Achieving Individually Fair Personalized Pricing

Thesis submitted in partial fulfillment of the requirements of the degree of

Master of Science in Computer Science and Engineering by Research

by

Shantanu Das 20171143 shantanu.das@research.iiit.ac.in



International Institute of Information Technology (Deemed to be University) Hyderabad - 500 032, INDIA June, 2023 Copyright © Shantanu Das, 2023 All Rights Reserved

International Institute of Information Technology Hyderabad, India

CERTIFICATE

It is certified that the work contained in this thesis, titled "Achieving Individually Fair Personalized Pricing" by Shantanu Das, has been carried out under my supervision and is not submitted elsewhere for a degree.

Date

Advisor: Prof. Sujit P Gujar

Date

Co-Advisor: Prof. Girish Varma

To my hard work and parents.

Acknowledgements

I would like to express my sincere gratitude to the following people for their invaluable support, guidance, and encouragement throughout the completion of my thesis:

Firstly, I would like to thank my advisor Prof. Sujit Gujar, and co-advisor Prof. Girish Varma for their unwavering guidance and support throughout the entire process. Their expertise and insights were instrumental in shaping the direction of my research and ensuring its successful completion.

I am also grateful to Prof. Shweta Jain, Prof. Ganesh Ghalme, and Dr. Swapnil Dhamal for their invaluable guidance throughout the duration of my research project. Their treasured comments and insights have helped greatly improve my understanding of the research topic and transform our efforts into an estimable acceptance at a prestigious conference.

I am also thankful to my research group at Machine Learning Lab. Braining our ideas and gaining invaluable insights into our common interests was a pleasure. I thank them for their encouragement, support, and valuable feedback throughout the entire process.

I would like to extend my sincere thanks to my friends and family for their unconditional love, support, and encouragement during this challenging time.

Lastly, I would like to thank everyone who participated in my study, without whom this research would not have been possible.

Thank you all for your support and guidance throughout this journey.

Abstract

The market is an indispensable aspect of our lives, enabling us to exchange valuable goods and services. Allocation and pricing are two significant challenges in determining the equilibrium dynamics of a market. With the advent of AI and technology, the way markets function has evolved greatly. With data analytics tools, sellers can engage in personalized pricing to maximize revenue. However, it is shown that personalized pricing is prone to unfairness among consumers. This thesis precisely studies the fairness challenges in personalized pricing. We aim to examine the concept of fairness in the context of monopoly markets that implement feature-based pricing. Our proposal introduces a new form of individual fairness, referred to as α -fairness, which ensures that individuals with similar characteristics are subjected to comparable pricing. To evaluate the loss in revenue to the seller in pursuit of fairness, we additionally introduce the notion of Cost of Fairness (CoF) – the ratio of the expected revenue in optimal feature-based pricing to the expected revenue in the given fair feature-based pricing.

First, we investigate the discrete valuation space and present an analytical solution for the most suitable fair feature-based pricing strategy. Our findings indicate that CoF, can be arbitrarily high for any fair pricing strategy. We observe that such valuation spaces are uncommon, so we focus on continuous valuation spaces that are well-studied in economics. With the standard assumption of the revenue function being continuous and concave with respect to the prices, we demonstrate that CoF is strictly less than two, regardless of the model parameters. Finally, we present a polynomial time algorithm that computes a fair feature-based pricing approach, successfully achieving CoF less than two.

Research Papers Based on the Thesis Work

Conference Papers

 Shantanu Das, Swapnil Dhamal, Ganesh Ghalme, Shweta Jain, and Sujit Gujar. Individual Fairness in Feature-Based Pricing for Monopoly Markets. Appeared in Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence, UAI 2022 [43].

Contents

Ch	apter		Page
1	Intro	duction	1
	1.1	Why Price Discrimination? (Discuss benefits through examples)	. 3
	1.2	Price Discrimination Markets	. 4
		1.2.1 Travel Fares	. 4
		1.2.2 Movies	. 5
	1.3	Unfair Price Discrimination	. 7
	1.4	Challenges and Problem Addressed	. 9
	1.5	Our Contributions	. 10
	1.6	Thesis Outline	. 11
2	Back	ground and Related Work	12
	2.1	Market Models	. 12
		2.1.1 Fisher's Model	. 13
		2.1.2 Arrow Debreu Model	. 15
	2.2	Competitive Markets	. 18
		2.2.1 Perfect Competition:	. 18
		2.2.2 Monopolistic Competition:	. 19
		2.2.3 Oligopoly:	. 20
		2.2.4 Monopoly:	. 20
	2.3	Pricing in Monopoly Markets	. 21
		2.3.1 Consumer Valuations	. 21
		2.3.2 Uniform Pricing	. 22
		2.3.2.1 Optimal Uniform Pricing	. 22
	2.4	Price Discrimination	. 22
		2.4.1 Conditions for Price Discrimination	. 23
		2.4.2 Types of Price Discrimination	. 24
		2.4.2.1 First-Degree Price Discrimination	. 24
		2.4.2.2 Second-Degree Price Discrimination	. 24
		2.4.2.3 Third-Degree Price Discrimination	. 24
		2.4.3 Feature-Based Pricing	. 25
		2.4.3.1 Market Segmentation in FP	. 25
		2.4.3.2 Consumer Features and Valuations	. 25
	2.5	Fairness	. 25
		2.5.1 Group fairness	. 25

		2.5.2	Individual Fairness
		2.5.3	Fairness in Pricing
	2.6	Import	ant Relevant Research
		2.6.1	Price Discrimination
		2.6.2	Fairness in Pricing 28
		2.6.3	The Limits of Price Discrimination
			2.6.3.1 Problem Description
			2.6.3.2 Solutions Proposed
			2.6.3.3 Claims
		2.6.4	Algorithmic Price Discrimination
			2.6.4.1 Problem Description
			2.6.4.2 Solutions Proposed
			2.6.4.3 Claims
		2.6.5	The Value of Personalised Pricing
			2.6.5.1 Solution proposed
			2.6.5.2 Claims
		2.6.6	Third-Degree Price Discrimination Versus Uniform Pricing
			2.6.6.1 Problem Description
			2.6.6.2 Solutions Proposed
			2.6.6.3 Claims
		2.6.7	The Value of Price Discrimination in Large Random Networks
			2.6.7.1 Problem Description :
			2.6.7.2 Solution proposed
			2.6.7.3 Claims
		2.6.8	Fairness, Welfare, and Equity in Personalized Pricing
			2.6.8.1 Problem Description
			2.6.8.2 Solutions Proposed
	2.7	Necess	sity for fairness in feature-based price discrimination
2	.		
3		/idual F	airness Model for Feature-Based Pricing
	5.1 2.2	Opuma	al FP
	3.2		1008
		5.2.1 2.2.2	α -faimess
	2.2	3.2.2 Ontinu	Cost of Fairness (CoF)
	3.3	Opum	al Fair Feature-Based Pricing (OP (FFP)
4	Disc	rete Val	uations
	4.1	Binary	^v Discrete Valuations
		4.1.1	Optimal Feature-based Pricing
		4.1.2	Optimal Fair Feature-based Pricing
		4.1.3	CoF Analysis
	4.2	CoF ca	an be arbitrarily bad

5	Cont	tinuous V	Valuations	.7
	5.1	Linear	methods to obtain Fair Feature-Based Pricing 4	8
		5.1.1	Pivot-Based linear pricing(LINP-FFP)	8
		5.1.2	CoF Bound	8
		5.1.3	Tightness of CoF bound:	1
	5.2	Propos	ed Algorithm	2
		5.2.1	Determining Optimal Pivot m	2
		5.2.2	Pricing Algorithm	2
6	Cond	clusions	and Future Work	6
	6.1	Conclu	sion	6
	6.2	Future	Work	7
Bił	oliogr	aphy .		7

List of Figures

Figure		Page
1.1	Classical Market (credits: [24] [©])	. 1
1.2	Online Market (credits: [45] [©])	. 2
1.3	Travel Fares (credits: [38] [©])	. 5
1.4	Uber Price Discrimination (credits: [9] [©])	. 6
1.5	Movie Tickets (credits: $[1]^{\odot}$)	. 7
1.6	Demand curves for binary market segments (credits: [34] [©])	. 7
1.7	Discriminatory pricing based on users' OS (credits: $[30]^{\degree}$)	. 8
2.1	Classification of Markets	. 13
5.1	Concave revenue function $\pi_i(\cdot)$ and its linear approximation $L_i(\cdot)$ (arrows show equations for $L_i(\cdot)$). Figure represents the case $\hat{p}_i - m \ge \alpha D_i/2$ for which LINP-FFP assigns $p_i = m + \alpha D_i/2$. The case $m - \hat{p}_i \ge \alpha D_i/2$ is similar.	. 50

List of Tables

Table	Page
2.1	Notations Table 1
3.1	Notations Table 2

Chapter 1

Introduction

"Pricing is an accurate and confident action that takes full advantage of the combination of customers' price sensitivity and alternative suppliers they have or could have."

-Miles

In this chapter, we describe the role of markets in our lives. We also discuss market pricing strategies and how they influence consumers' preferences.

Markets are systems of economic exchange that have been prevalent throughout human history. These systems are characterized by exchanging goods and services between *buyers* and *sellers*. The concept of classical markets is rooted in the principles of *supply* and *demand*, which are used to determine the *prices* of goods and services. The principles of markets are based on the idea of *free trade*. When the buyers



Figure 1.1: Classical Market (credits: [24][©])

and the sellers can negotiate prices without interference from external forces such as governments or other institutions, we call it free trade. In a free trading market, the prices rise when demand is high and fall when demand is low. Such markets are characterized by *competition* between buyers and sellers. In a competitive market, many buyers and sellers offer similar goods and services, and each party is free to negotiate prices to attract consumers or secure a sale. This competition helps ensure that prices remain reasonable and that goods and services are of high quality.

Markets have undergone significant changes over the years, particularly with the advent of technology. Today, many markets have moved online, with buyers and sellers conducting transactions through e-commerce platforms. These markets offer new opportunities for buyers and sellers to connect and conduct business. However, supply and demand still determine the prices in these online markets.



Figure 1.2: Online Market (credits: $[45]^{\odot}$)

As the internet and social media have become more widespread, personalized information about consumers has become more accessible. In modern markets, sellers utilize this data from online platforms or brokers to anticipate consumer preferences. This enables sellers to create more effective pricing strategies in order to generate higher revenues. One such pricing strategy sellers deploy is known as *price discrimination* or *discriminatory pricing*. In discriminatory pricing, the seller charges different groups differently for the same product or service.

This strategy is known to have been used for many years. Ideally, the seller would like to charge a consumer its valuation – the maximum amount a consumer is willing to pay. However, the seller might not be aware of such valuations. Another strategy could be to offer discounts to consumers who may vouch for the product/service. Even a seller might discount consumers willing to pay in cash or charge higher

prices to consumers in a hurry and less likely to shop around for a better deal. We can trace such price discrimination to the early 20th century when sellers began to recognize the potential benefits of charging different prices to different groups of consumers. One of the first examples of price discrimination occurred in the airline industry (Klein & Loebbecke [25]) when airlines began offering discounted prices for flights booked in advance or less popular flights. This allowed the airlines to increase their revenue by charging higher prices to consumers willing to pay more for a seat on a popular flight while offering lower prices to consumers ready to book a flight in advance or fly on a less popular route.

As sellers are typically unaware of the consumers' valuations, they rely on price-sensitive consumer data like purchase history, demographics, and other factors to forecast consumer preferences. In the modern, contemporary marketplace, with increased access to more informative consumer data, sellers use various techniques to learn more about consumers, leading to privacy and fairness concerns. This has made it necessary to comprehend price discrimination more thoroughly in theory and actuality.

The following section discusses a few examples of markets actively engaging in price discrimination to better understand this strategy in current marketplaces.

1.1 Why Price Discrimination? (Discuss benefits through examples)

As we saw, price discrimination is a popular strategy that sellers utilize in various industries and trades. However, to do so, sellers must spend significant resources collecting consumer data and making strong analytical inferences about their valuations to implement price discrimination efficiently. Given the overhead costs associated with price discrimination, it is natural to ask why firms often engage in this practice. Moreover, why is such a discriminatory practice considered legal by the authorities? We will try to answer these questions by discussing the benefits of price discrimination. There are several benefits of price discrimination for companies, consumers, and society as a whole. Some of the key benefits of this pricing strategy include the following:

- Increased revenue: Price discrimination can help a company increase revenue by charging higher
 prices to consumers willing to pay more for a product or service. For example, a movie theater may
 charge higher fees for tickets to popular movies during peak hours when demand is high, and consumers
 are willing to pay a premium to see the film. This allows the theater to capture more revenue from
 consumers willing to pay a higher price for a seat.
- Increased market share: Price discrimination can also help a company increase its market share by
 offering lower prices to specific groups of consumers. For example, a clothing retailer may provide
 discounts to students or senior citizens to attract these consumers and increase their market share. By

offering lower prices to these groups, the retailer can expand its customer base and gain a competitive advantage over other retailers.

- 3. More affordable prices for consumers: Price discrimination can also benefit consumers by allowing them to purchase products or services at more affordable prices. For example, a company may offer a discounted price for a subscription to its online service to consumers willing to pay for a longer-term subscription. This allows the customer to save money on the subscription, making the service more affordable.
- 4. Social Welfare: Price Discrimination is also beneficial due to its social value. If the producers in the market increase their output, price discrimination is proven to increase social welfare through the transactions in the market (Varian [44]). This can be understood by an example in Moriarty [33]. Regarding *medical drugs*, pharmaceutical companies charge lower prices in developing countries than in developed ones. However, if they were compelled to set the same price globally, they would opt for the higher price to make up for their marginal cost of production. As a result, they would likely stop supplying drugs at lower prices to people in developing countries, leaving them without access to essential medication.

Overall, price discrimination is a common pricing strategy used for many years and continues to be employed by sellers from many industries. While it can benefit both sellers and consumers in some cases, it can also be unfair to consumers if used to discriminate against certain groups and create confusion.

Now, in the following Section 1.2, we will discuss a few examples indicating the use of price discrimination in specific markets. These examples will demonstrate price discrimination's widespread utility and highlight genuine concerns regarding consumer fairness and privacy issues.

1.2 Price Discrimination Markets

1.2.1 Travel Fares

It is common to offer catered personalized prices while charging for travel fares. These pricing mechanisms are witnessed in all modes of transport, but price discrimination is most common in the airline industry. Consumers are typically divided into two groups based on the time between booking and travel dates: *leisure* and *business*. The unique behaviors between the two groups have led to this grouping method. Traveling for leisure consumers is more likely to be price sensitive; therefore, they often purchase tickets in advance to save money. When ticket costs increase as the departure date draws nearer, leisure travelers take advantage of *advance purchase* discounts. Leisure travelers would also have more elastic demand due to their price sensitivity. Those who travel for business require more flexibility than those who go for leisure; thus, they



Figure 1.3: Travel Fares (credits: [38][©])

would only make reservations a few days in advance. The presence of stricter demand implies a higher willingness to pay, and airlines may charge more for tickets. Similarly, airfares may also vary depending on the dates of travel. Travelers are usually willing to pay higher for traveling on holidays or weekends due to convenience and necessity. Airline industries identify these factors that affect consumer valuations and tend to extract most of the consumer surplus.

Deciding *cab fares* is another scenario where price discrimination is utilized extensively. Chang *et al.* [9] provide detailed empirical evidence on *Uber's price discrimination* strategy. They observe Uber charges higher cab fares to travelers who frequently stay at more expensive hotels. Based on their empirical evidence, for every \$100 increase in hotel room rates, the cab fare increased by \$0.10 to \$0.54. Cruise fares Namin *et al.* [36] have a similar pricing mechanism.

1.2.2 Movies

People go to movie theaters for leisure and typically prefer specific movie genres. Due to the higher demand for popular movies, people are willing to pay more for their preferred seats or convenient showtime. Since consumer preferences vary greatly, businesses can increase profits by using price discrimination to capture consumer surplus. Martin [29] explains market segmentation for movie theaters. Typically, the consumers are divided into two segments, viz. *nerdies* and *normies*.

Exhibit 1: Post About Discriminatory Uber Fares			
	Q		
Forum / United States / Hawaii / Oahu			
Uber - mark up if hotel is selected			
SimplyHuman · 1,512 forum posts New York Aug 18, 2021, 1:40 PM	\heartsuit		
I noticed a huge difference in price when my hotel is selected as opposed to picking an address right next to the hotel. As an example, Hyatt Regency to USS Arizona \$50, but using an address just outside the hotel, \$33. Same date, same time. this is just one example. This begs the question if Uber is building in a premium in its rates for hotel locations. I do understand that Uber has other pricing variable such as date and time of day, cab availability etc. my question is whether they are also using "hotel location" as a variable for pricing.			

