a 2-player game with 1 incumbent and 1 entrant. Consumers are played by automated bots in this experiment as we assume that consumers understand the idea of reputation system and treat firms in the same way. Bots are set to play the same strategy as defined in the model.

The game mainly focuses on the short-term impact of information barrier and entry problems. The game we use in the experiment only applies the first two periods of the repeated interaction as only the first two periods of price competition in the model affect players' total payoff. The simplified extensive-form game has 3 stages. At stage 1, incumbent and entrant sequentially enter the market and choose their product quality. At stage 2 and 3, both players begin price competitions after observing each other's quality and consumers choose whose product to purchase based on the predefined algorithm. Similar to the model, consumers' belief of player B's product quality (phrased as type in the experiment) may change at stage 3. The set of parameters $(u_H, u_L, \pi, c_I, c_E, \tau)$ are predetermined and are public information for all firms but consumers only learn their own information as assumed in the model. In the experiment, we will rename incumbent and entrant as "player A" and "player B". We will also use product type instead of quality, where firms produce either X(L) or Y(H). However, the instruction is still under a market environment as it is easier for subjects to understand the game without causing additional bias.

The experiment has one pair of main treatment. In the complete searching treatment (CC), consumers can find all firms which are currently in the market at stage 2 and 3, and compare the expected utility they provide. In the fractional searching treatment (FC), τ of the consumers can only find entrants at stage 2 when they start the automated repeated interaction (τ is public information to firms). We also add one control group. In oligopoly market (OL), all firms join the market at the beginning of the game and consumers know the quality of both firms' products at the beginning of stage 2. The game is the same as a Bertrand competition where both firms play as the incumbents.

Each experimental session consists of 2 practice periods and 15 paid periods with randomly matching between periods. Three treatment and control groups use between-subject design.

3.4.1 Treatment Variables

The major treatment variable is the comparison between complete searching and fractional searching. The OL treatment serves as the baseline and the stage 1 of the game is shown in Figure 3.1.

In the CC treatment, the incumbent and the entrant sequentially enter the market.

Figure 3.1: Game Interface (OL)

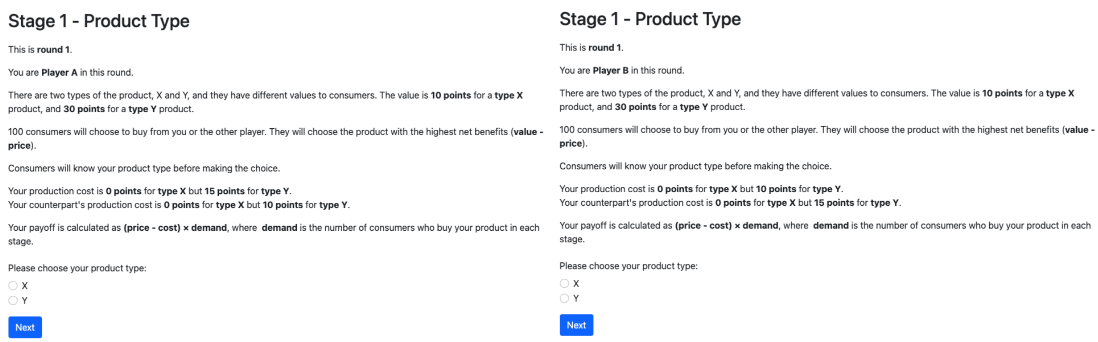

(i) Incumbent Quality Page (ii) Entrant Quality page

All consumers can choose to buy from either of them. At stage 1, the incumbent makes quality decision and the entrants observes that when making its quality decision. At stages 2 and 3, both players begin price competition after observing each other's quality. Figure 3.3 shows a screenshot of the game interface for the incumbent and the entrant in FC Treatment.

Figure 3.2: Game Interface (CC)

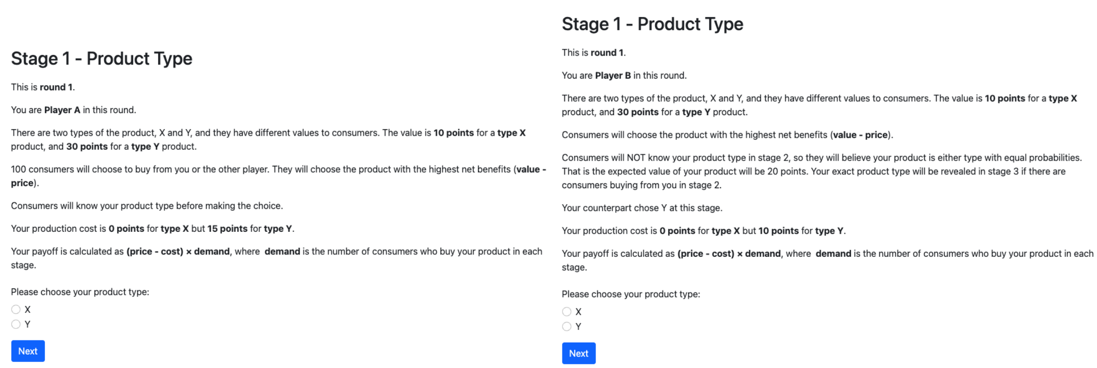

(i) Incumbent Quality Page (ii) Entrant Quality page

In the FC treatment, the incumbent and the entrant sequentially enter the market. However, a fraction of consumers can only choose to buy from the entrant whereas the

Figure 3.3: Game Interface (FC)

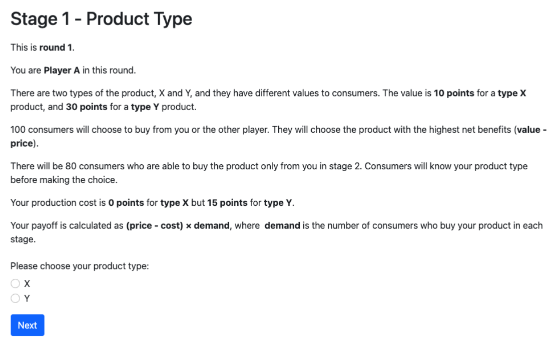

(i) Incumbent Quality Page

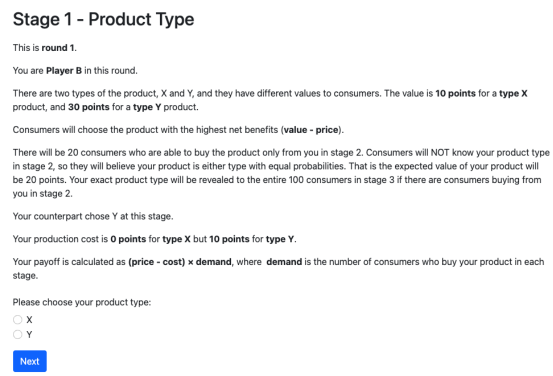

(ii) Entrant Quality page

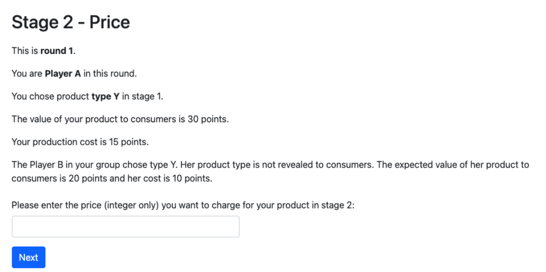

(iii) Incumbent Price Page

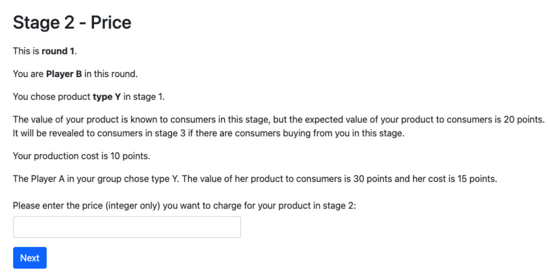

(iv) Entrant Price page

rest can only choose to buy from the incumbent. The same process follows as in the CC treatment. Figure 3.2 shows a screenshot of the game interface for the incumbent and the entrant in CC Treatment.

3.4.2 Hypotheses

In our experiment, we use two sets of parameters $(u_H, u_L, \pi, c_I, c_E, \tau) \in \{(30, 10, 0.5, 15, 10, 20), (30, 10, 0.5, 20, 7, 20)\}$. We expect to see that the behavior of most subjects in the experiment matches what our solutions in Section 3.3 predict. The details can be summarized in the following hypotheses.

Hypothesis 4. *In line with Proposition 1, under oligopoly, the incumbent and the entrant enter the market simultaneously. We expect to see that most entrant subjects in the experiment choose the high quality and set the price below the incumbent subjects' cost of producing high quality in both parameter settings.*

Hypothesis 5. *In line with Proposition 2, when the incumbent and the entrant enter the market sequentially, we expect to see the following. Under parameter set (1), we expect most entrant subjects to choose the low quality and set a price to acquire a positive demand where as most incumbent subjects to choose the high quality and set a price at their marginal cost in stage 2. Then most incumbent subjects will set a price at their marginal cost whereas most incumbents will set a price to acquire the entire market in stage 3. In comparison, under parameter set (2), we expect most entrant subjects to choose the high quality and set a price to acquire a positive demand whereas most incumbent subjects to choose the high quality and set a price at their marginal cost in stage 2. Then most incumbent subjects will set the same price whereas most entrant subjects will set a price lower than the incumbents' marginal cost in stage 3.*

Our primary research question is whether market entry problem can be alleviated by allowing fractional searching. Our model predicts that, when allowed, market entry can be easier for cost-efficient entrants and their reputation can be built up by offering high quality. This is summarized in the following hypotheses.

Hypothesis 6. *In line with Proposition 3, with fractional searching, we expect most* *entrant subjects to choose high quality and set a price to acquire a positive demand in their separate market whereas most incumbent subjects to do the same in stage 2. Then, in stage 3, we expect most entrant subjects to compete with the incumbent subjects by setting a price to acquire the entire market. The same conjecture holds in both parameter settings.*

Hypothesis 7. *With fractional searching, it is easier for entrant subjects to enter the market by producing High quality and easier for the market to reach Nash equilibrium.*

3.4.3 Session Information

Ten sessions were held between August 2020 and November 2020. In total, 102 subjects participated in our sessions. The subjected were paid the sum of all fifteen periods' payoffs, plus their show-up fee of four dollars. On average, subjects earned fourteen dollars during a forty-minute session. The experiments were developed on oTree **otree** and the subjects were recruited on Orsee **subject** which is operated by the Learning and Experimental Economics Projects of University of California at Santa Cruz (LEEPS Lab). Given the three between-subjects groups and two sets of parameters, the session information can be summarized as follows in Table 3.1.

Table 3.1: Summary of Sessions

Parameters	Group	#Sessions	#Subjects
	CC	2	22
$(30,10,0.5,15,10,20)$	FC	2	20
	OL	1	8
	CC	2	20
$(30,10,0.5,20,7,20)$	FC	2	20
	OL	1	12

3.5 Results

3.5.1 Data Overview

We have 50 subjects participating in our experiment with parameter set 1 and 52 subjects participating with parameter set 2. The main results are robust across the two parameter sets.

Table 3.2 presents statistics of subjects' quality and price choices in all experimental sessions. The first column indicate the treatment group. The second and third columns, quality^I and quality^E, indicate the quality[1] chosen by the incumbent subject and the entrant subject, respectively. The fourth column shows the total number of observation of that (quality^I, quality^E) combination. The fifth column shows the entry rate, calculated as the number of observations in which the entrant subjects' quality was revealed as a fraction of total number of observations within each quality combination. The last

[1] High quality is denoted by 1 and Low quality is denoted by 0. Same for the following tables.

four columns avg p_t^i (for $i \in \{I, E\}$ and $t \in \{1, 2\}$) shows the mean of prices chosen by subjects i at stage t of the experiment across all observations.

Result 1. *In the OL treatment group, most entrant subjects with lower cost chose the High quality and can almost always enter the market with High quality by undercutting the incumbent's price.*

As is shown in Table 3.2, subjects were able to get to Nash equilibrium in 48 out of 60 pairs under parameter set 1 and in 85 out of 90 pairs under parameter set 2 in the OL group. According to Figure 3.4[2], most entrant subjects were able to maximize their payoff by setting a price one tick size below the incumbent's cost of producing High quality and acquiring the entire market.

Result 2. *In the CC treatment group, most entrant subjects chose Low quality and cannot enter the market with High quality under parameter set 1. While under parameter set 2, when the difference in cost of producing High quality between incumbent and entrant is relatively larger, entrant subjects successfully entered the market around 80% to 96% of the time with High quality. Both support Hypothesis 5.*

Result 3. *In the FC treatment group, most entrant subjects chose the High quality and can always successfully enter the market under both parameter sets. Hypothesis 6 is supported.*

[2]The incumbent_price$t$ is the price chosen by the incumbent subject at stage t, and entrant_price$t$ is the price chosen by the entrant subject at stage t. One vertical red line indicates the predicted price in line with propositions.

In the CC group, 111 out of 165 and 87 out of 150 pairs were able to get to the Nash equilibrium under parameter set 1 and 2, respectively. Under parameter set 1, most incumbent subjects chose High quality and set a price to block the entry of entrant subjects, whereas most entrant subjects chose the fly-by-night strategy by producing Low quality. Under parameter set 2, though most pairs both chose High quality, around a third of the incumbent subjects chose Low quality instead and a few entrant subjects chose the fly-by-night strategy. In the FC group, 96 out of 150 and 125 out of 150 pairs reached the Nash equilibrium of both producing High quality under two parameter sets, respectively. There were 13% - 19% of entrant subjects playing the fly-by-night strategy. According to the distribution of prices in the Nash equilibria as is shown in Figure 3.5[3], most subjects set prices as are predicted by the model. In both Nash equilibria in the CC group, the lower entry rates were resulted from the higher prices set by the entrant subjects at stage 1. In Table 3.2, the avg p_1^Es are 4.74 and 10.02 in the two Nash equilibria, and are slightly higher than the model predictions at 4 and 9, respectively. This is also verified by looking at the bars to the red vertical lines in entrant_price1 of CC group in Figure 3.5a and Figure 3.5b. In comparison, the entry rate in the FC group is always at 1 and entrant subjects can always successfully enter the market. Both incumbent and entrant subjects set their prices in the equilibrium range at stage 1 and set some variable prices at stage 2.

[3] Q^i is the subject i's quality choice where $i \in \{I, E\}$. Two vertical lines indicate a predicted price range in line with propositions.

