

ON DYNAMIC PRICING AND ASSORTMENT PERSONALIZATION IN STRATEGIC SETTINGS

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by

Jiayang Gao

December 2017

© 2017 Jiayang Gao
ALL RIGHTS RESERVED

ON DYNAMIC PRICING AND ASSORTMENT PERSONALIZATION IN
STRATEGIC SETTINGS

Jiayang Gao, Ph.D.

Cornell University 2017

This thesis studies firm's operational decisions in strategic settings. In particular, it focuses on a firm's dynamic pricing or assortment optimization problems, in the presence of competitors or customers that respond to the firm's actions strategically. We study three problems in this area with different features. We optimize over the firm's strategies, and characterize the outcome of the game between the firm and its strategic counter-parties.

BIOGRAPHICAL SKETCH

Jiayang was born in Hangzhou, one of the most beautiful cities in east China, which many consider as the beginning of the Silk Road. She attended Zhejiang University, and transferred to New York University a year later and earned her bachelor's degree there with double major in Mathematics and Economics, and a minor in Computer Science. Jiayang was initially interested in pure math subjects like number theory and abstract algebra, but she started to realize the importance of real world application during her undergraduate research experience. She then joined the Operations Research department in Cornell. Upon graduating, she will join Laurion Capital Management in New York City, as a quantitative researcher. In her spare time, Jiayang enjoys eating good food, playing board games and traveling.

To my parents, Feifei Lu and Jianzhong Gao.

ACKNOWLEDGEMENTS

First, I would like to thank my advisors Huseyin Topaloglu and Krishnamurthy Iyer for being helpful, patient, supportive. During the past four years, I received many insightful suggestions from them, which made me a better researcher and writer.

I thank Yudong Chen, Bruce Turnbull, Shane Henderson, Robert Bland and again my advisor Krishnamurthy Iyer, who I worked with as a teaching assistant. I enjoy my time working with you and the students. I thank Yudong Chen in particular for being very supportive during my job hunting season.

I thank Wei Zhu, Xiaodong Yao, Matt Maxwell, Joseph Warfel and Kumar Abhishek for being great manager or mentor during my internships.

I thank Renxiong Wang, Yixuan Zhao, Bangrui Chen, Jie Xu, Wei Chen, Chaoxu Tong and Eric Guo for their help during my job hunting.

I thank Eric Ni, Shanshan Zhang, Yi Shen, Xiaoting Zhao and Wei Chen for giving me the warmest welcome when I came for campus visit 5 years ago. I am here because of you, and I do not regret it.

I thank Nanjing Jian, Yixuan Zhao and Jiayi Guo for making Rhodes 288 a home. I thank Wenxiong Zhang, Jinting Yang, Yuqi Lu and their families for bringing HFLS to Ithaca. I thank Yuan Chen for visiting Ithaca 7 times during my time here. I thank Audrey Shi for attending my graduation ceremony. You are like family to me.

I thank Miss Red for her loyalty and company.

I thank my best friend Eric Ni for sharing and contributing to my happiest years in Ithaca. Last but not least, I thank my parents for their love and understanding, and for being supportive to all decisions I made, however unconventional they may be.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
List of Figures	viii
1 Introduction	1
1.1 When Fixed Price Meets Priority Auctions: Competing Firms with Different Pricing and Service Rules	3
1.2 Price Competition under Linear Demand and Finite Inventories: Con- traction and Approximate Equilibria	6
1.3 Personalized Assortment Recommendation after Purchase	7
2 When Fixed Price Meets Priority Auctions: Competing Firms with Different Pricing and Service Rules	11
2.1 Introduction	11
2.1.1 Literature review	15
2.2 Model	19
2.3 Equilibrium structure of the customers' game	24
2.3.1 Structure of customers' equilibrium bidding function	24
2.3.2 Structure of the customers' equilibrium service decision	27
2.4 Existence and uniqueness of equilibrium in customers' game	33
2.4.1 System with mandatory service requirement	36
2.4.2 System with optional service requirement	40
2.4.3 System with positive reserve price	42
2.5 Competition between firms	44
2.5.1 High arrival rate and high capacity regime	46
2.5.2 Equilibria among firms in the limiting system	48
2.5.3 Price of stability of the limiting system	50
2.6 Conclusion	51
3 Price Competition under Linear Demand and Finite Inventories: Contraction and Approximate Equilibria	55
3.1 Equilibrium without Recourse	59
3.2 Equilibrium with Recourse	65
3.3 An Approximate Equilibrium	69
3.4 Future Research	73

4	Personalized Assortment Recommendation after Purchase	74
4.1	Introduction	74
4.1.1	Literature Review	78
4.2	Model	81
4.3	LP formulation for the offline problem	83
4.4	The balancing algorithm	86
4.4.1	The algorithm	86
4.4.2	The balancing step	88
4.4.3	Algorithm complexity	94
4.5	Feasibility and performance analysis	95
4.5.1	Feasibility of Algorithms	95
4.5.2	Performance guarantee	99
4.6	Asymptotic Optimality	101
4.7	The resolving algorithm	103
4.8	Numerical Analysis	108
4.8.1	Algorithm performance under simulation	108
4.8.2	An upper bound on the performance guarantee	113
4.9	Discussion and Conclusion	114
5	Conclusion	118
5.1	Strategic and forward looking customers	118
5.2	The role of information	119
A	Appendix for Chapter 2	121
A.1	Extensions to general distribution F	121
A.2	Proofs of the results in Section 2.3	121
A.3	Proofs of auxiliary results in Section 2.4	126
A.4	Proofs of Theorem 2 and Theorem 3	132
A.5	Proof of Theorem 4	142
A.6	Proofs in Section 2.5	147
A.7	Extension to non-preemptive queues for the bid-based firm	153
B	Appendix for Chapter 3	158
B.1	Omitted Results	158
B.2	Omitted Proofs	162
C	Appendix for Chapter 4	173
	Bibliography	179

LIST OF FIGURES

2.1	Expected payment in equilibrium of the customers' game	32
2.2	Expected waiting time in equilibrium of the customers' game	32
2.3	Total revenue per arrival in Nash equilibrium and under collusion, when $V = 5$	52
2.4	Ratio between the per arrival total revenues, when $V = 5$	52
2.5	Total revenue per arrival in Nash equilibrium and under collusion, when $V = 1.5$	53
2.6	Ratio between the per arrival total revenues, when $V = 1.5$	53
3.1	Best responses when equilibrium does not exist.	68
3.2	Revenue of the first firm as a function of its first period price.	68
3.3	Best responses when there are multiple equilibria.	69

CHAPTER 1

INTRODUCTION

Consider a revenue maximizing firm selling multiple products over finite time horizon. At each point in time, there are two important questions that it needs to consider. Which products should be made available to customers now? At what price? These two questions are fundamental components of the firm's operational decision. They are known as the dynamic assortment planning problem, and the dynamic pricing problem, respectively.

These questions are not trivial even under the simplest setting: a monopolist facing myopic customers. The decision at any point in time affects future revenue via remaining inventory. Thus, the problem is usually modeled using dynamic programming with time periods as decision epochs, and remaining inventory levels as states. The difficulty then arise due to potentially large state space and action space in the dynamic program, as well as the uncertainty in demand. To be more precise, the size of state space is exponential in the number of time periods, and the size of the action space in dynamic assortment problems is exponential in the number of products. This makes the dynamic program numerically intractable when the number of periods or the number of products is large. Approximation schemes have been developed under these settings. One widely used approach is to approximate the dynamic program using linear programming. On the other hand, since demand is uncertain, the firm has to act in response to the realized demand at any point in time, without knowing what future demand will be. Thus this problem is *online*, as

opposed to offline problems in which all future demand are deterministic and the firm's decision can be made at the beginning of the selling horizon. Under the same setting, online problems are more difficult than offline problems, because the firm in an online setting has less information about the future when making decisions. One resolution to this challenge is to consider the deterministic (offline) approximation of this online problem, and develop heuristics for the online problem based on the optimal solution to the deterministic problem.

In real world, firms often have to consider dynamic pricing and assortment planning problems under more complicated settings. For example, there often exists a strategic counter-party who alters its behavior based on the firm's action, and therefore affects the firm's payoff. The strategic counter-party can either be a competitor who changes the price and availability of its products based on the firm's decision, or a smart customer who calculates the expected payoff under each purchasing option and chooses the best option for herself. Understanding the behavior and decision processes of these strategic counter-parties is important for the firm. Therefore in this thesis, we consider dynamic pricing and assortment planning problems in strategic settings.

Apart from the difficulties discussed above, the strategic component adds another layer of complexity to the problem. When customers are strategic, and each customer's payoff depends on other customers' actions, each customer will have to account for other customers' actions. The firm has to analyze the outcome of the customers' interactions in order to evaluate its revenue under any policy, before it can optimize over all policies. When the firm's competitors are strategic in a dynamic problem, it

needs to analyze the competition among firms in future periods, in order to evaluate its future payoff and thus balance the trade-off between realizing more sales for the current period and saving more inventory for the future.

Due to the many difficulties discussed above, in each chapter of this thesis, we focus on modeling and resolving a subset of these challenges. In Chapter 2, we consider a service system with two competing firms with strategic customers, assuming a customer's payoff depends not only on her own action, but on other customers' actions as well. In Chapter 3, we consider a multi-period price competition problem among multiple firms with limited inventories of substitutable products. In Chapter 4, we consider a firm's dynamic assortment optimization problem with limited inventories of substitutable products, when it has the opportunity to recommend personalized assortments to customers after they make a purchase. In the remaining sections of this chapter, we briefly introduce each of these three problems, and summarize our key contributions.

1.1 When Fixed Price Meets Priority Auctions: Competing Firms with Different Pricing and Service Rules

In Chapter 2, we consider a service system with two competing firms that offer service via different pricing and service rules, the fixed-price firm and the bid-based firm. In the fixed-price firm, service is provided at a fixed rate, and in the bid-based firm,

customers submit a bid no less than a reserve price posted by the firm, and pay their bids. In the fixed-price firm, the expected waiting time is homogeneous across all customers, increasing in the arrival rate to this firm. In the bid-based firm, the expected waiting time for a customer is increasing in the arrival rate of customers with higher bids. Customers have different unit time waiting costs, and minimize the sum of the payment and the expected total waiting cost. Upon arrival, they strategically choose a firm to receive service from, and also choose a bid if they decide to receive service from the fixed-price firm. We assume that the customer cannot observe the real time congestion in the firms, or other customers' bids in the bid-based firm. Thus customers' decision are only based on steady state expectations.

This model has application in real world, when competing firms offer service via different pricing structures or service qualities. For example, in the cloud computing industry, Amazon EC2 offers "spot instances" in which the customers enter an auction to determine priority of service and payment. On the other hand, Microsoft Azure charges a fixed price for per hourly use of its service. Similarly in the local transportation industry, taxi service is usually offered at a fixed price for a specific original-destination pair, while Didi Dache, the largest car hailing company in China, gives customers the option to offer a customized amount of tip when they request for a car.

In order to analyze the competition between the firms, we have to first understand customers' decision. Since in both firms, the expected waiting time of a customer depends on other customers' choices, an equilibrium analysis is needed to analyze

the customers' strategies. We prove the existence and uniqueness of symmetric equilibrium for the game among customers, and show that the equilibrium has a simple threshold structure. In particular, customers with high and low waiting costs choose the bid-based firm, while customers with intermediate waiting costs choose the fixed-price firm. The bid-based firm serves two functions in equilibrium: customers with high waiting costs incur less waiting time than that in the fixed-price firm, by making a payment more than the fixed price, while customers with low waiting costs can reduce their payment by waiting for a longer period of time.

Using the characterization of the unique symmetric equilibrium for the customers' decision problem, we proceed to analyze the pricing game between the two firms. The bid-based firm chooses its reserve price, and the fixed-price firm chooses its fixed price, each maximizing its own revenue in the subsequent customers' equilibrium. We prove that there exists a mixed Nash equilibrium for the pricing game between firms. Moreover, in a limiting regime where the customer arrival rate and service capacities of the two firms increase proportionally, we characterize a set of approximate equilibria, under which the total expected revenue of the two competing firms is comparable to that under collusion.