Figure 1.4: Uber Price Discrimination (credits: [9][©])

Nerdies have specific preferences and are less likely to accept alternatives for the things they like. For example, they might prefer Marvel over DC or Sony over Philips. Typically, nerdies are ready to pay higher for their preferred choices. On the other hand, normies are more relaxed about their choices and do not pay as much attention to specific details. The sellers exploit this to gain higher revenue by offering different



Figure 1.5: Movie Tickets (credits: [1][©])

prices to the two types of consumers. However, to justify higher prices, they may offer additional services to the nerdies.



(a) Nerdie Demand





1.3 Unfair Price Discrimination

One might question whether charging different prices for the same product can cause consumers to distrust the companies and their products. Moreover, consumers may envy others who receive lower prices and perceive this as unfair. Consider the following scenario:

Example 1.1. A seller may offer products at low prices or for free to social media influencers who, in return, advertise the product. Now, consider two regular consumers, A and B, who have similar purchase histories, except consumer A regularly purchases products from the seller, but consumer B started buying from the seller recently. The seller offers a hefty discount to consumer A.

Consumer B will envy consumer A for the price and perceive their price as unfair compared to the price offered to consumer A, but may not feel the same towards the influencer. Price Discrimination affects our market or purchasing behavior and general social welfare; therefore, it is important to study pricing through the fairness lens.

Price discrimination has come under scrutiny recently due to its ethical concerns. We will now discuss a few examples of price discrimination that are considered unfair in the community.

Room Rates Orbitz is starting to show different results to users of Macs and PCs after finding Mac users spend more freely. In a recent search for hotels in Miami, Mac users saw more options over \$200. For El Paso, Texas, they saw more expensive options in the place of two cheaper ones.					
💣 Mac 🛛 Mi	ami 🥂 PC	💣 Mac	El Paso	🦊 PC	
1. Hyatt House \$118	1. Hyatt House \$118	1-5. Same for both	1-5. Same for bo	oth	
2. Design Suites \$124	2. Catalina Hotel \$209	6. Wyndham El Paso \$76	6. Travelodge \$	40	
3. Catalina Hotel \$209	3. Design Suites \$124	7. Studio Plus Deluxe \$54	7. Wyndham El	Paso \$76	
4. Churchill Suites \$189	4. The Richmond Hotel \$156	8. Hyatt Place El Paso \$76	8. Studio Plus D	eluxe \$54	
5. The Richmond Hotel \$156	5. Churchill Suites \$189	9. El Paso Marriott \$89	9. Days Inn \$55		
6. Eden Roc Renaissance \$212	6. Ocean Spray \$95	10. Radisson Hotel \$98	10. Hyatt Place	El Paso \$76	
7. The Palms Hotel & Spa \$224 7. South Seas Hotel \$175					
Source: WSJ searches of Orbitz that were performed at the same time for the same dates using a Mac with a Safari browser and a PC with Internet Explorer The Wall Street Journal					

Figure 1.7: Discriminatory pricing based on users' OS (credits: [30][©])

For instance, one notable example is *Orbitz*. It is an online travel agency showing higher prices to consumers visiting their page through Apple MacBook devices than the consumers using Windows devices. Mattioli [30], mentions that Mac users are shown 30% higher prices than Windows users. This sparked an outcry from both consumers and advocates alike.

Another example is *Uber's* strategy of charging personalized prices, which received heavy consumer backlash. Based on the consumers' *demand* for rides and the available *supply* of cab drivers, uber tends to apply price surges due to the natural interplay of supply and demand. Dholakia [14], describes this issue and proposes a few ideas Uber can adopt to regain their consumer's trust in the cab pricing. However, when booking a cab from the same location to the same destination at the same time, Uber shows different prices to different consumers (Mahadawi [27]). Consumers perceive this pricing strategy as unfair without having a say in it. A study conducted by Pandey & Caliskan [39] discovered that neighborhoods in Chicago with larger non-white populations, higher poverty rates, younger residents, and higher levels of education experienced higher taxi fares. This suggests that discriminatory practices may be present in the transportation industry.

Fine-grained data analysis in the e-commerce and retail industry has revealed several cases of price discrimination, as reported by Hinz *et al.* [21]. This has brought attention to the issue from regulatory bodies and the research community, as discussed by Michel [32]. While some economists have expressed concerns regarding the fairness of personalized pricing, laws and guidelines have been established to prohibit discrimination based on nationality or residence, as outlined by EU [17].

While price discrimination can sometimes be viewed as unethical, it is important to recognize that it can also serve several legitimate purposes Section 1.1, such as generating additional revenue, covering transportation and storage costs, expanding market reach, rewarding loyal customers, promoting a social cause, and others, as noted by Cassady [8]. However, it is crucial to conduct a thorough assessment to determine whether such practices may result in unfair treatment and prompt regulatory intervention, as discussed by Alan [3].

In this work, we focus on designing pricing strategies for a seller or monopolist who aims to maximize revenue through price discrimination while ensuring consumer fairness. By examining the different forms of price discrimination and their potential impact on consumers, we hope to shed light on this complex issue and contribute to a more comprehensive understanding of fair pricing practices in the marketplace.

1.4 Challenges and Problem Addressed

This section will briefly describe the main challenges we tackled and present an informal problem statement we solved in this work.

In Sections 1.2, 1.3, we saw that employing price discrimination not only requires an in-depth analysis of consumer valuation for efficiency but also demands careful consideration of consumer perceptions of prices, regulatory rules, and social welfare. Therefore, it is necessary to study price discrimination strategies under the scrutiny of the fairness perspective.

Challenge 1: The biggest challenge in fairness literature is defining fairness. There is no single definition of fairness that can be applied everywhere; instead, it is necessary to *formulate a fairness criterion* based on the task at hand.

Since a seller aims to maximize their revenue based on the prices, the corresponding pricing mechanism is typically modeled as an optimization problem to maximize the seller's revenue Bergemann *et al.* [4]. However, adding fairness constraints to the pricing problem inevitably leads to sub-optimal prices.

Challenge 2: The next challenge is finding the best sub-optimal prices which maximize the revenue under fairness constraints.

This work tackles the above challenges to ensure fairness in price discrimination strategies. Specifically, we work with *Feature-Based Pricing (FP)* in the monopoly market setting. This is an essential formulation in the literature on fair pricing strategies. A variant of fairness criteria and the price discrimination model has not been studied earlier. In the next section, we list our contribution briefly.

1.5 Our Contributions

- To capture the notion of individual fairness, proposed by Dwork *et al.* [15], we propose a fairness criterion α -fairness, which offers similar prices to similar individuals in feature-based pricing. We emphasize that fairness concerns may arise when individuals with comparable characteristics are separated into different segments and charged different prices, making it difficult to compare prices based on those characteristics. To compare the revenue generated by fair prices to the optimal revenue-maximizing prices, we define a quantity *CoF*: Cost of Fairness, the ratio of optimal revenue to fair revenue.
- To formulate the problem of maximizing revenue under α-fairness constraints, we propose a formulation *FFP*: Fair Feature-Based Pricing. We analyze optimal *Opt-FFP*: Optimal Fair Feature-Based Pricing and its revenue guarantees for a market with discrete valuations and binary segmentation. Next, we show that a constant lower bound on CoFis generally impossible.
- Next, we assume the revenue functions exhibit concavity with respect to the offered prices ¹, and demonstrate that a constant upper bound on the cost of fairness can be attained. In this context, we show that the seller can calculate the *Opt-FFP* prices using a convex program. This is possible provided the seller has access to distributional information, including knowledge of the valuation distribution functions of all consumers.
- Next, we identify a class of *FFP* strategies, namely *LINP-FFP*: Linear Fair Feature-Based Pricing, that satisfies α -fairness. With the help of these pricing strategies, we then show that the CoF is strictly less than 2 irrespective of model parameters used to predict consumer valuations from features.
- Finally, we propose an algorithm *OPT-LINP-FFP*: Optimal Linear Fair Feature-Based Pricing. Given the optimal revenue-maximizing prices from the solutions of feature-based pricing, *OPT-LINP-FFP*

¹A commonly accepted assumption in economics given that a significant number of probability distributions follow this pattern, as pointed out by Bergemann *et al.* [5]

returns a list of optimal LINP-FFPprices which satisfy α -fairness. The runtime complexity for this algorithm is $O(k \log(k))$, where k is the number of segments. Without access to complete distributional information, it computes α -fair pricing and achieves the aforementioned CoF.

1.6 Thesis Outline

The contents of my thesis are briefly described below.

- In Chapter 2, we will begin our discussion by providing the necessary background on *general equilibrium theory* of markets and describe two fundamental market models, viz. *Fisher Market* and *Arrow-Debreu Market*. We will emphasize the hardness of equilibrium computation in the presence of indivisible goods. Further, we will discuss *competitive markets* and their classification based on the degree of competition. Next, we will present the theoretical framework for understanding pricing strategies in monopoly markets. Next, we will introduce important fairness criteria in the pricing literature. Finally, we will provide a brief review of papers that compare the performance of uniform pricing vs. personalized pricing under specific assumptions on valuations distribution.
- In Chapter 3, we will formally introduce our fairness criteria α -fairness and a performance metric CoF to compare the revenue of potentially unfair prices vs. fair prices. Next, we will formulate the fair pricing problem OPT_{FFP} that maximizes revenue and satisfies α -fairness.
- In Chapter 4, we will analyze the solutions to OPT_{FFP} for discrete prices and consumer valuations. Then, we will show that it is necessary to assume standard restrictions on consumer valuations else, the CoF may turn out to be arbitrarily bad.
- In Chapter 5, we will solve OPT_{FFP} for continuous consumer valuations. Under the standard concavity assumptions on revenue functions, we will analyze CoF and provide a constant upper bound of 2. Next, we will identify a class of prices LINP-FFP, that achieve this bound and provide an algorithm to solve for LINP-FFP.

Chapter 2

Background and Related Work

"The market is like a lake agitated by the wind, where the water is incessantly seeking its level without ever reaching it."

- Leon Walras

In this Chapter, we provide the necessary background to understand the problem formulation. We also provide a non-exhaustive literature review and discuss a few relevant papers. Finally, we present an informal description of our problem statement.

This chapter will introduce markets and describe two main market models in literature Section 2.1, with the corresponding equilibrium conditions. Next, in Section 2.2, we discuss the classification of markets based on the number of sellers and the type of interactions between them. Then in Section 2.3, Section 2.4, we introduce the pricing strategies for these market models and formulate the revenue maximization problem in each case. In Section 2.5, we provide background on fair learning and pricing literature. In Section 2.6, we discuss the important relevant papers followed by a brief review of a few selected papers on related topics. Finally, in Section 2.7, we address the need for individual fairness in pricing and provide an informal description of the problem addressed in this work.

2.1 Market Models

Designing a market model and studying the problem of market equilibria involves determining a set of prices and allocating goods to consumers that satisfy their budget constraints while maximizing their total valuation. Moreover, the sellers should empty their stock, and the market should be cleared. Economists have been studying this problem since the 19th century, developing various models to capture the concept of market equilibrium. One of the earliest models was introduced by *Walras* in his 1874 publication, *"Elements of Pure Economics,"* which describes the equilibrium of an economic system in terms of *supply* and *demand*. This model expressed the equilibrium conditions where supply equals demand. The existence

of an equilibrium for the Walrasian system was first proved by Wald in 1936, though under severe restrictions. In 1954, Arrow and Debreu proved the existence of equilibrium under much milder assumptions and were awarded the Nobel Prize for their contribution.

The following section will discuss two celebrated market models in the economics literature and provide relevant notations.



Figure 2.1: Classification of Markets

2.1.1 Fisher's Model

The *Fisher market model*, given by Irving Fisher (see Brainard & Scarf [7]) is one of the first fundamental market models in the mathematical economic theory of markets. The model assumes a market where consumers aim to purchase goods while maximizing their total valuation within their budget constraints. The market equilibrium is reached when the demand for goods equals the supply. In this equilibrium, consumers exhaust their budgets, all goods are sold, and each consumer's purchased bundle maximizes their valuation given their budget and the equilibrium prices. Fisher's model is significant in the literature because desired computational properties exist during the computation of equilibrium prices for the goods under mild realistic assumptions on the consumers' valuation functions.

Formally, a Fisher market model \mathcal{M} consists of a set of agents called *consumers* denoted by $\mathbf{B} = \{1, ..., m\}$ and the set of *goods* denoted by $\mathbf{A} = \{1, ..., n\}$. The model defines each consumer *i*:

- an initial *budget* $e_i > 0$, which may be thought of as cash that can be used to purchase items but has no inherent worth to the consumer, and
- a non-negative *valuation* function v_i : ℝⁿ → ℝ₊ that converts a bundle of m goods representing the quantity of each good to a real-valued valuation. Upon obtaining y_i ∈ ℝ^m things, the consumer's valuation is represented by v_i(y_i).

Without loss of generality, the initial budgets assigned to the set of consumers are normalized as $\sum_{i}^{m} e_i = 1$, and the quantity of each good is assumed to be 1. These are typical assumptions used to study and analyze the Fisher market model for convenience. In Fisher's market, every consumer aims to exhaust their budget money, and in exchange, they wish to procure a set of goods with maximum valuation according to their valuation function. Let $\mathbf{p} = (p_1, \dots, p_n)$ be the *price* vector for goods, where p_j represents the price for the good *j*, and let $\mathbf{y} = (y_1, \dots, y_m)$ where $y_i \in \mathbb{R}^n_+$, $\forall i \in \mathbf{B}$ be the set of *allocations* of goods to the set of consumers where y_{ij} is the amount of good *j* allocated to consumer *i*.

Definition 2.1 (*Market Outcome*). A tuple of price vector for n goods and the corresponding allocation vectors $\langle \mathbf{p}, \mathbf{y} \rangle$ is defined as the market outcome.

An allocation **y** is valid when the assigned bundles of goods y_i are within the market's budget constraints for each consumer *i*. Formally, for a consumer *i*,

$$\sum_{j}^{n} p_j \ y_{ij} \le e_i \tag{2.1}$$

When the market outcome satisfies the consumers' objective, i.e., maximizes their individual utilities by allocating the optimal bundle of goods given their budget constraints, and all the goods are *sold*, or the market is cleared, then we say that the market equilibrium is achieved. The formal definition of market equilibrium is given by (Nisan *et al.* [37]):

Definition 2.2 (*Equilibrium (Fisher Market)*). A market outcome $\langle \mathbf{p}, \mathbf{y} \rangle$ is in equilibrium if and only *if*,

- For all consumers i the allocated bundle y_i maximizes their valuation for a given price vector **p** i.e., y_i ∈ argmax v_i(y_i), ∀i ∈ B
- Every good is either sold completely or is assigned a price 0, i.e., $(\sum_{i}^{m} y_{ij}-1)p_{j} = 0, \forall items j \in \mathbf{A}.$
- The allocation **y** is such that the initial budgets for all consumers are spent completely, i.e., $\sum_{i}^{n} y_{ij} p_j = e_i, \forall i \in \mathbf{B}$

According to the theory of general equilibrium of markets (Maxfield [31]), such a market equilibrium exists if at least one consumer seeks each item and each consumer desires at least one good in the market. Under some standard assumptions on the valuation functions of the consumers, this equilibrium can be computed using the renowned Einsenberg-Gale convex program.

The Fisher Market model assumes the existence of an expendable currency, or money in the form of an initial budget endowment, i.e., e_i for all consumers $i \in \mathbf{B}$, for an equilibrium. This is a strong assumption that restricts the modeling of more general markets. To resolve this, Arrow and Debreu proposed a more general setting which we discuss below.

2.1.2 Arrow Debreu Model

The *Arrow-Debreu model* provides a framework for analyzing a market economy with a finite number of goods, consumers, and sellers. The model assumes that consumers have well-defined valuations for goods: continuous, non-satiated, and convex. Additionally, each consumer has an initial allocation of the goods, with a positive amount of at least one item.

Remark: This is the key difference between the Fisher and Arrow-Debreu models. The condition of *initial budget*, e_i is dropped, and consumers can have any arbitrary endowment to participate in the Arrow-Debreu market.

Sellers use non-increasing returns to scale or constant returns to scale technology to transform inputs into outputs. Each seller aims to optimize their revenue, while each consumer maximizes their valuation within budgetary constraints.

Similar to the Fisher market model, in the Arrow-Debreu model, there are *n* goods indexed by $j \in \{1, ..., n\}$, and a set of *m* consumers, indexed by $i \in \{1, ..., m\}$. Each seller/consumer *j* has an endowment

of the goods, denoted by e_{ij} , and a valuation function $v_i(\mathbf{y})$, where $\mathbf{y} = (y_1, \dots, y_n)$ is a bundle of goods that *j* can consume.