Table 3.2: Summary Statistics

treatment	quality^I	quality^E	obs	entry rate	avg p_1^I	avg p_1^E	avg p_2^I	avg p_2^E
(a) *param. set 1*								
CC	1	1	40	0.00	17.58	12.42	18.40	12.38
	1	0	111	0.85	16.02	4.74	18.94	1.86
	0	1	6	0.67	4.67	11.67	3.33	12.50
	0	0	8	1.00	0.88	7.38	2.00	1.75
FC	1	1	96	1.00	27.88	18.58	17.52	16.42
	1	0	28	1.00	26.96	17.82	21.21	4.32
	0	1	22	1.00	8.27	18.05	4.91	19.95
	0	0	4	1.00	7.75	13.00	5.75	3.75
OL	1	1	48	0.92	17.85	14.62	16.50	14.77
	1	0	1	0.00	20.00	5.00	20.00	5.00
	0	1	10	1.00	3.40	16.70	5.60	18.50
	0	0	1	1.00	1.00	1.00	1.00	1.00
(b) *param. set 2*								
CC	1	1	87	0.78	20.60	10.02	20.43	16.77
	1	0	11	0.73	20.18	8.55	20.09	4.36
	0	1	48	0.96	2.00	9.62	1.69	16.19
	0	0	4	1.00	0.75	7.75	0.50	1.25
FC	1	1	125	1.00	29.07	19.17	19.54	19.07
	1	0	19	1.00	28.16	17.84	20.53	0.21
	0	1	3	1.00	10.00	20.00	0.00	20.00
	0	0	3	1.00	10.00	11.67	0.67	0.00
OL	1	1	85	0.96	20.72	18.65	20.12	18.76
	1	0	3	0.00	20.00	0.67	20.00	0.00
	0	1	1	0.00	1.00	27.00	0.00	19.00
	0	0	1	1.00	0.00	0.00	0.00	0.00

Figure 3.4: Baseline Price Distribution

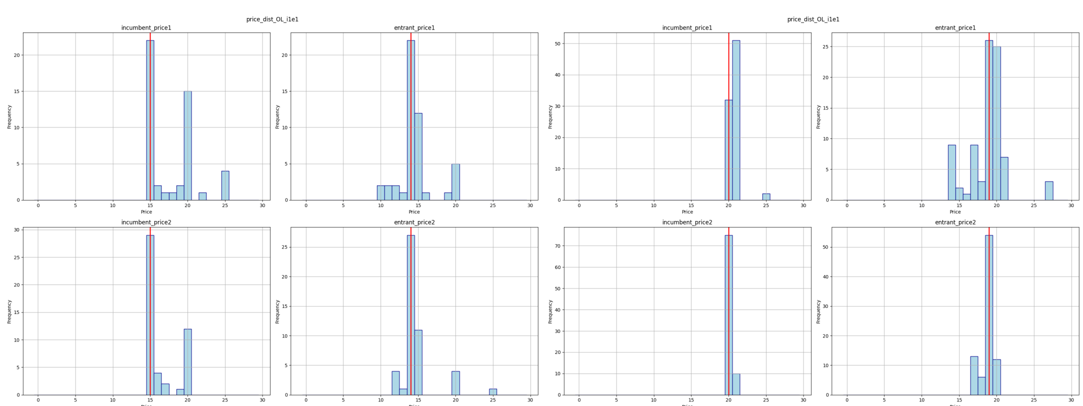

(a) param. set 1 (b) param. set 2

Figure 3.5: Price Distribution

(a) param. set 1

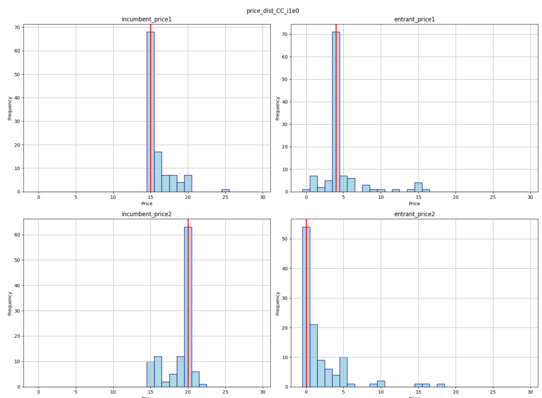

(i) $Q^I = 1, Q^E = 0$ (CC)

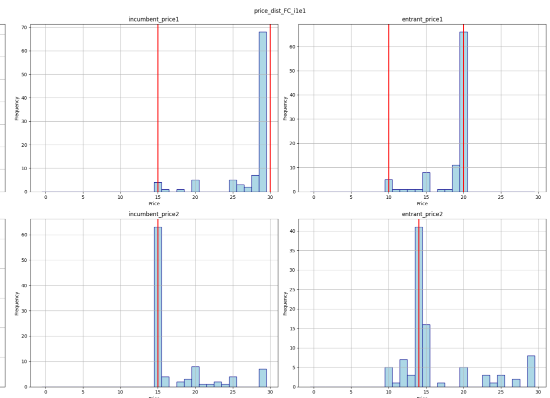

(ii) $Q^I = 1, Q^E = 1$ (FC)

(b) param. set 2

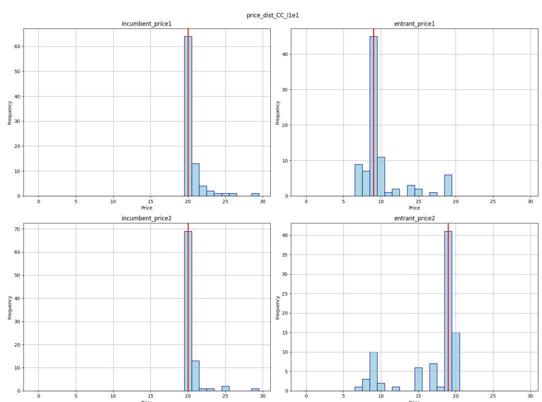

(i) $Q^I = 1, Q^E = 1$ (CC)

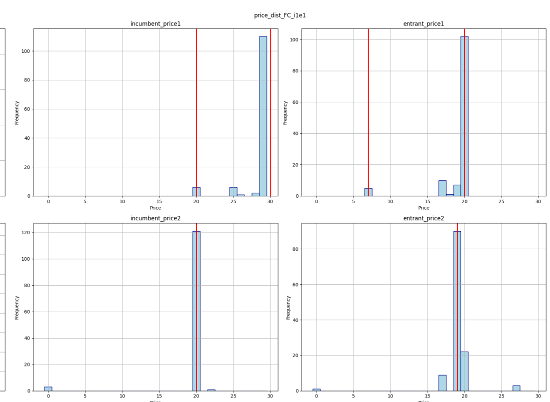

(ii) $Q^I = 1, Q^E = 1$ (FC)

3.5.2 Average Quality and Entry Rate

Figure 3.6 and Figure 3.7 show how the average quality of incumbent subjects' and entrant subjects' changes in different treatment groups. As is shown in Figure 3.6, the average quality of entrants' is converging to 1 pretty quickly during round 2 or 3. The average quality of incumbents' shows a less obvious pattern during first four rounds under parameter set 1 but converge later to 1 under both parameter sets. In Figure 3.7a, we can see a clear difference between CC and FC treatment groups. Though the average quality of incumbents' converges to 1, the average quality of entrants' converges to 0 in the CC group compared to 1 in the FC group. In Figure 3.7b, the entrants' average quality converges to 1 in both groups but the pattern of the incumbents' average quality is unclear in the CC group while also converging to 1 in the FC group. Hypotheses 4 - 6 are supported.

Table 3.3 reports the results of the estimation of the logistic regressions of incumbent's quality (columns (1) and (5)), entrant's quality (columns (2) and (6)), entry status (columns (3) and (7)), and NE outcome status (columns (4) and (8)). The standard errors are cluster-robust at the outcome level, shown in parentheses. The independent variable round is the round number in the experimental session, and FC and OL are the indicator dummies for FC and OL treatment groups. Columns (1)-(4) apply parameter set 1 and columns (5)-(8) apply parameter set 2. The coefficient estimates support most of our hypotheses. Under parameter set 1, columns (1) and (2) show that incumbent subjects' quality is not different across three treatment groups, whereas entrant subjects are more likely to producing High quality in FC and OL treatment groups, which is statistically significant at the $p = 0.01$ level. Column (3) also shows that it is easier for entrants to enter the market in the FC group at the $p = 0.01$ significance level. Under parameter set 2, column (5) shows that incumbent subjects are more likely to produce High quality in the FC and OL groups whereas the odds of entrant subjects producing High quality is lower in the FC treatment, though it is not statistically significant as is shown in column (6). Column (7) suggests it is easier for entrants to enter the market in the FC group at the $p = 0.01$ level. Lastly, columns (4) and (8) show that FC treatment group does not increase the odds of reaching Nash equilibrium but it shows some evidence that subjects are learning to converge to Nash equilibrium under parameter set 1 while OL treatment group increases the odds of achieving NE under parameter set 2, both at the $p = 0.01$ level.

Result 4. *It is easier for entrant subjects to enter the market with High quality in the FC treatment group than in the CC treatment group. However, our experiment did not show evidence that subjects converge to Nash equilibrium better in the FC treatment than in the CC treatment group. Hypothesis 7 is partially rejected.*

Figure 3.6: Baseline Average Quality

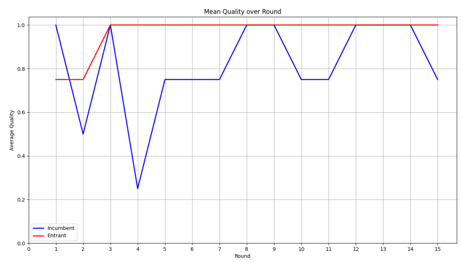

(a) param. set 1

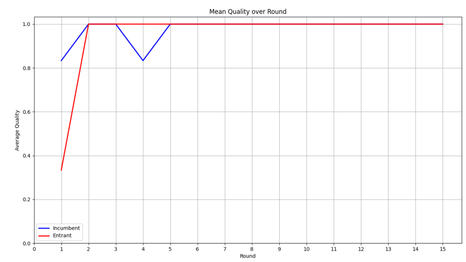

(b) param. set 2

Figure 3.7: Average Quality

(a) param. set 1

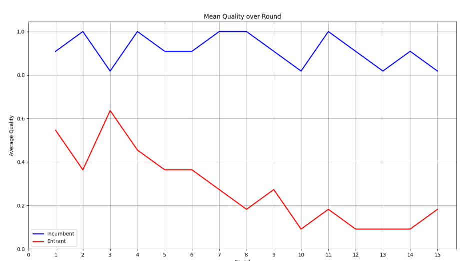

(i) CC

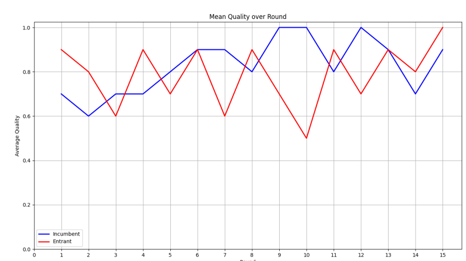

(ii) FC

(b) param. set 2

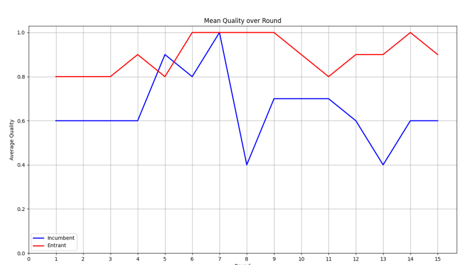

(i) CC

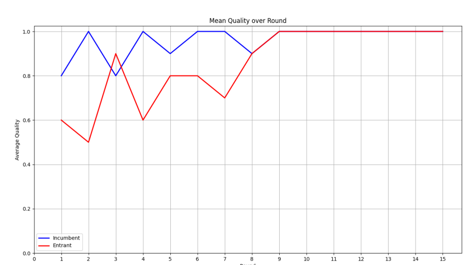

(ii) FC

Table 3.3: Logistic Regression

	parameter set 1				parameter set 2			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
round	0.058	-0.070*	0.179***	0.104***	0.022	0.252***	0.063	0.129
	(0.053)	(0.037)	(0.036)	(0.028)	(0.032)	(0.054)	(0.109)	(0.081)
FC	-0.821	2.302***	13.275***	-0.152	2.547***	-0.484	25.582***	1.364
	(1.047)	(0.124)	(1.506)	(2.697)	(0.693)	(0.557)	(1.248)	(0.980)
OL	-0.890	4.381***	1.966*	0.691	3.153***	0.927*	0.821	2.621***
	(1.374)	(0.696)	(1.107)	(2.488)	(0.701)	(0.537)	(1.108)	(0.776)
Observations	375	375	375	375	390	390	390	390
Note:								*p<0.1; **p<0.05; ***p<0.01

3.6 Conclusion

We examine problems associated with informational entry barriers when there is presence of reputation systems. We first propose a model to capture the existence of reputation systems and show how this results in inefficiencies that cost-effective entrants cannot enter the market where there are incumbents with well-established reputations. In order to mitigate such problems, we propose the fractional searching mechanism in which it allows entrants to not compete against incumbents during their initial entry into the market. This provides an opportunity for entrants to establish some reputation for their possibly superior quality. Additionally, we propose an experiment design simulating such markets and implement a simplified fractional searching algorithm. Our results suggest that the existence of such reputation systems creates an informational barrier for cost-effective entrants to enter the market but this problem can be alleviated by introducing fractional searching. On the one hand, cost-effective entrants are more likely to produce high quality when the market allows fractional searching. On the other hand, the market entry rate is significantly higher with fractional searching. This paper adds new insights into both theoretical models and empirical solutions to informational barrier problems. Furthermore, our model and experiment assumes exogenous beliefs held by consumers when entrant's quality is unknown, it is worth exploring how fractional searching could affect the market if beliefs are endogenous instead.

Appendix A

Supplement to Chapter One

A.1 Summary statistics for profits

Figure A.1: Profit Distribution

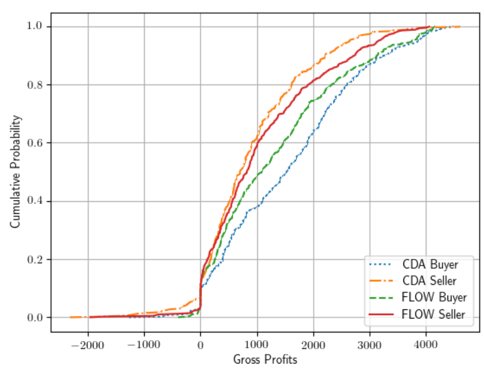

(i) Gross CDF

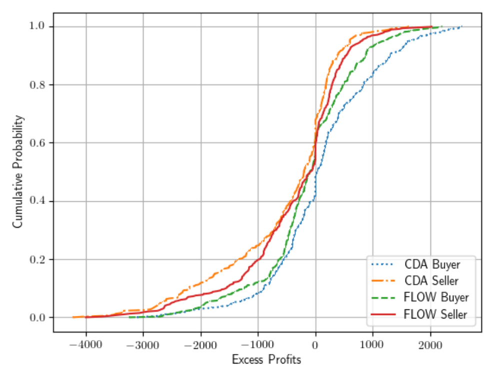

(ii) Excess CDF

Table A.1: Summary of Profits

	T1 - T20		T1 - T10		T11 - T20	
	CDA	FLOW	CDA	FLOW	CDA	FLOW
(a) *Gross Profits*						
Overall	1223	1184	1158	1120	1288	1249
	(1128)	(1106)	(1090)	(1082)	(1161)	(1127)
Sellers	891	1053	936	1039	845	1067
	(936)	(1036)	(997)	(992)	(871)	(1080)
Buyers	1555	1316	1380	1200	1730	1431
	(1203)	(1159)	(1136)	(1162)	(1246)	(1146)
(b) *Excess Profits*						
Overall	-227	-266	-292	-330	-162	-201
	(1031)	(910)	(1043)	(975)	(1016)	(836)
Sellers	-544	-382	-499	-396	-590	-368
	(1028)	(951)	(1069)	(1066)	(985)	(822)
Buyers	90	-149	-85	-265	265	-34
	(933)	(853)	(976)	(871)	(854)	(819)

Note: Gross profit is the average of individual trader's end-of-period profit across 5 groups over all 20 trading periods. Excess profit is the average of the difference between individual trader's end-of-period profit and the individual CE profit across 5 groups over all 20 trading periods.