1.2 Price Competition under Linear Demand and Finite Inventories: Contraction and Approximate Equilibria

In Chapter 3, we consider a multi-period price competition among multiple firms. Each firm sells a perishable good with limited inventory. The demand that a firm sees at any period is assumed to be deterministic and linear in prices of all firms in that period. In particular, demand at a firm is decreasing in its own price, while increasing in all other firms' prices. Each firm chooses its price at each period of the selling horizon, in order to maximize its own revenue.

This model is applicable to real world settings where multiple firms sell substitutable and perishable products with limited inventory. For example, in the airline industry, multiple airlines may operate between the same origin and destination, with similar departure time. Seats on these flights are highly substitutable, thus airlines set price competitively. Also, many electronic products have similar features, thus firms taking the prices of their competitors into consideration when setting their prices.

We study two equilibrium concepts for this problem. In the equilibrium without recourse, each firm chooses a price trajectory for the whole selling horizon before the first period and commits to it, assuming that all other firms do the same. In the equilibrium with recourse, in each period, each firm observes the remaining inventory levels of all firms, and then chooses a price for that period, assuming that all other

firms do the same. Although the demand of each firm is a deterministic function of the prices, so that there is no uncertainty in the responses of the firms, equilibria with and without recourse can still be very different.

We show that the equilibrium without recourse uniquely exists, by proving that the best response of each firm to the price trajectories of all other firms is a contraction mapping. As a result, equilibrium without recourse can be computed by best response iterations.

For the game with recourse, we give examples under which equilibrium does not exist, or is not unique. We thus study approximate equilibria for the game with recourse. We focus on a low influence regime, in which the price charged by a firm affects its demand much more than it affects the demand of any other firm. We show that the equilibrium without recourse can be used to construct an approximate equilibrium with recourse, which has the same price trajectory as the equilibrium without recourse.

1.3 Personalized Assortment Recommendation after Purchase

In Chapter 4, we consider a firm's multi-period assortment optimization problem with limited inventories of substitutable products, assuming that it has the opportunity to offer a personalized assortment to each customer after she makes a purchase. Each

period has two stages. A customer arrives to stage 1 and is shown an assortment, from which she can choose to make a purchase, or to leave without a purchase. If she makes a purchase in stage 1, she then enters stage 2 in which she is offered another assortment for her to consider. Customer behavior follows general choice models. The firm maximizes its expected revenue, by choosing an assortment to offer at each stage of each time period, based on remaining time and inventory. Moreover, the assortment offered in stage 2 is personalized based on the customer's purchase in stage 1.

This model is applicable for many web-based sellers, who have the opportunity to make personalized recommendations to customers after they make purchase. For example, after a customer books a hotel via Priceline.com, she can find rental car deals recommended to her in the confirmation email that she receives. Similarly, Amazon.com also make recommendations in the confirmation email that it sends to customers after they make a purchase. In both cases above, recommendations are made based on the item purchased by the customer.

Although we do not explicitly model the customer's strategic behavior in this chapter, the choice probabilities can be seen as an aggregate outcome of customers' strategies. We can also extend this model to allow more complicated strategic behavior. A possible extension is to consider forward looking customers, who may act suboptimally in the first stage, in order to "disguise" their types and get a better recommendation in the second stage.

The problem studied in this chapter is an online problem, and we first study an

offline version of this problem, which can be considered as a fluid approximation of the original online problem. The offline problem can be modeled using a linear program, using the probability of offering each assortment at each stage as a decision variable. This linear program is exponential in the number of products, but it can be solved efficiently using column generation.

Using the solutions to the offline problem, we propose a balancing algorithm for the original online problem. The algorithm solves the offline LP once at the beginning of the selling horizon, and updates the solutions at each time period so that they remain feasible for the offline LP with the current time and inventory information. In each time period, the algorithm uses a policy suggested by these feasible solutions.

We prove that the balancing algorithm attains an expected revenue of at least $1/3$ of the optimal offline revenue, which is an upper bound on the optimal online revenue. We also give an example, under which the optimal online revenue is only 47% of the optimal offline revenue. This suggests that the best constant performance guarantee that any algorithm can achieve is no more than 47%, as long as we do the analysis by comparison with the offline LP. In a limiting regime, in which the number of time periods and inventory levels increase proportionally and both tend to infinity, we prove that the revenue under the balancing algorithm converges in probability to the optimal offline revenue.

We propose a resolving algorithm as a natural extension of the balancing algorithm. It follows the balancing algorithm, but periodically resolves the offline LP to adjust for stochastic realization of sales. The resolving algorithm also has a performance

guarantee of $1/3$ of the optimal offline revenue, and it is also asymptotically optimal. Numerical experiments suggest that resolving the LP can significantly improve the algorithm performance when inventory is scarce, even resolving is performed only once in the middle of the selling horizon.

CHAPTER 2

WHEN FIXED PRICE MEETS PRIORITY AUCTIONS: COMPETING FIRMS WITH DIFFERENT PRICING AND SERVICE RULES

2.1 Introduction

Competing firms often offer their products and services through various modes differing in their pricing structure and service quality. One main motivation behind this service differentiation among competing firms is to target heterogeneous customers, differing in their preferences over the quality and urgency of service, their tolerance for uncertainty, etc. For example, in the cloud computing service Amazon EC2, the customers can choose to obtain service by bidding for a computing resource in a quasi-auction market (“spot instances”, see [7]), whereas in the competing service Microsoft Azure, customers pay for a fixed-price per hourly use [49]. Similarly, in the local transportation industry, a regular taxi service company offers a fixed price for a specific original-destination pair, while Didi Dache, a major competitor of the taxi companies in China, gives customers the option to offer a customized amount of tip to the driver when they request for a car, and the drivers are more likely to accept a request as the amount of tip increases. (For more information about Didi Dache and tipping, see [14] and [57].) In this case, the Didi customers are essentially bidding in a priority queue.

With an assortment of options to obtain services, customers have to make strategic trade-offs among the cost, quality and priority of service. This choice is further complicated by the fact that, owing to the presence of resource constraints on the firms' part, each customer's choice among the different firms influences and is influenced by how other customers make the same decision. This suggests that the characteristics of the set of customers availing service from a particular firm is determined endogenously, and an equilibrium analysis is needed to understand the customers' decision. Further, each firm has to model the customers' equilibrium response in order to evaluate and optimize the design of their service in order to target customers with a specific set of characteristics.

To understand the equilibrium behavior of customers and the resulting competition among firms offering different pricing and service rules, in this chapter, we consider a setting consisting of customers with heterogeneous delay-sensitivity, and two competing service providers who offer service via two different pricing and service rules. In the fixed-price firm, the service is offered at a fixed price, and the service quality (waiting time) is the same for everyone who chooses to obtain service from this firm. On the other hand, in the bid-based firm, the service is provided according to a first-price priority auction, where each customer on arrival submits a bid equal to the price they will pay for service, and the expected waiting time is decreasing in the bids. Each customer chooses the firm (and a bid, if she chooses the bid-based firm) that maximizes her total expected utility, given by the difference between the utility of receiving service, and her expected cost. This cost is comprised of the customer's payment for service and the delay cost she incurs by waiting until service completion.

A customer's strategy in this context consists of her choice of the firm to obtain service from (or to balk without obtaining service), and her bid in the bid-based firm if she chooses to obtain service there. We analyze a symmetric equilibrium of the preceding system, where all customers adopt the same strategy, and where each customer is making a best response to others' strategies.

For given fixed price in the fixed-price firm, we show by explicit construction that there exists a *unique* symmetric equilibrium. More importantly, we characterize the structure of the equilibrium strategy, and show that it has a multi-threshold structure: the bid-based firm is visited by customers with very high or very low delay-sensitivity, whereas the fixed-price firm is visited by customers with moderate delay-sensitivity. This result can be extended to the case where the bid-based firm charges a reserve price.

This structural characterization of the symmetric equilibrium of the customers' game provides the main insight of this chapter, namely that the bid-based firm simultaneously enables two sets of customers, with markedly different characteristics, to make optimal trade-offs between the delay costs and the cost of obtaining service. On one hand, for those customers with very low unit waiting cost, it acts as a means to reduce their total cost by submitting low bids and waiting longer for service. On the other hand, for customers with very high unit waiting costs, it acts as a venue to demand high priority (and lower waiting time) for service. This insight coincides with what we see in the China's local transportation service industry, in which Didi Dache has a reserve price that is much lower than the price charged by the taxi company

(usually about 30% lower, see [66]), but some delay-sensitive customers choose to add a considerable amount of tip to avoid waiting (see [64]).

Finally, using the uniqueness of the customers' equilibrium strategy, we study the price competition game between the two firms, where the fixed-price firm sets the fixed-price and the bid-based firm sets the reserve price. We show existence of a mixed Nash equilibrium. Moreover, by analyzing the game in a limiting regime where the arrival rate and the firms' service capacity increase proportionally, we show the total expected revenue in equilibrium achieved by the two firms under competition is close to the revenue obtained by firms under collusion.

We summarize our main contributions as follows:

(1) *Characterization of structure of customers' equilibrium strategy:* We show that, in any symmetric equilibrium, the customers' strategy has a simple multi-threshold structure: among all customers who choose to obtain service from the system, the customers with relatively high per unit time waiting cost and those with relatively low per unit time waiting costs choose to obtain service from the bid-based firm. On the other hand, the customers who obtain service from the fixed-price firm have relatively moderate per unit time waiting cost. Consequently, in this market, the fixed-price firm only attracts customers with moderate waiting costs.

(2) *Existence and uniqueness of customers' equilibrium:* Using the structure of the equilibrium strategy, we show the existence and uniqueness of a symmetric equilibrium of the customers' game. Our proof proceeds by obtaining necessary and sufficient

conditions on a threshold strategy to constitute a symmetric equilibrium. These conditions are obtained by imposing the continuity of the expected waiting times and total cost of the customers with unit waiting costs at the thresholds. Finally, we show the existence (and uniqueness) by explicitly constructing a solution satisfying these conditions.

(3) *Competition between firms:* Finally, we investigate how competing firms set their prices in order to maximize their revenue in equilibrium. Using the uniqueness of the customers' equilibrium strategy, we study the resulting price competition between the two firms, where the fixed-price firm sets the fixed-price and the bid-based firm sets the reserve price, and show existence of a mixed Nash equilibrium. Moreover, we analyze this competition under a limiting regime where the arrival rate and the firms' service capacity increase proportionally. For the resulting limiting game, we show that the total expected revenue in equilibrium achieved by the two firms under competition is close to the revenue obtained by firms under collusion.

2.1.1 Literature review

There are multiple strands of literature from game theory and auctions, revenue management, and queueing theory that are related to our work.

Our work builds on existing work on revenue maximizing pricing policies of a monopolist, in the presence of strategic customers. [47] consider a queueing system with a single server and finite number of customer classes each with a different

per unit time waiting cost and service time distribution. The authors devise an incentive compatible pricing policy, such that the social welfare is maximized when each customer class endogenously choose their arrival rate and priority. Similar models have been studied in the context of a revenue maximizing service provider [2, 70, 28].

Although we do not specify a queueing model in this chapter, our results is applicable to many commonly used queueing models (e.g., $G/M/k$ priority queues), since the expected waiting time function under those models satisfy our assumptions. Thus, the analysis of customers' behavior in the bid-based firm is closely related to priority queues. One of the earliest work in priority queues is by [36]. He considers a model of an $M/M/1$ queue where service is provided in the decreasing order of the customers' bids, and where the customers are non-strategic with an exogenously specified bid distribution. In this setting, he obtains expressions for the expected waiting time as a function of the bid in both preemptive and non-preemptive settings. [8, 44, 29, 32, 35] build on this work by considering strategic customers that determine their priorities endogenously through their payments. Our analysis of the bid-based queue further makes use of results from auction theory [53, 38]. More broadly, our work contributes to the literature at the intersection of game theory and queueing theory. See [33] for a comprehensive survey of various models of queueing systems with strategic customers and servers.