The model assumes a set of *competitive markets*, one for each good *j*, in which sellers can trade the good at a price p_j . The initial endowment of a consumer *i* is given by $\mathbf{q_i} = (q_{i1}, \dots, q_{in})$, where q_{ij} is the amount of good *j* the consumer *i* owns initially. Similar to the Fisher market model, we assume that the total quantity of each good available is 1, i.e., $\sum_{i \in [m]} q_{ij} = 1$ is true for all $j \in [n]$.

The Arrow-Debreu model assumes that all consumers have complete information about the market and that they maximize their valuation subject to their budget constraints. The *budget constraint* for agent i is given by:

$$p_1 y_{i1} + p_2 y_{i2} + \ldots + p_n y_{in} \le p_1 e_{i1} + p_2 e_{i2} + \ldots + p_n e_{in}$$

$$(2.2)$$

 y_{ij} is the amount of good j consumed by agent i, and the right-hand side is the agent's initial endowment.

The model seeks to find an equilibrium set of prices \mathbf{p} at which there is no longer any surplus demand for any good and no profit for sellers. In other words, the market clears, and all goods are sold at a price the consumers are willing to pay while sellers break even. This equilibrium concept is formally captured in Definition 2.3. Note for this equilibrium:

- *Price taking:* Each agent is assumed to be a price taker, meaning they do not influence the market price for any good.
- *Efficiency:* The market equilibrium is efficient if no other allocation of goods can make any agent better off without making another agent worse off.

The Arrow-Debreu model is flexible enough to be applied to different contexts. For example, it can be extended to spatial or inter-temporal settings by properly defining the goods depending on their location or delivery time. The model can also incorporate expectations and uncertainty in the analysis by treating goods as conditional on distinct global situations.

The model has been used to evaluate financial and monetary markets, international trade, and the effects of policy changes on resource allocation, among other topics. It provides a general equilibrium structure that captures a market economy's complex interactions between consumers, sellers, and goods.

Definition 2.3 (*Equilibrium (Arrow-Debreu Market)*). A market outcome $\langle \mathbf{p}, \mathbf{y} \rangle$ is in equilibrium if and only if,

- Market clearing: for all j = 1, ..., n, the total quantity of good j demanded by all agents must equal the total quantity of good j supplied, i.e., $\sum_{i=1}^{n} y_{ij} = 1$
- Individual optimization: each agent *i* chooses a consumption bundle y_i^* that maximizes their valuation $v_i(y)$ subject to their budget constraint, i.e.,

$$y_i^* \in \operatorname{argmax} v_i(y)$$
, subject to $\mathbf{p}^{\mathsf{T}} \mathbf{y_i} \leq \mathbf{p}^{\mathsf{T}} \mathbf{e_i} \quad \forall i$

where e_i is the agent's initial endowment.

Indivisible Goods (Special Case)

It is a well-known fact in the general equilibrium theory of markets, an equilibrium (similar to the ones described in Definition 2.2 and Definition 2.3) does not necessarily always exist. The market equilibrium with indivisible goods induces a combinatorial optimization problem which makes the computation of market equilibrium hard (if it exists). To understand this better, let us consider the set of all feasible allocations in the market for both divisible and indivisible goods. For a market with *n* divisible goods and *m* consumers, the set of feasible allocations is an *m*-tuple of *n* sized real vectors, i.e. $\mathcal{A}_{div} := \{\mathbf{y} | \mathbf{y} = \{y_1, \dots, y_m\}$ where $y_i \in \mathbb{R}_+, \forall i \in [m]\}$. Note that the search space is O(mn). However, for *n* indivisible goods with *m* consumers, the feasible allocation set is given by $\mathcal{A}_{indiv} := \{0, 1\}^{n \times m}$. The search space is $O(2^{mn})$.

It is shown by Bikhchandani & Mamer [6] (and many others in this line of work) that if there exists a single divisible good ^{*a*} alongside other indivisible goods, market equilibrium is guaranteed to exist under the assumption of *quasilinear consumer valuations* later relaxed by Fujishige & Yang [18].

In the earlier sections of this chapter, we witnessed markets from a third-eye astute perspective where everything about the market is known, viz., the number of consumers & sellers, the valuation functions of the consumers, the production quantities, etc. The existence of a market equilibrium is discussed under the assumption that the market characteristics are common knowledge. What happens to the market prices if these characteristics are unknown? In practice, the sellers set the prices of goods with the objective of *revenue maximization*. Given this real-world situation, we must address some important questions: How is the market equilibrium affected under the pricing strategies with the objective of revenue maximization? How do sellers

 $^{^{}a}$ Typically, money is assumed to be an indivisible good in the economics literature. In such settings, money is the divisible good available in the market unless stated otherwise.

respond to the *competitive pricing* offered by their market rivals? These queries provide information about a specific aspect of the markets, such as *competition* in the market or *competitive markets*. According to available literature, with underpricing strategies for revenue maximization, the market equilibrium may shift towards a higher price and lower quantity traded. Sellers may respond to competitive pricing offered by their market rivals by matching the price or offering a differentiated product to maintain their market share. In the remainder of this chapter, we will discuss competitive markets in detail.

2.2 Competitive Markets

Market competition exists from a seller's perspective because multiple sellers offer similar products or services to many consumers. As a result, sellers are forced to compete with one another to attract consumers and sell their products or services.

Competition is driven by the desire to capture market share and increase profits. With more sellers in a market, each seller's market share becomes smaller. To maintain or increase their market share, sellers must offer competitive prices, better quality products and services, or unique value propositions that set them apart.

In a competitive market, sellers must be more efficient and innovative. To remain competitive, they must find ways to produce their products or services at lower costs or with higher quality. This leads to increased productivity and efficiency, which can benefit the seller, the consumer, and the overall economy.

Ultimately, competition benefits the consumer, resulting in lower prices, better quality products or services, and more choices. However, it also forces sellers to continuously improve and innovate to stay ahead. The competition also encourages sellers to be more efficient in their operations and to seek out new markets and consumers, which can lead to economic growth and job creation. In the economics literature, the competitive markets are classified into four main categories viz., *perfect competition markets, monopolistic competition markets, nonopoly markets, monopoly markets*. These markets will be discussed briefly below.

2.2.1 Perfect Competition:

A competitive market is a type of marketplace characterized by the absence of any single entity that has the power to influence the market price. The market comprises multiple consumers and sellers, typically small sellers that produce similar goods. This can be formally characterized as $|\mathbf{B}| >> |\mathbf{A}|$ Consumers have full access to market information, which plays a crucial role in determining prices through the interaction of supply and demand. In a competitive market, no individual seller has significant market power. The output level is optimal as sellers have limited control over market prices. The products in such a market are homogeneous or identical to competitors', and there are no significant barriers to entering or exiting the market.

This market structure is considered highly efficient for allocating resources and *maximizing social welfare*. However, it is rare in practice and is mainly a theoretical model that provides a reference point for other market structures. It can also be used as a benchmark for evaluating the performance of other markets.

Examples of pure competition market structures include *stock markets*, where numerous consumers and sellers interact to determine the price of stocks, *agricultural markets*, where farmers sell their produce to consumers at a prevailing market price; and *craft markets*, where artisans sell their handmade goods to consumers at market prices. These examples highlight the features of a competitive market, such as the absence of market power, identical products, and the interaction of supply and demand forces.

2.2.2 Monopolistic Competition:

Monopolistic competition, which is comparable to pure competition, describes a market structure in which several tiny sellers compete with one another. Such small sellers engaged in monopolistic competition offer very similar but distinctive products. The key difference between monopolistic and pure competitive markets is in the former, it is not expected that manufacturing will be at the lowest feasible cost, producing the best results in a market, as is the case with pure competition.

These factors give sellers in a monopolistic competition market power to charge higher prices within a certain range. The products are remarkably similar, but small differences are the basis for sellers' marketing and advertising. Differentiation can include styles, brand names, locations, packaging, advertisements, pricing strategies, etc.

Examples include fast food restaurants, clothing stores, breakfast cereal sellers, service and repair markets, tutors, beauty salons, and spas. Products and services at a beauty salon are similar, but these sellers will use certain value propositions, such as quality of services and appealing pricing, to attract more consumers. They may even advertise brand-name beauty products that are themselves in monopolistic competition; there is little that separates makeup and hair products as far as what constitutes these products and their use. Another example is the market for fast-food restaurants, where there are numerous sellers offering similar types of food, but each seller has some control over the price they charge due to differences in quality, service, or location.

In monopolistic competition, sellers have some market power but are still price-takers in the long run. This means they cannot sustainably charge higher prices than their competitors without losing consumers. Businesses can freely enter the market when profits are attractive and there is easy entry and exit in monopolistic competition.

2.2.3 Oligopoly:

Few sellers control an oligopoly, which has little competition. They can work together or compete against one another to increase prices and profits by using their combined market power. Getting into an oligopoly is challenging. The most dominant businesses have control over raw materials, patents, and financial and physical resources, which disadvantages future rivals. This contributes to high pricing. However, consumers will resort to product alternatives available on the market if prices are too high. This sheds light on a *fundamental trade-off* that sellers in a market frequently face: setting a competitive price low enough to attract consumers and maintain market share, but, at the same time, the prices need to be high enough to generate high revenue and make a profit.

In such a market, products can be of two types: *uniform* in nature, i.e., all sellers in the oligopoly offer the same product, or the sellers can regulate various other factors such as quality, promotion offers, discounts, and so on to differentiate their product from the products their market competitors offer.

For instance, Microsoft, Sony, and Nintendo are the three sellers controlling the video game systems market. The auto and petroleum industries serve as other examples of oligopolies. When the dynamic relationship between sellers and consumers evolves, so do pricing, profitability, and output levels.

2.2.4 Monopoly:

In a monopoly market, there is only one seller of a specific product or service and no direct competitors. The monopolist has complete control over the product's supply. As a result, the seller commands substantial market power and the ability to control the cost of the goods or services they offer.

Monopoly markets often arise due to barriers to entry, which prevent other sellers from entering the market and competing with the monopolist. These barriers can be natural, such as a unique natural resource or economies of scale. They can also be created through legal or regulatory means, such as patents or exclusive licenses.

A monopolist can set higher pricing for its goods than in a competitive market due to its market strength, leading to bigger profits. Consumers may have to pay more for the goods, and the market is likely to produce less and operate less efficiently than it would if it were competitive. Due to the absence of competition, monopolies are not motivated to innovate or improve their products. As a result, consumers may be stuck with outdated or low-quality goods at higher prices. Due to these factors, governments usually regulate monopoly markets to prevent the misuse of market dominance and safeguard the interests of consumers.

In a monopoly market, equilibrium is characterized by the intersection of the monopolist's *marginal revenue (MR)* and *marginal cost (MC)* curves. The marginal revenue curve shows how much additional

revenue the monopolist would earn by selling one additional unit of output, while the marginal cost curve shows how much additional cost the monopolist would incur by producing one additional unit of output.

The monopolist chooses the quantity of output where MR equals MC, which maximizes its profit. At this quantity, the monopolist sets a price higher than the marginal cost of producing a unit of output, resulting in a higher price and lower output than in a competitive market.

The monopolist limits output to raise prices, which increases the overall revenue generated but also impacts the general welfare of market consumers.

2.3 Pricing in Monopoly Markets

In a monopoly market, there is a single seller for the products or items present in the market. Consumers can buy products only from a single monopolistic seller. In such a setting, the seller controls the entire market share and decides the prices for the products in the market so that the revenue gains the seller obtains are maximized. The monopolist can employ many pricing strategies to achieve the maximum possible revenue return.

A simple and naive approach to setting prices is to assign the same price of the product to all the consumers (or consumer types which we will describe in later sections) irrespective of their valuations for the product, this strategy is called *uniform pricing*. A monopolist typically resorts to uniform pricing in the absence of any information about consumer valuations.

The seller could charge different prices to different consumers based on their willingness to pay for the product. Consumers who perceive a higher value for the product will agree to pay a higher price than others, generating higher revenue for the seller. The pricing strategy then boils down to correctly identifying the consumer valuations and offering a price equal to their valuations. This develops a sense of discrimination of prices offered to consumers, hence, this kind of pricing strategy is termed as *price discrimination*.

Before getting into the monopoly pricing strategies, we will introduce some notations for consumer valuation functions.

2.3.1 Consumer Valuations

Every consumer has a valuation function $v_i : \mathbb{R}^n_+ \to \mathbb{R}_+$ that maps the bundle of *n* goods to the value perceived by the individual consumers. However, this work assumes n = 1, i.e., a single good market with money. The valuation function signifies the amount consumers will pay against the purchase. Particularly, valuation is the maximum amount of money a consumer is willing to pay for a given quantity of the product.

Let **B** = {1,..., *m*} be the set of consumers and let v_i be the valuation of the unit amount of product corresponding to the *i*th consumer. The valuations are drawn i.i.d. from a fixed support set $D_{\mathcal{F}}$ according

to the valuation distributions with pdf f_i , and cdf \mathcal{F}_i for each consumer $i \in \mathbf{B}$. We further assume some standard regularity conditions on the valuation distributions popular in literature (Myerson [35]).

2.3.2 Uniform Pricing

The monopolist offers a single price from the support set of consumer valuations $p \in D_{\mathcal{F}}$ for the single product offered in his monopoly market to all the consumers $i \in \mathbf{B}$, and the consumers choose to either make a purchase of a unit amount of the product at the set price p (if $p \leq v_i$), or they do not make the purchase at all (if $p > v_i$). Therefore, for a given consumer i, the probability of making the purchase at a price p is given by, $\mathbb{P}(p \leq v_i) = 1 - \mathcal{F}_i(p)$. The expected revenue generated from each consumer i,

$$\pi_i(p) = p(1 - \mathcal{F}_i(p))$$

2.3.2.1 Optimal Uniform Pricing

The optimal price p_{\star} , which maximizes the total expected revenue for the seller in the monopoly market in the uniform pricing mechanism, is given by the following optimization formulation.

$$p_{\star} \in \underset{p \in D_{\mathcal{F}}}{\operatorname{argmax}} \sum_{i \in \mathbf{B}} \pi_i(p) \tag{2.4}$$

(2.3)

In the following section, we formally introduce the price discrimination strategy. We also discuss the challenges faced by sellers to implement price discrimination efficiently.

2.4 Price Discrimination

In this pricing strategy, the seller divides the consumers into *market segments* based on consumer characteristics. For example, in student/senior citizen discounts, the market is split into two segments, while in electronic retail pricing, the sellers can split the market based on multiple features. Let $\mathbf{M} := \{1, ..., k\}$ be the set of market segments, and let $\mathbf{p} = \{p_1, ..., p_k\}$ be the price vector such that p_i is the price assigned to the *i*th market segment by the seller. The revenue generated by *i*th segment is given by:

$$\pi_i(p_i) = p_i(1 - \mathcal{F}_i(p_i)) \tag{2.5}$$

where \mathcal{F}_i is the cdf for valuation distribution in the *i*th. Note that the price terms differ in Eq. (2.3) and in Eq. (2.5).

The revenue-maximizing prices $\hat{\mathbf{p}}$ are computed by solving the following optimization problem:

$$\widehat{\mathbf{p}} \in \underset{\mathbf{p} \in \mathbb{R}^k_+}{\operatorname{argmax}} \sum_{i}^k \pi_i(p_i)$$
(2.6)

It is important to note that after the market segments are fixed, the optimal price for a given segment is independent of the optimal price for any other segment in the market as the seller aims to maximize his revenue from each segment. Therefore, optimal prices for all segments $i \in \mathbf{M}$ can be computed independently of other segments. Formally, for all segments i, the optimal price \hat{p}_i is given by:

$$\widehat{p}_i \in \operatorname*{argmax}_{p \in \mathbb{R}_+} \pi_i(p) \tag{2.7}$$

Now that we have a formal understanding of the price discrimination strategy, in the following subsections we will discuss market conditions for price discrimination (Section 2.4.1) and its types (Section 2.4.2).

2.4.1 Conditions for Price Discrimination

For the price discrimination strategy to be viable, the market should satisfy three conditions (Varian [44]). First, the seller should be able to exert some market power. In price discrimination, the prices offered could be greater than the marginal cost of production, usually offered to consumers willing to pay a higher amount in exchange for the purchase. Moreover, to increase sales, the seller might also offer lower prices, sometimes even below the cost of production. However, considering the objective of maximizing revenue, lower prices can only be offered to marginal consumers or a group of similar consumers. This is possible when the seller has significant influence in the market.