A.2 Summary statistics with unweighted prices

Table A.2: Summary Statistics of Experimental Sessions

	T1 - T20		T1 - T10		T11 - T20	
	CDA	FLOW	CDA	FLOW	CDA	FLOW
Average price	9.04	9.57	8.21	9.18	8.62	9.37
$\lvert P_t - P_{CE} \rvert$	-0.83	-0.21	-1.59	-0.54	-1.21	-0.37
$\lvert P_t - P_{t-1} \rvert$	0.14	0.15	0.14	0.15	0.14	0.15
Std$(P_t - P_{t-1})$	0.08	0.05	0.08	0.05	0.07	0.05

Note: Summary statistics using prices from each second across 5 groups and 120 seconds over 20 trading periods

A.3 Regression with interaction term

Table A.3: Regression Summary with interactions

	$\lvert$ price deviation $\rvert$	$\lvert$ price change $\rvert$	price volatility	order number	order size	traded volume	filled contract/Q_{CE}	realized surplus	surplus$_{buy-sell}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	7.661***	2.788***	0.357***	51.2***	6.6	810.6***	0.675***	0.527***	-2.311***
	(0.902)	(0.410)	(0.079)	(6.4)	(5.6)	(127.8)	(0.120)	(0.120)	(0.713)
FLOW	0.136	-1.287***	-0.094**	-16.5***	3.9	-306.8***	-0.163**	-0.036	0.266
	(0.629)	(0.316)	(0.044)	(4.1)	(4.5)	(78.2)	(0.080)	(0.065)	(0.585)
FLOW × round	-0.023	0.051**	0.002	-0.4*	0.9***	12.2*	0.008*	0.001	0.015
	(0.060)	(0.024)	(0.004)	(0.3)	(0.3)	(7.3)	(0.004)	(0.004)	(0.041)
round	-0.213***	-0.094***	-0.012**	-0.1	0.7***	12.8*	0.012**	0.019***	0.012
	(0.059)	(0.021)	(0.005)	(0.4)	(0.3)	(7.7)	(0.005)	(0.006)	(0.045)
Observations	4,800	4,600	200	200	200	200	200	200	200
R^2	0.319	0.155	0.337	0.745	0.552	0.495	0.310	0.313	0.739
Adjusted R^2	0.318	0.153	0.313	0.736	0.536	0.477	0.285	0.288	0.730

Note: *p<0.1; **p<0.05; ***p<0.01

Table A.4: Hypothesis Test for treatment effect

H_0: $\beta_2 + 20\beta_4 = 0$	$\lvert$ price deviation $\rvert$	$\lvert$ price change $\rvert$	price volatility	order number	order size	traded volume	filled contract/Q_{CE}	realized surplus	surplus$_{buy-sell}$
— test-stat —	0.151	5.636	3.55	7.624	3.681	5.704	1.976	0.667	1.057
p-value	0.88	0.0	0.0	0.0	0.0	0.0	0.048	0.505	0.29

A.4 Experiment Instructions (CDA)

Welcome, and thank you for participating!

From now until the end of the experiment, please do not communicate with other participants. Please turn off your cell phones and stay focused on the task. If you have any questions, please raise your hand or let the experimenters know; they will answer your questions.

Please pay careful attention to the instructions, as real money is at stake. During the experiment, you will earn Experimental Currency Units (ECUs). At the end of the experiment, we will convert your earnings into US Dollars (USD) at one dollar for every 1000 ECUs. You are guaranteed a show-up fee of USD 6 but can earn considerably more.

Basic Idea

In this experiment, you and seven other participants will be traders in a simple automated financial market. Using the information displayed on your screen, you will make trades to earn as many ECUs as possible. As explained below, your earnings will depend on the trade offers you and the other traders make in your group. There will be 22 trading periods, each lasting 2 minutes.

The two first rounds are practice rounds (they don't count for your actual payment).

Your final earnings for the session will be the sum of your show-up fee and the cumulative profits from all trading periods, converted into US Dollars.

Buy Orders and Sell Orders (Treat CDA)

All traders can submit or update *buy* and/or *sell orders*. Each order consists of two elements or pieces of information:

- The ***total quantity*** stating the number of shares that the agent wants to trade.
- The ***limit price*** of the order. In buy orders, the limit price indicates the maximum price the trader is willing to pay per share. In sell orders, the limit price indicates the minimum price the trader is willing to accept for each share.

Example

The blue line and blue square in Figure A.2 show that the trader submitted a buy order for a total of 200 shares at a limit price of 13 ECUs. If someone is willing to sell for 13 ECUs or less, the former trader's buy order will be filled. If that counterpart wants to trade less than 200 shares, then the order will be filled only partially.

Figure 1 also shows (in red) that the trader submitted a sell order for 100 shares. The limit price of this order is 15 ECUs. That is, if someone is willing to buy at 15 ECUs or a higher price, the sell order will be filled. If said counterpart wants to trade less than 100 shares, then the order will be filled only partially.

Figure A.2: CDA Your Input

Submitting orders in the graphical interface (Treat CDA)

To set a sell order, you drag the red square anywhere inside the “your input” grid to set the limit price and the total quantity (volume). You submit your order by clicking on the “Send Sell” red button, as is shown in Figure A.3.

Similarly, to set a buy order, you can drag the blue square anywhere inside the grid to set the limit price and the total quantity (volume). Then you can submit your order by clicking on the “Send Buy” button.

In case your order is not completely executed, your order will remain active with the remaining quantity. For example, if your order is for buying 200 shares and you only obtain 150 shares when trading with another participant. Then, your order remains active, but now only for 50 shares.

You can submit a new buy (sell) order, or cancel it, at any time. However, if you have an active order, you first need to cancel it by pressing the “cancel buy” (“cancel sell”) button. Then you can submit a new order.

Figure A.3: CDA Your Input

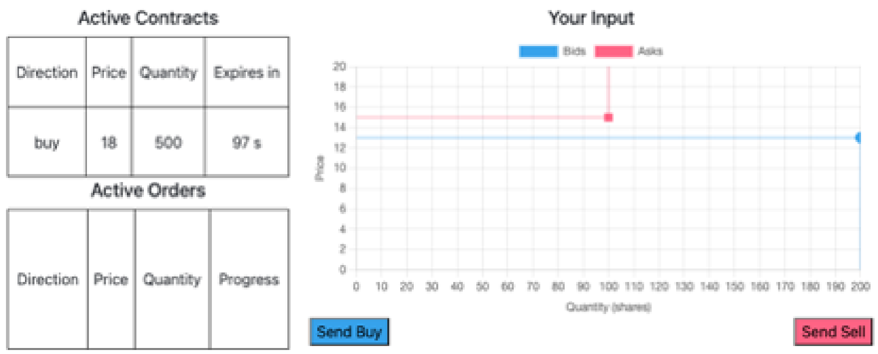

Contracts

Why do you want to trade? What is a share worth to you? The answer to both questions is that the computer will give you a contract that enables you to profit from trading. You will receive your contract near the beginning of the trading period, and it will appear in the “Active Contracts” table on the left of the screen, as shown in Figure A.2.

For example, your contract might read **“Buy 500 shares at a price of 18 ECUs in 97 seconds.”** This means that 97 seconds from now, the computer will buy up to 500 shares from your inventory at a price of 18. To take advantage of this opportunity, you need to buy some shares to build up your inventory. Your profit on each share (up to the 500th share) is the contract price (here 18) minus the price you paid when you bought that share. Of course, for this contract it would not be worthwhile to buy shares at a higher price than 18 or to buy more than 500 shares.

Similarly, a contract that says **"Sell 150 shares at a price of 13 ECUs in 80 seconds"** means that, in 80 seconds, the computer will sell you up to 150 shares at the price of 13 ECUs per share. Here you make a profit by selling up to 150 shares at prices above 13 ECUs per share. For this contract, it would not be worthwhile to sell shares at a lower price than 13 or to sell more than 150 shares.

Keep in mind that you can buy more shares than you have cash for and thus have a negative cash balance, or sell more shares than you own and thus have a negative inventory (this is called shorting). If you have a BUY contract for 400 shares, your inventory should be between zero and your contract quantity, 400 shares. If you have a SELL contract for 500 shares, your inventory should be between -500 (minus 500) shares and zero.

How the market works (Treat CDA)

The market price at which you can buy or sell shares depends on other traders' buy and sell orders. Consider the example illustrated in Figure A.4. All traders, except Trader 1, have active buy or sell orders. In particular, the four sellers want to sell at prices of 15 or more, respectively (see red steps). These four sellers conform to the market supply (the sum of their four sell orders) shown in red. The four buyers want to buy at prices of 13, 12 and 11 or less, respectively (see blue steps). These three buyers conform to the market demand (the sum of their three buy orders) shown in blue.

The green dots indicate the most recent trades. The latest trade is represented by the largest dot.

As you can see, the demand and supply do not cross each other inside the box but along the y-axis. That is there cannot be trade at that point.

Let us see what happens when Trader 1 decides to submit a buy order. There are two cases.

- First, if her buy order has a price below the lowest sell order price (say 15), then Trader 1 will only add to the demand, and no transaction will occur.

- Second, if Trader 1's buy order comes at a price above the lowest selling price (say 13), then she will transact immediately with the seller(s) with the lowest price. Suppose the seller with the lowest price does not offer enough volume (quantity) to fill Trader's 1 order. In that case, the buy order of Trader 1 will be filled only partially or the remaining quantity will be purchased from the second lowest selling price, and so on.

Figure A.4: CDA Market

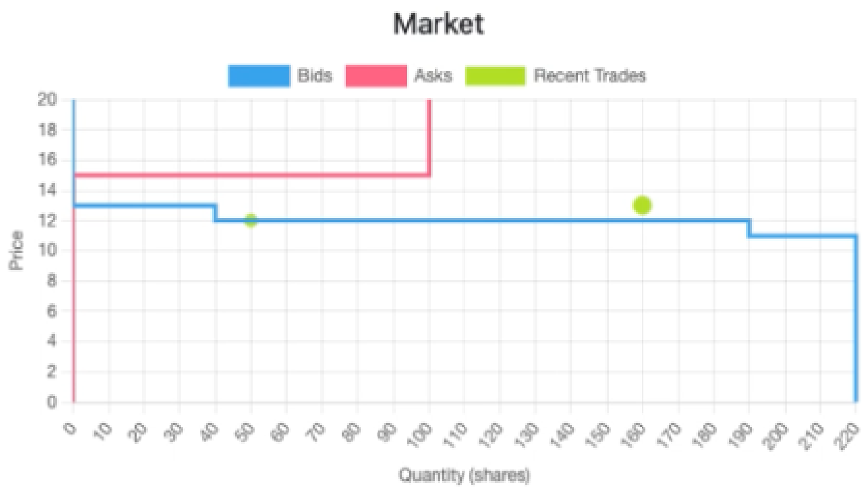

Your Earnings

You make a profit by selling shares at higher prices than you bought them.

However, notice that you do not need to buy a share in order to sell it. Even if you have zero inventory you can still sell shares (making your inventory negative). Similarly, even if you have no cash, you can still buy shares (making your cash balance negative).

Also, notice that you can buy and sell from and to either the computer and other participant traders, or both. Your profits from transactions with other participants will appear immediately in your cash account, whereas your profits from contracts with the computer will be added to the cash account when the contract expires (aka contract deadline).

Examples: You lose money when you buy at a higher price than you sold, or sell at a lower price than you bought. For example, suppose you pay 23 ECUs for one share that you sell to the computer at a contract price of 17 ECUs. Then your profit is 17 - 23 = - 6 ECUs. That is, you incurred a loss of 6 ECUs on that share. If, instead, you bought that share at price of 14 ECUs before the deadline, you would get a profit of 17 - 14 = 3 ECUs on that share.

Projected Profits

The middle bottom box shows your *projected profits* which indicate your period profits if the active contract (if any) expires immediately and the period ends right away. This provides a summary of your current state.

Session Dollar Earnings: After the last trading round, the computer will add up all your period earnings to determine your actual compensation. These earnings will be converted into US dollars at one dollar for every 1000 ECUs.

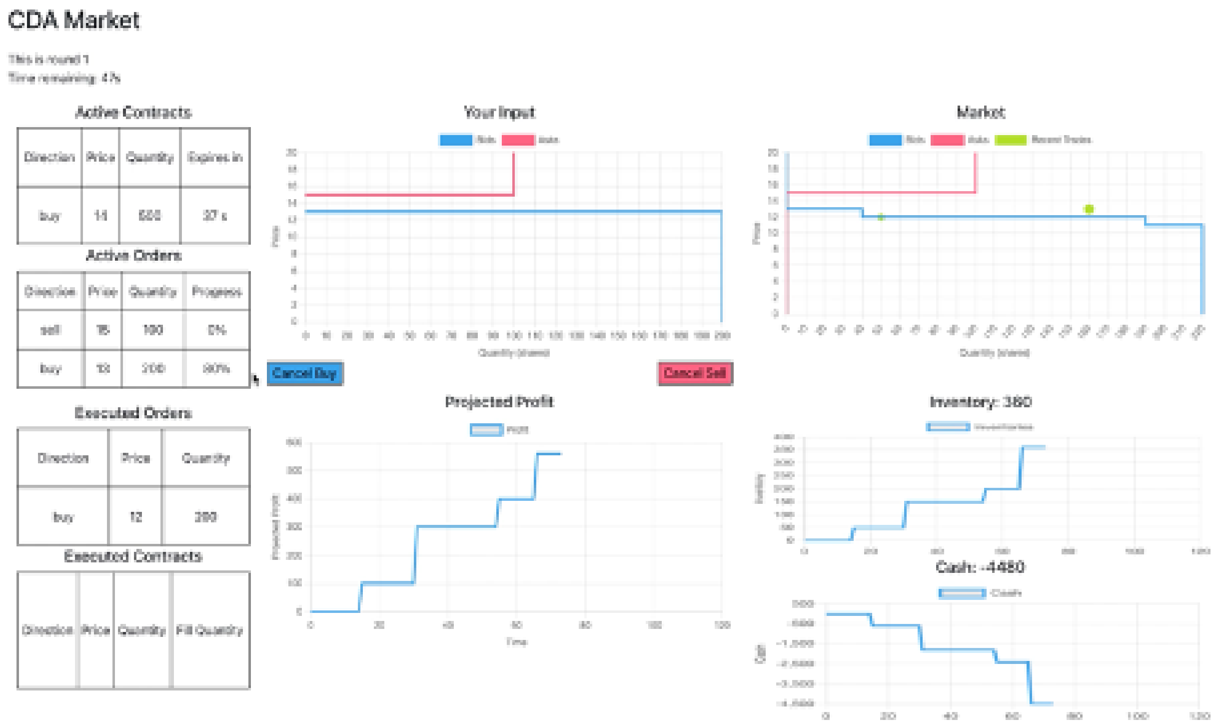

Figure A.5: CDA UI

A.5 Experiment Instructions (FLOW)

Welcome, and thank you for participating!

From now until the end of the experiment, please do not communicate with other participants. Please turn off your cell phones and stay focused on the task. If you have any questions, please raise your hand or let the experimenters know; they will answer them.

Please pay careful attention to the instructions, as real money is at stake. During the experiment, you will earn Experimental Currency Units (ECUs). At the end of the experiment, we will convert your earnings into US Dollars (USD) at one dollar for every 1000 ECUs. You are guaranteed a show-up fee of USD 6 but can earn considerably more.

Basic Idea

In this experiment, you and seven other participants will be traders in a simple automated financial market. Using the information displayed on your screen, you will make trades to earn as many ECUs as possible. As explained below, your earnings will depend on the trade offers you and the other traders make in your group. There will be 20 trading periods, each lasting 2 minutes.