There are a number of papers that study a market with firms offering a multitude of price-quality combinations (offered by either a monopolist or competing firms) for

customers to choose from. [52] consider a monopolist choosing a price-quality schedule to maximize its revenue. [4] consider revenue maximizing price/lead-time menu for a single server system where customers' value for service completion is a monotone continuous function of their per unit time waiting cost. [55] study a similar model where the per unit time waiting costs are a linear or sublinear function of the value of service, and they consider the asymptotic regime with large arrival rates and service capacity. See also [67] and [45] for related models studying pricing and scheduling policy in the asymptotic large system regime. [34] consider a competition on price and quality in the cloud computing market. In these papers, the customers' strategies in equilibrium have similar structure as that in our model, such that customers who prefer higher quality pay higher price, which also leads to threshold structure of customers' strategy when the price-quality menu is discrete. However, our analysis of the customers' strategy is fundamentally different and more complex. This is mainly because for a fixed price-quality menu, a customer's payoff is a function of her own action only, and independent of other customers' actions. Thus there is no game among customers, and the customers' strategy can be solved by an optimization problem. However, in our model, the quality of service, as measured by the expected waiting time, is a function of a customer's own strategy as well as those of others. Consequently, the customers' decision is the outcome of a game among customers.

There are some papers which consider the game among customers in a market with multiple pricing and service policies [17, 12, 69, 1]. [3] analyze a queueing system under different pricing and service policies, where customers are strategic and have a value for service that depends multiplicatively on their delay cost. [1] model the

Amazon EC2 cloud computing service as a hybrid system where customers can choose to enter a bid-based priority queue, or to obtain service from a fixed-price queue with infinite capacity. The authors show that any equilibrium has a single threshold structure, with all customers below the threshold entering the bid-based priority queue, and the rest the fixed-price queue. [12] prove the existence of similar single threshold strategy equilibrium in the context of a seller adopting a fixed-price or an auction mechanism to sell their products. In contrast, in our model, the fixed-price firm has finite capacity, implying that customers obtaining service from the fixed-price firm may have to wait for service. This waiting for service in the fixed-price firm induces customers with high unit waiting cost to choose the bid-based firm in equilibrium, resulting in an equilibrium strategy with a multi-threshold structure.

In our model, the customers with intermediate types (defined by their unit time waiting costs) choose the fixed-price firm, while the customers on the two extremes choose the bid-based firm. Similar intermediate-versus-extremes equilibrium structure was recognized in several other papers under different circumstances. [71] consider trading position on a FIFO queue with an intermediary at a fee, when customers have different unit time waiting costs. They show that in equilibrium, customers with intermediate waiting costs do not participate in the trade and remain in their FIFO position, while customers with lower or higher waiting costs trade their priorities. [4] study the design of a price/lead-time menu in order to maximize revenue, when customers differ in patience levels. They conclude that pricing out the customers with intermediate patience levels while serving the most patient and impatient ones may increase revenue.

The rest of this chapter is organized as follows. In Section 2.2, we describe our model of a market with a fixed-price firm and a bid-based firm. In Section 2.3, we study the customers' game, and characterize the structure of the strategy in any symmetric equilibrium. Using this structure, in Section 2.4 we obtain necessary and sufficient conditions for a strategy to constitute an equilibrium, and prove our main result, namely the existence and uniqueness of a symmetric equilibrium of the customers' game. Finally, in Section 2.5, we study the firms' game, and prove the existence of mixed Nash equilibrium. We also study the price of stability of the firms' game in a limiting regime, and characterize the conditions under which it is small.

2.2 Model

Consider a setting with two competing firms, a fixed-price firm and a bid-based firm, offering service to a Poisson stream of customers with rate $\lambda > 0$. The fixed-price firm charges a fixed price $P > 0$ to offer a service where all customers incur uniform expected waiting times until service completion. This uniform expected waiting time increases with the arrival rate of customers. We assume that the fixed-price firm has a finite service capacity of n , such that if the arrival rate of customers to the firm is greater than or equal to n , then the firm can no longer ensure finite expected waiting times. For example, the fixed-price firm may use an $G/G/n$ FIFO queue to provide service to its customers.

On the other hand, the bid-based firm provides service via a first-price auction

[38] with reserve price r . More precisely, customers arriving to the bid-based firm submit a bid no less than r on arrival that denotes the price they are willing to pay for service. Customers are then served in the decreasing order of their bids (with ties broken uniformly at random), and are charged their bid upon service completion. We assume that the bid-based firm has a service capacity of k , implying that if the total arrival rate of customers to the firm is greater than or equal to k , then the firm cannot ensure finite expected waiting time to all of its customers. For most of our analysis, we assume that a customer of the bid-based firm experiences an expected waiting time that depends only on the arrival rates of customers with higher bids. For example, the bid-based firm may provide service to its customers using a $G/M/k$ preemptive priority queue. (Note that for a non-preemptive priority queue, a customer's expected waiting time also depends on the arrival rate of customers with lower bids. Our analysis in Section 2.3 and Section 2.4 can be reproduced for non-preemptive queues; see Appendix A.7 for details.)

Each arriving customer is characterized by three features: their service requirement, their value for service completion and their cost for waiting until service completion. We address each one separately below.

We assume that each customer's service requirement is exponentially (and independently) distributed with mean 1 (equal to the service rate $\mu = 1$). The homogeneity of service requirement is a restrictive assumption, and a more general model will allow for non-homogeneous distribution for the service requirement. On the other hand, our model can easily accommodate more general distributions for service requirement,

and the independence assumption can be expected to hold in many service systems where the customer base is fairly large and diverse.

Next, each customer obtains a value $V > 0$ upon service completion. As mentioned earlier, we assume that this value is uniform across customers. Although this assumption may seem restrictive, one justification arises out of the interpretation of the value V as the opportunity cost faced by the customer. For example, one may consider V to be fixed price charged by a competitor to the two firms who offers service with negligible waiting times. In such settings, if each customer has access to this competitor, then our assumption of uniform value of service completion holds. However, if not all customers can access this competitor, then a more appropriate model would require different values of service completion for different (classes of) customers.

Finally, we assume that the customers incur a heterogeneous cost for waiting until service completion. More precisely, each customer incurs a disutility proportional to the total time she spends in the system until service completion, and we refer to the proportionality constant as the customer's unit waiting cost. We assume that each customer's unit waiting cost c is drawn independently and identically from a continuously differentiable bounded distribution F . For ease of notation, we assume that this distribution F is the uniform distribution on $[0, 1]$ in this chapter. All of our analytical results extend directly to the general case, as we discuss in Appendix A.1.

We assume the arrival rate λ , the price P in the fixed-price firm, the reserve price r in the bid-based firm, the distribution of the service requirement, the value

of service completion V , the service capacities k and n , and the distribution F of customers' unit waiting cost are common knowledge among the customers and the two service providers.

On arrival to the system, each customer decides based on her unit waiting cost, whether to obtain service and if so, from which firm. If she decides to obtain service from the bid-based firm, then she further chooses a bid to submit. A customer choosing not to obtain service leaves the system never to return, and obtains zero utility. A customer with unit waiting cost c waiting for a time W until service completion, and making a payment m obtains a utility equal to $V - c \cdot W - m$. (We refer to the quantity $c \cdot W + m$ as the total cost incurred by the customer.) We assume that the customers are strategic and seek to maximize their total expected utility, where the expectation is with respect to the steady state distribution of the system. Implicitly, this steady state expectation entails assuming that the customers cannot observe the state of the system, such as the queue lengths in each firm or the existing bids in the bid-based firm, before making their decision. This assumption is valid in many settings, especially when the queue is not physical, e.g., in call centers, online service industry, etc.

Consequently, we represent a customer's strategy by a pair of functions $x(\cdot)$ and $b(\cdot)$ of her unit waiting cost c , where $x(c) \in \{\text{LEAVE, FIX, BID}\}$ denotes her decision about whether to obtain service and if so, from which firm, and $b(c) \geq 0$ denotes her bid upon joining the bid-based firm. We refer to the function $x(\cdot)$ as the customer's service decision and the function $b(\cdot)$ as her bid function. For the sake of completeness,

we define $b(c) = P$ if $x(c) = \text{FIX}$ and $b(c) = 0$ if $x(c) = \text{LEAVE}$.

We focus on the symmetric setting, where all customers follow the same strategy (x, b) . In this scenario, we let $w_F(x, b)$ denote the expected waiting time in the fixed-price firm in steady state. For a customer with unit waiting cost c , the expected total cost on receiving service from the fixed-price firm is then given by $cw_F(x, b) + P$. Similarly, we let $w_B(b'|x, b)$ denote the expected waiting time in steady state for a customer joining the bid-based firm and making a bid b' . The total expected cost for such a customer with unit waiting cost c is then given by $cw_B(b'|x, b) + b'$.

We say a strategy (x, b) forms a symmetric equilibrium if, assuming all other customers act according to the strategy (x, b) , each customer's expected utility is maximized by following the same strategy (x, b) . Formally, we require that (x, b) satisfy the following conditions:

$$x(c) = \begin{cases} \text{LEAVE,} & \text{only if } V \leq \min\{\min_{b'}\{cw_B(b'|x, b) + b'\}, cw_F(x, b) + P\}; \\ \text{FIX,} & \text{only if } cw_F(x, b) + P \leq \min\{V, \min_{b'}\{cw_B(b'|x, b) + b'\}\}; \\ \text{BID,} & \text{only if } \min_{b'}\{cw_B(b'|x, b) + b'\} \leq \min\{cw_F(x, b) + P, V\}, \end{cases}$$

and

$$b(c) \in \arg \min_{b'} \{cw_B(b'|x, b) + b'\}, \quad \text{if } x(c) = \text{BID}.$$

Here, we break ties arbitrarily. The first condition specifies that the customer will choose to obtain the service only if the total expected cost is less than or equal to the value of service completion. In this case, the customer will choose the fixed-price firm if the total expected cost therein is no more than that in the bid-based firm under the best possible bid. Otherwise, the customer will choose the bid-based firm. The

second condition requires that upon choosing to obtain service from the bid-based firm, the customer will enter a bid that minimizes her total expected cost.

2.3 Equilibrium structure of the customers' game

In this section, we characterize the structure of the equilibrium strategy of the customers' game. This structure is used later in our proof of the existence of customers' equilibrium. For the ease of notation, we assume in Sections 2.3 and 2.4 that $r = 0$. We will extend all results to the case where $r > 0$ in Subsection 2.4.3.

2.3.1 Structure of customers' equilibrium bidding function

We begin our analysis of the symmetric equilibrium by focusing on the equilibrium bidding function. We show that in a symmetric equilibrium (x, b) , the bidding function $b(\cdot)$ is completely specified once the service decision $x(\cdot)$ is known. Towards this, for $c \in [0, 1]$, define $B(c|(x, b))$ as the proportion of customers in the bid-based firm with unit waiting cost below c : $B(c|(x, b)) \triangleq \int_0^c \mathbf{I}\{x(u) = \text{BID}\} du$. We have the following lemma showing the monotonicity of the expected waiting time, payment, and the total cost in any symmetric equilibrium. The proof relies on the fact that a customer does not gain by unilaterally deviating from the equilibrium strategy. We provide the details in Appendix A.2.

Lemma 1. *In any symmetric equilibrium (x, b) of the customers' game, for all*

customers that choose to obtain service, the expected waiting time is non-increasing, the expected payment is non-decreasing, and the total expected cost is strictly increasing in the unit waiting cost. Moreover, the bidding function is strictly increasing at all unit waiting cost c where $B(c|(x, b))$ is strictly increasing.

The implication of the lemma is as follows. In a symmetric equilibrium, the bidding function $b(c)$ is strictly increasing whenever $B(c|(x, b))$, the proportion of customers in the bid-based firm with unit waiting cost lower than c , is strictly increasing. And since the service in the bid-based firm is in the decreasing order of bids, this implies that in equilibrium, the customers in the bid-based firm are served in decreasing order of the unit waiting cost. Thus, the expected waiting time of a customer in the bid-based firm is solely a function of the service decision $x(\cdot)$ and their unit waiting cost c . We use $w(c|x)$ to denote the expected waiting time of a customer with unit waiting cost c when everyone uses the service decision function $x(\cdot)$. Using this result, we obtain the following characterization of the expected waiting time function and the bidding function in terms of the service decision in equilibrium.

Lemma 2. *In any symmetric equilibrium (x, b) , the expected waiting time function $w(\cdot|x)$ and the bidding function $b(\cdot)$ are completely determined by the customers' service decision $x(\cdot)$. In particular, the bidding function satisfies*

$$b(c) = \int_0^c w(t|x)dt - cw(c|x). \quad (2.1)$$

for all c such that $x(c) \neq \text{LEAVE}$.