Secondly, to successfully offer discriminated prices to the consumers, the seller should be able to sort the consumers into segments with similar valuations and prevent the consumers from strategically shifting across segments. The seller should be able to prevent strategic *self-selection* by the consumers into segments with a lower price.

Third, the seller should be able to prevent *resaling* in the market. This is important because if the product purchased in lower price segments is sold at higher prices by the consumers, the seller will experience a significant loss in revenue.

¹With notational abuse, we denote both the valuation functions for i^{th} segment in price discrimination and i^{th} consumer in market models with cdf \mathcal{F}_i , and pdf f_i

2.4.2 Types of Price Discrimination

The price discrimination strategy is divided into three categories based on how consumers are divided into market segments (Pigou [40], Varian [44]).

2.4.2.1 First-Degree Price Discrimination

First-degree or perfect price discrimination, also known as *personalized pricing*, is the most extreme form of price discrimination. In this case, the seller charges each consumer *i* exactly their valuation v_i , extracting the entire consumer surplus. In personalized pricing, the set of market segments coincides with the set of consumers, i.e., $\mathbf{M} = \mathbf{B}$, and the number of market segments equals the number of consumers.

First-degree price discrimination is usually more challenging to implement in practice, as the seller must access accurate information about each consumer's valuation (i.e., v_i for all $i \in \mathbf{B}$). It may also result in some consumers being priced out of the market, leading to a potential loss of revenue for the seller.

2.4.2.2 Second-Degree Price Discrimination

Second-degree price discrimination is a pricing strategy where a seller charges different prices based on the quantity q_i of a product or service purchased by a consumer *i*. In this strategy, consumers who purchase larger quantities of a product or service will receive a lower price per unit than those who purchase smaller quantities, i.e., if $q_i > q_j$, then $\frac{p_i}{q_i} < \frac{p_j}{q_j}$. The aim of quantity discrimination is to incentivize consumers to purchase larger quantities of a product or service, which can lead to increased revenue and profits for the seller. This strategy can also help the seller to clear excess inventory or sell perishable goods before they expire.

2.4.2.3 Third-Degree Price Discrimination

Third-degree price discrimination, also known as *group pricing* is a pricing strategy where a seller charges different prices to different groups of consumers based on their characteristics. The market is divided into k different consumer types typically identified by their features $x_i \in X$. The aim of this strategy is to maximize profits by catering to different market segments, each with its own valuation. The main challenge is the optimal segmentation of consumers into groups based on characteristics such as age, income, location, or other demographic factors. However, discriminatory prices based on groups may develop a sense of unfair treatment of the consumers.
2.4.3 Feature-Based Pricing

One of the conditions for price discrimination is that the seller should be able to segment the consumers in the market and offer different prices to each segment. One way of segmenting the market is based on collecting consumer data. Sellers can collect consumer data in many forms like mode of payment, purchase history, etc. Price Discrimination involving market segmentation based on consumer data/features is usually termed as *Feature-Based Pricing (FP)*. We will explore this pricing model in detail in Chapter 3.

2.4.3.1 Market Segmentation in FP

Consider a monopoly market where the seller has an infinite supply of goods². Every consumer in the market is associated with a feature vector $x_i \in \mathbb{R}_+$, where the index *i* represents the *i*th unique feature observed by the seller. The market is divided into finite number of segments $X = \{x_1, x_2, ..., x_k\}$, where x_i represents the *i*th segment. The seller, given access to X, can price discriminate across segments to extract maximum revenue. Note that multiple consumers may have the same feature vector, and all the consumers with identical features belong to the same market segment.

2.4.3.2 Consumer Features and Valuations

Elmachtoub *et al.* [16] discusses an important aspect of feature-based pricing. The goal of market segmentation in price discrimination is to segregate consumers based on their valuations, i.e., consumers in the same segment should ideally have the same or similar valuations for the product. For this to be valid in feature-based pricing, the consumer features should be informative of the corresponding consumer valuations. This notion is formally captured in (Elmachtoub *et al.* [16]).

2.5 Fairness

A crucial challenge in fair learning tasks is formulating a definition of fairness and relevant fairness measures. The fairness criteria in literature are primarily categorized into *group fairness* and *individual fairness*.

2.5.1 Group fairness

Group fairness measures are defined to mitigate bias in decision-making among different groups. These groups are typically identified based on sensitive attributes such as age, race, gender, etc. Examples of group fairness criteria are *statistical parity* and *equalized odds*. Consider a prediction function $f : X \to Y$

²We assume that the seller never runs short of goods and has enough supply to match the market demand

where X is the set of input features and Y is the set of labels. A is a set of binary protected attribute where A = 0, (A = 1) correspond to majority (minority) classes. Statistical parity is expressed as:

$$\mathbb{P}(f(x) = y | x, A = 0) = \mathbb{P}(f(x) = y | x, A = 1)$$

whereas equalized odds are given by:

$$\mathbb{P}(f(x) = 1 | Y = 1, A = 0) = \mathbb{P}(f(x) = 1 | Y = 1, A = 1)$$
, and
 $\mathbb{P}(f(x) = 1 | Y = 0, A = 0) = \mathbb{P}(f(x) = 1 | Y = 0, A = 1)$

2.5.2 Individual Fairness

On the other hand, Individual Fairness focuses on treating similar individuals in a similar manner, regardless of their membership in a protected group. The idea is that if two individuals are similar in terms of their characteristics, they should receive similar treatment and predictions from the model, irrespective of their sensitive attributes. Consider two individuals with feature vectors $x, x' \in X$ and d, a metric defined on feature space X. Then individual fairness is defined as (Dwork *et al.* [15]):

$$|f(x) - f(x')| \le d(x, x')$$

2.5.3 Fairness in Pricing

The fairness criteria examined in pricing literature are defined from the perspective of the seller whose goal is to satisfy the given fairness constraints in the process of maximizing revenue. Cohen *et al.* [10] formulated four fairness measures from the seller's perspective.

Consider the price discrimination in a monopoly market described in Section 2.4 with binary market segmentation (k = 2). Let β_i denote the fraction of consumers in the i^{th} segment. V_i is the valuation sampled from a distribution with cdf \mathcal{F}_i . We start with price fairness to enforce that the prices of two segments are similar.

Definition 2.4 (*Price Fairness*). Let \hat{p}_i denote the optimal revenue-maximizing prices for each segment *i* and let p_i be some suboptimal set of prices. Then, p_i satisfies price fairness when,

$$|p_1 - p_0| \le (1 - \alpha)|\widehat{p}_1 - \widehat{p}_0|$$

Next, we have demand fairness.

Definition 2.5 (*Demand Fairness*). Let $\overline{\mathcal{F}}_i(p_i) := \mathbb{P}(v_i \ge p_i)$ be the market share and $\beta_i \overline{\mathcal{F}}_i$ the demand associated with segment *i*. Following the price notations in Definition 2.4, p_i satisfies demand fairness when,

$$|\bar{\mathcal{F}}_1(p_1) - \bar{\mathcal{F}}_0(p_0)| \le (1 - \alpha)|\bar{\mathcal{F}}_1(\widehat{p}_1) - \bar{\mathcal{F}}_0(\widehat{p}_0)|$$

Since the consumers typically pay a price less than their actual valuation, this disparity is termed *consumer surplus*. Next, we have a fairness criterion that involves consumer surplus.

Definition 2.6 (*Surplus Fairness*). Let $S_i(p_i) := \mathbb{E}[(V_i - p_i)]$ be the consumer surplus for segment *i*. Following the price notations in Definition 2.4, p_i satisfies surplus fairness when,

 $|S_1(p_1) - S_0(p_0)| \le (1 - \alpha)|S_1(\widehat{p}_1) - S_0(\widehat{p}_0)|$

Consumers with a valuation less than the allocated price do not make the purchase. This is termed as *no-purchase valuation*. The fairness criterion based on no-purchase valuation is defined as:

Definition 2.7 (*No-purchase valuation Fairness*). Let $N_i(p_i) := \mathbb{E}[V_i|V_i < p_i]$ be the no-purchase valuation for segment *i*. Following the price notations in Definition 2.4, p_i satisfies no-purchase valuation fairness when,

$$|N_1(p_1) - N_0(p_0)| \le (1 - \alpha)|N_1(\hat{p}_1) - N_0(\hat{p}_0)|$$

Having defined all necessary notation, terminology, and fairness, we now brief a few important, relevant works in the following subsection.

2.6 Important Relevant Research

In this section, we present a brief literature review on a comparative study of uniform pricing versus price discrimination followed by fairness in pricing.

2.6.1 Price Discrimination

The concept of price discrimination is first discussed in Pigou [40] with a market of two segments and linear valuations. After this, a stream of literature is available on the effects of price discrimination on social welfare. Aguirre *et al.* [2] identify sufficient market conditions for price discrimination to either increase or decrease social welfare. Next, Cowan [11] and Bergemann *et al.* [4] discuss the sufficient conditions for an increase in consumer surplus under third-degree price discrimination.

A second line of work on price discrimination explores the revenue-maximizing objective from the sellers' perspectives. Schmalensee [42] and Varian [44] connect the two lines of work and discuss the conditions on valuation functions that cause an in-social surplus when the seller's revenue increases.

Bergemann *et al.* [4] also shows that under complete information about the consumers' valuation distributions, there exists a set of price vectors and market segmentation achieving any linear combination of consumer surplus and seller surplus under certain necessary constraints. Cummings *et al.* [12] extend this work under access to only partial information on consumers' valuations. They consider the market segmentation for maximum social surplus from an intermediary's perspective and provide approximate guarantees. Malueg & Snyder [28], Bergemann *et al.* [5], Elmachtoub *et al.* [16] discuss the revenue benefits of price discrimination over uniform pricing based on the available information on consumers' valuations.

When the valuations of the consumers are not known, Elmachtoub *et al.* [16] propose feature-based pricing and provide bounds on the value generated using idealized personalized pricing and Feature-based pricing over Uniform pricing. The value of feature-based pricing depends on the correlation between valuations and consumer features. Huang *et al.* [22] consider the first-degree price discrimination over the social network where the centrality measures in social networks determine the characteristics of the consumers. They provide bounds on the value of network-based personalized pricing in large random social networks with varying edge densities.

2.6.2 Fairness in Pricing

Recently, many questions have been raised on the ethical side of price discrimination methods. Moriarty [33] strongly criticizes online personalized pricing and suggests that personalized prices compete unfairly for the social surplus the transactions create. Gerlick & Liozu [19] points out the need to design personalized pricing with ethical considerations, which can provide win-win outcomes for organizations and consumers. Richards *et al.* [41] discusses that discriminatory pricing leads to the perception of unfairness amongst consumers, undermining retail platforms' stability. They discuss that when consumers are involved in forming the prices, this leads to improved fairness perception, thus leading to better retentivity. Levy & Barocas [26] discusses that web-based platforms typically use many private features of user profiles to connect consumers and sellers. Users interacting on such media leads to discrimination regarding race, gender, and other protected characteristics. All these studies lead to understanding the optimal price discriminatory strategies under the fairness constraint, which is the focus of our work.

Finally, Kallus & Zhou [23] presents a list of metrics like *price disparity*, *equal access*, *allocative efficiency fairness* to measure and analyze fairness in feature-based pricing and study its interplay with welfare. The metrics discussed are mainly the group fairness notions entirely different from the α -fairness discussed in this paper. Though the above papers discuss the ethical issues in price discrimination, we emphasize that none provide a systematic approach to designing the pricing strategy that maximizes the revenue and ensures a fairness guarantee.

We will review a few selected papers from the literature comparing revenues generated by uniform pricing Eq. (2.4) and price discrimination strategies Eq. (2.6). Bergemann *et al.* [4] present a study regarding revenue comparison of third-degree price discrimination with uniform pricing when sellers have access to consumer valuations which was later relaxed by Cummings *et al.* [12]. Elmachtoub *et al.* [16] introduced feature-based pricing and compared its revenue with uniform pricing. Next, [5] show that appropriate market segmentation can achieve any combination of consumer and seller surplus. Finally, Huang *et al.* [22] considers a variant of price discrimination for consumers embedded in social networks.

2.6.3 The Limits of Price Discrimination

This paper Bergemann *et al.* [4] explores the benefits of using third-degree price discrimination (and first-degree price discrimination, a special case of third-degree price discrimination) in monopoly markets over uniform pricing by characterizing the effects on both consumer and seller surplus. They further analyze the impact of market segmentation on social welfare in price discrimination under complete knowledge of consumer valuations.

2.6.3.1 Problem Description

A monopolist sells a good to a continuum of consumers, each demanding one unit. The consumer valuations are sampled from the set $\mathcal{V} = \{v_1, \dots, v_r\}$ and a market is represented as a distribution over \mathcal{V} as a simplex of (r-1) dimensions. The goal is to design a mechanism that returns optimal market segmentation that maximizes social welfare and revenue.

2.6.3.2 Solutions Proposed

They present a special class of market segmentation called *Extremal Markets* with certain useful properties. Further, they prove that any segmentation can be written as a linear combination of *Extremal Markets*. To solve the problem, they provide welfare and revenue bounds based on certain assumptions.

2.6.3.3 Claims

• The first major claim (result) is given by:

Surplus Triangle: There exists a segmentation and optimal pricing rule with consumer surplus u and seller surplus s if and only if $u \ge 0$, $s \ge s^*$, and $u + s \le w^*$.

Extremal and Direct Segmentation: For any segmentation and optimal pricing rule (σ, **p**), there exist:
(i) an extremal segmentation and an optimal pricing rule (σ', **p**') and (ii) a direct segmentation σ''
(and associated direct pricing rule **p**'') that achieve the same joint distribution over valuations and prices. As such, they achieve the same seller surplus, consumer surplus, total surplus, and output.

2.6.4 Algorithmic Price Discrimination

In this paper, Cummings *et al.* [12] consider a generalization of the third-degree price discrimination problem studied in Bergemann *et al.* [4], where an intermediary between the consumer and the seller can design market segments to maximize any linear combination of consumer surplus and seller revenue. They consider three models of information about the consumers' valuations and provide optimal market segmentations.

2.6.4.1 Problem Description

They consider a third-degree price discrimination model with finite consumers $\mathbf{B} = \{1, ..., m\}$ and valuations $\mathcal{V} = \{v_1, ..., v_r\}$. Next, they consider three different price discrimination models of consumer valuations.

- (*Bayesian Model*) The intermediary and the seller know the value-type distributions: \mathcal{F}_i for all $i \in \mathbf{B}$, and distribution over consumers β .
- (Sample Complexity Model) The intermediary and the seller observe a batch of signal-value pairs sampled from the underlying distribution. The idea is to estimate the number of samples required to get an *ε* approximation.
- (*Online Learning Model*) In this model, the intermediary attempts to learn the segmentation online using only bandit feedback from the consumers' decision to purchase or not at the seller's chosen price.

2.6.4.2 Solutions Proposed

A summary of model-wise solutions in this paper are:

- *(Bayesian Model)* They show that the optimal segmentation can be computed using a linear program (LP).
- (*Sample Complexity Model*) In this model, since the intermediary is unaware of the valuations' distributions, they estimate the underlying distributions based on the available samples under standard conditions like boundedness and regularity. They construct MHR-like distributions and provide sample complexity bounds for the same.

• (*Online Learning Model*) Optimal segmentation in this model is achieved by employing *no-regret* algorithms. Here the intermediary can only observe the purchase decision instead of the consumer valuation (in the above model).

2.6.4.3 Claims

- In the Bayesian Model, the range of achievable values for consumer surplus and revenue depends on the distribution, and one may not always be able to achieve the full consumer surplus.
- In the Sample Complexity model, under the standard assumptions on underlying distributions for each consumer type we need a polynomial number of samples for ϵ approximation of optimal segmentation.
- Online Learning model can be reduced to a sample complexity model by replacing the observed valuations with binary responses observed by online bandit corresponding to consumer purchases.

2.6.5 The Value of Personalised Pricing

In this paper, Elmachtoub *et al.* [16] considers three different pricing strategies and introduces *Feature-Based Pricing* for the first time. They compare the revenue generated through the proposed pricing strategies by taking a ratio of the corresponding revenues and providing theoretical bounds.

2.6.5.1 Solution proposed

To prove lower and upper bounds on the quantities $\frac{\Pi(\hat{\mathbf{p}})}{\Pi(p_{\star})}$ and $\frac{\Pi(\hat{\mathbf{p}})}{\Pi(\hat{\mathbf{p}_{\star}})}$. Note that $\hat{\mathbf{p}}_{\mathbf{x}}$ corresponds to the optimal feature-based price vector. The consumer valuations, $v_i \sim \mathcal{F}_i$ govern the above ratios the authors aim to bound. To do so, the authors use the following estimates on consumer valuations.

$$S := \frac{\inf (v_i | \mathcal{F}_j(v_i) = 1)}{\mu}, M := 1 - \frac{c}{\mu}, \mathcal{D} := \frac{\mathbb{E}[|v - \mu|]}{2\mu}$$

These three statistics are re-scaled unit-less quantities associated with the valuations V that represent *maximal valuation, per-unit cost, and mean absolute deviation* respectively.