The first two rounds are practice rounds (they don't count for your actual payment).

Your final earnings for the session will be the sum of your show-up fee and the cumulative profits from all trading periods, converted into US Dollars.

Buy Orders and Sell Orders (Treat FLOW)

All traders can submit or update *buy* and/or *sell orders*. Each order consists of four elements:

- The ***total quantity*** stating the number of shares that the agent wants to trade,

- The ***maximum rate*** of shares to trade per second.
- The ***low price*** of the order. In buy orders, the low price indicates the price under which the order is filled at the *maximum rate*. In sell orders, the low price indicates the price under which the order is not filled.
- The ***high price*** of the order. In buy orders, the high price indicates the price above which the order is not filled. In sell orders, the high price indicates the price above which the order is filled at the *maximum rate*.

Example

The blue line and blue box in Figure A.6 show that the trader submitted a buy order for a total of 500 shares at the maximum rate of 15 shares per second at prices between 6 (low price) and 12 (high price). When the market price is at or below 6 ECUs per share, her order will be filled at her *maximum rate* (15 shares per second). Her buy order will not be executed when the market price is at or above 12 ECUs per share. When the market price is between 6 to 12 ECUs per share, her order will be filled at a rate between zero and her maximum rate. The higher the market price, the lower the rate of execution.

Figure A.6 also shows (in red) that the trader submitted a sell order for 50 shares. Her *maximum rate* is 19 shares per second, her *low price* is 14, and her *high price* is 18. When the market price is at or above 18 ECUs per share, her order will be filled at her *maximum rate* of 19 shares per second. Her sell order will not be executed when the market price is at or below 14 ECUs per share. When the market price is between 14 to 18 ECUs per share, her order will be filled at a rate between zero and her maximum rate. The higher the market price, the higher the rate of execution.

Figure A.6: FLOW Your Input

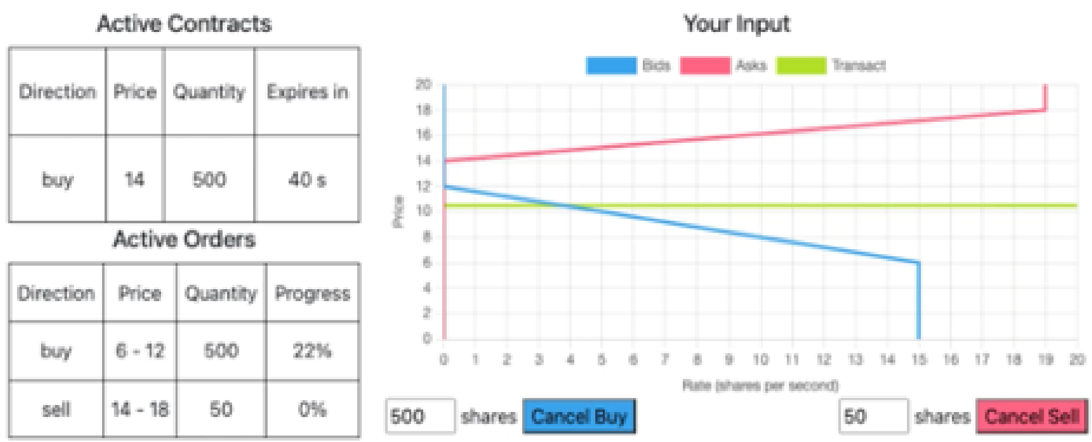

Submitting orders in the graphical interface (Treat FLOW)

To set a sell order, you move the red dot along the vertical axis to set the low price, and you drag the red square anywhere inside the grid to set the high price and maximum rate. You enter the **total number of shares** you want to sell in the box next to the “Send Sell” red button. When you have finished setting the order, submit your order by clicking on the “Send Sell” button.

Similarly, to set a buy order, you move the blue dot along the vertical axis to set the high price, and you drag the blue square anywhere inside the grid to set the low price and the maximum rate. You enter the **total number of shares** you want to buy in the box next to the "Send Buy" blue button. When you have finished setting the order, submit your order by clicking on the "Send Buy" button.

You can submit a new buy (sell) order, by first canceling your active buy (sell) orders and then submitting a new buy (sell) order. You do not need to wait for an order to be fully executed to cancel it.

Contracts

Why do you want to trade? What is a share worth to you? The answer to both questions is that the computer will give you a contract that enables you to profit from trading. You will receive your contract near the beginning of the trading period, and it will appear in the "Active Contracts" table on the left of the screen, as shown in Figure A.6.

For example, your contract might read **"Buy 500 shares at a price of 14 ECUs in 40 seconds."** This means that 40 seconds from now, the computer will buy up to 500 shares from your inventory at a price of 14. To take advantage of this opportunity, you need to buy some shares to build up your inventory. Your profit on each share (up to the 500th share) is the contract price (here 14) minus the price you paid when you bought that share. Of course, for this contract it would not be worthwhile to buy shares at a higher price than 14 or to buy more than 500 shares.

Similarly, a contract that says **“Sell 150 shares at a price of 13 ECUs in 80 seconds”** means that, in 80 seconds, the computer will sell you up to 150 shares at the price of 13 ECUs per share. Here you make a profit by selling up to 150 shares at prices above 13 ECUs per share. For this contract, it would not be worthwhile to sell shares at a lower price than 13 or to sell more than 150 shares.

Keep in mind that you can buy more shares than you have cash for and thus have a negative cash balance, or sell more shares than you own and thus have a negative inventory (this is called shorting). If you have a BUY contract for 400 shares, your inventory should be between zero and your contract quantity, 400 shares. If you have a SELL contract for 500 shares, your inventory should be between -500 (minus 500) shares and zero.

How the market works (Treat FLOW)

The market price and the rate at which you can buy or sell shares depends on the buy and sell orders of all traders, as illustrated in the top chart of Figure A.7. Here, one trader submitted the buy order and sell order shown in Figure A.6. Together with the buy and sell orders submitted by the other 7 traders, they resulted in the market demand (sum of buy orders) shown in blue and market supply (summed sell orders) shown in red. The supply and demand intersect at a price of 10.48 and a Rate of 3.8. Therefore the trader from Figure A.6 is buying at a price of 10.48 ECU to some other trader(s) at a rate of 3.8 shares per second (where the price crosses the agents' demand curve). Thus, her cash position is decreasing at a rate of 10.48 x 3.8 = 19.82 ECU per second while her inventory is increasing at a rate of 3.8 shares per second. This will continue until she buys her total quantity of 500 shares, or until the other traders change their supply or demand, e.g., when they reach their requested total quantity or when they place a new order.

Figure A.7: FLOW Market

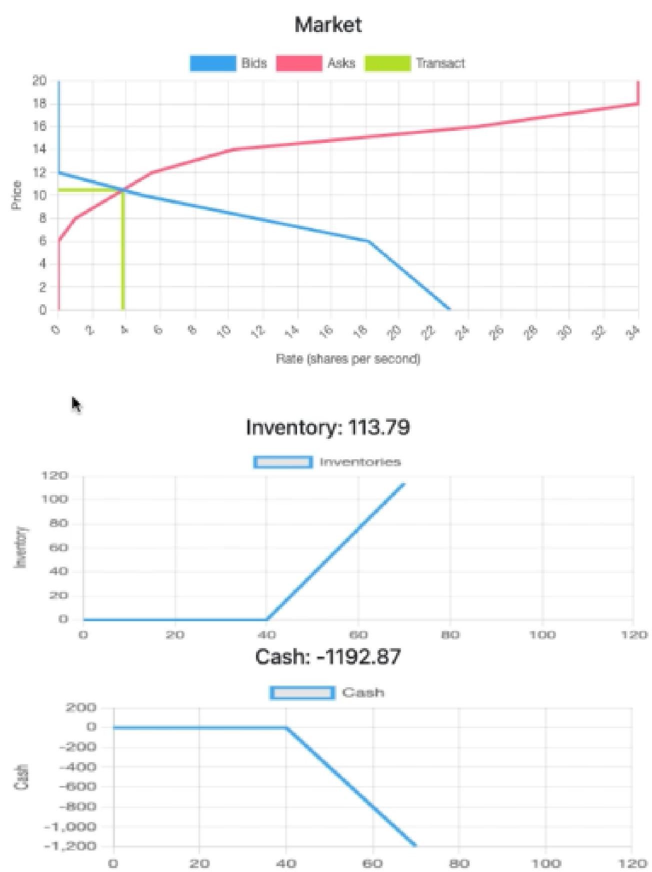

Your Earnings

You make a profit by selling shares at higher prices than you bought them.

However, notice that you do not need to buy a share in order to sell it. Even if you have zero inventory you can still sell shares (making your inventory negative). Similarly, even if you have no cash, you can still buy shares (making your cash balance negative).

Also, notice that you can buy and sell from and to either the computer and other participant traders, or both. Your profits from transactions with other participants will appear immediately in your cash account, whereas your profits from contracts with the computer will be added to the cash account when the contract expires (aka contract deadline).

Examples: You lose money when you buy at a higher price than you sold, or sell at a lower price than you bought. For example, suppose you pay 23 ECUs for one share that you sell to the computer at a contract price of 17 ECUs. Then your profit is 17 - 23 = - 6 ECUs. That is, you incurred in a loss of 6 ECU on that share. If, instead, you bought that share at price of 14 ECUs before the deadline, you would get a profit of 17 - 14 = 3 ECUs on that share.

Projected Profits

The middle bottom box shows your *projected profits* which indicate your period profits if the active contract (if any) expires immediately and the period ends right away. This provides a summary of your current state.

Session Earnings: After the last trading period, the computer will add up all your period earnings to determine your actual compensation. These earnings will be converted into US Dollars at a rate of one US Dollar for every 1000 ECUs.

Figure A.8: FLOW UI

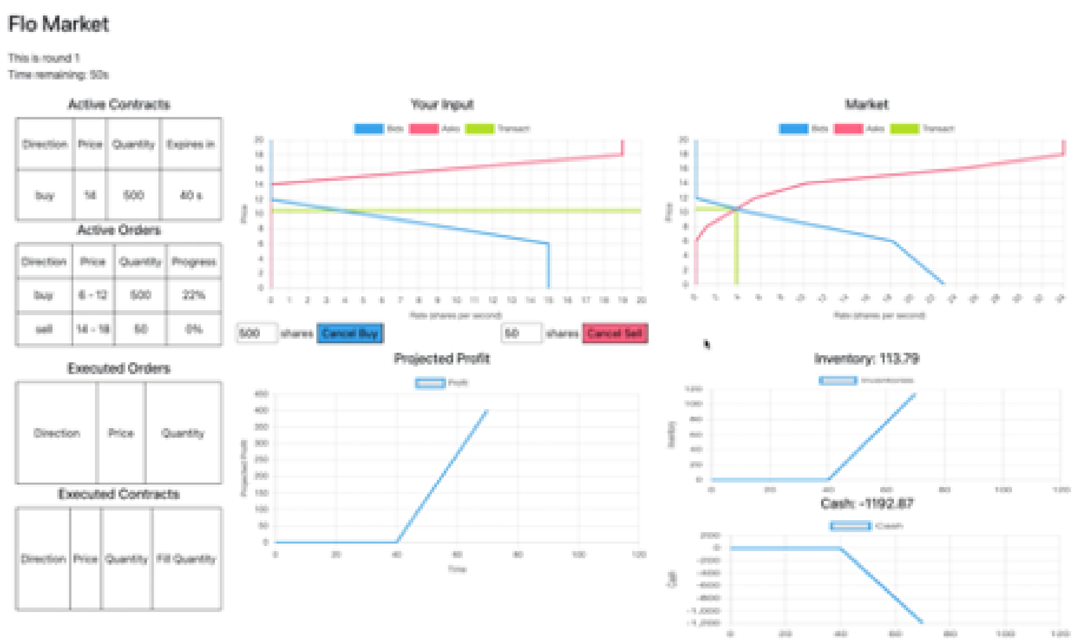

Appendix B

Supplement to Chapter Two

B.1 Examples of comparative statics

Following eq (2.5), we provide an example of sufficient conditions in this section for $\frac{\partial z^*}{\partial a} > 0$. Recall that

$$\begin{aligned}\frac{\partial z^*}{\partial a} &= \frac{\frac{1}{2}(v_H - v_L + 2c)(z^* - 1)^2 - c(z^{*2} + \frac{1}{2}) - \frac{(v_H - v_L + 2c)c}{2(v_H - v_L)}}{2(z^* - 1) + a(1 - z^*)(v_H - v_L + 2c) + 2acz^*} \\ &= \frac{\frac{v_H - v_L}{2} z^{*2} - (v_H - v_L + 2c)z^* + \frac{v_H - v_L}{2} - \frac{c^2}{v_H - v_L}}{[2 - a(v_H - v_L + 2c) + 2ac]z^* - [2 - a(v_H - v_L + 2c)]}\end{aligned}$$

Both the numerator and the denominator can be considered functions of z^*, which can be used to find conditions that determine the sign of $\frac{\partial z^*}{\partial a}$. We consider the sufficient conditions where the numerator and the denominator share the same sign with the following conditions.

$$0 \leq z^* \leq \frac{(v_H - v_L + 2c) - \sqrt{(v_H - v_L + 2c)^2 - 2(v_H - v_L)(\frac{v_H - v_L}{2} - \frac{c^2}{v_H - v_L})}}{v_H - v_L}$$

$$2 - a(v_H - v_L + 2c) < 0$$

or

$$2 - a(v_H - v_L + 2c) > 0$$

$$z > 1 - \frac{2ac}{2 - a(v_H - v_L + 2c) + 2ac}$$

In our experiment with $(v_H, v_L, c) = (500, 100, 5)$, we would have $\frac{\partial z^*}{\partial a} > 0$ when $0 \leq z^* \leq 0.865$ and $a > 0.005$, or when $1 - \frac{10a}{2-400a} \leq z^* \leq 0.865$ and $a < 0.005$.

B.2 Simulations

In the simulations, we apply the CRRA utility function $u(x) = x^{\alpha}$, which is consistent with Crosetto and Filippin (2013). As in our experiment, we set the baseline $(v_H, v_L, c) = (500, 100, 5)$ and vary these parameters. Given the minor theoretical increase, it is difficult to observe the comparative statistics in the model without ex post uncertainty (C). However, we observe trends where the reservation probabilities decrease with the level of risk tolerance in our modified model with ex post uncertainty

(UC).

Figure B.1: Heatmaps of reservation probabilities.