Proof sketch. From Lemma 1, we obtain that in equilibrium, a customer's expected waiting time in the bid-based firm depends only on the proportion of customers in

the bid-based firm with lower unit waiting costs, which in turn depends only on the service decision $x(\cdot)$ and the customer's unit waiting cost. Since the expected waiting time in the fixed-price firm depends only on the proportion of customers in the fixed-price firm, this holds true also for a customer in the fixed-price firm. Together, this implies that, in a symmetric equilibrium (x, b) , the expected waiting time of a customer is given by a function $w(c|x)$ only of their unit waiting cost c and the service decision $x(\cdot)$.

To obtain (2.1), note that in equilibrium, for any customer with unit waiting cost c , the marginal decrease in the expected waiting cost resulting from a marginal increase in the bid, must equal the marginal increase in the resulting payment. Assuming differentiability, this implies $b'(c) = -cw'(t|x)$, which on integrating yields (2.1). See Appendix A.2 for details. \square

As a consequence of Lemma 2, in order to find a symmetric equilibrium, it suffices to focus only on the customers' service decision $x(\cdot)$, and use (2.1) to obtain the bidding function $b(\cdot)$. Moreover, for any given service decision $x(\cdot)$, if customers bid according to the bidding function given by (2.1) in the bid-based firm, we observe that the total expected cost of a customer is given by

$$TC(c|x) \triangleq cw(c|x) + b(c) = \int_0^c w(t|x)dt, \quad (2.2)$$

for all $x(c) \neq \text{LEAVE}$.

2.3.2 Structure of the customers' equilibrium service decision

Having characterized the bidding function in a symmetric equilibrium of the customers' game, we now focus on the service decision $x(\cdot)$ in a symmetric equilibrium. Before we proceed, we note that starting from an equilibrium, if we alter the actions of a measure zero set of customers from their current action to a different best response action, the resulting service decision continues to be a symmetric equilibrium (without a specific tie-breaking rule). Thus, in order to avoid unnecessary technicalities, in the rest of the chapter, we focus only on those symmetric equilibria where each action is either employed by a set of customers of positive measure, or never adopted by any customers. (In particular, we ignore equilibria where one of the firms services a set of customers with measure zero.) Then, the following theorem, the main result of this section, states that the service decision $x(\cdot)$ in any symmetric equilibrium has a simple multi-threshold structure.

Theorem 1. *Let $V > 0$ and $P > 0$. In any symmetric equilibrium (x, b) , the service decision function $x(\cdot)$ has multi-threshold structure. Specifically, there exists thresholds $0 < c_1 \leq c_2 \leq c_\ell \leq 1$, such that*

$$x(c) = \begin{cases} \text{LEAVE,} & \text{if } c \in (c_\ell, 1]; \\ \text{FIX,} & \text{if } c \in (c_1, c_2); \\ \text{BID,} & \text{if } c \in [0, c_1] \cup [c_2, c_\ell]. \end{cases}$$

For a symmetric equilibrium with no customer obtaining service from the fixed-price

firm, the thresholds satisfy $c_1 = c_2 = c_\ell \in (0, 1]$. For a symmetric equilibrium with some customers obtaining service from the fixed-price firm, the thresholds satisfy $0 < c_1 < c_2 < c_\ell \leq 1$.

Proof. Fix a symmetric equilibrium (x, b) . Consider customer with unit waiting cost c , who obtains service in equilibrium. The total expected cost of such a customer in equilibrium is no more than V , i.e., $cw(c|x) + b(c) \leq V$. Thus, the total expected cost of any customer with unit waiting cost $c' < c$ using action $(x(c), b(c))$ is less than V . Consequently, the expected total cost of such a customer using action $(x(c'), b(c'))$ is also less than V , and such a customer also chooses to obtain service. Thus, there exists a threshold $c_\ell \in [0, 1]$ such that all customers with unit waiting cost below c_ℓ choose to obtain service (i.e., $x(c) \in \{\text{FIX}, \text{BID}\}$), and the rest choose to leave the system without obtaining service (i.e., $x(c) = \text{LEAVE}$). Observe that for $V > 0$, we have $c_\ell > 0$, for if no other customers obtain service from the system, then it is optimal for small enough $\epsilon > 0$, a customer with unit waiting cost ϵ strictly prefers to obtain service from the bid-based firm by making a small enough bid. (Note that it may be the case that $c_\ell = 1$, implying that no customers choose to leave the system without obtaining service.)

Now, we consider how customers with unit waiting cost below c_ℓ choose between the fixed-price firm and the bid-based firm. Let C_F denote the set of unit waiting costs for which the customer's equilibrium choice is to obtain service from the fixed-price firm. We begin by showing that the set C_F is convex.

Let $c, \hat{c} \in C_F$ with $\hat{c} < c$. Since $x(c) = x(\hat{c}) = \text{FIX}$, we obtain in equilibrium,

$$cw_F(x) + P \leq \min_{b'}\{cw_B(b'|x) + b'\}, \quad cw_F(x) + P \leq V,$$

$$\hat{c}w_F(x) + P \leq \min_{b'}\{\hat{c}w_B(b'|x) + b'\}, \quad \hat{c}w_F(x) + P \leq V.$$

Let $\beta \in (0, 1)$, and $\tilde{c} = \beta c + (1 - \beta)\hat{c}$. Taking convex combination of the two inequalities on the right side above, we obtain $\tilde{c}w_F(x) + P \leq V$, implying the customer with unit waiting cost \tilde{c} would prefer obtaining service from the fixed-price firm over leaving without obtaining service. Similarly, we obtain

$$\begin{aligned} \tilde{c}w_F(x) + P &\leq \beta \min_{b'}\{cw_B(b'|x) + b'\} + (1 - \beta) \min_{b'}\{\hat{c}w_B(b'|x) + b'\} \\ &< \min_{b'}\{\tilde{c}w_B(b'|x) + b'\}. \end{aligned}$$

Thus, the customer with unit waiting cost \tilde{c} would strictly prefer obtaining service from the fixed-price firm over obtaining service from the bid-based firm. Taken together, this implies $x(\tilde{c}) = \text{FIX}$, and hence C_F is convex.

If C_F is empty, then all customers with unit waiting cost below c_ℓ choose to obtain service from the bid-based firm, and the service decision function is given by $x(c) = \text{BID}$ for $c \leq c_\ell$, and $x(c) = \text{LEAVE}$ for $c > c_\ell$. This fits the representation of the service decision function in the theorem with $c_1 \triangleq c_\ell$ and $c_2 \triangleq c_\ell$.

Hence, for the rest of the proof, suppose C_F is non-empty. Define $c_1 \triangleq \inf C_F$ and $c_2 \triangleq \sup C_F$, such that $c_1 < c_2$. We now show that $0 < c_1$ and $c_2 < c_\ell$.

Suppose $c_1 = 0$. By convexity of C_F , this implies that there exists an $\bar{\epsilon} > 0$, such that $x(c) = \text{FIX}$ for all $c < \bar{\epsilon}$. For a customer with unit waiting cost $\epsilon < \bar{\epsilon}$, his total expected cost is $\epsilon w_F(x) + P$. Suppose instead the customer chooses to obtain service

from the bid-based firm with a zero bid. This ensures that his expected waiting time is equal to that of a customer with the lowest priority in the bid-based firm. By Lemma 1, the expected waiting time in equilibrium is non-increasing in unit waiting cost, and hence the expected waiting time of a customer with the lowest priority in the bid-based firm is at most $w(\epsilon|x) = w_F(x)$. Thus, the total expected cost of the customer with unit waiting cost ϵ on choosing to obtain service from the bid-based firm with a zero bid is at most $\epsilon w_F(x) < \epsilon w_F(x) + P$. This contradicts the assumption that (x, b) is an equilibrium. Hence, we obtain that in any symmetric equilibrium, $c_1 > 0$.

Finally, suppose $c_2 = c_\ell$. Since C_F is convex (and non-empty), this implies that all customers with unit waiting cost below c_1 obtain service from the bid-based firm, all customers with unit waiting cost between c_1 and c_ℓ obtain service from the fixed-price firm, and the rest choose to leave without obtaining service. Consequently, a customer with unit waiting cost c_1 has the highest priority in the bid-based firm, and has an expected waiting time of 1. From Lemma 1, we know that the expected waiting time in an equilibrium is non-increasing in the unit waiting cost, implying that the expected waiting time of a customer in the fixed-price firm, $w_F(x)$, is less than or equal 1. However, since C_F is non-empty, the expected waiting time in the fixed-price firm has to be strictly greater than the service completion time 1. Thus we obtain a contradiction.

Summing up, we obtain that any symmetric equilibrium (x, b) where at least some customers choose to obtain service from the fixed-price firm, there exists thresholds

$0 < c_1 < c_2 < c_\ell \leq 1$, such that for all $c \in [0, c_1] \cup [c_2, c_\ell]$, we have $x(c) = \text{BID}$, for all $c \in (c_1, c_2)$, we have $x(c) = \text{FIX}$, and for all $c \in (c_\ell, 1]$, we have $x(c) = \text{LEAVE}$. \square

The preceding theorem states that in a symmetric equilibrium, a customer's decision about which firm to obtain service from has a simple threshold structure: customers with very low and very high unit waiting cost choose to obtain service from the bid-based firm, and those with intermediate unit waiting cost prefer to obtain service via the fixed-price firm. This structure suggests that the bid-based firm serves two different functions in the system. For those customers with very high unit waiting cost, the bid-based firm provides means to obtain high priority and get service immediately. On the other hand, for customers with very low unit waiting cost, the bid-based firm allows them to obtain service at low costs, albeit after longer waiting times.

In Fig. 2.1 and Fig. 2.2 we illustrate the equilibrium strategy, the expected waiting time in equilibrium and the equilibrium payment for the following parameter values: $\lambda = 40$, $P = 3$, and V is sufficiently high that every customer obtains service in equilibrium. The bid-based firm is an $M/M/20$ preemptive priority queue, and the fixed-price firm is an $M/M/30$ FIFO queue. From these plots, we observe the threshold structure of equilibrium, and the monotonicity of the expected waiting time and the payment with respect to unit time waiting cost. Note also that the low unit waiting cost customers in the bid-based firm pay less than the price in the fixed-price firm, and wait longer for service completion, while the high unit waiting cost customers in the bid-based firm pay more than the price in the fixed-price firm,

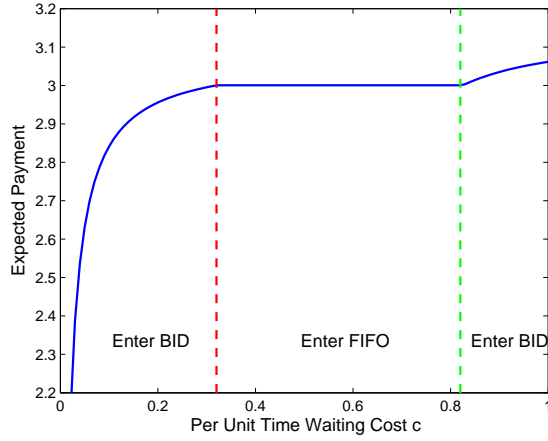


Figure 2.1: Expected payment in equilibrium of the customers' game

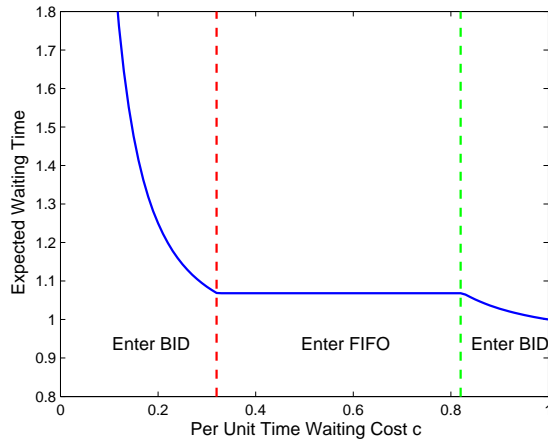


Figure 2.2: Expected waiting time in equilibrium of the customers' game

and have lower waiting times.