2.6.5.2 Claims

The paper provides bounds on the value generated using *personalized pricing* and *feature-based pricing* over *uniform pricing*. The bounds are partitioned into three cases based on the heterogeneity of the market which is captured by the Mean Absolute Deviation \mathcal{D} of the valuation distribution F, the three categories are:

• (L) Low Heterogeneity: $0 \le \mathcal{D} \le \delta_L$

- (M) Medium Heterogeneity: $\delta_L \leq \mathcal{D} \leq \delta_M$
- (H) High Heterogeneity: $\delta_M \leq \mathcal{D} \leq \delta_H$,

where $\delta_L, \delta_M, \delta_H$ are constants that depend on M and S :

$$\delta_L := -\frac{M \log\left(\frac{S+M-1}{M}\right)}{W_{-1}\left(\frac{-1}{e^{\frac{S+M-1}{M}}}\right)}, \quad \delta_M := \frac{M \log\left(\frac{S+M-1}{M}\right)}{1 + \log\left(\frac{S+M-1}{M}\right)}, \quad \delta_H := \frac{M(S-1)}{S+M-1}$$

The bounds on the ratio of revenues for each market category discussed above are:

Closed-form, tight upper bounds are provided for the value of personalized pricing over uniform pricing (i.e. Π(p̂)/Π(p⋆)) when the scale, margin, and coefficient of deviation of the valuation distribution are known. When these upper bounds are small, suggests that the value of any personalized pricing strategy is limited.

Upper Bounding $\frac{\Pi(\hat{p})}{\Pi(p_{\star})}$ using *S*, *M*, and *D*). For any *F* with scale *S* > 1, margin *M* > 0, and coefficient of deviation *D*, we have the following: (Note, -W() is the Lambert W-function)

1. If $0 \leq \mathcal{D} \leq \delta_L$, then

$$\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \leq \frac{-W_{-1}\left(\frac{\bar{M}-1}{e}\right)}{1-\frac{\mathcal{D}}{M}}$$

(Low Heterogeneity)

2. If $\delta_L \leq \mathcal{D} \leq \delta_M$, then

$$\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \le \frac{M \log\left(\frac{S+M-1}{M}\right)}{\mathcal{D}}$$

(Medium Heterogeneity)

3. If $\delta_M \leq \mathcal{D} \leq \delta_H$, then

$$\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \le -W_{-1}\left(\frac{-1}{e\left(\frac{S+M-1}{M}\right)\left(1-\frac{\mathcal{D}}{M}\right)}\right)$$

(High Heterogeneity)

Moreover, for any S, M, D there exists a valuation distribution F with scale S, margin M and coefficient of deviation D such that the corresponding bound is tight.

• Closed-form lower bounds on the value of idealized personalized pricing that rely on necessary shape assumptions such as uni-modality or uni-modality and symmetry. The bound is tight for any specified

coefficient of deviation(\mathcal{D}). These lower bounds provide guarantees on how much value personalized pricing can provide as a function of the market heterogeneity.

Lower Bounding $\frac{\Pi(\hat{p})}{\Pi(p_{\star})}$ using \mathcal{D} . Consider a valuation distribution $V \sim F$, with margin M > 0 and coefficient of deviation \mathcal{D} .

1. If V is unimodal and symmetric, then

$$\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \ge \frac{1}{1 - 2\frac{\mathcal{D}}{M}}$$

Moreover, for every value of $\frac{\mathcal{D}}{M}$ there exists a unimodal and symmetric distribution such that this bound is tight.

2. If V is unimodal and

$$-0 \le \frac{\mathcal{D}}{M} \le \frac{1}{3}$$
, then

$$- 0 \le \frac{\mathcal{D}}{M} \le \frac{1}{3}$$
, then

$$\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \geq \frac{8\frac{\mathcal{D}}{M}}{\left(1 + \frac{\mathcal{D}}{M}\right)^2}$$

 $\frac{\Pi(\widehat{p})}{\Pi(p_{\star})} \ge \frac{1}{1 - \frac{\mathcal{D}}{M}}$

Moreover, if $\frac{\mathcal{D}}{M} = 0$, this bound is tight, and, as $\frac{\mathcal{D}}{M}$ tends to 1, there exists a family of unimodal valuation distributions such that this bound is asymptotically tight.

• The lower and upper bounds for feature-based pricing are generated on its value compared to personalized and uniform pricing. These bounds make explicit the relationship between the informational value of the features and the value of feature-based pricing in a market.

2.6.6 Third-Degree Price Discrimination Versus Uniform Pricing

Bergemann *et al.* [5] compare the seller's profit generated by optimal uniform pricing and third-degree price discrimination in a monopoly market.

2.6.6.1 Problem Description

They consider a monopoly market with k types of consumers where β_i is the fraction of consumers in the *i*th segment. The main contribution of this paper is a systematic comparison of revenues generated by personalized pricing versus uniform pricing. They also provide a tight bound on the ratio of revenues generated through each pricing approach.

2.6.6.2 Solutions Proposed

Under the standard assumptions on the revenue functions $\pi_i()$ described by:

- All the revenue functions $\pi_i()$ share a common support \mathcal{V} .
- Revenue functions are concave.

They prove a constant lower bound on the *profit-ratio* given by $\frac{\Pi(p_{\star})}{\Pi(\hat{\mathbf{p}})}$.

Note: We will use a similar quantity to compare performance of fair versus unfair feature-based pricing schemes.

2.6.6.3 Claims

They prove the optimal uniform price performs at least half as good as the optimal price discrimination policy under the standard assumptions mentioned above, and the bound is tight. The claim is formally stated as:

Proposition 2.1. Suppose that the profit functions $\pi_i(p)$ are concave and defined in the same bounded interval $[\underline{v}, \overline{v}]$ for all $i \in \{1, ..., k\}$. Then the uniform price $p_s = \underline{v} + \overline{v}/2$ delivers a 1/2-approximation for the monopolist's profits.

2.6.7 The Value of Price Discrimination in Large Random Networks

In this paper, Huang *et al.* [22] extends the existing literature on *network-based* pricing. They evaluate the benefits of price discrimination versus uniform pricing where consumers can influence other consumers' valuations for the product or service offered in the monopoly market.

2.6.7.1 Problem Description :

A monopolist sells a divisible product to *n* consumers embedded in a social network *G*. The social network *G* is modeled as an Erdos Renyi random graph with *n* vertices and a given non-zero edge probability $\rho(n)$.

The optimal pricing $\hat{\mathbf{p}}$ and profit earned for optimal pricing $\Pi(\hat{\mathbf{p}})$ are:

$$\widehat{p_i} = \left(\frac{a+c}{2}\right)\mathbf{1} + \left(\frac{a-c}{2}\right)\frac{\rho}{||G+G^T||}(G-G^T)\,\kappa\left(G+G^T,\frac{\rho}{||G+G^T||}\right)$$

where $\kappa \left(G + G^T, \frac{\rho}{||G + G^T||}\right)$ is the Bonacich Centrality, and *a* is the standalone consumer valuation.

$$\Pi(\widehat{\mathbf{p}}) = \frac{1}{2} \left(\frac{a-c}{2}\right)^2 \mathbf{1}^T \left(I - \frac{\rho}{||G+G^T||}(G+G^T)\right) \mathbf{1}$$

Bonacich Centrality ($\kappa(G, \alpha)$): It is the discounted sum of weighted walks of all lengths ending at node *i* for all nodes in *G*. The discount factor for a walk of length *l* is α^{l} .

2.6.7.2 Solution proposed

They classify the network-based problem into three major cases arising in random graphs based on the network density $\rho(n)$. They compare personalized pricing with uniform prices for each case.

To perform this comparison, they perform asymptotic analysis and estimate the value of positive network effects, which is given by

$$VoN(n) = \mathbf{E}_{G}[\Pi(\widehat{\mathbf{p}})] - \frac{n}{2} \left(\frac{a-c}{2}\right)^{2}$$

They estimate $\mathbf{E}_G[\Pi(\widehat{\mathbf{p}})]$ by estimating the spectral norm of $||G + G^T||$ and the number of weighted walks of length *l* in graph *G* and multi-graph $G + G^T$.

The authors prove and use the estimates for large random graphs in three cases viz. *very sparse graph*, *sparse graphs*, and *dense graphs* for computing the Value of Price Discrimination in asymptotic settings.

2.6.7.3 Claims

The authors provide the necessary and sufficient conditions for a graph to be price discrimination-free is given by:

Theorem 2.1. *G* is a price discrimination-free network if and only if $G^k 1 = (G^T)^k 1$ for each positive integer *k*.

Further, they analyze regret growth rates for social network densities.

2.6.8 Fairness, Welfare, and Equity in Personalized Pricing

In this paper, Kallus & Zhou [23] presents a list of metrics to measure and analyze fairness in First Degree Price discrimination or Personalized Pricing and study its interplay with welfare.

2.6.8.1 Problem Description

The paper discusses some problems around fairness and welfare in the personalized pricing model. They attempt to analyze the multiple objectives that pricing decisions affect with fairness and welfare implications.

2.6.8.2 Solutions Proposed

This paper considers a feature-based pricing model associated with a demand curve that depends on the price set by sellers. More importantly, the consumer features contain a binary sensitive attribute A. Let P be the price variables for consumers, and $d : \mathbb{R}_+ \to \mathbb{R}_+$ be the demand function. The authors propose the following *group fairness* criteria in the context of personalized pricing and dynamic pricing.

• *Price Parity:* This metric allocates equalized prices across consumer groups separated by the protected attributes. Formally, the fairness condition is given as

$$\mathbb{E}[P|A=a] = \mathbb{E}[P|A=b]$$

• *Model Error Fairness:* A major part of the personalized pricing model is to predict the valuations of the consumers, given their feature vectors. This raises the question of whether this price suboptimality affects the groups equally. Hence the suboptimal price function due to a predicted demand curve is adopted and is given as:

$$\widehat{\mathbf{p}} \in \underset{\mathbf{p}}{\operatorname{argmax}} \mathbb{E}[p \cdot d(p)]$$

• *Access and Equal Access:* A Personalized Pricing strategy may affect the demand curve based on the availability of the good. Therefore the paper describes a fairness metric to deal with the accessibility of the good. Access parity is given as:

$$\mathbb{E}[d(P)|A = a] = \mathbb{E}[d(P)|A = b]$$

• Allocative Efficiency and Fairness: This applies to the settings where we must sort the consumers using prices based on their valuations. Assume monotonicity in demand curve $p > p' \implies D(p) \ge D(p')$, then the fairness metric, concordance is given as

$$P(v_1 > v_2 | p_1 > p_2)$$

The following section describes an informal problem statement we address in this work.

2.7 Necessity for fairness in feature-based price discrimination

As we saw in the previous discussions, price discrimination benefits the seller as it generates increased revenue. However, it has many fairness concerns. Our work aims to formulate and solve individual fairness problems in feature-based price discrimination. Let the price vector $\hat{\mathbf{p}}$ be the optimal solution to the revenue-maximizing feature-based pricing. Consider two segments *i*, *j* identified by the features $x_i, x_j \in X$ respectively. However, the prices \hat{p}_i, \hat{p}_j may be unfair. Let **p** be another price vector that satisfies the fairness constraints such that the revenue generated using the fair prices is maximized. Specifically, maximize the objective defined in Eq. (2.5) under the additional fairness constraints.

In the next chapter, we discuss our formulation of the fairness constraints required to obtain fair prices and provide an algorithm to convert given optimal prices into fair prices. Moreover, in doing so, we analyze the loss in revenue and provide a constant upper bound on the *Cost of Fairness* (Section 3.2), which is a measure to compare revenues generated by different prices.

Notation	Description	
\mathcal{M}	Market Model	
B, A	Set of buyers & goods in a market respectively	
e_i, v_i	Initial budget & valuation of consumer <i>i</i> respectively	
q_i	initial endowment of consumer <i>i</i>	
р	Price vector for <i>n</i> goods	
	Amount of good <i>j</i> allocated to consumer <i>i</i>	
$\pi_i(p_i)$	Expected revenue generated in the i^{th} segment	
p_{\star}	Optimal Uniform Pricing	
p	Price vector for OPT _{FP}	
$\mathcal{F}_i, f_i()$	Valuations CDF, PDF for i^{th} consumer segment respectively	
X	Set of all consumer features/types	
$D_{\mathcal{V}}$	V Support set of consumers' valuations	
x _i	<i>i</i> Consumer feature of the <i>i</i> th segment	
β_i	The fraction of consumers in the i^{th} segment	
П(р)	Total expected revenue generated by p	

Table 2.1: Notations Table 1

Chapter 3

Individual Fairness Model for Feature-Based Pricing

This chapter will introduce α -fairness – the fairness criteria we propose in this work, and *CoF* – a ratio of revenues to measure loss in revenue due to fairness. We further formulate *OPT_{FP}* – Optimal Feature-Based Pricing, and *OPT_{FFP}* – Optimal Fair Feature-Based Pricing.

Group fairness strategies aim to achieve fairness across groups categorized according to sensitive attributes such as race or gender. Conversely, individual fairness ensures fairness between every pair of individuals in the feature space. Our proposed approach, α -fairness in feature-based pricing, is inspired by individual fairness in the Machine Learning literature. It emphasizes that similar individuals should receive similar treatment. Our alpha fairness criterion can be applied at individual and group levels, depending on the granularity of the available features in feature-based pricing. Specifically, alpha fairness ensures fairness among population segments that are distinguished from each other based on their feature vectors. Before presenting our fairness definition, we will formally define the feature-based pricing strategy.

3.1 Optimal FP

Every consumer in the monopoly market is associated with a feature vector x. The set of all features is denoted by $X := \{x_1, \ldots, x_k\}$ where i^{th} market segment is identified by the feature vector $x_i \in X$. As mentioned earlier in Section 2.4.3, multiple consumers may have the same feature vector and all the consumers with identical features belong to the same market segment. In other words, the set of consumers is partitioned by unique features. A feature-based price function $p : X \to \mathbb{R}_+$ maps the consumer features to their prices. We will interchangeably use $p(x_i)$ and p_i to simplify the notations. The Feature-Based pricing is modeled as an optimization problem denoted by a 4-tuple OPT_{FP}($\mathcal{V}, X, \mathcal{F}, \beta$), where consumer valuations v_i are sampled from common support $\mathcal{V}, \mathcal{F} := \{\mathcal{F}_1, \ldots, \mathcal{F}_k\}$ is the set of cdfs for valuation distributions, and $\beta := \{\beta_1, \ldots, \beta_k\}$ is the set of fractions of consumers in each segment. We denote the expected revenue per consumer in segment *i* by $\pi_i(p_i)$ Eq. (2.3), and expected revenue over all segments $\Pi(\mathbf{p}) := \sum_{x_i \in X} \beta_i \pi_i(p_i).$ We assume that β_i s are known to the seller.

In the absence of fairness constraints, $OPT_{FP}(\cdot)$ reduces to charging each segment separately and optimal FP strategy $\hat{\mathbf{p}}$ consisting \hat{p}_i for segment *i* is given by $\hat{p}_i \in \underset{p \in \mathbb{R}}{\operatorname{argmax}} \pi_i(p_i)$.

3.2 Definitions

This section provides two important definitions regarding fairness in feature-based pricing (Shantanu *et al.* [43]). The first definition is one of the major contributions of our work, i.e., the definition of the fairness criterion we propose. Meanwhile, the second definition is to quantify the revenue loss incurred due to the introduction of fairness constraints in OPT_{FP} . The second definition is important for the theoretical analysis provided in the remainder of this thesis.

3.2.1 α -fairness

Let $d : X \times X \to \mathbb{R}_+$ be a distance function over X. We assume that such a function exists and is well defined in X, i.e., (X, d) is a metric space. The distance function quantifies the dissimilarity between feature vectors of individuals belonging to market segments. For simplicity we write $d(x_i, x_j) := d_{ij}$. Individual fairness in FP strategy is defined as:

Definition 3.1 (α -fairness). A price function $\mathbf{p} : X \to \mathbb{R}_+$ is α -fair with respect to d iff for all $x_i, x_j \in X$, we have

 $|p_i - p_j| \le \alpha \cdot d_{ij}. \tag{3.1}$

We call a pricing strategy *Fair Feature-based Pricing* (α -*FFP*) which satisfies Eq. (3.1) for a given value of α . It is easy to see from the definition that any α -FFP is also α' -FFP for any $\alpha' \ge \alpha$. We will drop the quantifier α and call it FFP when it is clear from the context.