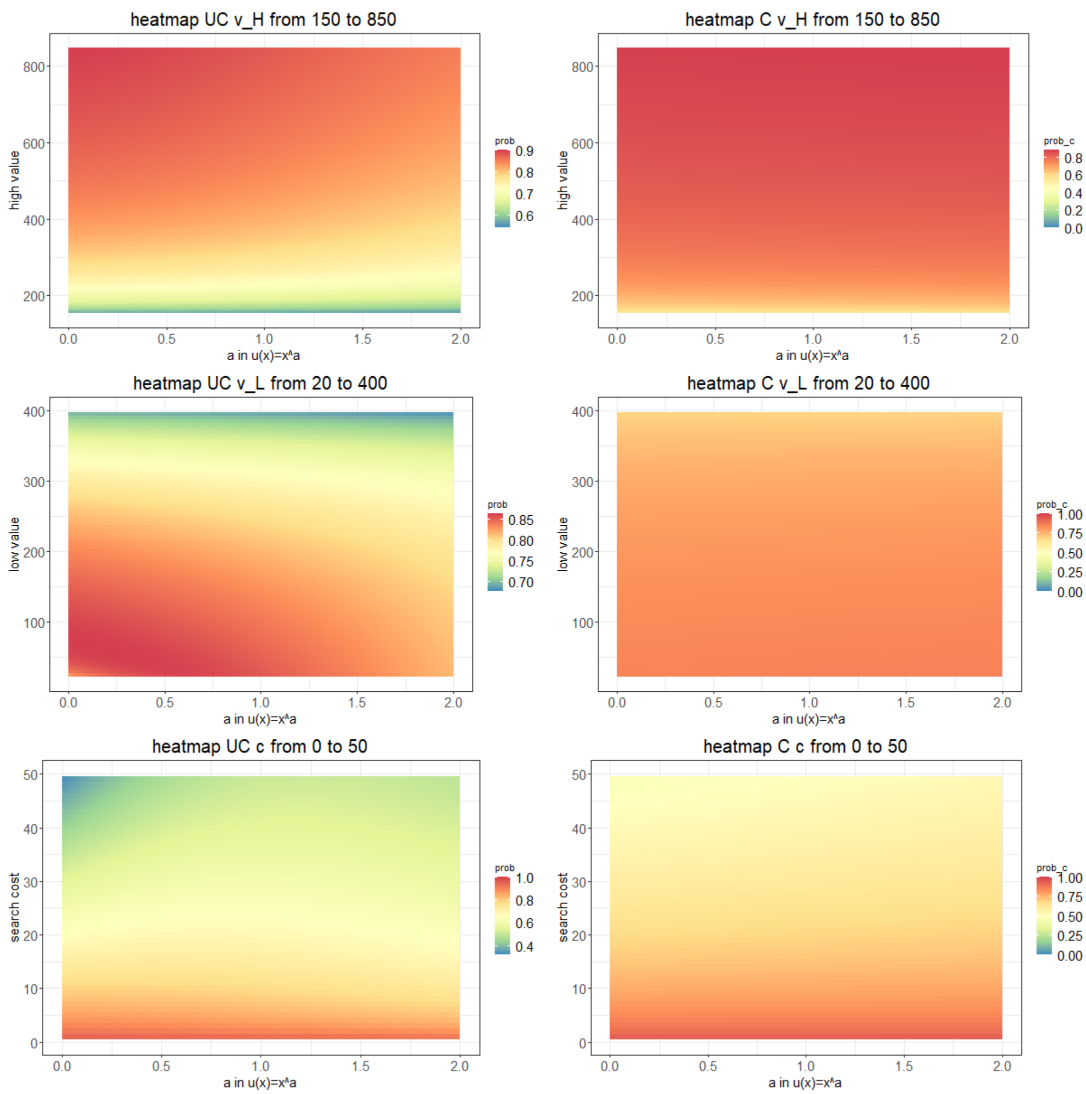

(Note: The x-axis refers to the α in $u(x) = x^{\alpha}$. The y-axis refers to an exogenous parameter. The color reflects the z^* value given the parameters. We compare the model predictions with ex post uncertainty (left panel) to the model predictions without ex post uncertainty (right panel).)

B.3 Estimates of α for the BRET task

Table B.1: Estimates of α for the BRET, assuming CRRA $u(x) = x^{\alpha}$.

K	α	K	α	K	α
1	$0 \le \alpha \le 0.014$	36	$0.551 \le \alpha \le 0.574$	71	$2.39 \le \alpha \le 2.508$
2	$0.015 \le \alpha \le 0.025$	37	$0.575 \le \alpha \le 0.599$	72	$2.509 \le \alpha \le 2.636$
3	$0.026 \le \alpha \le 0.036$	38	$0.6 \le \alpha \le 0.625$	73	$2.637 \le \alpha \le 2.773$
4	$0.037 \le \alpha \le 0.046$	39	$0.626 \le \alpha \le 0.652$	74	$2.774 \le \alpha \le 2.921$
5	$0.047 \le \alpha \le 0.058$	40	$0.653 \le \alpha \le 0.68$	75	$2.922 \le \alpha \le 3.081$
6	$0.059 \le \alpha \le 0.069$	41	$0.681 \le \alpha \le 0.709$	76	$3.082 \le \alpha \le 3.255$
7	$0.07 \le \alpha \le 0.08$	42	$0.71 \le \alpha \le 0.739$	77	$3.256 \le \alpha \le 3.444$
8	$0.081 \le \alpha \le 0.092$	43	$0.74 \le \alpha \le 0.769$	78	$3.445 \le \alpha \le 3.651$
9	$0.093 \le \alpha \le 0.104$	44	$0.77 \le \alpha \le 0.801$	79	$3.652 \le \alpha \le 3.878$
10	$0.105 \le \alpha \le 0.117$	45	$0.802 \le \alpha \le 0.834$	80	$3.879 \le \alpha \le 4.129$
11	$0.118 \le \alpha \le 0.129$	46	$0.835 \le \alpha \le 0.869$	81	$4.13 \le \alpha \le 4.406$
12	$0.13 \le \alpha \le 0.142$	47	$0.87 \le \alpha \le 0.904$	82	$4.407 \le \alpha \le 4.715$
13	$0.143 \le \alpha \le 0.155$	48	$0.905 \le \alpha \le 0.941$	83	$4.716 \le \alpha \le 5.062$
14	$0.156 \le \alpha \le 0.169$	49	$0.942 \le \alpha \le 0.98$	84	$5.063 \le \alpha \le 5.453$
15	$0.17 \le \alpha \le 0.183$	50	$0.981 \le \alpha \le 1.02$	85	$5.454 \le \alpha \le 5.898$
16	$0.184 \le \alpha \le 0.197$	51	$1.021 \le \alpha \le 1.061$	86	$5.899 \le \alpha \le 6.41$
17	$0.198 \le \alpha \le 0.212$	52	$1.062 \le \alpha \le 1.105$	87	$6.411 \le \alpha \le 7.003$
18	$0.213 \le \alpha \le 0.226$	53	$1.106 \le \alpha \le 1.15$	88	$7.004 \le \alpha \le 7.7$
19	$0.227 \le \alpha \le 0.242$	54	$1.151 \le \alpha \le 1.197$	89	$7.701 \le \alpha \le 8.53$
20	$0.243 \le \alpha \le 0.257$	55	$1.198 \le \alpha \le 1.247$	90	$8.531 \le \alpha \le 9.534$
21	$0.258 \le \alpha \le 0.273$	56	$1.248 \le \alpha \le 1.298$	91	$9.535 \le \alpha \le 10.776$
22	$0.274 \le \alpha \le 0.29$	57	$1.299 \le \alpha \le 1.352$	92	$10.777 \le \alpha \le 12.351$
23	$0.291 \le \alpha \le 0.307$	58	$1.353 \le \alpha \le 1.409$	93	$12.352 \le \alpha \le 14.412$
24	$0.308 \le \alpha \le 0.324$	59	$1.41 \le \alpha \le 1.469$	94	$14.413 \le \alpha \le 17.229$
25	$0.325 \le \alpha \le 0.342$	60	$1.47 \le \alpha \le 1.531$	95	$17.23 \le \alpha \le 21.3$
26	$0.343 \le \alpha \le 0.36$	61	$1.532 \le \alpha \le 1.597$	96	$21.31 \le \alpha \le 27.76$
27	$0.361 \le \alpha \le 0.379$	62	$1.598 \le \alpha \le 1.666$	97	$27.761 \le \alpha \le 39.532$
28	$0.38 \le \alpha \le 0.398$	63	$1.667 \le \alpha \le 1.739$	98	$39.533 \le \alpha \le 68.274$
29	$0.399 \le \alpha \le 0.418$	64	$1.74 \le \alpha \le 1.816$	99	$\alpha \ge 68.275$
30	$0.419 \le \alpha \le 0.438$	65	$1.817 \le \alpha \le 1.898$		
31	$0.439 \le \alpha \le 0.459$	66	$1.899 \le \alpha \le 1.985$		
32	$0.46 \le \alpha \le 0.481$	67	$1.986 \le \alpha \le 2.077$		
33	$0.482 \le \alpha \le 0.503$	68	$2.078 \le \alpha \le 2.174$		
34	$0.504 \le \alpha \le 0.526$	69	$2.175 \le \alpha \le 2.278$		
35	$0.527 \le \alpha \le 0.55$	70	$2.279 \le \alpha \le 2.389$		

Note: K is the number of boxes collected.

B.4 Main results with MPL first crossing point

Figure B.2: Search duration vs risk tolerance.

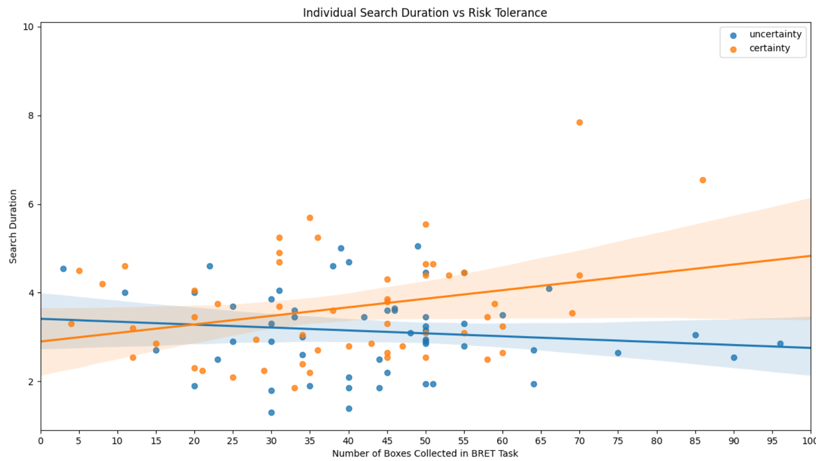

(Note: The plot shows how search duration changes with the level of risk tolerance when using the average behavior in the last 10 paid rounds and number of boxes collected in BRET.)

Figure B.3: Search duration vs risk tolerance.

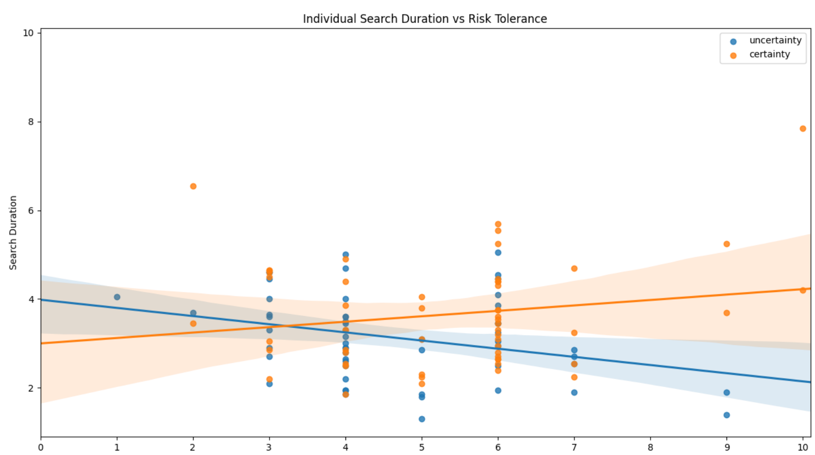

(Note: The plot shows how search duration changes with the level of risk tolerance when using the average behavior in the last 10 paid rounds and number of risky choices in MPL.)

Table B.2: Regression summary.

	Dependent Variable: Average Search Duration			
	BRET	BRET	MPL	MPL
	1-20 rounds	11-20 rounds	1-20 rounds	11-20 rounds
Risk Tolerance	2.173^{**}	2.241^{*}	1.261	0.560
	(0.835)	(1.189)	(0.837)	(1.185)
Risk Tolerance x UC	-2.815^{**}	-3.043^{*}	-3.789^{***}	-4.161^{**}
	(1.172)	(1.669)	(1.290)	(1.826)
UC	0.583	1.018	1.035^{*}	1.476^{*}
	(0.527)	(0.750)	(0.586)	(0.828)
const	2.965^{***}	2.878^{***}	3.235^{***}	3.497^{***}
	(0.359)	(0.512)	(0.413)	(0.584)
Observations	109	109	109	109
R^2	0.118	0.041	0.130	0.066

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Table B.3: Tests of regression coefficients.

	BRET		BRET		MPL		MPL	
	1-20 rounds		11-20 rounds		1-20 rounds		11-20 rounds	
H_0	t-stat	p-value	t-stat	p-value	t-stat	p-value	t-stat	p-value
$\beta_1 > 0$	2.603	0.005	1.885	0.031	1.506	0.068	0.473	0.319
$\beta_1 + \beta_2 < 0$	-0.781	0.218	-0.684	0.248	-2.575	0.006	-2.592	0.005
$0.5\beta_2 + \beta_3 = 0$	-3.484	0.001	-1.826	0.071	-3.641	0.000	-2.306	0.023

Appendix C

Supplement to Chapter Three

C.1 Proof of Proposition 2

The game can be solved by backward induction. From the last decision nodes for entrant, we have the following conditions.

1) Both firms choose L:

Entrant will gain an advantage at the beginning since consumers will expect them to produce H with a positive probability. Consumers update Eu_{it} based on entrant's quality to u_L for $t > 1$ and both firms will share the market after the first period.

	Incumbent	Entrant
p_{i1}	0	$\pi(u_H - u_L) - \varepsilon$
$p_{it}, t > 1$	0	0
payoff	0	$\pi(u_H - u_L) - \varepsilon$

2) Incumbent chooses L and entrant chooses H:

The results depend on the consumers' choices. Although consumers prefer H to L, they have a low prior belief towards entrant which plan to provide H.

If $c_E > \pi(u_H - u_L)$, consumers will choose to purchase from incumbent and entrant cannot enter the market.

	Incumbent	Entrant
$p_{it}, t > 0$	$c_E - \pi(u_H - u_L) - \varepsilon$	c_E
profit	$\frac{1}{1-\delta}[c_E - \pi(u_H - u_L) - \varepsilon]$	0

If $c_E \leq \pi(u_H - u_L)$, consumers will choose to purchase from entrant instead. Entrants successfully enter the market and take the entire market share.

	Incumbent	Entrant
p_{i1}	0	$\pi(u_H - u_L) - \varepsilon$
$p_{it}, t > 1$	0	$u_H - u_L - \varepsilon$
profit	0	$\pi(u_H - u_L) - c_E - \varepsilon + \frac{\delta}{1-\delta}(u_H - u_L - c_E - \varepsilon)$

3) Incumbent chooses H and entrant chooses L:

Entrant plays a fly-by-night strategy since consumer strictly prefer H to L with complete information. However, it could also be profitable to play such strategy since consumers expected entrants to play H with a positive probability, which leave space for profit.

If $c_I \leq (1-\pi)(u_H - u_L)$, consumer will prefer to purchase from incumbent and entrant cannot enter.

	Incumbent	Entrant
$p_{it}, t > 0$	$(1-\pi)(u_H - u_L) - \varepsilon$	0
profit	$\frac{1}{1-\delta}[(1-\pi)(u_H - u_L) - c_I - \varepsilon]$	0

If $c_I > (1-\pi)(u_H - u_L)$, consumer will prefer to purchase from entrant but will go back to incumbent starting from $t > 1$ once the entrant's quality is revealed.