From the perspective of equilibrium analysis, the preceding theorem is important. From Lemma 2, in any symmetric equilibrium (x, b) , the customers' bidding function $b(\cdot)$ is fully specified from the service decision $x(\cdot)$. Combining this result with

the preceding theorem, we obtain that any symmetric equilibrium (x, b) can be characterized by three thresholds $c_1, c_2, c_\ell \in [0, 1]$. Thus from this point on, we only need to consider symmetric strategies characterized by c_1, c_2, c_ℓ . This greatly simplifies the analysis for showing existence of the equilibrium, as we see in the following section.

2.4 Existence and uniqueness of equilibrium in customers' game

Having determined the structure of a symmetric equilibrium of customers' game, we are now ready to present our main results regarding the existence and uniqueness of customers' equilibrium. We present our results in two steps: first, in Section 2.4.1, we consider a system where obtaining service is mandatory for all customers, denoted by SYS_{man} . Although this restriction is impracticable, we use the results for this system to show existence (and uniqueness) in the original system where customers have the option to leave the system without obtaining service, which we denote by SYS_{op} . This is achieved by carefully relating an equilibrium of the SYS_{man} system with a corresponding equilibrium of the SYS_{op} system. We provide the details in Section 2.4.2. We then extend the results to systems with positive reserve price in Section 2.4.3.

To state our results, we need two primitives. Let $\Gamma(x)$ denote the expected waiting

time of a customer in the bid-based firm, if the arrival rate of customers to that firm with higher priority is equal to $x \in [0, k)$. And, let $\Phi(x)$ denote the expected waiting time of a customer in the fixed-price firm, if the arrival rate of the customers to that firm is $x \in [0, n)$.

We make the following assumptions on $\Gamma(x)$ and $\Phi(x)$.

- Assumption 1.**
1. *The function Γ is finite, strictly increasing and continuous in $[0, k)$, and infinite after k .*
 2. *The function Φ is finite, strictly increasing and continuous in $[0, n)$ and infinite after n .*
 3. $\Gamma(0) = \Phi(0) = 1$.
 4. $\int_0^y \Gamma(t)dt \rightarrow \infty$, as $y \rightarrow k$.

The following lemma shows that Assumption 1 holds under many commonly studied queueing system. The proof is given in Appendix A.3.

Lemma 3. *Suppose the fixed-price firm operates a $G/G/n$ queue with service rate 1, and the bid-based firm operates a $G/M/k$ preemptive queue with service rate 1. Then the system satisfies Assumption 1.*

Note that implicitly, we have used the fact that the queue in the bid-based firm is preemptive. Without this assumption, the expected waiting time of a customer will be a function not only of the arrival rate of the customers with higher priority, but also of the total arrival rate to the queue. However all our analysis in Sections

2.3 and 2.4 can be reproduced for non-preemptive queues. Please refer to Appendix A.7 for details.

Following Theorem 1, we represent a customer's strategy by a vector of thresholds $\bar{c} = (c_1, c_2, c_\ell)$. Assuming all customers follow the strategy \bar{c} , let $w_F(\bar{c})$ be the expected waiting time in the fixed-price firm. Similarly, let $w_B(c|\bar{c})$ be the expected waiting time of a customer of unit waiting cost c if she chooses to obtain service from the bid-based firm and subsequently makes the optimal bid from (2.1). Letting $\alpha = c_2 - c_1$ for simplicity of notation, we have the following expressions:

$$w_F(\bar{c}) = \Phi(\lambda\alpha), \quad (2.3)$$

$$w_B(t|\bar{c}) = \begin{cases} \Gamma(\lambda c_\ell - \lambda t) & \text{if } t \in [c_2, c_\ell]; \\ \Gamma(\lambda c_\ell - \lambda\alpha - \lambda t) & \text{if } t \in [0, c_1]. \end{cases} \quad (2.4)$$

The first equation follows from the fact that, under strategy \bar{c} , the arrival rate of customers to the fixed-price firm equals $\lambda(c_2 - c_1) = \lambda\alpha$. The second equation follows from the fact that for $t \in [c_2, c_\ell]$, the arrival rate of customers to the bid-based firm with higher priority than the customer with unit waiting cost t is equal to $\lambda(c_\ell - t)$, whereas it equals $\lambda(c_\ell - \alpha - t)$ for $t \in [0, c_1]$, since all customers with unit waiting cost in $[c_1, c_2]$ obtain service from the fixed-price firm. (Note that these expressions make use of our assumption that the unit waiting costs are uniformly distributed. We briefly describe in Appendix A.1 how these expressions differ for non-uniform distributions.)

2.4.1 System with mandatory service requirement

We first consider the setting where the customers do not have the option to leave the system without obtaining service. In other words, we assume that the customers' value for service completion V is sufficiently high to render obtaining service mandatory for all customers. Denote such a system by $\text{SYS}_{\text{man}}(\lambda, P)$ if the arrival rate is λ and the fixed-price price is P . In such a setting, customers arrive to the system and make the choice between obtaining service from the fixed-price firm or the bid-based firm, and a symmetric equilibrium for this system can be represented as $\bar{c} = (c_1, c_2, 1)$. In the following, we show that for all $\lambda \in (0, n + k)$ and for all $P \geq 0$, there exists a unique symmetric equilibrium. Towards that goal, define $P_{\text{max}}(\lambda)$ as follows:

$$P_{\text{max}}(\lambda) \triangleq \begin{cases} \frac{1}{\lambda} \int_0^\lambda \Gamma(t) dt - 1 & \text{if } \lambda < k; \\ \infty & \text{otherwise.} \end{cases} \quad (2.5)$$

As we show below, for each arrival rate λ , $P_{\text{max}}(\lambda)$ imposes a bound on the value of the price P in the fixed-price firm, so that for $P \geq P_{\text{max}}(\lambda)$, no customer obtains service from the fixed-price firm in a symmetric equilibrium of the customers' game.

Suppose the strategy $\bar{c} = (1, 1, 1)$ constitutes a symmetric equilibrium, where no customers choose to obtain service from the fixed-price firm. In such an equilibrium, if the customer with unit waiting cost equal to one obtains service from the fixed-price firm, her expected waiting time to service completion is 1 and her expected payment is P . Since in equilibrium such a customer prefers obtaining service from the bid-based

firm, it must be the case that

$$\int_0^1 w_B(t|\bar{c}) dt \leq 1 + P. \quad (\text{Pref-BID})$$

Here, the left hand side denotes the total expected cost of the customer with unit waiting cost c_ℓ , as per (2.2). The necessary condition (Pref-BID) requires that this be less than that of obtaining service from the fixed-price firm. Note that, using (2.4), it is straightforward to show that this implies $P \geq P_{\max}(\lambda)$. Thus, for $\bar{c} = (1, 1, 1)$ to be a symmetric equilibrium for $\text{SYS}_{\text{man}}(\lambda, P)$, a necessary condition is $P \geq P_{\max}(\lambda)$. The following theorem shows that it is also sufficient. The proof of this theorem is given in Appendix A.4.

Theorem 2. *There exists a unique symmetric customers' equilibrium of the system $\text{SYS}_{\text{man}}(\lambda, P)$ of the form $\bar{c} = (1, 1, 1)$ if and only if $P \geq P_{\max}(\lambda)$.*

The preceding theorem states that for any $\lambda \in (0, n + k)$, $P_{\max}(\lambda)$ is the highest price in the fixed-price firm for which one may expect customers to choose that firm for obtaining service. For values of P greater than $P_{\max}(\lambda)$, all customers prefer the bid-based firm over the fixed-price firm. The following theorem, our main result, shows that for values of P less than $P_{\max}(\lambda)$, there exists a unique symmetric equilibrium, in which there is a positive arrival rate of customers to the fixed-price firm.

Theorem 3. *For all $0 < P < P_{\max}(\lambda)$, there exists a unique symmetric customers' equilibrium $\bar{c} = (c_1, c_2, 1)$, with $0 < c_1 < c_2 < 1$. In other words, in equilibrium, the arrival rate of customers to the fixed-price firm is positive.*

Our proof proceeds by first identifying a set of necessary and sufficient conditions

for a strategy $\bar{c} = (c_1, c_2, 1)$ to be a symmetric equilibrium for $\text{SYS}_{\text{man}}(\lambda, P)$, and second showing that these conditions have a unique solution by explicit construction. In the following, we provide the intuition behind these necessary and sufficient conditions.

Suppose $0 < P < P_{\text{max}}(\lambda)$, and consider a symmetric equilibrium $\bar{c} = (c_1, c_2, 1)$ with $0 < c_1 < c_2 < 1$, where the arrival rate of customers to the fixed-price firm is positive. In this equilibrium, a customer with unit waiting cost c_1 must be indifferent between obtaining service in the bid-based firm and the fixed-price firm. For otherwise, a customer with unit waiting costs slightly higher than c_1 would strictly prefer to mimic the behavior customer with unit waiting cost c_1 . Equating the customer's total expected cost, we have

$$\int_0^{c_1} w_B(t|\bar{c})dt = c_1 w_F(\bar{c}) + P. \quad (\text{ContT-P})$$

Note that since all customers with unit waiting cost below c_1 obtain service from the bid-based firm, we have $w(t|\bar{c}) = w_B(t|\bar{c})$ and consequently, as per (2.2), the left hand side denotes the total expected cost of a customer with unit waiting cost c_1 . On the other hand, the right hand side denotes the total expected cost for such a customer on obtaining service from the fixed-price firm.

Similarly, consider a customer with unit waiting cost c_2 . Since all customers with unit waiting costs between c_1 and c_2 obtain service from the fixed-price firm, the expected waiting time of the customer with unit waiting cost c_2 must equal that of a customer with unit waiting cost c_1 . Since from Lemma 1, we know that the expected waiting time in equilibrium is non-increasing in the unit waiting cost, and since all

customers in the fixed-price firm have waiting times $w_F(\bar{c})$, it follows that

$$w_B(c_2|\bar{c}) = w_B(c_1|\bar{c}) = w_F(\bar{c}). \quad (\text{ContW-P})$$

The necessary condition (ContW-P) thus states that the expected waiting time in equilibrium must be continuous at c_1 .

The following proposition formalizes the preceding discussion, and shows that the aforementioned necessary conditions are also sufficient for a strategy to be an equilibrium. We provide the proof in Appendix A.3.

Proposition 1. *A strategy $\bar{c} = (c_1, c_2, 1)$, with $0 < c_1 < c_2 < 1$, is a symmetric customers' equilibrium for the system $\text{SYS}_{\text{man}}(\lambda, P)$ with $0 < P < P_{\text{max}}(\lambda)$ if and only if the conditions (ContT-P) and (ContW-P) hold.*

Using (2.4) and (2.3), we can summarize the conditions (ContT-P) and (ContW-P) as follows:

$$\begin{aligned} \int_0^{c_1} \Gamma(\lambda(1 - \alpha - t))dt &= c_1\Phi(\lambda\alpha) + P \\ \Gamma(\lambda(1 - \alpha - c_1)) &= \Phi(\lambda\alpha). \end{aligned} \quad (2.6)$$

Thus, following Proposition 1, showing the existence (and uniqueness) of a symmetric equilibrium for the system $\text{SYS}_{\text{man}}(\lambda, P)$ with $P \in (0, P_{\text{max}}(\lambda))$ requires showing that there exists a (unique) $c_1 \in [0, 1)$ and $\alpha \in (0, 1 - c_1]$ that satisfy the set of equations (2.6). We show that this is indeed the case in the proof of Theorem 3. We provide the details of the proof in Appendix A.4.

2.4.2 System with optional service requirement

In this section, we extend our existence and uniqueness result to systems where customers may choose not to obtain service. We denote the system with optional service requirement with arrival rate λ , fixed price P , and the value of service completion V by $\text{SYS}_{\text{op}}(\lambda, P, V)$.