3.2.2 Cost of Fairness (CoF)

Next, we define CoF as the deviation from optimality due to fairness constraints given in Eq. (3.1). It is defined as the ratio of expected revenue generated by optimal and fair feature-based pricing.

Definition 3.2 (*Cost of Fairness* (CoF)). Cost of fairness for an FFP strategy **p** is defined as the ratio of overall expected revenue generated by $\hat{\mathbf{p}}$ to the expected revenue generated by **p**

$$CoF = \frac{\Pi(\widehat{\mathbf{p}})}{\Pi(\mathbf{p})}.$$
(3.2)

Notation	Description
$OPT_{FP}((\mathcal{V},\mathcal{X},\mathcal{F},\beta))$	Optimal Feature-Based Pricing problem
$OPT_{FFP}(\mathcal{V}, \mathcal{X}, d, \mathcal{F}, \beta, \alpha)$	Optimal Fair Feature-Based Pricing problem
$d_{ij} := d(x_i, x_j)$	A real-valued metric on the consumer feature space X
$D_i := \min_{i \neq j} d_{ij}$	Minimum distance of another consumer from consumer i in (X, d)
α	Fairness parameter
CoF	Cost of Fairness
<u>v</u>	Marginal cost of production
\overline{v}	Maximum valuation in continuous case
p*	Price vector for OPT _{FFP}
L_m	Linear approximation of concave revenue curve with m as parameter
$\Pi_m(\mathbf{L})$	Total expected revenue generated by OPT-LINP-FFP with pivot m
p	Price vector for OPT-LINP-FFP

Table 3.1: Notations Table 2

To clarify the above definitions, we analyze the CoF of the network-based pricing model given by [22].

The network-based pricing model accurately captures consumer valuations embedded in a social network. However, the network-based model strongly assumes that the consumers' valuations are identical outside the social network. In the real world, consumer valuations depend on many other personalized traits like age, gender, location, etc. Hence, next, we study fair pricing for general features.

Fairness in Network-Based Pricing

We now discuss a pricing strategy based on consumers' social network proposed by Huang *et al.* [22]. The model is summarised as follows: There are *n* consumers embedded in a directed social network with non-empty adjacency matrix $G \in \mathbb{R}^{n \times n}_+$ where g_{ij} is the influence of consumer *j* on consumer *i*. The seller sets a price vector $\mathbf{p} \in \mathbb{R}^n_+$ where p_i is the price for i^{th} consumer. It is assumed that consumer valuation varies only with their position in the network. Hence, the only consumer feature is their respective network externalities. The seller faces a marginal cost of production $c \ge 0$ for the product, and every consumer has a standalone utility $a > c \ge 0$. $\rho \in (0, 1)$ is a positive network externalities coefficient representing the strength of the network effect. The optimal price vector proposed by Huang *et al.* [22] is then given by

$$p^* = \left(\frac{a+c}{2}\right)\mathbf{1} + \left(\frac{a-c}{2}\right)\frac{\rho}{||G+G^T||}(G-G^T) K$$
(3.3)

where $K = \kappa \left(G + G^T, \frac{\rho}{||G + G^T||} \right)$ is the Bonacich Centrality.

Network-Based pricing stated above is α -fair with

$$\alpha = \left(\frac{a-c}{2}\right)\frac{\rho}{||G+G^T||}$$

Proof. The optimal price (Eq. (3.3)) for consumer *i* is given by

$$p_i = \left(\frac{a+c}{2}\right) + \left(\frac{a-c}{2}\right) \frac{\rho}{||G+G^T||} \kappa_i$$

where κ_i is the *i*th entry of the vector $(G - G^T) K$, the associated consumer feature in network-based pricing model. Consider two consumers *i*, *j* in the network model with features κ_i and κ_j respectively then difference in price $p_i - p_j$ is:

$$p_i - p_j = \left(\frac{a - c}{2}\right) \frac{\rho}{||G + G^T||} (\kappa_i - \kappa_j)$$

Taking the metric for consumers i, j to be the absolute difference between κ_i and κ_j then prices follow

$$|p_i - p_j| = \left(\frac{a - c}{2}\right) \frac{\rho}{||G + G^T||} |\kappa_i - \kappa_j|$$

From Definition 3.1, price vector p is α -fair with $\alpha = \left(\frac{a-c}{2}\right) \frac{\rho}{||G+G^T||}$

3.3 Optimal Fair Feature-Based Pricing (OPT_{FFP})

In this section, we introduce our feature-based pricing strategy FPP which is both α -fair and can price discriminate based on consumer features. Given the optimal FP, we want to ensure α -fairness in the pricing strategy. α -fairness is achieved by maximizing revenue while satisfying the fairness constraints. We will divide our discussions into two cases, viz. *discrete consumer valuations* and *continuous consumer valuations*.

Let (X, d) be a metric space. We model the Optimal fair feature-based pricing (FFP) problem as an optimization problem that maximizes $\Pi(\mathbf{p})$ with α -fairness constraints described in Eq.(3.1). We denote this problem as $OPT_{FFP}(\mathcal{V}, X, d, \mathcal{F}, \beta, \alpha)$ and the corresponding optimal price vector obtained using FFP strategy is denoted by \mathbf{p}^* . Let the feature-based market segments be given by $X = \{x_1, \ldots, x_k\}$, their valuations are drawn from a set \mathcal{V} . Then the OPT_{FFP} is given by the following optimization problem:

Definition 3.3 (*OPT*_{*FFP*} with discrete valuation). *OPT*_{*FFP*} for discrete valuations is given by:

$$\mathbf{p}^{\star} \in \underset{\mathbf{p} \in \mathcal{W}^{k}, \forall i \in [k]}{\operatorname{argmax}} \sum_{i=1}^{k} \beta_{i} \pi_{i}(p_{i})$$
(3.4a)

subject to,
$$|p_i - p_j| \le \alpha d(x_i, x_j), \forall i \ne j$$
 (3.4b)

 $p_i \ge 0, \forall i \in [k] \tag{3.4c}$

The distinction between discrete and continuous valuations is obtained by changing the valuations' support set \mathcal{V} . For discrete valuations, $\mathcal{V} := \{v_1, \ldots, v_r\}$; here, the valuations can take *r* discrete values, and the seller can pick one value from \mathcal{V} as a price offered to a consumer segment. Similarly, for continuous valuations, the support set is given by $\mathcal{V} := [\underline{v}, \overline{v}]$. In this case, \underline{v} signifies the marginal production cost faced by the seller, and \overline{v} is an upper limit on the price.

In the following two chapters, we will discuss discrete and continuous consumer valuations and provide optimal prices for each case.

Chapter 4

Discrete Valuations

In this chapter, we assume the consumer valuations are sampled from a discrete set, and the monopolist can select prices only from the valuations' support set, i.e., $v_i, p_i \in \mathcal{V}$. In Section 4.1.1, we derive OPT_{FP} then, we propose how to achieve α -fairness in Section 4.1.2 and finally, we analyze the CoF for discrete valuation setting in Section 4.1.3.

4.1 **Binary Discrete Valuations**

 OPT_{FFP} with discrete valuations is a combinatorial optimization problem; therefore, to begin our discussion, we consider the simplest setting as follows:

Let the consumer segments be given by $X = \{x_1, x_2\}$ and their valuations are drawn from a discrete set $\mathcal{V} = \{v_1, v_2\}$, we assume $v_1 < v_2$ without loss of generality. Let $\beta_1 = \beta$ and $\beta_2 = 1 - \beta$. Further, let the pmf for discrete valuations be defined as $f_i(v_1) = q_i$, $f_i(v_i) = 1 - q_i$, for $i \in \{1, 2\}$.

Then total expected revenue generated by **p** is given by:

$$\Pi(\mathbf{p}) = \beta \cdot p_1 \cdot [q_1 \mathbb{1}(v_1 \ge p_1) + (1 - q_1) \mathbb{1}(v_2 \ge p_1)] + (1 - \beta) \cdot p_2 \cdot [q_2 \mathbb{1}(v_1 \ge p_2) + (1 - q_2) \mathbb{1}(v_2 \ge p_2)]$$
(4.1)

4.1.1 Optimal Feature-based Pricing

As discussed earlier, $\Pi(\mathbf{p})$ can be maximized by solving for optimal $\pi_i(p_i)$ for each market segment independently if there are no fairness constraints. This problem is an integer program with the price for each consumer type being a discrete variable. The revenue generated depends on β_i and $f_i(\cdot)$ (β , q_1 , q_2 in the current simplest case). The optimal FP is then given as

For
$$i \in \{1, 2\}$$
: $\widehat{p}_i = \begin{cases} v_1 & \text{if } q_i \ge 1 - \frac{v_1}{v_2} \\ v_2 & \text{otherwise} \end{cases}$ (4.2)

Proof. For a market segment i, $\pi_i(v_1) = v_1$ and $\pi_i(v_2) = v_2(1 - q_i)$. So, $\widehat{p}_i = v_1$ if

$$\pi_i(v_1) \ge \pi_i(v_2) \implies v_1 \ge v_2(1-q_i) \implies q_i \ge 1 - \frac{v_1}{v_2}$$

otherwise, $\widehat{p}_i = v_2$.

Next, we analyze the fairness aspects of the above pricing strategy.

4.1.2 Optimal Fair Feature-based Pricing

Let (X, d) be a metric space. We model the Optimal fair feature-based pricing (OPT_{FFP}) problem as an integer program that maximizes $\Pi(\mathbf{p})$ with α -fairness constraints described in Eq.(3.1). We denote the corresponding optimal price vector as \mathbf{p}^* . First, we make an interesting and very useful claim for binary valuations.

Lemma 4.1. When $\mathcal{V} = \{v_1, v_2\}$, and if $\widehat{\mathbf{p}}$ is not α -fair, OPT $_{FFP}(\mathcal{V}, \mathcal{X}, d, \mathcal{F}, \beta, \alpha)$ reduces to $OPT_{FP}(\widetilde{\mathcal{V}}, \mathcal{X}, \mathcal{F}, \beta)$ where $\widetilde{\mathcal{V}}$ is either $\{v_1\}$, or $\{v_2\}$, or $\{v_1, v_1 + \alpha d_{12}\}$.

Proof. Let (p_1, p_2) be the tuple of offered prices. Note that if $v_2 - v_1 \le \alpha d_{12}$ or $\hat{p}_1 = \hat{p}_2$, then the optimal $\mathbf{p}^* = \hat{\mathbf{p}}$ with support $\{v_1, v_2\}$ and $\hat{\mathbf{p}}$ will be trivially fair. We consider a more interesting case when $v_2 - v_1 > \alpha d_{12}$ and $\hat{p}_1 \neq \hat{p}_2$. In this case, the only candidate support sets for optimal fair pricing strategy are $\{v_1\}, \{v_2\}, \{v_1, v_1 + \alpha d_{12}\}, \{v_2 - \alpha d_{12}, v_2\}$. The optimal FFP does not take values from the set $\{v_2 - \alpha d_{12}, v_2\}$ as the consumers with valuation v_1 would not make any purchase. Hence, the expected revenue with support $\{v_2 - \alpha d_{12}, v_2\}$ will be less than or equal to the expected revenue with support $\{v_2\}$.

We now relax the constraint of binary valuation and analyze the optimal fair pricing scheme for *n* valuations. The consumer segments are $X = \{x_1, x_2\}$ with $\beta_1 = \beta$ and $\beta_2 = 1 - \beta$, the valuations are drawn from the set $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$, and $f_1(v_i) = q_{i,1}$ and $f_2(v_i) = q_{i,2}$. This is a simple extension of the pricing problem, $OPT_{FP}(\mathcal{V}, X, \mathcal{F}, \beta)$ modeled as an integer program where the prices are drawn from the set \mathcal{V} . If $\hat{\mathbf{p}}$ is not α -fair then, the corresponding $OPT_{FFP}(\mathcal{V}, X, d, \mathcal{F}, \beta, \alpha)$ can be solved by reducing it to $OPT_{FP}(\tilde{\mathcal{V}}, X, \mathcal{F}, \beta)$ with $\tilde{\mathcal{V}}$ given by:

$$\widetilde{\mathcal{V}} = \begin{cases} \{v_i\}, v_i \in \mathcal{V} & \text{if } p_1^{\star} = p_2^{\star} \\ \{v_j, v_j + \alpha d_{12}, v_j - \alpha d_{12}\}, v_j \in \mathcal{V} & \text{if } p_1^{\star} \neq p_2^{\star} \end{cases}$$
(4.3)

Given the set $\widehat{\mathcal{V}}$, the pricing problem $OPT_{FP}(\widetilde{\mathcal{V}}, \mathcal{X}, \mathcal{F}, \beta)$ can be solved in constant time. It is easy to see that computing $\widehat{\mathcal{V}}$ takes $O(n^2)$ time for *n* valuations and 2 consumer types. Therefore, the fair pricing problem $OPT_{FFP}(\mathcal{V}, \mathcal{X}, d, \mathcal{F}, \beta, \alpha)$ can be solved in $O(n^2)$ time.

4.1.3 CoF Analysis

For n = 2, based on the values of q_1, q_2 we have the following cases:

1. $p_1^{\star} = p_2^{\star} = v_1$ 2. $p_1^{\star} = p_2^{\star} = v_2$ 3. $p_1^{\star} = v_1 + \alpha d_{12}, p_2^{\star} = v_1$ 4. $p_1^{\star} = v_1, p_2^{\star} = v_1 + \alpha d_{12}$

In cases 1 and 2, optimal fair pricing is equivalent to uniform pricing and therefore 'trivially' fair with CoF = 1, i.e., $\Pi(\widehat{\mathbf{p}}) = \Pi(\mathbf{p}^*)$. For case 3, $\Pi(\widehat{\mathbf{p}})$ and $\Pi(\mathbf{p}^*)$ are given as:

$$\Pi(\hat{\mathbf{p}}) = \beta(v_2)(1 - q_1) + (1 - \beta)v_1$$
$$\Pi(\mathbf{p}^{\star}) = \beta(v_1 + \alpha d_{12})(1 - q_1) + (1 - \beta)v_1$$

Then the cost of fairness for case 3 is given as:

$$CoF = \frac{\Pi(\widehat{\mathbf{p}})}{\Pi(\mathbf{p}^{\star})} = \frac{\beta(v_2)(1-q_1) + (1-\beta)v_1}{\beta(v_1 + \alpha d_{12})(1-q_1) + (1-\beta)v_1}$$

= $\frac{\beta(v_2 - v_1) + v_1 - \beta v_2 q_1}{\beta \alpha d_{12}(1-q_1) - \beta v_1 q_1 + v_1}$
= $\frac{\beta\left(1 - \frac{v_1}{v_2}\right) + \frac{v_1}{v_2} - \beta q_1}{\beta\left(\frac{\alpha d_{12}}{v_2}\right)(1-q_1) - \beta\left(\frac{v_1}{v_2}\right)q_1 + \frac{v_1}{v_2}}$ (4.4)

Replacing β with $(1 - \beta)$ and q_1 with q_2 in the above expression, we get a similar approximation of CoF for case 4.

4.2 CoF can be arbitrarily bad

Observation 4.1. Cost of fairness with discrete valuations can go arbitrarily bad.

Proof. From Eq. (4.4) when $\frac{v_1}{v_2} \rightarrow 0$, we have $\text{CoF} = \frac{v_2}{\alpha d_{12}}$. The CoF (in Case 3 and/or Case 4) is arbitrarily bad if $d_{12} > 0$ when there is a large difference between v_1 and v_2 . Note that $d_{12} = 0$ is uninteresting as the seller cannot distinguish between two segments.

Note that v_2 being arbitrarily large need not be a typical setting. Hence, we work with bounded support valuations in the backdrop of the above negative results.

Example 4.1 (Bike). Consider a seller selling motorbikes in a monopoly market. The minimum age to own a driving license and ride a motorbike is 18 years. Assume two consumers with features x_1, x_2 , and valuations v_1, v_2 respectively where x_i represents the age of the consumers (in years). Assume $x_1 = 17$, and $x_2 = 18$. Note, v_1 is negligible compared to v_2 , as the 18 years old consumer can legally ride a bike. Let p_1, p_2 be the respective feature-based prices offered to the consumers.

In a market described in Example 4.1, OPT_{FFP} prices will face a significant CoF due to a drastic shift in consumer valuations corresponding to a unit change in one of the consumer features.

The above example demonstrates the motivation for certain regularity assumptions on the consumer valuations to ensure minimal CoF for FFP strategies. Based on the economics literature, We will make some standard assumptions on the revenues $\pi_i(\cdot)$, viz., *concave revenues* and *common support* for valuations (Bergemann *et al.* [5]).