	Incumbent	Entrant
p_{i1}	c_I	$c_I - (1-\pi)(u_H - u_L) - \varepsilon$
$p_{it}, t > 1$	$u_H - u_L - \varepsilon$	0
profit	$\frac{\delta}{1-\delta}(u_H - u_L - c_I - \varepsilon)$	$c_I - (1-\pi)(u_H - u_L) - \varepsilon$

4) If both firms choose H:

Incumbent will have an advantage given its high reputation and can undercut the price of entrant. However, if entrant has a lower cost, she can undercut incumbent's price if consumers has a high prior belief that entrant will produce H, or if entrant has a strong cost advantage over incumbent.

If $c_I - c_E > (1-\pi)(u_H - u_L)$, entrant has an advantage over incumbent and can take the market.

	Incumbent	Entrant
p_{i1}	c_I	$c_I-(1-\pi)(u_H-u_L)-\varepsilon$
$p_{it}, t>1$	c_I	$c_I-\varepsilon$
profit	0	$(c_I-c_E)-(1-\pi)(u_H-u_L)-\varepsilon+\frac{\delta}{1-\delta}(c_I-c_E-\varepsilon)$

If $c_I-c_E \leq (1-\pi)(u_H-u_L)$, incumbent gets the advantage and can set a lower price to permanently prevent entrant from entering the market.

	Incumbent	Entrant
$p_{it}, t>0$	$c_E+(1-\pi)(u_H-u_L)-\varepsilon$	c_E
profit	$\frac{1}{1-\delta}[c_E-c_I+(1-\pi)(u_H-u_L)-\varepsilon]$	0

Go back to the initial node for incumbent. When incumbent chooses L,

- If $c_E \leq \pi(u_H-u_L)$, entrant will choose H and take over the market when $-c_E+\frac{\delta}{1-\delta}(u_H-u_L-c_E-\varepsilon)\geq 0$.

- Entrant will choose L and share the market with incumbent if

  $c_E>\pi(u_H-u_L)$, or

  $c_E\leq\pi(u_H-u_L)$ and $-c_E+\frac{\delta}{1-\delta}(u_H-u_L-c_E-\varepsilon)<0$.

There is no entry barrier when incumbent produces L. Entrant can always enter the market with either L or H. As a results, incumbent always earn 0 when they choose L.

When incumbent chooses H,

- If $c_I\leq(1-\pi)(u_H-u_L)$, we will also have $c_I-c_E<(1-\pi)(u_H-u_L)$. Entrant cannot enter the market and incumbent earns a positive payoff.

- If $c_I>(1-\pi)(u_H-u_L)$ and $c_I-c_E>(1-\pi)(u_H-u_L)$, entrant compares the fly-by-night payoff and long-term high quality payoff. When $\frac{\delta}{1-\delta}(c_I-c_E-\varepsilon)-c_E\geq 0$, entrant produce H and take over the entire market.

- Entrant will produce L and earn a fly-by-night payoff if

  $c_I>(1-\pi)(u_H-u_L)$ and $c_I-c_E\leq(1-\pi)(u_H-u_L)$,

  or

  $c_I>(1-\pi)(u_H-u_L)$, $c_I-c_E>(1-\pi)(u_H-u_L)$, and $\frac{\delta}{1-\delta}(c_I-c_E-\varepsilon)-c_E<0$.

For incumbent, choosing L is always weakly dominated by choosing H given entrant's choices and in many conditions strictly dominated by choosing H. Incumbent sets up the entry barrier with H and entrants' response to the incumbent's strategy conditional on the market conditions. The market equilibrium can be summarized as follows.

- If $c_I \leq (1-\pi)(u_H - u_L)$, incumbent set up a strong entry barrier by producing H.
- Entrant will choose a fly-by-night strategy by producing L when
  $c_I > (1-\pi)(u_H - u_L)$ and $c_I - c_E \leq (1-\pi)(u_H - u_L)$, or
  $c_I > (1-\pi)(u_H - u_L)$, $c_I - c_E > (1-\pi)(u_H - u_L)$ and $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - c_E < 0$.
  Incumbent get 0 at $t = 1$ but can still defend the market with H.
- If $c_I - c_E > (1-\pi)(u_H - u_L)$ and $\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - c_E \geq 0$, entrant successful enters the market with H.

C.2 Proof of Proposition 3

At $t = 1$, both incumbent and entrant play as monopolist in their own niche market. Both firms set the price equal to consumers' expected utility and set the quality for the long-run consideration. At $t > 1$, both players' compete in the same market and entrant's quality is revealed from her choice at first period.

To motivate entrant to produce H, entrant should be able to produce H and take the market if she does so, which requires $c_I > c_E$ and $\pi u_H + (1-\pi)u_L \geq c_E$. Another comparison is between the fly-by-night strategy and producing H for reputation. Entrant produces H when she earns more long-run benefit from H than short-run benefit from L.

$$\tau(\pi u_H + (1-\pi)u_L - c_E) + \frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) \geq \tau(\pi u_H + (1-\pi)u_L) \tag{C.1}$$

$$\frac{\delta}{1-\delta}(c_I - c_E - \varepsilon) - \tau c_E \geq 0 \tag{C.2}$$

The third part can be simply proved by comparing the conditions for entrant to enter the market with H with and without fractional searching.

C.3 Price Distribution (Param. set 1)

Figure C.1: CC treatment group

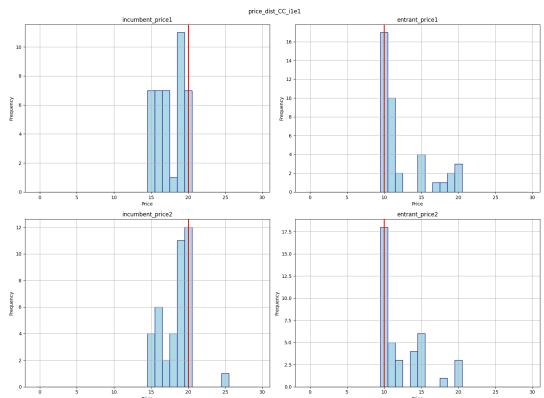

(i) $Q^I = 1, Q^E = 1$

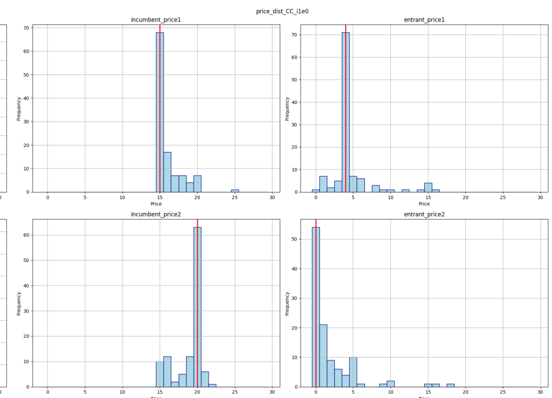

(ii) $Q^I = 1, Q^E = 0$

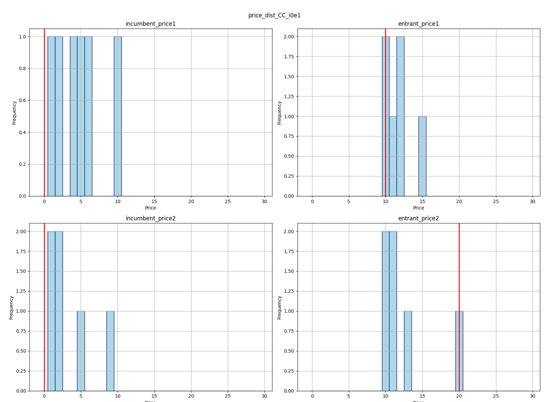

(iii) $Q^I = 0, Q^E = 1$

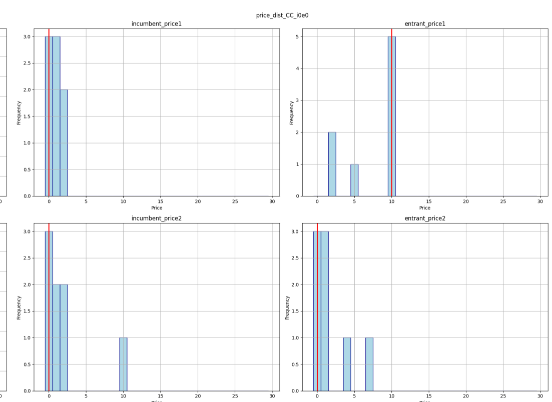

(iv) $Q^I = 0, Q^E = 0$

C.4 Price Distribution (Param. set 1)

Figure C.2: FC treatment group

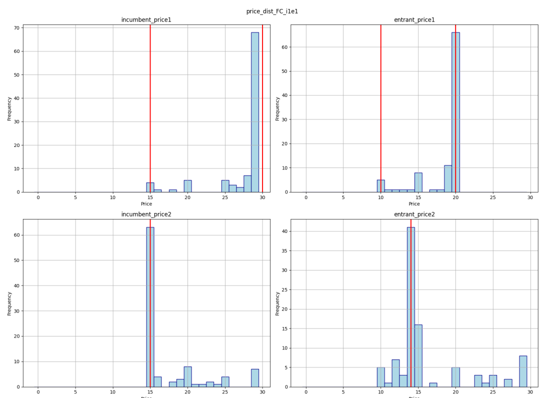

(i) $Q^I = 1, Q^E = 1$

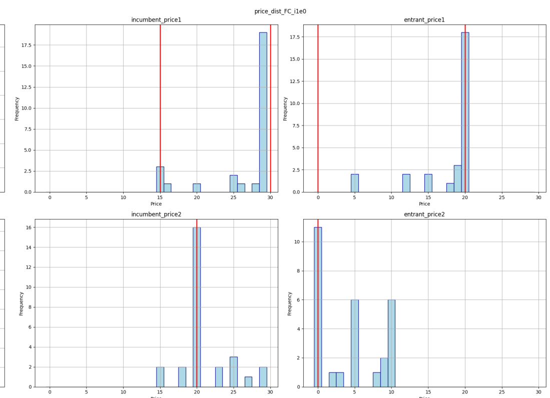

(ii) $Q^I = 1, Q^E = 0$

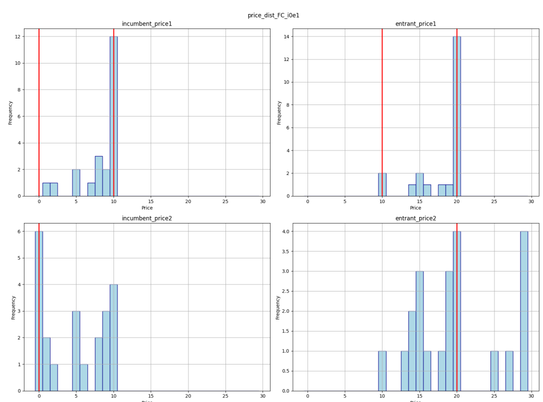

(iii) $Q^I = 0, Q^E = 1$

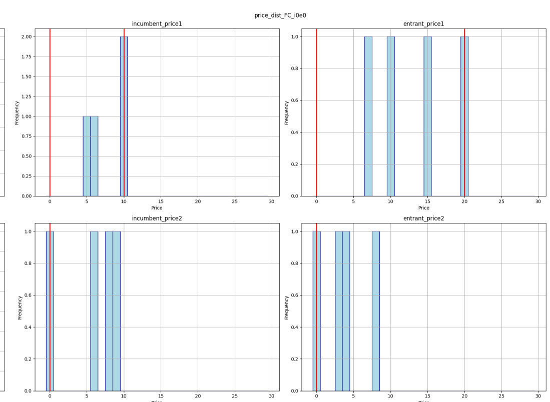

(iv) $Q^I = 0, Q^E = 0$

C.5 Price Distribution (Param. set 2)

Figure C.3: CC treatment group

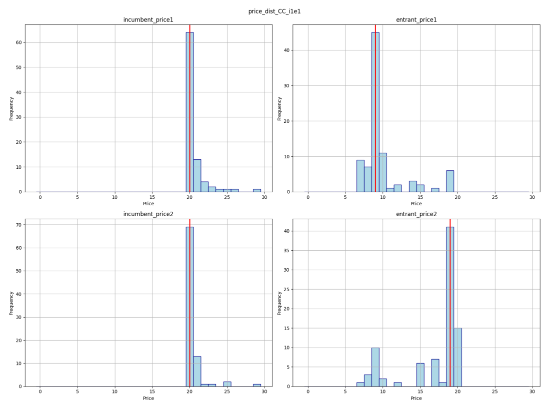

(i) $Q^I = 1, Q^E = 1$

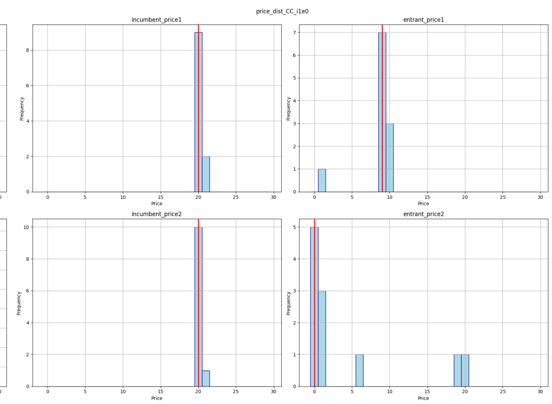

(ii) $Q^I = 1, Q^E = 0$

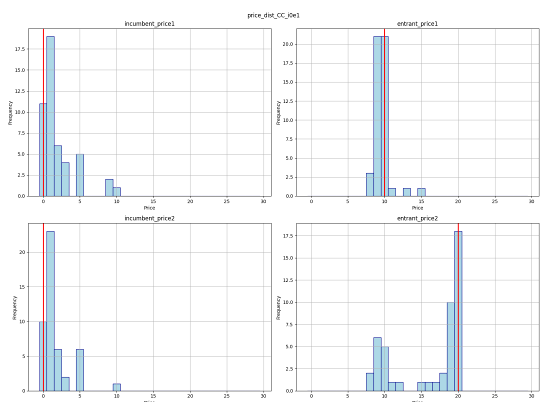

(iii) $Q^I = 0, Q^E = 1$

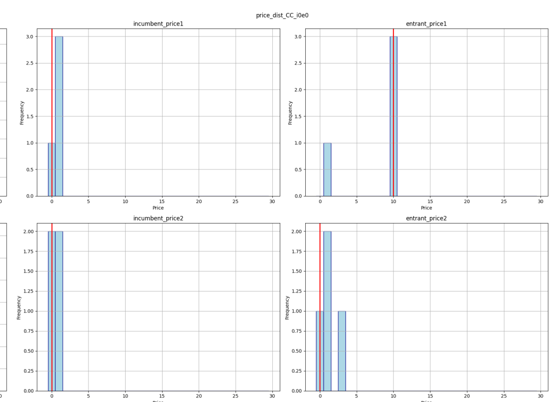

(iv) $Q^I = 0, Q^E = 0$

C.6 Price Distribution (Param. set 2)