For this system, consider a symmetric equilibrium $\bar{c} = (c_1, c_2, c_\ell)$. In equilibrium, each customer choosing to obtain service must have a total expected cost that is less than or equal to the value of service completion V . In particular, this holds for a customer with unit waiting cost c_ℓ . Moreover, if c_ℓ is strictly less than one, the total expected cost of such a customer must exactly equal to V . For if this were not true, a customer with unit waiting cost slightly greater than c_ℓ would find it preferable to obtain service from the system. Thus, we obtain the following necessary condition on an equilibrium:

$$c_\ell < 1, \quad \int_0^{c_\ell} w(t|\bar{c})dt = V, \quad \text{OR} \quad c_\ell = 1, \quad \int_0^{c_\ell} w(t|\bar{c})dt \leq V. \quad (\text{IND})$$

Here, as per (2.2), the integral denotes the total expected cost of a customer with unit waiting cost c_ℓ . Note that we have $w(t|\bar{c}) = w_F(\bar{c})$ for $t \in (c_1, c_2)$ and $w(t|\bar{c}) = w_B(t|\bar{c})$ for $t \in [0, c_1] \cup [c_2, c_\ell]$. Using (IND), we can now relate the symmetric equilibria of the system SYS_{op} with those of the system SYS_{man} . The following lemmas formalize this argument. The proof is provided in Appendix A.3.

Lemma 4. *If the strategy $(c_1, c_2, 1)$ is a symmetric equilibrium for the system $\text{SYS}_{\text{man}}(\lambda, P)$ with $\lambda \in (0, n + k)$, then the strategy $\bar{c}(u) = (c_1 u, c_2 u, u)$ for $u \in (0, 1]$*

is a symmetric equilibrium for the system $\text{SYS}_{\text{op}}(\frac{\lambda}{u}, Pu, V)$ if and only if the condition (IND) holds for the strategy $\bar{c}(u)$.

The preceding lemma states that a symmetric equilibrium of the system SYS_{man} can be used to construct a symmetric equilibrium of a related SYS_{op} system, as long as one can ensure that the condition (IND) is satisfied. Conversely, the following lemma constructs a symmetric equilibrium for a SYS_{man} system using the symmetric equilibrium of a related SYS_{op} system.

Lemma 5. *If the strategy (c_1, c_2, u) is a symmetric equilibrium of the system $\text{SYS}_{\text{op}}(\lambda, P, V)$ then the strategy $(\frac{c_1}{u}, \frac{c_2}{u}, 1)$ is a symmetric equilibrium of the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$.*

Let $U_\lambda = (0, 1] \cap (0, \frac{n+k}{\lambda})$. The preceding lemmas, together with Theorem 3, imply that there exist functions $\mathcal{C}_i : U_\lambda \rightarrow [0, 1]$ for $i = 1, 2$, such that for each $u \in U_\lambda$, we have $\mathcal{C}_1(u) \leq \mathcal{C}_2(u) \leq u$, and the strategy $(\frac{\mathcal{C}_1(u)}{u}, \frac{\mathcal{C}_2(u)}{u}, 1)$ is the unique symmetric equilibrium of the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$. Using this, we now state the existence (and uniqueness) result for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$.

Theorem 4. *For each $\lambda > 0$, $P > 0$ and $V > 0$, there exists $u = u(\lambda, P, V) \in U_\lambda$ such that the strategy $(\mathcal{C}_1(u), \mathcal{C}_2(u), u)$ constitutes the unique symmetric customers' equilibrium for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$. Further, for each $\lambda > 0$ and $V > 0$, there exists a threshold $P(\lambda, V)$ such that for all $P \geq P(\lambda, V)$, in the symmetric equilibrium, the arrival rate of customers to the fixed-price firm is zero, whereas for all $P \in (0, P(\lambda, V))$, the arrival rate of customers to the fixed-price firm is positive.*

To see the intuition behind the result, observe that, by using the expressions (2.4) and (2.3), we can write the expected waiting time function for the strategy as

$$\bar{\mathcal{C}}(u) = (\mathcal{C}_1(u), \mathcal{C}_2(u), u),$$

$$w(t|\bar{\mathcal{C}}(u)) = \begin{cases} \Gamma(\lambda u - \lambda t) & \text{if } t \in [\mathcal{C}_2(u), u]; \\ \Phi(\lambda \mathcal{A}(u)) & \text{if } t \in (\mathcal{C}_1(u), \mathcal{C}_2(u)); \\ \Gamma(\lambda u - \lambda \mathcal{A}(u) - \lambda t) & \text{if } t \in [0, \mathcal{C}_1(u)], \end{cases}$$

where $\mathcal{A}(u) = \mathcal{C}_2(u) - \mathcal{C}_1(u)$. By straightforward algebra, the condition (IND) then reduces to

$$\begin{aligned} \int_0^{u-\mathcal{A}(u)} \Gamma(\lambda t) dt + \mathcal{A}(u) \Phi(\lambda \mathcal{A}(u)) &\leq V, \quad \text{if } u < 1, \\ \int_0^{u-\mathcal{A}(u)} \Gamma(\lambda t) dt + \mathcal{A}(u) \Phi(\lambda \mathcal{A}(u)) &= V, \quad \text{if } u = 1. \end{aligned} \quad (2.7)$$

Thus, showing the existence (and uniqueness) of a symmetric equilibrium for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$ requires showing that the preceding equation has a unique solution $u = u(\lambda, P, V)$. This is obtained by showing that both $\mathcal{A}(u)$ and $u - \mathcal{A}(u)$ are continuous and non-decreasing in $u \in U_\lambda$. We provide the details in Appendix A.5.

2.4.3 System with positive reserve price

The existence and uniqueness of customers' equilibrium can be easily extended to the setting where reserve price in the bid-based firm r is positive. When $P > r$, there is an straightforward one-to-one mapping between the systems with fixed price P , reserve price r , value V , and the system with fixed price $P - r$, reserve price 0, value $V - r$. In particular, the two systems have the same equilibrium thresholds, and any

customer who receives service in equilibrium pays an additional r in the former system in addition to her payment in the latter system. To see this, note that in the former system, all customers who receive service pays at least r , thus subtracting r from both the prices and the value does not change the customers' service decision. As a result, we can obtain the service decisions in the former system with reserve price, by calculating that in the latter system, which has zero reserve price. Customers' bids in the former system is then given by r plus their bids in the latter system.

When $P \leq r$, the system cannot be mapped directly to any system that we have studied in the earlier parts of the chapter. However, we can follow the same analysis procedures, to show that there exists a unique customers' equilibrium. In particular, c_1 is equal to 0 in equilibrium, and (c_2, c_ℓ) satisfies condition (IND), as well as a condition which suggests that the customer with unit waiting cost c_2 is indifferent between the two firms (or $c_2 = c_\ell$ and the customer with unit waiting cost c_2 prefers fixed-price over bid-based). Namely,

$$c_\ell < 1, \quad P + \int_0^{c_\ell} w(t|\bar{c})dt = V, \quad \text{OR} \quad c_\ell = 1, \quad P + \int_0^{c_\ell} w(t|\bar{c})dt \leq V, \quad (\text{IND})$$

and

$$c_2 < c_\ell, \quad P + c_2 w_F(\bar{c}) = r + c_2 w_B(c_2|\bar{c}),$$

$$\text{OR} \quad c_2 = c_\ell, \quad P + c_2 w_F(\bar{c}) \leq r + c_2 w_B(c_2|\bar{c}).$$

The following theorem summarizes the discussions above.

Theorem 5. *For each $\lambda > 0$, $P > 0$, $r > 0$ and $V > 0$, there exists a strategy characterized by thresholds $\bar{c} = (c_1, c_2, c_\ell)$ that constitutes the unique symmetric customers' equilibrium for the system with optional service requirements. In particular,*

when $r \geq P$, we have $c_1 = 0$ and the equilibrium strategy is a two-thresholds strategy.

2.5 Competition between firms

Next, using the unique equilibrium among the customers for any given fixed price P and reserve price r , we study the competition between the two firms in setting their prices. In particular, we consider a game between the two firms, where the firm operating fixed-price service sets P and the firm operating bid-based service sets r so as to maximize their respective expected revenue. In this resulting game, we investigate the existence of Nash equilibria.

Without loss of generality, we assume that the fixed-price firm chooses $P \in [0, V]$ and the bid-based firm chooses $r \in [0, V]$, since otherwise no customer would obtain service from the respective firm, resulting in zero profit for that firm. Furthermore, we assume that $V \geq 1$. This is because, if $V < 1$, some customers would never choose to obtain service from either firm, even if they were to be served immediately upon arrival, and such customers can be ignored for the sake of analysis. The resulting system with $V < 1$ can then be mapped to a system with $V \geq 1$ and lower λ , using techniques as in the proof of Lemma 4.

Given these assumptions, we now state the following result regarding the existence of mixed Nash equilibria for the game between the two firms.

Theorem 6. *For $V > 1$, the game between the two firms has a mixed Nash equilibrium.*

This result follows from the continuity of the expected revenue of the two firms in P and r , and from the fact that they take values in a compact set. We prove the former statement by showing that \bar{c} is continuous in both P and r . The details are given in Appendix A.6.

Since we only obtain existence of equilibrium in *mixed* strategies for the general case, we next analyze the system under a specific assumption for the expected waiting times and in a regime where the arrival rate and the service capacity is large. In Section 2.5.2, we analyze this system in the limit where the arrival rate goes to infinity, and fully characterize the set of *pure* equilibria of the limiting game. Using standard arguments, we show in Theorem 7 that these pure equilibria are ϵ -equilibria for the system for finite, but large enough, arrival rate. Finally, in Section 2.5.3, we use the characterization of the set of pure equilibria of the limiting game to study its revenue properties, in comparison to the setting where the fixed-price and the bid-based service were operated by a single firm. In particular, we show that the total revenue in the equilibria of the limiting game is fairly close to that obtained in the system with a single firm. These results together suggest that for large arrival rates and large capacity, the competition between the firms does not significantly affect the revenues in equilibrium.

2.5.1 High arrival rate and high capacity regime

In this subsection, we introduce the limiting regime and the waiting time expressions under which we perform our analysis. In particular, we consider the limiting regime where $\lambda \rightarrow \infty$, with $n = q_F \lambda$ and $k = q_B \lambda$, for some $q_F > 0$ and $q_B > 0$. Since n and k denote the capacities of the two firms, this regime applies to settings where the arrival rate and the capacities of the two firms are proportionally large. We restrict our analysis in the following parameter regime.

Assumption 2 (Duopoly regime). 1. $0 < q_F < 1$ and $0 < q_B < 1$

2. $q_F + q_B \geq 1$.

In particular, the first condition ensures that no firm has enough capacity to serve all customers in the market, while the second condition ensures that the two firms together have enough capacity to serve everyone.

To obtain analytical expressions for the thresholds, we focus on a setting with explicit expected waiting time expressions $w_F(\cdot)$ and $w_B(\cdot|\bar{c})$, motivated by $M/M/1$ FIFO and (preemptive) priority queues. In particular, suppose the strategy adopted by the customers is given by $\bar{c} = (c_1, c_2, c_\ell)$. Let $\rho_F = \lambda(c_2 - c_1)/n$ denote the congestion (see [27]) at the fixed-price firm, and let $\rho_B = \lambda(c_\ell - c_2 + c_1)/k$ denote the congestion at the bid-based firm.

We assume that the expected waiting time at the fixed-price firm is given explicitly

by

$$w_F(\bar{c}) = \frac{\rho_F}{1 - \rho_F} \frac{1}{n} + 1.$$

Essentially, making this assumption implies that when a customer waiting for service in the fixed-price firm, the service requirement ahead of her decreases at rate n . The first term then denotes the waiting time until the beginning of service in an $M/M/1$ FIFO queue with service rate n and arrival rate $\lambda(c_2 - c_1)$. Once the customer starts her service, her service is processed at rate 1, and incurs an additional waiting time of one, as in the second term. As an example of an instance where this waiting time expression holds, consider a airport taxi service, where customers are riders waiting for taxi, each of which arrive at rate n . Thus, the time waiting for arrival of a taxi is given by the first term in the expression. Once a taxi arrives, the time it takes to reach the destination is fixed and independent of the service rate n .

Similarly, we assume the expected waiting time of a customer with unit waiting cost c in the bid-based firm is given by

$$w_B(c; \bar{c}) = \frac{1}{k(1 - \rho_B + \rho_B B(c; \bar{c}))^2} - \frac{1}{k} + 1,$$

where $B(c; \bar{c})$ is the fraction of customers that join the bid-based firm with unit cost less than c . Namely, we have

$$B(x; \bar{c}) = \begin{cases} \frac{x}{c_\ell - c_2 + c_1} & \text{if } x \leq c_1; \\ \frac{x + c_1 - c_2}{c_\ell - c_2 + c_1} & \text{if } x \geq c_2. \end{cases}$$

To obtain the expression for $w_B(\cdot; \bar{c})$, we again make the assumption that for a customer waiting for service in the bid-based firm, the service requirement ahead of her decreases at rate k . The first two terms then denote the waiting time until the beginning of for service in a $M/M/1$ priority queue with service rate κ and total

arrival rate $\lambda(c_\ell - c_2 + c_1)$, where we have used the waiting time expression for priority queues as in [36]. The last term then denotes the time in service.