As argued in Section 3 of Dhangwatnotai *et al.* [13], valuation distributions satisfying the Monotone Hazard Rate (MHR) satisfy the abovementioned assumptions regarding revenue functions. It is also observed that the revenue functions are concave for another commonly analyzed family of distributions in literature called the regular distributions, in which the virtual valuation is non-decreasing (Section 4.3 of Bergemann *et al.* [5]). MHR is a common assumption in Econ-CS (Hartline & Roughgarden [20], Myerson [35]). Therefore, in the following section, we analyze the cost of fairness for such valuation distributions and the associated concave revenue functions.

Chapter 5

Continuous Valuations

This chapter will discuss the *OPT_{FFP}* strategy for continuous valuations under the standard concavity assumptions and provide an upper bound on CoF. We propose *LINP-FFP* – an FFP strategy that achieves the CoF bound. We also provide an algorithm *OPT-LINP-FFP* to obtain α -fair prices.

Assumptions: We make the following assumptions on the valuation functions and revenue curves (Bergemann *et al.* [5]):

- *Common Support:* For a consumer segment corresponding to feature vector x_i, the valuations are drawn from a continuous set D_F := [v, v̄]. The support is the same across all consumer segments x_i ∈ X.
- *Concave Revenues:* The expected revenue π_i(p) generated by the consumer segment x_i, is a concave function over prices p ∈ D_F.

The lower bound on the common support for valuations \underline{v} is the *marginal cost* defined as a minimum feasible valuation for which a seller is willing to sell the product. The marginal cost may include the cost of production, transportation, etc. On the other hand, the upper bound, \overline{v} is the maximum consumer valuation. Without loss of generality, we consider maximum consumer valuation greater than marginal cost; i.e., trade occurs.

In Section 3.3, we introduced the problem formulation for the Optimal FFP strategy. With reference to the problem formulation in Eq. (3.4), we will now explain a few theoretical contributions regarding OPT_{FFP} prices, bound the CoF, and identify a class of pricing strategies that satisfy the bound.

We begin with a tight upper bound on the CoF under conditions as mentioned above followed by a pricing scheme based on the available information about the revenue functions (Section 5.1), and finally, we present an algorithm that achieves the CoF bound in Section 5.2.

5.1 Linear methods to obtain Fair Feature-Based Pricing

The consumer feature space X, is equipped with a metric $d : X \times X \to \mathbb{R}_+$. Recall that we denote the distance between two feature vectors x_i, x_j as $d_{ij} := d(x_i, x_j)$. Let $D_i := \min_{j \neq i} d_{ij}$. We propose a class of α -fair pricing strategies with the following proposition.

Proposition 5.1. For a given $m \in [\underline{v}, \overline{v}]$, if the price function satisfies $|p_i - m| \leq \frac{\alpha}{2} D_i$ for all $i \in [K]$ then it satisfies α -fairness.

Proof. From triangle inequality, we have $|p_i - p_j| \le |p_i - m| + |p_j - m| \le \frac{\alpha}{2}D_i + \frac{\alpha}{2}D_j \le \alpha d_{ij}$. The last inequality results from the fact that $D_i = \min_{k \ne i} d_{ik} \le d_{ij}$ and $D_j = \min_{k \ne j} d_{ik} \le d_{ji} = d_{ij}$.

The intuition behind Proposition 5.1 explains, to ensure that the prices for different segments are not too different, it is enough to ensure that the pricing for each segment is not too distant from some common point m. The pricing for all the segments would hence be around this point and could be determined with respect to this point. We term this point as *pivot*.

The following subsection will discuss the pricing mechanism developed as one major contribution in this work.

5.1.1 Pivot-Based linear pricing(LINP-FFP)

We now present the our FFP model, an α -fair pricing strategy that is pivot-based and satisfies the condition in Proposition 5.1, with access to only \hat{p}_i for a given *m*.

$$p_{i} = \begin{cases} m + \alpha D_{i}/2 & \text{if } \widehat{p}_{i} - m \ge \alpha D_{i}/2 \\ m - \alpha D_{i}/2 & \text{if } m - \widehat{p}_{i} \ge \alpha D_{i}/2 \\ \widehat{p}_{i} & \text{otherwise} \end{cases}$$
(5.1)

We call this pricing scheme *Linearized Pivot-Based Fair Feature-Based Pricing (LINP-FFP)*. It is easy to see that the above pricing strategy is α -fair.

We now present the CoF bound for LINP-FFP.

5.1.2 CoF Bound

In this subsection, we will present the second major contribution of our work, a constant upper bound on the CoF under assumptions described in 5.

In the following Theorem 5.1, we propose the CoF upper bound for LINP-FFP. Later in Section 5.1.3, we formally argue that the bound is tight and also applies to OPT_{FFP} . To prove Theorem 5.1, we also present Lemma 5.1 along with its formal proof.

Theorem 5.1. The Cost of Fairness for optimal fair price discrimination with concave revenue functions satisfies

$$CoF \leq \frac{2}{1 + \min\left\{\alpha \frac{\min_i D_i}{\bar{v} - \underline{v}}, 1\right\}}$$

Proof. We prove that the above CoF is satisfied by LINP-FFP and hence the theorem. Let $m \in [\underline{v}, \overline{v}]$ be a pivot point (See Figure Fig. 5.1). Let

$$\gamma_{i} := \begin{cases} \frac{(m-\underline{v})+\alpha D_{i}/2}{\widehat{p}_{i}-\underline{v}} & \text{if } \widehat{p}_{i}-m \ge \alpha D_{i}/2\\ \frac{(\bar{v}-m)+\alpha D_{i}/2}{\bar{v}-\widehat{p}_{i}} & \text{if } m-\widehat{p}_{i} \ge \alpha D_{i}/2\\ 1 & \text{otherwise} \end{cases}$$
(5.2)

Let $\widehat{\pi}_i := \pi_i(\widehat{p}_i)$ be the expected revenue generated from the *i*th segment under the OPT_{FFP} prices $\widehat{\mathbf{p}}$.

We will present a linear approximation of revenue functions that satisfy the Assumptions 5. The following lemma presents an important relation between linear and optimal revenues.

Linear Approximation of revenue curves

Lemma 5.1. The pricing strategy given in Eq. (5.1) guarantees at-least γ_i fraction of optimal revenue from segment *i*, *i.e.*, $\pi_i \geq \gamma_i \widehat{\pi}_i$.

Proof. A lower bound to the concave revenue functions $\pi_i(\cdot)$ for any segment *i* is the piecewise linear approximation L_i , given by (see Figure Fig. 5.1):

$$L_{i}(p) = \begin{cases} \frac{\widehat{\pi}_{i}}{\widehat{p}_{i} - \underline{v}}(p - \underline{v}), & p \leq \widehat{p}_{i} \\ \frac{-\widehat{\pi}_{i}}{\overline{v} - \widehat{p}_{i}}(p - \overline{v}), & p > \widehat{p}_{i} \end{cases}$$
(5.3)

So, for each consumer segment *i* we have,

$$L_i(p) \le \pi_i(p), \ \forall p \in [\underline{v}, \overline{v}]$$



Figure 5.1: Concave revenue function $\pi_i(\cdot)$ and its linear approximation $L_i(\cdot)$ (arrows show equations for $L_i(\cdot)$). Figure represents the case $\hat{p}_i - m \ge \alpha D_i/2$ for which LINP-FFP assigns $p_i = m + \alpha D_i/2$. The case $m - \hat{p}_i \ge \alpha D_i/2$ is similar.

Expected revenues generated per consumer in segment *i* by pricing rule in Eq. (5.1) for $\hat{p}_i - m \ge \alpha D_i/2$, $m - \hat{p}_i \ge \alpha D_i/2$, and remaining cases are given below in the respective order

$$\pi_i(p_i) \ge L_i(p_i) = \frac{\widehat{\pi}_i}{\widehat{p}_i - \underline{v}} (m + \alpha D_i/2 - \underline{v}) = \widehat{\pi}_i \gamma_i$$

$$\pi_i(p_i) \ge L_i(p_i) = \frac{-\widehat{\pi}_i}{\overline{v} - \widehat{p}_i} (m - \alpha D_i/2 - \overline{v}) = \widehat{\pi}_i \gamma_i$$

$$\pi_i(p_i) = L_i(\widehat{p}_i) = \widehat{\pi}_i$$

This proves the lemma.

Let π_i^* denote the expected revenue generated from the *i*th segment by **p**^{*}. So, CoF for optimal FPP is given by:

$$\operatorname{CoF} = \frac{\sum_{i \in [K]} \beta_i \widehat{\pi}_i}{\sum_{i \in [K]} \beta_i \pi_i^{\star}} \leq \frac{\sum_{i \in [K]} \beta_i \widehat{\pi}_i}{\sum_{i \in [K]} \beta_i \pi_i}$$
(Optimality of π_i^{\star})
$$\leq \frac{\sum_{i \in [K]} \beta_i \widehat{\pi}_i}{\sum_{i \in [K]} \beta_i \gamma_i \widehat{\pi}_i}$$
(Lemma 5.1)

In order to prove the said CoF bound, it suffices to show that there exists an *m* (and hence a corresponding pricing strategy using Eq. (5.1)) for which the said CoF bound is satisfied. It can be seen that for $m = (\underline{v} + \overline{v})/2$, and replacing denominators in Eq. (5.2) by $\overline{v} - \underline{v}$, we have that

$$\operatorname{CoF} \leq \frac{\sum_{i \in [K]} \beta_i \widehat{\pi}_i}{\sum_{i \in [K]} \beta_i \widehat{\pi}_i \left(\frac{1}{2} + \min\{\frac{\alpha D_i}{2(\bar{v} - \underline{v})}, 1\}\right)} \\ \leq \frac{\sum_{i \in [K]} \beta_i \widehat{\pi}_i}{\left(\sum_{i \in [K]} \beta_i \widehat{\pi}_i\right) \left(\frac{1}{2} + \min\{\frac{\alpha \min_j D_j}{2(\bar{v} - \underline{v})}, 1\}\right)} \\ = \frac{2}{1 + \min\left\{\alpha \frac{\min_j D_j}{\bar{v} - \underline{v}}, 1\right\}}$$

-	
_	

It is worth noting here that the cost of fairness does not depend on the number of segments and the population distribution among these segments. So, if the segments are well separated in terms of the distance between features of consumers across segments, the number of segments and the distribution of consumer population in these segments do not affect revenue guarantee. Also, if the admissible prices are supported over a large interval, the fairness guarantee becomes weaker. This insight discourages pricing schemes with wildly varying prices across segments. Finally, if $\alpha = 0$, i.e., without any fairness constraints, we recover the bound of 2 proved in Bergemann *et al.* [5].

We emphasize that the bound is strictly less than 2 because, under fairness constraints, $\alpha \neq 0$ and typically, the consumer types are well separated in the feature space according to the metric *d* else, the consumer types are indistinguishable for the seller hence, $d_{ij} \neq 0$ for all $i, j \in [K]$.

5.1.3 Tightness of CoF bound:

We claim that the CoF bound presented above is tight. In the following example, equality holds and proves the tightness of the bound.

Example 5.1 (Tightness of the CoF bound). Consider K = 2 where $\beta_1 = \beta_2 = \frac{1}{2}$. Consider \mathcal{F}_i be such that $\pi_i(\cdot) = L_i(\cdot)$ with $\hat{p}_1 = \underline{v} + \epsilon$, $\hat{p}_2 = \overline{v} - \epsilon$, where $\epsilon \to 0$, and $\hat{\pi}_1 = \hat{\pi}_2$. It can be seen that if α is such that $\alpha d_{12} < \overline{v} - \underline{v}$, any FP satisfying $p_2 - p_1 = \alpha d_{12}$ and $p_1, p_2 \in [\hat{p}_1, \hat{p}_2]$ is an optimal FFP (fair FP), and the corresponding $CoF = \frac{2}{1 + \frac{\alpha d_{12}}{\overline{v} - \underline{v}}}$. If $\alpha d_{12} \ge \overline{v} - \underline{v}$, the optimal FP is α -fair and so, CoF = 1. Hence, for this example, $CoF = \frac{2}{1 + \frac{\alpha d_{12}}{\overline{v} - \underline{v}}}$. This shows the tightness of the CoF bound derived in Theorem 5.1.

We now present an algorithm, OPT-LINP-FFP, to find the optimal pivot m^* in the above LINP-FFP strategy when only \hat{p} and $\hat{\pi}_i$ s are known.

5.2 **Proposed Algorithm**

As LINP-FFP satisfies α -fairness (Proposition 5.1), and also achieves CoF bounds in Theorem 5.1, we look for a pricing strategy optimal within the class of LINP-FFP. It reduces to finding an *optimal pivot* that maximizes revenue. This section proposes a *binary-search-based algorithm* for the same. For pricing **p**, the expected revenue generated per consumer is given by $\Pi(\mathbf{p}) = \sum_{i=1}^{K} \beta_i \pi_i(p_i)$. Let $\tau_i := \frac{\alpha}{2} D_i$. Observe from Lemma 5.1 that $\Pi(\mathbf{p})$ is lower bounded as:

$$\Pi(\mathbf{p}) \ge \Pi_m(\mathbf{L}) = \sum_{i=1}^K \beta_i \gamma_i \widehat{\pi}_i = \sum_{i:|\widehat{p}_i - m| < \tau_i} \beta_i \widehat{\pi}_i + \sum_{i:\widehat{p}_i - m \ge \tau_i} \beta_i \widehat{\pi}_i \frac{m + \tau_i - \nu}{\widehat{p}_i - \nu} + \sum_{i:m - \widehat{p}_i \ge \tau_i} \beta_i \widehat{\pi}_i \frac{\overline{\nu} - m + \tau_i}{\overline{\nu} - \widehat{p}_i}$$
(Eq. (5.6))

5.2.1 Determining Optimal Pivot m

As we can see, the revenue generated by $L_{IN}P$ -FFP is lower bounded by a piecewise linear function in m. With the aim of achieving a better lower bound, we now address the problem of determining an optimal pivot.

$$m^{\star} \in \operatorname*{argmax}_{m \in [\nu, \bar{\nu}]} \Pi_m(\mathbf{L}) \tag{5.7}$$

where $\Pi_m(\mathbf{L})$ is the total expected revenue generated by LINP-FFP at pivot *m*.

5.2.2 Pricing Algorithm

In what follows, we call the candidate points *m* for the optimal pivot, i.e., for maximizing $\Pi_m(\mathbf{L})$, as *critical points*. We denote the set of these critical points as *C*.

Lemma 5.2. $\Pi_m(\mathbf{L})$ as a function of *m* is concave and piecewise linear with the set of critical points $C = \left(\{\widehat{p}_i - \frac{\alpha}{2}D_i, \widehat{p}_i + \frac{\alpha}{2}D_i\}_{i \in [K]} \cap [\underline{v}, \overline{v}]\right) \cup \{\underline{v}, \overline{v}\}.$

Proof. It is easy to see that for a segment *i*, γ_i as a function of *m* is continuous and piecewise linear with breakpoints (i.e., points at which piecewise linear function changes slope): $\hat{p}_i - \frac{\alpha}{2}D_i$ and $\hat{p}_i + \frac{\alpha}{2}D_i$ provided

they are in the range $[\underline{v}, \overline{v}]$. The set of breakpoints is hence $\{\widehat{p}_i - \frac{\alpha}{2}D_i, \widehat{p}_i + \frac{\alpha}{2}D_i\} \cap [\underline{v}, \overline{v}]$. Also, the slope monotonically decreases at the breakpoints, i.e., γ_i is a concave function of m.

From Eq. (5.6), we can see that $\Pi_m(\mathbf{L})$ is a weighted sum over all segments, of γ_i 's with constant weights $\beta_i \hat{\pi}_i$. So, $\Pi_m(\mathbf{L})$ as a function of *m* is concave and piecewise linear with breakpoints belonging to the following set: $\{\hat{p}_i - \frac{\alpha}{2}D_i, \hat{p}_i + \frac{\alpha}{2}D_i\}_{i \in [K]} \cap [\underline{v}, \overline{v}]$. Hence, a point *m* that maximizes $\Pi_m(\mathbf{L})$ belongs to either the aforementioned set of breakpoints, or the set of its boundary points $\{\underline{v}, \overline{v}\}$. Thus, the set of critical points $C = (\{\widehat{p}_i - \frac{\alpha}{2}D_i, \widehat{p}_i + \frac{\alpha}{2}D_i\}_{i \in [K]} \cap [\underline{v}, \overline{v}]) \cup \{\underline{v}, \overline{v}\}$.