Figure C.4: FC treatment group

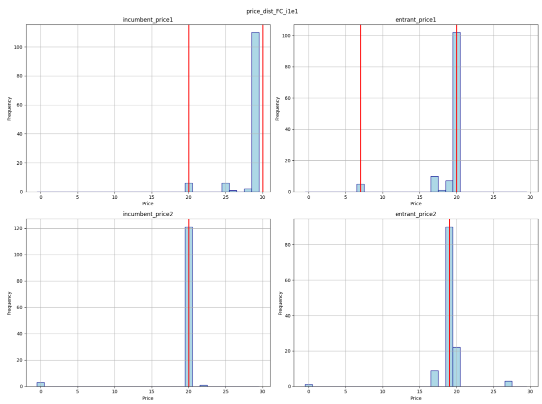

(i) $Q^I = 1, Q^E = 1$

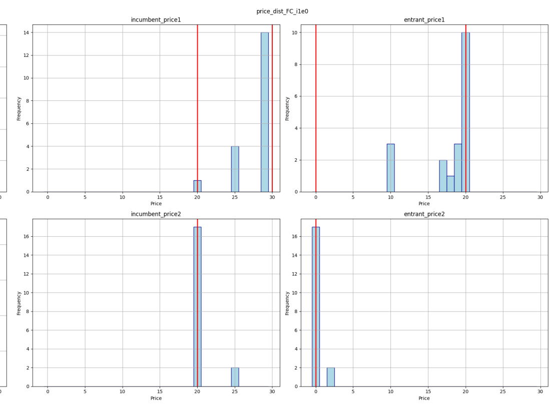

(ii) $Q^I = 1, Q^E = 0$

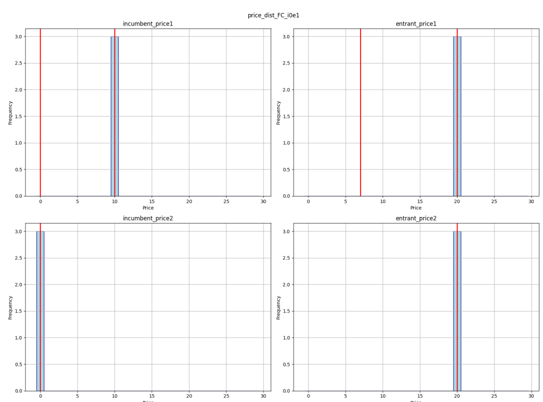

(iii) $Q^I = 0, Q^E = 1$

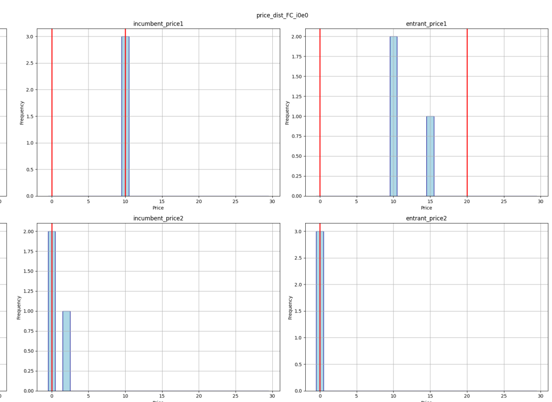

(iv) $Q^I = 0, Q^E = 0$

C.7 Experiment Instructions (OL)

General Ideas In the following rounds, you will be matched with another participant as your counterpart in each round.

- The match is anonymous: you do not know who you are matched with.

- The experiment randomly rematches between rounds so your counterpart changes every round.
- There are two roles: player A and player B. One of you will be randomly assigned to player A and the other one will be player B.
- Each round is a new game and is not correlated with other rounds.

In each round, the game has 3 stages. You and your counterpart choose the type of your products at stage 1 and compete with price and product type at stages 2 and 3.

Stage 1: At stage 1, you and your counterpart simultaneously choose the type of your products.

- Production is costly. The costs of producing X and producing Y are different.
- You can observe both players' costs before you make the decision.
- When you choose your product type, you cannot observe your counterpart's decision.
- You cannot change your product type at stage 2 or 3.

Stage 2: At stage 2, you and your counterpart simultaneously choose the price of your products.

- You can observe both players' product types before you choose your price.
- Your price is required to be between the cost of your product and the value consumers get from your product.

After you choose your prices, 100 bot consumers choose whose product they will purchase at this stage. For consumers, their net benefit = product value - product price.

- Consumers choose to purchase from the player that provides them a higher net benefit.
- Consumers observe both players' product types and prices before they make decisions.
- If both of you provide the same net benefit, consumers will prefer the one with a higher product value.
- If both of you provide the same net benefit and the same product value, both of you will equally share the market.

Stage 3: At stage 3, you and your counterpart repeat what you have done at stage 2: both of you simultaneously choose your price and consumers choose whose product they will purchase. You can choose a different price at stage 3 based on what you learn from the result of stage 2. However, what you have chosen at stage 2 does not affect the result of stage 3.

Your payoff in this round

At stage 2 and 3, your stage payoff is (your product price - your product cost) $\times$ your demand.

- Your demand is the number of consumers who purchase your product. It can be either the whole market (100), 0, or half of the market (50).

Your payoff in this round is the sum of your payoffs at stages 2 and 3. If you finish reading, please raise your hand in the Zoom meeting room with the blue icon under ”participants” and wait for the experimenter to advance the page.

C.8 Experiment Instructions (CC)

General Ideas In the following rounds, you will be matched with another participant as your counterpart in each round.

- The match is anonymous: you do not know who you are matched with.
- The experiment randomly rematches between rounds so your counterpart changes every round.
- There are two roles: player A and player B. One of you will be randomly assigned to player A and the other one will be player B.
- Each round is a new game and is not correlated with other rounds.

In each round, the game has 3 stages. You and your counterpart choose the type of your products at stage 1 and compete with price and product type at stages 2 and 3.

Stage 1: At stage 1, you and your counterpart sequentially choose the type of your products, X or Y.

- Production is costly. The costs of producing X and producing Y are different.
- You can observe both players' costs before you make the decision.
- Player A first choose their product type. Then, player B observe player A's product type and choose their type.
- Your product type affects your actions and payoffs at stage 2 and 3. You cannot change your product type at stage 2 or 3.

Stage 2: At stage 2, you and your counterpart first simultaneously choose the price of your products.

- You can observe both players' product types before you choose your price.

- Your price is required to be between the cost of your product and the value consumers get from your product, so none of the participants in this game gets negative payoff.

After you choose your prices, 100 bot consumers choose whose product they will purchase at this stage. For consumers, their net benefit = product value - product price.

- Consumers choose to purchase from the player that provides them a higher net benefit.

- If both of you provide the same net benefit, consumers will prefer the one with a higher product value.

- If both of you provide the same net benefit and the same product value, both of you will equally share the market.

However, consumers only observe part of product information:

- Consumers only observe player A's product type at stage 2. Player B's type is unknown to consumers.

- Besides value from X and Y, consumers have a third value when the product type is unknown. As player B's product type is unknown to consumers at stage 2, consumers will use the value of unknown product when calculating the net benefit from player B.

- Both players know the value of unknown product when they make decisions.

- Consumers know both players' prices when they make purchase decisions.

Stage 3: At stage 3, you and your counterpart repeat what you have done at stage 2: both of you simultaneously choose your price and consumers choose whose product they will purchase. You can choose a different price at stage 3 based on what you learn from the result of stage 2. Whether consumers know player B's type at stage 3 depends on player B's demand at stage 2:

- If player B gets positive demand at stage 2, their product type becomes known to consumers at stage 3. Consumers will use the value of player B's product type when calculating the net benefit.

- If player B gets 0 demand at stage 2, their product type is still unknown to consumers at stage 3.

Your payoff in this round At stage 2 and 3, your stage payoff is (your product price - your product cost) $\times$ your demand.

- Your demand is the number of consumers who purchase your product. It can be either the whole market (100), 0, or half of the market (50).

Your payoff in this round is the sum of your payoffs at stages 2 and 3. If you finish reading, please raise your hand in the Zoom meeting room with the blue icon under "participants" and wait for the experimenter to advance the page.

C.9 Experiment Instructions (FC)

General Ideas In the following rounds, you will be matched with another participant as your counterpart in each round.

- The match is anonymous: you do not know who you are matched with.

- The experiment randomly rematches between rounds so your counterpart changes every round.

- There are two roles: player A and player B. One of you will be randomly assigned to player A and the other one will be player B.
- Each round is a new game and is not correlated with other rounds.

In each round, the game has 3 stages. You and your counterpart choose the type of your products at stage 1 and compete with price and product type at stages 2 and 3.

Stage 1: At stage 1, you and your counterpart sequentially choose the type of your products, X or Y.

- Production is costly. The costs of producing X and producing Y are different.
- You can observe both players' costs before you make the decision.
- Player A first choose their product type. Then, player B observe player A's product type and choose their type.
- Your product type affects your actions and payoffs at stage 2 and 3. You cannot change your product type at stage 2 or 3.

Stage 2: At stage 2, you and your counterpart first simultaneously choose the price of your products.

- You can observe both players' product types before you choose your price.
- Your price is required to be between the cost of your product and the value consumers get from your product, so none of the participants in this game gets negative payoff.

After you choose your prices, 100 bot consumers choose whose product they will purchase at this stage. For consumers, their net benefit = product value - product price.

- At stage 2, you and your counterpart will be placed in two separate markets, where 20 consumers can only purchase from player B and 80 consumers can only purchase from player A.

- You and your counterpart cannot affect each other at stage 2. Consumers will purchase when their net benefit is non-negative.

However, consumers only observe part of product information:

- Consumers only observe player A's product type at stage 2. Player B's type is unknown to consumers.

- Besides value from X and Y, consumers have a third value when the product type is unknown. As player B's product type is unknown to consumers at stage 2, consumers will use the value of unknown product when calculating the net benefit from player B.

- Both players know the value of unknown product when they make decisions.

- Consumers know both players' prices when they make purchase decisions.

Stage 3: At stage 3, all 100 consumers return to the same market and share the information of products they purchased at stage 2. Whether consumers know player B's type at stage 3 depends on player B's demand at stage 2.

- If player B gets positive demand at stage 2, their product type becomes known to all consumers at stage 3.

- If player B gets 0 demand at stage 2, their product type is still unknown to all consumers at stage 3.

You and your counterpart then interact in the same market. Both of you simultaneously choose your price and consumers choose whose product they will purchase. You can choose a different price at stage 3 based on what you learn from the result of stage 2. Consumers' decision is as the following:

- Consumers choose to purchase from the player that provides them a higher net benefit.
- If both of you provide the same net benefit, consumers will prefer the one with a higher product value.
- If both of you provide the same net benefit and the same product value, both of you will equally share the market.

Your payoff in this round

At stage 2 and 3, your stage payoff is (your product price - your product cost) $\times$ your demand.

- Your demand is the number of consumers who purchase your product. It can be either your own market or 0 at stage 2, or the whole market (100), 0, or half of the market (50) at stage 3.

Your payoff in this round is the sum of your payoffs at stages 2 and 3. If you finish reading, please raise your hand in the Zoom meeting room with the blue icon under "participants" and wait for the experimenter to advance the page.

C.10 Sequential Entry with Endogenous Prior Belief

In the previous sections, π is set as an exogenous parameter and consumers are set to believe that firms produce H only when they always produce H in the past. What if consumers endogenously decide their prior belief π? Note that consumers are not defined as players in the extensive form game and thus do not have optimization problems but follow a pre-defined behavioral rule. From the previous propositions, we have learned that incumbent will always choose H to set up the entry barrier. For consumers, given that they can observe incumbent's quality, they have no incentive to motivate entrant to produce H. However, consumers' belief affects entrant's strategies and consumers may get a higher expected utility by adjusting π. To make π an endogenous parameter, consumers' belief need to be consistent with firm's strategy given the belief.

Consumers need to learn the distribution of firm's cost to generate the prior belief, which requires us to release Assumption 6. Let us assume that consumers know c_I as incumbent is defined to be in the market before entrant tries to enter. We also assume that consumers believe that $c_E \sim U[\underline{c}, \bar{c}]$, where $0 < \underline{c} < c_I < \bar{c}$. That is, consumers belief that entrant's cost is somewhere around incumbent's cost. Note that the equilibrium in the section highly depends on consumers' prior belief and we just provide a simple assumption where consumers are not quite sure whether entrant is cost efficient or not.

In the sequential entry without fractional searching, we can transform the conditions in Proposition 2 with respect to c_E and thus calculate the probability of entrant choosing a specific quality. For example, let $\varepsilon \longrightarrow 0$, we can transform the conditions as follows.

$$c_I - c_E > (1-\pi)(u_H - u_L) \longrightarrow c_E < c_I - (1-\pi)(u_H - u_L) \tag{C.3}$$

$$\frac{\delta}{1-\delta}(c_I - c_E) - c_E \geq 0 \longrightarrow c_E \leq \delta c_I \tag{C.4}$$

In this way, we can rephrase Proposition 2 from consumers' perspective as follows.

- From consumers perspective, entrant produces L when:
  (1) $c_E \geq c_I - (1-\pi)(u_H - u_L)$ (2) $\delta c_I < c_E < c_I - (1-\pi)(u_H - u_L)$
- Similarly, entrant produces H when $c_E < c_I - (1-\pi)(u_H - u_L)$ and $c_E \leq \delta c_I$.

Note that consumers' prior belief needs to be consistent with firms' strategies. Let $\hat{c} = \min\{\delta c_I, c_I - (1-\pi^*)(u_H - u_L)\}$. Consumers' endogenous belief π^* can be defined as follows.

- If $\hat{c} \geq \bar{c}$, consumers believe that entrant will always produce H and consumers set $\pi^* = 1$. However, this contradicts our assumption that $c_I < \bar{c}$ so consumers never set $\pi^* = 1$. Plug $\pi^* = 1$ back to the condition, we have $\pi^* = 1$ when $\delta c_I \geq \bar{c}$.

- If $\hat{c} \leq \underline{c}$, consumers believe that entrant will always produce L and consumers set $\pi^* = 0$. Plug $\pi^* = 0$ back to the condition, we have $\pi^* = 0$ when $\min\{\delta c_I, c_I - (u_H - u_L)\} \leq \underline{c}$. Based on our assumption, we have $c_I - (u_H - u_L) < 0$ so the condition above is always satisfied. That is to say, in any circumstances, consumers' belief of entrant producing L is an equilibrium belief.

- If $\hat{c} \in (\underline{c}, \bar{c})$, consumers believe that it is possible for entrant to produce either L or H and $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$. Plug $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$ back to the condition, we have $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$ when $\min\{\delta c_I, c_I - \pi^*(u_H - u_L)\} \in (\underline{c}, \bar{c})$.

Although not mathematically complete, the equilibrium belief π^* reveals an important aspect of consumers belief that consumers never believe that entrant produces H with 100 percent certainty but the belief of entrant producing L is always an equilibrium belief regardless of the exogenous parameters. The endogenous belief explains the entry barrier from consumers' perspective.

In the sequential entry with fractional searching, applying similar methods, we can get new conditions from consumers' perspective.

- If $c_E < c_I$, $c_E \leq \pi u_H + (1-\pi)u_L$, and $c_E \leq \frac{\delta}{\delta+(1-\delta)\tau} c_I$, entrant produces H.

- Otherwise, entrant produces L.

Let $\hat{c} = \min\{c_I, \pi u_H + (1-\pi)u_L, \frac{\delta}{\delta+(1-\delta)\tau} c_I\}$. Consumers' endogenous belief π^* can be defined as follows.

- If $\hat{c} \geq \bar{c}$, entrant will always produce H and consumers set $\pi^* = 1$. Plug $\pi^* = 1$ back to the condition, we have $\pi^* = 1$ when $\frac{\delta}{\delta+(1-\delta)\tau} c_I \geq \bar{c}$. Similar to the previous conditions, this never holds.

- If $\hat{c} \leq \underline{c}$, entrant will always produce L and consumers set $\pi^* = 0$. Plug $\pi^* = 0$ back to the condition, we have $\pi^* = 0$ when $\min\{u_L, \frac{\delta}{\delta+(1-\delta)\tau}c_I\} \leq \underline{c}$.