2.5.2 Equilibria among firms in the limiting system

Substituting the above expected waiting time expressions into equations (ContT-P), (ContW-P) and (IND), let $(c_1^\lambda, c_2^\lambda, c_\ell^\lambda)$ denote the equilibrium thresholds for any fixed $\lambda > 0$. Recall that $\alpha^\lambda = c_2^\lambda - c_1^\lambda$. The following lemma characterizes the limiting thresholds and the expected revenue per arrival as λ approaches infinity under Assumption 2. The proof is given in Appendix A.6.

Lemma 6. *Under Assumption 2, as λ goes to $+\infty$, the asymptotic properties of the system are given by:*

1. *If $P > r$ and $P < V - q_B$, we have $c_\ell^\lambda \rightarrow c_\ell^\infty(P, r) = \min\{V - P, 1\}$, $c_1^\lambda \rightarrow c_1^\infty(P, r) = c_\ell^\infty(P, r) \sqrt{q_B / \left(\frac{1}{q_F + q_B - c_\ell^\infty(P, r)} - \frac{1}{q_F} + \frac{1}{q_B} \right)}$, and $\alpha \rightarrow \alpha^\infty(P, r) = c_\ell^\infty(P, r) - q_B$. Moreover, the limiting expected revenue of the fixed-price firm per arrival is $R_F^\infty(P, r) = (c_\ell^\infty(P, r) - q_B) P$, and the limiting expected revenue of the bid-based firm per arrival is $R_B^\infty(P, r) = q_B P$.*
2. *If $P \leq r$ and $r < V - q_F$, we have $c_\ell^\lambda \rightarrow c_\ell^\infty(P, r) = \min\{V - r, 1\}$, $c_1^\lambda \rightarrow 0$, and $\alpha^\lambda \rightarrow \alpha^\infty(P, r) = q_F$. Moreover, the limiting expected revenue of the fixed-price firm per arrival is $R_F^\infty(P, r) = q_F P$, and the limiting expected revenue of the bid-based firm per arrival is $R_B^\infty(P, r) = (c_\ell^\infty(P, r) - q_F) r$.*

3. If $P > r$ and $P \geq V - q_B$, we have that c_1^λ, c_2^λ and c_i^λ all tend to $\min\{V - r, q_B\}$ as $\lambda \rightarrow \infty$. The limiting expected revenue of the fixed-price firm per arrival is $R_F^\infty(P, r) = 0$, and the limiting expected revenue of the bid-based firm per arrival is $R_B^\infty(P, r) = q_B r$.
4. If $P \leq r$ and $r \geq V - q_F$, we have $c_1^\lambda \rightarrow 0$, whereas c_2^λ and c_i^λ both tend to $\min\{V - P, q_F\}$ as $\lambda \rightarrow \infty$. The limiting expected revenue of the fixed-price firm per arrival is $R_F^\infty(P, r) = q_F P$, and the limiting expected revenue of the bid-based firm per arrival is $R_B^\infty(P, r) = 0$.

From Lemma 6, we obtain that the firms' limiting expected revenue per arrival has simple closed form. We thus study a “limiting game”, where the fixed-price and bid-based firms choose P and r respectively to maximize their payoffs $R_F^\infty(P, r)$ and $R_B^\infty(P, r)$. First, we justify this analysis of the limiting game via the following theorem, which states that any Nash equilibrium of the limiting game is an ϵ -equilibrium of the original game.

Theorem 7. *For any $\epsilon > 0$, there exists a λ_0 , such that for any $\lambda > \lambda_0$, every equilibrium (P, r) of the limiting game is an ϵ -equilibrium in the original game with arrival rate λ .*

This theorem follows from the continuity of the firms' revenue in λ , which follows directly from the continuity of the equilibrium thresholds in λ . The details are given in Appendix A.6.

With this theorem in place, we now characterize the set of Nash equilibria of

the limiting game in the following lemma. The details of the proof are provided in Appendix A.6.

Lemma 7. *The following holds under Assumption 2.*

1. *If $V \geq 2 - q_B$, the set of Nash equilibria of the limiting game is given by $\{(P, r) : P = V - 1, r \leq \frac{1 - q_B}{q_F}(V - 1)\}$. Under any equilibrium (P, r) , $R_B^\infty(P, r) = q_B(V - 1)$ and $R_F^\infty(P, r) = (1 - q_B)(V - 1)$.*
2. *If $V < 2 - q_B$, the set of Nash equilibria of the limiting game is given by $\{(P, r) : P = \frac{V - q_B}{2}, r \leq V - 1\}$. Under any equilibrium (P, r) , $R_B^\infty(P, r) = q_B \frac{V - q_B}{2}$ and $R_F^\infty(P, r) = \left(\frac{V - q_B}{2}\right)^2$.*

2.5.3 Price of stability of the limiting system

From the preceding lemma, we obtain that for fixed parameters q_B, q_F and V , the firms' payoffs R_B^∞ and R_F^∞ are the same under all Nash equilibria. In the following, we compare the total revenue obtained by the two firms in any such equilibrium to the maximum total revenue that can be obtained if the two firms colluded to set P and r . More precisely, we define the *price of stability* (PoS) (see [56]) to be the ratio of the maximum revenue obtained by the two firms under collusion to the total revenue of the two firms in equilibrium. We have the following result for the price of stability in the limiting game.

Theorem 8. *For $V \geq 2$, we have $\text{PoS} = 1$. For $1 \leq V < 2$, PoS is strictly decreasing in V , with $\text{PoS} \leq \frac{9}{8}$ if $1.5 \leq V < 2$ and $\text{PoS} \leq \frac{1}{1 - q_B^2}$ if $1 \leq V < 1.5$.*

The proof of this theorem is given in Appendix A.6. From this theorem, we see that in the firms' limiting game, for small q_B or large V , the price of stability is relatively low. Moreover, for sufficiently large V , the total payoff under Nash equilibria is the same as the total payoff under the collusion. This suggests that in the limiting game, the competition between the two firms does not significantly affect the revenues in equilibrium. By Theorem 7, we conclude the same holds for ϵ -equilibrium in the finite system for large enough arrival rates.

We conclude this section by two numerical examples that illustrate our discussion above. Figure 2.3 shows the total revenue per arrival in a Nash equilibrium and under collusion, as k, n and λ increase proportionally, assuming $q_B = q_F = 0.5$, and $V = 5$. Figure 2.4 shows the price of stability¹ for the same set of parameters. Observe that the ratio is very close to 1 for large λ , which coincides with the first case of Theorem 8. Similarly, Figure 2.5 and 2.6 show the revenues and price of stability as k, n and λ increase proportionally, assuming $q_B = q_F = 0.5$, and $V = 1.5$. The price of stability fluctuates around $9/8$ for large λ , which coincides with the second case of Theorem 8.

2.6 Conclusion

In this chapter, we analyze a model of two competing service providers offering service which differ in pricing rules and priority of service. We show under general settings the existence and uniqueness of a symmetric equilibrium of the customers' game,

¹Note that we do not have any monotonicity results for the price of stability as arrival rate increases. Thus we have zig zag patterns for price of stability.

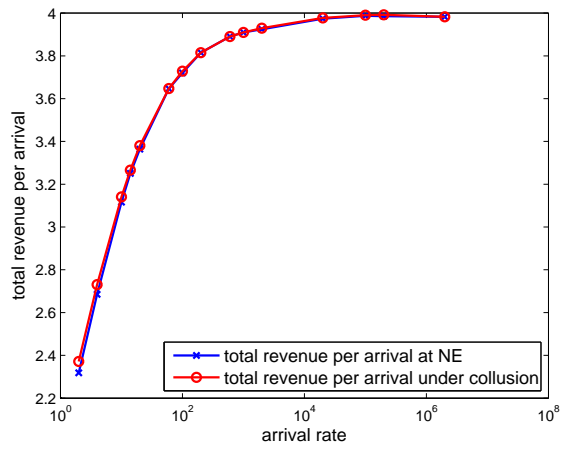


Figure 2.3: Total revenue per arrival in Nash equilibrium and under collusion, when $V = 5$




Figure 2.4: Ratio between the per arrival total revenues, when $V = 5$




Figure 2.5: Total revenue per arrival in Nash equilibrium and under collusion, when $V = 1.5$

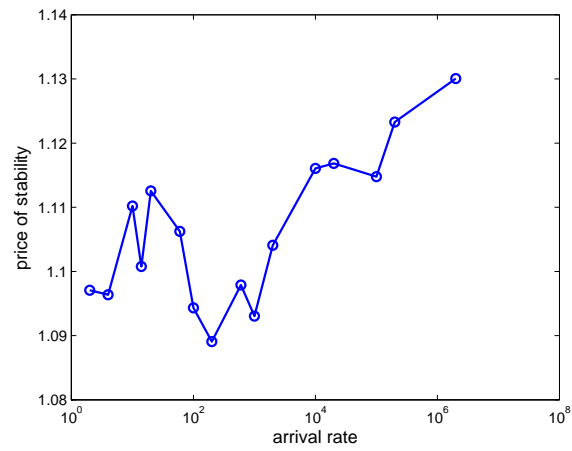


Figure 2.6: Ratio between the per arrival total revenues, when $V = 1.5$

where the customers' strategy has a multi-threshold structure. In particular, we show the customers endogenously segregate between the two firms, with customer with very low and high waiting costs choosing to obtain service from the bid-based firm, whereas customers with moderate waiting costs choosing to obtain service from the fixed-price firm. With the characterization of the unique equilibrium among the customers, we study the price competition between the two firms, and show that under high arrival rate and explicit waiting time assumptions, the competition between the two firms does not substantially affect the total revenue per arrival.

There are many avenues for future research. In many practical applications, one needs to consider service abandonments and dynamic pricing of the firms in response to the real-time state of the system. Incorporating these practical considerations into our model is an interesting area for future research.

CHAPTER 3
PRICE COMPETITION UNDER LINEAR DEMAND AND FINITE
INVENTORIES: CONTRACTION AND APPROXIMATE
EQUILIBRIA

In many practical situations, multiple firms selling substitutable products set their prices competitively to sell limited inventories over a finite selling horizon, given that the demand of each firm jointly depends on the prices charged by all firms. For example, airlines competitively set the prices for their limited seat inventories in a particular market. Firms selling electronic products take the prices of their competitors into consideration when setting their prices. In this chapter, we consider multiple firms with limited inventories of substitutable products. Each firm chooses the prices that it charges for its product over a finite selling horizon. The demand that each firm faces is a deterministic function of the prices charged by all of the firms, where the demand of a firm is linearly decreasing in its price and linearly increasing in the prices of the other firms. Each firm chooses its prices over a finite selling horizon to maximize its total revenue.

MAIN CONTRIBUTIONS. We study two types of equilibrium for the competitive pricing setting described above. In an equilibrium without recourse, at the beginning of the selling horizon, each firm selects and commits to the prices it charges over the whole selling horizon, assuming that the other firms do the same. In an equilibrium with recourse, at each time period in the selling horizon, each firm observes the inventories of all of the firms and chooses its price at the current time period, again

under the assumption that the other firms do the same. Essentially, an equilibrium without recourse corresponds to an open-loop equilibrium [19], whereas an equilibrium with recourse corresponds to a Markov perfect equilibrium (MPE) [20] in the dynamic game among the firms. Despite the fact that the demand of each firm is a deterministic function of the prices so that there is no uncertainty in the firms' responses, we show a clear contrast between the two equilibrium notions.