Three main steps of OPT-LINP-FFP

- 1. Computing the set of all critical points C.
- 2. Binary search for optimal pivot m^* in the set of all critical points C.
- 3. After obtaining the optimal pivot point m^* , compute the prices for all consumer segments using Eq. (5.1).

Our algorithm OPT-LINP-FFP (Optimal Linearized Pivot-based Fair Feature-based Pricing) which determines an optimal pivot m^* and provides an α -fair pricing strategy ($\tilde{\mathbf{p}}$) is presented in Algorithm 1. This is formalised in the following Theorem 5.2

Theorem 5.2. The OPT-LinP-FFP algorithm (a) returns optimal pivot point m^* and runs in $O(K \log(K))$ time, and (b) achieves the CoF bound given in Theorem 5.1.

Proof. (a) The first module is the creation and sorting of the set of critical points *C*, which takes $O(K \log(K))$ time. Owing to Lemma 5.2, we can find an optimal pivot m^* using binary search over *C*. Here, the number of critical points are at most 2K + 2, i.e., $|C| \le 2K + 2$. So, in the second module that finds an optimal pivot, the binary search in the outer (*while*) loop runs for $O(\log(|C|))$ iterations, and the inner (*for*) loops run for O(K) iterations overall. Thus, the running time of the second module is $O(K \log(K))$. The third module that computes pricing for the different segments runs in O(K) time. So, the total running time of Algorithm 1 is $O(K \log(K))$.

(b) From Theorem 5.1, for $m = (\underline{v} + \overline{v})/2$, the CoF bound holds. Also, $\Pi_{m^*}(\mathbf{L}) \ge \Pi_m(\mathbf{L})$ for all $m \neq m^*$. We have:

$$\operatorname{CoF} = \frac{\Pi(\widehat{\mathbf{p}})}{\Pi(\widetilde{\mathbf{p}})} \le \frac{\Pi(\widehat{\mathbf{p}})}{\Pi_m \star(\mathbf{L})} \le \frac{\Pi(\widehat{\mathbf{p}})}{\Pi_m(\mathbf{L})}$$

Algorithm 1: OPT-LINP-FFP

Input: α , $\{(\widehat{p}_i, \widehat{\pi}_i, \beta_i, D_i)\}_{i=1}^K$

Output: $m^*, \tilde{\mathbf{p}}$

/*Creating and sorting the set of critical points

;

 $C \leftarrow \{\underline{v}, \overline{v}\};$ for $i \in [K]$ do

 $\tau_i \leftarrow \frac{\alpha}{2} D_i$;

if
$$\widehat{p}_i - \tau_i > \underline{v}$$
 then
 $\[C \leftarrow C \cup \{\widehat{p}_i - \tau_i\} ;\]$
if $\widehat{p}_i + \tau_i < \overline{v}$ then

$$C \leftarrow C \cup \{\widehat{p}_i + \tau_i\}$$

sort(C);

/*Binary search for optimal pivot

$$\ell \leftarrow 0, r \leftarrow |\mathcal{C}| - 1;$$

while $\ell \leq r$ do

 $z \leftarrow \lfloor \frac{\ell + r}{2} \rfloor //*$

C[z] is the current pivot /*Computing the expression in Eq 5.6 at current and adjacent

*/

*/

*/

critical points

$$\Pi_{\mathcal{C}[z-1]} \leftarrow 0, \Pi_{\mathcal{C}[z]} \leftarrow 0, \Pi_{\mathcal{C}[z+1]} \leftarrow 0;$$

for
$$y \leftarrow \{z - 1, z, z + 1\}$$
 do

for
$$i \leftarrow 1$$
 to K do

if
$$\widehat{p}_i \ge C[y] + \tau_i$$
 then
 $\downarrow \gamma_i \leftarrow \frac{C[y] - \underline{v} + \tau_i}{\widehat{p}_i - \underline{v}}$

else if
$$\widehat{p}_i \leq C[y] - \tau_i$$
 then

$$\sum_{i \neq j} \gamma_i \leftarrow \frac{\overline{v} - C[y] + \tau_i}{\overline{v} - \widehat{p}_i}$$
else

if $\Pi_{C[z-1]} \leq \Pi_{C[z]} \leq \Pi_{C[z+1]}$ then $\lfloor \ell \leftarrow z+1;$ else if $\Pi_{C[z-1]} \geq \Pi_{C[z]} \geq \Pi_{C[z+1]}$ then $\lfloor r \leftarrow z-1$

else

 $m^{\star} \leftarrow C[z];$ break;

 Algorithm 1: OPT-LINP-FFP (continued...)

 /*Pricing for the different segments
 */

 for $i \in [K]$ do
 */

 if $\hat{p}_i \ge m^* + \tau_i$ then
 $\Box \quad \tilde{p}_i \leftarrow m^* + \tau_i$

 else if $\hat{p}_i \le m^* - \tau_i$ then
 $\Box \quad \tilde{p}_i \leftarrow m^* - \tau_i$

 else
 $\Box \quad \tilde{p}_i \leftarrow \hat{p}_i$

Experiments on Synthetic Data: We validate our theoretical claims on synthetically generated consumer valuations. The data is generated by approximating valuation distributions $f_k(\cdot)$ as triangular functions over a chosen common support by generating random peaks for each consumer segment k. We find the revenue peaks $\hat{\pi}_k$ and the corresponding \hat{p}_k values for OPT-LINP-FFP. The consumer features are *m*-dimensional random vectors where each entry is in the range [0, 1] and the distance metric used is Euclidean 2-norm. We assume that $\beta_k = 1/n$ for all consumer types k where n is the number of consumer types. On simulated data, OPT-LINP-FFP achieves CoF= 1.0806 (worst case CoF = 1.1834, average case CoF = 1.0806) with coefficient of variation = 0.027 for 500 iterations.

With this, we end our discussions on our contributions to the literature on fairness in pricing. In the next chapter, we conclude our work by summarizing the contents of this thesis. We also suggest possible directions with immense potential for theoretical and applied contributions.

Chapter 6

Conclusions and Future Work

6.1 Conclusion

This thesis work explored the importance of markets and their emergence through time. We saw that markets are a system of profitable exchange of goods and services which eventually attains an equilibrium governed by the forces of *supply* and *demand*. In Chapter 2, we provided the necessary background on the general equilibrium theory of markets and described two fundamental market models, *Fisher Market*, and *Arrow-Debreu Market*. Further, we saw that available information on consumer preferences plays a crucial role in market dynamics. Incomplete information and imbalanced market share lead to *competition* among the sellers, which tends to benefit the consumers and social welfare. We presented the theoretical framework for understanding pricing strategies in monopoly markets and introduced important fairness criteria in the pricing literature. Finally, we briefly reviewed papers that compared the performance of uniform pricing versus personalized pricing under specific assumptions on valuations distribution. Further, we saw that available information on consumer preferences plays a crucial role in market dynamics.

Next, we direct the focus of our discussion toward *monopoly markets* and different pricing strategies in such markets to maximize revenue returns. Through various examples, we observe that such pricing practices have raised ethical and fairness concerns among consumers and regulatory bodies. We first presented a brief literature review on fairness in pricing strategies to address these concerns.

In Chapter 3, we formally introduced our fairness criteria, named α -fairness, and a performance metric called CoF, which compares the revenue of potentially unfair and fair prices. Then, we formulated the fair pricing problem, OPT_{FFP}, that maximizes revenue while satisfying α -fairness.

In Chapter 4, we analyzed the solutions to OPT_{FFP} for discrete prices and consumer valuations. We showed that it was necessary to assume standard restrictions on consumer valuations. Otherwise, the CoF could turn out to be arbitrarily bad. Whereas, Chapter 5, solved OPT_{FFP} for continuous consumer valuations. Under the standard concavity assumptions on revenue functions, we analyzed CoF and provided a constant

upper bound of 2. Next, we identified a class of prices, named LINP-FFP, that achieved this bound and provided an algorithm to solve for LINP-FFP.

6.2 Future Work

We must empirically verify our framework as we present significant theoretical results regarding individual fairness in monopoly markets. Towards this, we need to collect relevant real-world data and perform empirical analysis to draw valuable inferences. The following research direction could be to analyze our framework for markets with more sellers, for instance, in a duopoly or oligopoly market setting. Furthermore, we assume the market segmentation is based on consumer features, but a monopolist can choose to segment a market differently to obtain optimal revenues under α -fairness constraints.

Bibliography

- Movie Ticket Offers. https://cdn.grabon.in/gograbon/images/web-images/uploads/ 1618571888087/movie-ticket-offers.jpg.
- [2] Aguirre, Iñaki, Cowan, Simon, & Vickers, John. 2010. Monopoly Price Discrimination and Demand Curvature. *The American Economic Review*, **100**(4), 1601–1615.
- [3] Alan, Kaplan. 2020. Hospital Price Discrimination Is Deepening Racial Health Inequity. https://catalyst.nejm.org/doi/full/10.1056/CAT.20.0593.
- [4] Bergemann, Dirk, Brooks, Benjamin, & Morris, Stephen. 2015. The Limits of Price Discrimination. *American Economic Review*, **105**(3), 921–57.
- [5] Bergemann, Dirk, Castro, Francisco, & Weintraub, Gabriel. 2020. Third-degree price discrimination versus uniform pricing. *Games and Economic Behavior*, 131, 275–291.
- [6] Bikhchandani, Sushil, & Mamer, John W. 1997. Competitive Equilibrium in an Exchange Economy with Indivisibilities. *Journal of Economic Theory*, **74**(2), 385–413.
- [7] Brainard, William C., & Scarf, Herbert E. 2005. How to Compute Equilibrium Prices in 1891. *The American Journal of Economics and Sociology*, 64(1), 57–83.
- [8] Cassady, Ralph. 1946. Techniques and Purposes of Price Discrimination. *Journal of Marketing (pre-1986)*, 11(000002), 135.
- [9] Chang, Yuqian, Winston, Clifford, & Yan, Jia. 2022. Does Uber Benefit Travelers by Price Discrimination? *The Journal of Law and Economics*, 65(1), 1–28.
- [10] Cohen, Maxime C., Elmachtoub, Adam N., & Lei, Xiao. 2022. Price Discrimination with Fairness Constraints. *Management Science*, 68(12), 8536–8552.
- [11] Cowan, Simon. 2016. Welfare-increasing third-degree price discrimination. The RAND Journal of Economics, 47(2), 326–340.

- [12] Cummings, Rachel, Devanur, Nikhil R., Huang, Zhiyi, & Wang, Xiangning. 2020. Algorithmic Price Discrimination. *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2432–2451.
- [13] Dhangwatnotai, Peerapong, Roughgarden, Tim, & Yan, Qiqi. 2015. Revenue maximization with a single sample. *Games and Economic Behavior*, **91**, 318–333.
- [14] Dholakia, Utpal. 2015. Everyone Hates Uber's Surge Pricing Here's How to Fix It. https: //hbr.org/2015/12/everyone-hates-ubers-surge-pricing-heres-how-to-fix-it.
- [15] Dwork, Cynthia, Hardt, Moritz, Pitassi, Toniann, Reingold, Omer, & Zemel, Richard. 2012. Fairness through Awareness. Page 214–226 of: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. ITCS '12. New York, NY, USA: Association for Computing Machinery.
- [16] Elmachtoub, Adam N., Gupta, Vishal, & Hamilton, Michael L. 2021. The Value of Personalized Pricing. *Manage. Sci.*, 67(10), 6055–6070.
- [17] EU, Communications Authority. 2020. Pricing and payments. https://europa.eu/youreurope/ citizens/consumers/shopping/pricing-payments/index_en.htm.
- [18] Fujishige, Satoru, & Yang, Zaifu. 2000. Equilibrium in an exchange economy with multiple indivisible commodities and money. *Journal of Mathematical Economics*, **33**(3), 353–365.
- [19] Gerlick, Joshua A., & Liozu, Stephan M. 2020. Ethical and Legal Considerations of Artificial Intelligence and Algorithmic Decision-making in Personalized Pricing. *Journal of Revenue and Pricing Management*, 19(2).
- [20] Hartline, Jason D, & Roughgarden, Tim. 2009. Simple versus optimal mechanisms. *Pages 225–234* of: Proceedings of the 10th ACM conference on Electronic commerce.
- [21] Hinz, Oliver, Hann, II-Horn, & Spann, Martin. 2011. Price Discrimination in E-Commerce? An Examination of Dynamic Pricing in Name-Your-Own Price Markets. *MIS Quarterly*, 35(1).
- [22] Huang, Jiali, Mani, Ankur, & Wang, Zizhuo. 2019. The Value of Price Discrimination in Large Random Networks. In: Proceedings of the 2019 ACM Conference on Economics and Computation. EC '19. New York, NY, USA: Association for Computing Machinery.
- [23] Kallus, Nathan, & Zhou, Angela. 2021. Fairness, Welfare, and Equity in Personalized Pricing. Page 296–314 of: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT '21. New York, NY, USA: Association for Computing Machinery.

- [24] Kamei, Precious. 2019. India: Street Markets For The Mad Shopper In You. https://www.outlookindia.com/outlooktraveller/explore/story/69635/ street-markets-in-india-you-need-to-explore.
- [25] Klein, Stefan, & Loebbecke, Claudia. 2000. The transformation of pricing models on the web: examples from the airline industry. *Pages 19–21 of: 13th International Bled Electronic Commerce Conference*.
- [26] Levy, Karen, & Barocas, Solon. 2017. Designing Against Discrimination in Online Markets. *Berkeley Technology Law Journal*, **32**(3).
- [27] Mahadawi, Arwa. 2018. Is your friend getting a cheaper Uber fare than you are? https://www.theguardian.com/commentisfree/2018/apr/13/uber-lyft-prices-personalized-data.
- [28] Malueg, David A., & Snyder, Christopher M. 2006. Bounding the relative profitability of price discrimination. *International Journal of Industrial Organization*, 24(5), 995–1011.
- [29] Martin, Adam. 2020. Movies, Markets, and Economic Growth. Library of Economics and Liberty.
- [30] Mattioli, Dana. 2012. On Orbitz, Mac Users Steered to Pricier Hotels. https://www.wsj.com/ articles/SB10001424052702304458604577488822667325882.
- [31] Maxfield, Robert R. 1997. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, **27**(1), 23–51.
- [32] Michel, Stefan. 2016. Is personalized pricing fair? https://www.imd.org/research-knowledge/ articles/is-personalized-pricing-fair/.
- [33] Moriarty, Jeffrey. 2021. Why online personalized pricing is unfair. *Ethics and Information Technology*, 23(3), 495–503.
- [34] MyEconomics. 2020. Theory of Economic Behavior. https://www.myeconomics.info/ plus-two-economics-chapter-2-note-in-english-price-elasticity-of-demand.html.
- [35] Myerson, Roger. 1981. Optimal Auction Design. *Mathematics of Operations Research*, 6(1), 58–73.
- [36] Namin, Aidin, Gauri, Dinesh K., & Kwortnik, Robert J. 2020. Improving revenue performance with third-degree price discrimination in the cruise industry. *International Journal of Hospitality Management*, 89.
- [37] Nisan, Noam, Roughgarden, Tim, Tardos, Eva, & Vazirani, Vijay V. 2007. Algorithmic Game Theory. New York, NY, USA: Cambridge University Press.
- [38] Olena, Ostapenko. Online purchase or booking of tickets vector image. https://www.vectorstock. com/royalty-free-vector/online-purchase-or-booking-of-tickets-vector-14847023.
- [39] Pandey, Akshat, & Caliskan, Aylin. 2021. Disparate Impact of Artificial Intelligence Bias in Ridehailing Economy's Price Discrimination Algorithms. New York, NY, USA: Association for Computing Machinery. Page 822–833.
- [40] Pigou, Arthur. 1920. The Economics of Welfare. Macmillan.
- [41] Richards, Timothy J., Liaukonyte, Jura, & Streletskaya, Nadia A. 2016. Personalized Pricing and Price Fairness. *International Journal of Industrial Organization*, 44.
- [42] Schmalensee, Richard. 1981. Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination. *American Economic Review*, **71**(1), 242–47.
- [43] Shantanu, Das, Swapnil, Dhamal, Ganesh, Ghalme, Shweta, Jain, & Sujit, Gujar. 2022. Individual Fairness in Feature-Based Pricing for Monopoly Markets. In: The 38th Conference on Uncertainty in Artificial Intelligence.
- [44] Varian, Hal R. 1985. Price discrimination and social welfare. *The American Economic Review*, **75**(4), 870–875.
- [45] Yacoby, Rakefet. 2022. Online Marketplaces: The Best eCommerce Guide on How to Sell Your Product On The Web in 2023. https://www.mayple.com/blog/online-marketplaces.