- If $\hat{c} \in (\underline{c}, \bar{c})$, it is possible for entrant to produce either L or H and $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$. Plug $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$ back to the condition, we have $\pi^* = \frac{\hat{c}-\underline{c}}{\bar{c}-\underline{c}}$ when $\min\{\pi^* u_H + \pi^* u_L, \frac{\delta}{\delta+(1-\delta)\tau}c_I\} \in (\underline{c}, \bar{c})$.

Compared to sequential entry without fractional searching, $\pi^* = 0$ is not guaranteed as an equilibrium belief when there exists fractional searching. Also, as $\frac{\delta}{\delta+(1-\delta)\tau} > \delta$, consumers would have a higher belief of entrant producing H when fractional searching exists.

Bibliography

Abdellaoui, Mohammed (2000). “Parameter-free elicitation of utility and probability weighting functions”. In: *Management science* 46(11), pp. 1497–1512.

Abdellaoui, Mohammed, Peter Klibanoff, and Laetitia Placido (2011). *Ambiguity and compound risk attitudes: an experiment*. Tech. rep. Working paper. Northwestern University.

Aldrich, Eric M. and Daniel Friedman (2022). “Order Protection Through Delayed Messaging”. In: *Management Science* 69(2), pp. 774–790.

Aldrich, Eric M. and Kristian López Vargas (Mar. 2019). “Experiments in high-frequency trading: comparing two market institutions”. In: *Experimental Economics* 23(2), pp. 322–352. ISSN: 1573-6938. DOI: 10.1007/s10683-019-09605-2. URL: http://dx.doi.org/10.1007/s10683-019-09605-2.

Anderson, Lisa R and Jennifer M Mellor (2008). “Predicting health behaviors with an experimental measure of risk preference”. In: *Journal of health economics* 27(5), pp. 1260–1274.

Asano, Takao, Hiroko Okudaira, and Masaru Sasaki (2015). “An experimental test of a search model under ambiguity”. In: *Theory and Decision* 79, pp. 627–637.

Atkeson, Andrew, Christian Hellwig, and Guillermo Ordoñez (Nov. 2014). “Optimal Regulation in the Presence of Reputation Concerns *”. In: *The Quarterly Journal of Economics* 130(1), pp. 415–464.

Bagwell, Kyle (1990). “Informational product differentiation as a barrier to entry”. In: *International Journal of Industrial Organization* 8(2), pp. 207–223.

Barasinska, Nataliya, Dorothea Schäfer, and Andreas Stephan (2008). “Financial risk aversion and household asset diversification”. In.

Barsky, Robert B et al. (1997). “Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study”. In: *The Quarterly Journal of Economics* 112(2), pp. 537–579.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak (1964). “Measuring utility by a single-response sequential method”. In: *Behavioral science* 9(3), pp. 226–232.

Board, Simon and Moritz Meyer-ter-Vehn (2013). “REPUTATION FOR QUALITY”. In: *Econometrica* 81(6), pp. 2381–2462.

Bohnet, Iris and Steffen Huck (May 2004). “Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change”. In: *American Economic Review* 94(2), pp. 362–366.

Bolton, Gary, Ben Greiner, and Axel Ockenfels (2013). “Engineering Trust: Reciprocity in the Production of Reputation Information”. In: *Management Science* 59(2), pp. 265–285.

Bolton, Gary E., Elena Katok, and Axel Ockenfels (2005). “Cooperation among strangers with limited information about reputation”. In: *Journal of Public Economics* 89(8). The Experimental Approaches to Public Economics, pp. 1457–1468.

Brandts, Jordi, Paul Pezanis-Christou, and Arthur Schram (2008). “Competition with Forward Contracts: A Laboratory Analysis Motivated by Electricity Market Design”. In: *The Economic Journal* 118(525), pp. 192–214.

Brolley, Michael and Marius Zoican (June 2023). “Liquid speed: A micro-burst fee for low-latency exchanges”. In: *Journal of Financial Markets* 64, p. 100785. ISSN: 13864181. DOI: 10.1016/j.finmar.2022.100785.

Brown, Meta, Christopher J. Flinn, and Andrew Schotter (2011). “Real-Time Search in the Laboratory and the Market”. In: *American Economic Review* 101(2), pp. 948–974.

Budish, Eric, Peter Cramton, Albert Kyle, et al. (Apr. 2023). “Flow Trading”. In.

Budish, Eric, Peter Cramton, and John Shim (July 2015). “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response”. In: *The Quarterly Journal of Economics* 130(4), pp. 1547–1621. ISSN: 0033-5533. URL: https://doi.org/10.1093/qje/qjv027.

Caplin, Andrew, Mark Dean, and Daniel Martin (2011). “Search and Satisficing”. In: *American Economic Review* 101(7), pp. 2899–2922.

Charness, Gary, Ninghua Du, and Chun-Lei Yang (2011). “Trust and trustworthiness reputations in an investment game”. In: *Games and Economic Behavior* 72(2), pp. 361–375.

Che, Yeon-Koo and Johannes Hörner (Dec. 2017). “Recommender Systems as Mechanisms for Social Learning”. In: *The Quarterly Journal of Economics* 133(2), pp. 871–925.

Chen, Daniel L, Martin Schonger, and Chris Wickens (2016). “oTree—An open-source platform for laboratory, online, and field experiments”. In: *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.

Chetty, Raj and Adam Szeidl (2007). “Consumption commitments and risk preferences”. In: *The Quarterly Journal of Economics* 122(2), pp. 831–877.

Cho, Junghoo and Sourashis Roy (2004). “Impact of Search Engines on Page Popularity”. In: *Proceedings of the 13th International Conference on World Wide Web*. WWW ’04. Association for Computing Machinery: New York, NY, USA, pp. 20–29.

Cho, Junghoo, Sourashis Roy, and Robert E Adams (2005). “Page quality: In search of an unbiased web ranking”. In: *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pp. 551–562.

Choi, Syngjoo et al. (2007). “Consistency and heterogeneity of individual behavior under uncertainty”. In: *American economic review* 97(5), pp. 1921–1938.

Cooper, David J., Susan Garvin, and John H. Kagel (1997). “Signalling and Adaptive Learning in an Entry Limit Pricing Game”. In: *The RAND Journal of Economics* 28(4), pp. 662–683.

Cox, James C. and Ronald L. Oaxaca (1989). “Laboratory experiments with a finite-horizon job-search model”. In: *Journal of Risk and Uncertainty* 2, pp. 301–329.

Crosetto, Paolo and Antonio Filippin (2013). “The “bomb” risk elicitation task”. In: *Journal of Risk and Uncertainty* 47(1), pp. 31–65.

Dana, James D. and Yuk-Fai Fong (2011). “PRODUCT QUALITY, REPUTATION, AND MARKET STRUCTURE”. In: *International Economic Review* 52(4), pp. 1059–1076.

Degryse, Hans and Nikolaos Karagiannis (2022). “Priority Rules”. In: URL: https://ssrn.com/abstract=3186009.

Dinerstein, Michael et al. (July 2018). “Consumer Price Search and Platform Design in Internet Commerce”. In: *American Economic Review* 108(7), pp. 1820–59.

Dohmen, Thomas et al. (2011). “Individual risk attitudes: Measurement, determinants, and behavioral consequences”. In: *Journal of the european economic association* 9(3), pp. 522–550.

Eckel, Catherine C and Philip J Grossman (2002). "Sex differences and statistical stereotyping in attitudes toward financial risk". In: *Evolution and human behavior* 23(4), pp. 281–295.

Einav, Liran, Chiara Farronato, and Jonathan Levin (2016). "Peer-to-Peer Markets". In: *Annual Review of Economics* 8(1), pp. 615–635.

Einav, Liran, Theresa Kuchler, et al. (Apr. 2015). "Assessing Sale Strategies in Online Markets Using Matched Listings". In: *American Economic Journal: Microeconomics* 7(2), pp. 215–47.

Farrell, Joseph (1986). "Moral Hazard as an Entry Barrier". In: *The RAND Journal of Economics* 17(3), pp. 440–449.

Fradkin, Andrey (Aug. 2015). "Search Frictions and the Design of Online Marketplaces". In: ACM.

Fradkin, Andrey et al. (2015). "Bias and Reciprocity in Online Reviews: Evidence From Field Experiments on Airbnb". In: Association for Computing Machinery: New York, NY, USA, p. 641.

Friedman, Daniel, Sameh Habib, et al. (2022). "Varieties of Risk Preference Elicitation". In: *Games and Economic Behavior*.

Friedman, Daniel, R Mark Isaac, et al. (2014). *Risky curves: On the empirical failure of expected utility*. Routledge.

Ghose, Anindya, Panagiotis G. Ipeirotis, and Beibei Li (2012). “Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowdsourced Content”. In: *Marketing Science* 31(3), pp. 493–510.

Gneezy, Uri and Jan Potters (1997). “An experiment on risk taking and evaluation periods”. In: *The quarterly journal of economics* 112(2), pp. 631–645.

Greiner, Ben (2015). “Subject pool recruitment procedures: organizing experiments with ORSEE”. In: *Journal of the Economic Science Association* 1(1), pp. 114–125.

Haas, Marlene, Mariana Khapko, and Marius Zoican (2021). “Speed and learning in high-frequency auctions”. In: *Journal of Financial Markets* 54(C).

Heflebower, R. B. (1957). “Barriers to New Competition”. In: *The American Economic Review* 47(3), pp. 363–371.

Hey, John D and Chris Orme (1994). “Investigating generalizations of expected utility theory using experimental data”. In: *Econometrica: Journal of the Econometric Society*, pp. 1291–1326.

Hey, John D. (1987). “Still searching”. In: *Journal of Economic Behavior & Organization* 8(1), pp. 137–144.

Holt, Charles A. and Susan K. Laury (2002). “Risk Aversion and Incentive Effects”. In: *American Economic Review* 92(5), pp. 1644–1655.

Hörner, Johannes (June 2002). “Reputation and Competition”. In: *American Economic Review* 92(3), pp. 644–663.

Houser, Daniel and John Wooders (2006). “Reputation in Auctions: Theory, and Evidence from eBay”. In: *Journal of Economics & Management Strategy* 15(2), pp. 353–369.

Huck, Steffen, Gabriele K. Lünser, and Jean-Robert Tyran (2012). “Competition fosters trust”. In: *Games and Economic Behavior* 76(1), pp. 195–209.

Huck, Steffen, Gabriele K. Lünser, and Jean-Robert Tyran (2016). “Price competition and reputation in markets for experience goods: an experimental study”. In: *The RAND Journal of Economics* 47(1), pp. 99–117.

Jeon, Doh-Shin and Stefano Lovo (2012). “Reputation as an Entry Barrier in the Credit Rating Industry”. In: *IDEI Working Papers* 675.

Kittaka, Yuta and Ryo Mikami (2020). “Consumer Search and Stock-out: A Laboratory Experiment”. In: *ISER DP*(1104).

Kogut, Carl A. (1990). “Consumer search behavior and sunk costs”. In: *Journal of Economic Behavior & Organization* 14(3), pp. 381–392.

Kremer, Ilan, Yishay Mansour, and Motty Perry (2014). “Implementing the “Wisdom of the Crowd””. In: *Journal of Political Economy* 122(5), pp. 988–1012.

Krouse, Clement G. (1984). “Brand Name as a Barrier to Entry: The Rea Lemon Case”. In: *Southern Economic Journal* 51(2), pp. 495–502.

Kyle, Albert S and Jeongmin Lee (Nov. 2017). “Toward a Fully Continuous Exchange”. In: *Oxford Review of Economic Policy* 33(4), pp. 650–675. ISSN: 0266-903X.

Li, Lingfang (Ivy) and Erte Xiao (2014). “Money talks: rebate mechanisms in reputation system design”. In: *Management Science* 60(8), pp. 2054–2072.

Lippman, Steven A. and John J. McCall (1976). “The Economics of Job Search: A Survey”. In: *Economic Inquiry* 14(3), pp. 347–368.

Lynch, Michael et al. (1986). “Product Quality, Consumer Information and ”Lemons” in Experimental Markets”. In: *Empirical Approaches to Consumer Protection Economics*, pp. 251–306.

Mailath, George J. and Larry Samuelson (Apr. 2001). “Who Wants a Good Reputation?” In: *The Review of Economic Studies* 68(2), pp. 415–441.

Mason, Charles F. and Clifford Nowell (1992). “Entry, Collusion, and Capacity Constraints”. In: *Southern Economic Journal* 58(4), pp. 1002–1014.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier (Aug. 2014). “Promotional Reviews: An Empirical Investigation of Online Review Manipulation”. In: *American Economic Review* 104(8), pp. 2421–55.

Miller, Ross M. and Charles R. Plott (1985). “Product Quality Signaling in Experimental Markets”. In: *Econometrica* 53(4), pp. 837–872.

Miura, Takahiro, Keigo Inukai, and Masaru Sasaki (2017). “The effect of risk attitudes on search behavior: A laboratory search experiment”. In: *Discussion Papers In Economics And Business* 17(17).

Müller, Stephan and Georg Götz (2017). "Quality competition and entry deterrence: When to launch a second brand". In: *Managerial and Decision Economics* 38(8), pp. 1201–1215.

Nachman, David C (1972). *On risk aversion and optimal stopping*. Tech. rep. Discussion paper.

Palley, Asa B. and Mirko Kremer (2014). "Sequential Search and Learning from Rank Feedback: Theory and Experimental Evidence". In: *Management Science* 60(10), pp. 2525–2542.

Pandey, Sandeep et al. (2005). "Shuffling a stacked deck: the case for partially randomized ranking of search engine results". In: *arXiv preprint cs/0503011*.

Picone, Gabriel, Frank Sloan, and Donald Taylor (2004). "Effects of risk and time preference and expected longevity on demand for medical tests". In: *Journal of Risk and Uncertainty* 28(1), pp. 39–53.

Resnick, Paul et al. (Dec. 2000). "Reputation Systems". In: *Commun. ACM* 43(12), pp. 45–48.

Rouviere, Elodie and Raphael Soubeyran (2008). "Collective Reputation, Entry and Minimum Quality Standard". In: *FEEM Working Paper*(7).

Schmalensee, Richard (1982). "Product Differentiation Advantages of Pioneering Brands". In: *The American Economic Review* 72(3), pp. 349–365.

Schotter, Andrew and Yale M. Braunstein (1981). “Economic search: an experimental study”. In: *Economic inquiry* 19(1), pp. 1–25.

Schunk, Daniel (2009). “Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence”. In: *Journal of Economic Dynamics and Control* 33(9), pp. 1719–1738.

Schunk, Daniel and Joachim Winter (2009). “The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment”. In: *Journal of Economic Behavior & Organization* 71(2), pp. 347–360.

Sonnemans, Joep (1998). “Strategies of search”. In: *Journal of Economic Behavior & Organization* 35(3), pp. 309–332.

Ursu, Raluca M. (2018). “The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions”. In: *Marketing Science* 37(4), pp. 530–552.

Vellodi, Nikhil (Oct. 2018). “Ratings Design and Barriers to Entry”. In: *NET Institute Working Paper*(18-13).

Villas-Boas, J. Miguel (2004). “Consumer Learning, Brand Loyalty, and Competition”. In: *Marketing Science* 23(1), pp. 134–145.

Weitzman, Martin L. (1979). “Optimal Search for the Best Alternative”. In: *Econometrica* 47(3), pp. 641–654.

Zwick, Rami et al. (2003). “Consumer Sequential Search: Not Enough or Too Much?” In: *Marketing Science* 22(4), pp. 503–519.