We consider the diagonal dominant regime, where the price charged by each firm affects its demand more than the prices charged by the other firms. In other words, if all of the competitors of a firm decrease their prices by a certain amount, then the firm can decrease its price by the same amount to ensure that its demand does not decrease. This regime is rather standard in the existing literature and it is used in, for example, [6] and [21]. Focusing on the equilibrium without recourse, we show in Section 3.1 that the best response of each firm to the price trajectories of the other firms is a contraction mapping, when viewed as a function of the prices of the other firms. In this case, it immediately follows that the equilibrium without recourse always exists and it is unique (see [68, Section 2.5]).

We give counterexamples in Section 3.2 to show that an equilibrium with recourse may not exist or may not be unique. Motivated by this observation, we look for an approximate equilibrium that is guaranteed to exist. We call a strategy profile for the firms an ϵ -equilibrium with recourse if no firm can improve its total revenue by more than ϵ by deviating from its strategy profile. We consider a low influence regime, where the effect of the price of a firm on the demand of another firm is diminishing,

which naturally holds when the number of firms is large. We show in Section 3.3 that the equilibrium without recourse can be used to construct an ϵ -equilibrium with recourse that has the same price trajectory as the equilibrium without recourse. So, intuitively speaking, an ϵ -equilibrium with recourse is expected to exist when the number of firms is large.

Our results fill a gap in a fundamental class of revenue management problems. Although there is no uncertainty in the firms' responses, the equilibria with and without recourse are not the same concept and can be qualitatively quite different. While the equilibrium without recourse uniquely exists, the same need not hold for the equilibrium with recourse. Also, our contraction argument for showing the existence and uniqueness of the equilibrium without recourse uses the Karush-Kuhn-Tucker (KKT) conditions for the firm's problem. Though contraction arguments are standard for showing existence and uniqueness of equilibrium [68], to the best of our knowledge, this duality-based contraction argument is new for price competition under limited inventories. This argument becomes surprisingly effective when dealing with linear demand functions, but it is an open question whether similar arguments hold for other demand functions. Lastly, our results indicate that in a low influence regime the equilibrium without recourse can be used to construct an ϵ -equilibrium with recourse with the same price trajectory as the equilibrium without recourse.

LITERATURE REVIEW. Similar to us, [21] considers price competition among multiple firms with limited inventories over a finite selling horizon. There are three key differences between their work and ours. First, they focus on a continuous-time setting,

whereas we study a discrete-time formulation. Second, they consider a generalized Nash game [60] where each firm considers all firms' capacity constraints while setting their prices, whereas in our model, each firm only considers its own capacity constraints. Most importantly, they focus on open-loop and closed-loop equilibria, and show that in the diagonally dominant regime a unique open-loop equilibrium exists and coincides with a closed-loop equilibrium. Although an equilibrium without recourse in our setting is the same as an open-loop equilibrium, our equilibrium with recourse is more restrictive than their closed-loop equilibrium. In particular, their closed-loop equilibria need not be *perfect*, whereas our equilibrium with recourse is a Markov perfect equilibrium. Thus, we show that the equilibrium with recourse can be different from the equilibrium without recourse. More precisely, although the former equilibrium need not exist or be unique, the latter is an approximate equilibrium with recourse in the low influence regime.

There are a number of papers that study price competition over a single period. [50] shows that pure Nash equilibrium (NE) exists for a wide class of supermodular demand models. [22] provides sufficient conditions for uniqueness of equilibrium in the Bertrand game when the demands of the firms are nonlinear functions of the prices, there is a non-linear cost associated with satisfying a certain volume of demand and each firm seeks to maximize its expected profit. [58] identifies the conditions for existence and uniqueness of pure NE when the demands are characterized by a mixture of multinomial logit models and the cost of satisfying a certain volume of demand is linear in the demand volume. [26] considers price competition among multiple firms when the relationship between demand and price is characterized by

the nested logit model and provides conditions to ensure the existence and uniqueness of the equilibrium. [54] proves the existence of pure strategy equilibrium in a price competition between two suppliers when capacity is private information.

Considering the papers on price competition over multiple time periods, [39] studies a stochastic game when there are strategic consumers choosing the time to purchase. [41] studies a competitive pricing problem when the relationship between demand and price is captured by the multinomial logit model and inventory levels are public information. [46] studies the pricing game between two firms with limited inventories facing stochastic demand. The authors characterize the unique subgame perfect Nash equilibrium. [43] shows the existence of a unique pure MPE in a pricing game between two firms offering vertically differentiated products.

3.1 Equilibrium without Recourse

There are n firms indexed by $N = \{1, \dots, n\}$. Firm i has c_i units of initial inventory, which cannot be replenished over the selling horizon. There are τ time periods in the selling horizon indexed by $T = \{1, \dots, \tau\}$. We use p_i^t to denote the price charged by firm i at time period t . Using $\mathbf{p}^t = (p_1^t, \dots, p_n^t)$ to denote the prices charged by all of the firms at time period t , the demand faced by firm i at time period t is given by $D_i^t(\mathbf{p}^t) = \alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t$, where $\alpha_i^t > 0$, $\beta_i^t > 0$ and $\gamma_{i,j}^t > 0$. We assume that the price charged by each firm affects its demand more than the prices charged by the other firms, in the sense that $\sum_{j \neq i} \gamma_{i,j}^t < \beta_i^t$ for all $i \in N$, $t \in T$. Also, using $\mathbf{p}_{-i}^t = (p_1^t, \dots, p_{i-1}^t, p_{i+1}^t, \dots, p_n^t)$ to denote the prices charged by firms other than

firm i at time period t , to avoid negative demand quantities, we restrict the strategy space of the firms such that each firm i charges the price p_i^t at time period t that satisfies $\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \geq 0$, given the prices \mathbf{p}_{-i}^t charged by the other firms. If the firms other than firm i commit to the price trajectories $\mathbf{p}_{-i} = \{\mathbf{p}_{-i}^t : t \in T\}$, then we can obtain the best response of firm i by solving the problem

$$\begin{aligned} \max \left\{ \sum_{t \in T} \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) p_i^t : \right. \\ \sum_{t \in T} \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) \leq c_i, \\ \left. \alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \geq 0 \quad \forall t \in T, \quad p_i^t \geq 0 \quad \forall t \in T \right\}. \end{aligned} \quad (3.1)$$

Since $\beta_i^t > 0$, problem (3.1) has a strictly concave objective function and linear constraints, which implies that the best response of firm i is unique.

Using the non-negative dual multipliers v_i and $\{u_i^t : t \in T\}$ for the first and second constraint in problem (3.1), the KKT conditions for this problem are

$$\begin{aligned} \left(\sum_{t \in T} \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) - c_i \right) v_i &= 0, \\ \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) u_i^t &= 0 \quad \forall t \in T, \\ \alpha_i^t - 2\beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t + \beta_i^t (v_i - u_i^t) &= 0 \quad \forall t \in T. \end{aligned} \quad (3.2)$$

Since problem (3.1) has a concave objective function and linear constraints, the KKT conditions above are necessary and sufficient at optimality; see [10]. In other

words, for a feasible solution $\{p_i^t : t \in T\}$ to problem (3.1), there exist corresponding non-negative dual multipliers v_i and $\{u_i^t : t \in T\}$ that satisfy the KKT conditions in (3.2) if and only if $\{p_i^t : t \in T\}$ is the optimal solution to problem (3.1). Note that we do not associate dual multipliers with the constraints $p_i^t \geq 0$ for all $t \in T$ in problem (3.1) since it is never optimal for firm i to charge a negative price. Therefore, we can actually view the constraints $p_i^t \geq 0$ for all $t \in T$ as redundant constraints. We use the KKT conditions in (3.2) extensively to characterize the best response of firm i to the price trajectories \mathbf{p}_{-i} of the other firms. In the rest of this section, we exclusively focus on the *strategies without recourse*, where each firm i commits to a price trajectory $\{p_i^t : t \in T\}$ at the beginning of the selling horizon and does not adjust these prices during the course of the selling horizon. If the price trajectory $\{p_i^t : t \in T\}$ chosen by each firm i is the best response to the price trajectories \mathbf{p}_{-i} chosen by the other firms, then we say that the price trajectories $\{\mathbf{p}^t : t \in T\}$ chosen by the firms is an *equilibrium without recourse*. We show that there exists a unique equilibrium without recourse. Furthermore, if we start with any price trajectory $\{\mathbf{p}^t : t \in T\}$ for the firms and successively compute the best response of each firm to the price trajectories of the other firms, then the best response of each firm forms a contraction mapping when viewed as a function of the prices charged by the other firms. Using this result, we show that there exists a unique equilibrium without recourse. To capture the best response of firm i to the prices charged by the other firms, we define for each $\nu \geq 0$, the set of time periods

$$\mathcal{T}_i(\nu, \mathbf{p}_{-i}) = \left\{ t \in T : \frac{\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t}{\beta_i^t} > \nu \right\}.$$

In the next lemma, we use $\mathcal{T}_i(\cdot, \mathbf{p}_{-i})$ to give a succinct characterization of the solution $\{p_i^t : t \in T\}$ and the corresponding dual multipliers v_i and $\{u_i^t : t \in T\}$ that satisfy the KKT conditions.

Lemma 8. *If a feasible solution $\{p_i^t : t \in T\}$ to problem (3.1) and the corresponding non-negative dual multipliers v_i and $\{u_i^t : t \in T\}$ satisfy the KKT conditions in (3.2), then we have*

$$p_i^t = \begin{cases} (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) / (2\beta_i^t) + v_i/2 & \text{if } t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i}) \\ (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) / \beta_i^t & \text{if } t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i}), \end{cases}$$

$$u_i^t = \begin{cases} 0 & \text{if } t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i}) \\ v_i - (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) / \beta_i^t & \text{if } t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i}). \end{cases}$$

Proofs of all lemmas are in Appendix B.1. By Lemma 8, we can characterize the solution $\{p_i^t : t \in T\}$ and the dual multipliers v_i and $\{u_i^t : t \in T\}$ that satisfy the KKT conditions in (3.2) only by using the value of v_i . If we know the value of v_i , then we can compute the set of time periods $\mathcal{T}_i(v_i, \mathbf{p}_{-i})$, in which case, we can choose the values of $\{p_i^t : t \in T\}$ and $\{u_i^t : t \in \mathcal{T}_i\}$ as given in Lemma 8. Throughout the rest of this section, we indeed choose the values of $\{p_i^t : t \in T\}$ and $\{u_i^t : t \in \mathcal{T}_i\}$ as given in Lemma 8, since we are interested in solutions that satisfy the KKT

conditions. Naturally, we do not know the value of v_i that allows us to obtain an optimal solution $\{p_i^t : t \in T\}$ to problem (3.1). In the next lemma, we give a characterization of the value of v_i that corresponds to the solution $\{p_i^t : t \in T\}$ and the dual multipliers v_i and $\{u_i^t : t \in T\}$ satisfying the KKT conditions in (3.2). In particular, we consider the function

$$G_i(\nu, \mathbf{p}_{-i}) = \begin{cases} \sum_{t \in \mathcal{T}_i(\nu, \mathbf{p}_{-i})} (\alpha_i^t - \beta_i^t \nu + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i & \text{if } \nu > 0 \\ \left[\sum_{t \in \mathcal{T}_i(\nu, \mathbf{p}_{-i})} (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i \right]^+ & \text{if } \nu = 0. \end{cases}$$

Lemma 24 in Appendix B.1 shows that $G_i(\cdot, \mathbf{p}_{-i})$ is strictly decreasing over some $[0, \nu^*]$ and has a unique root. In the next lemma, we use its root to characterize a solution to the KKT conditions.

Lemma 9. *If a feasible solution $\{p_i^t : t \in T\}$ to problem (3.1) and the corresponding non-negative dual multipliers v_i and $\{u_i^t : t \in T\}$ satisfy the KKT conditions in (3.2), then we have $G_i(v_i, \mathbf{p}_{-i}) = 0$.*

By Lemma 9, if a feasible solution $\{p_i^t : t \in T\}$ to problem (3.1) and the corresponding non-negative dual multipliers v_i and $\{u_i^t : t \in T\}$ satisfy the KKT conditions in (3.2), then v_i must be the unique root of $G_i(\cdot, \mathbf{p}_{-i})$. Also, by Lemma 8, the values of $\{p_i^t : t \in T\}$ and $\{u_i^t : t \in T\}$ must be given as in Lemma 8. In the next theorem, we use these results to show that the best response of firm i is a contraction mapping when viewed as a function of the prices of the other firms.

Theorem 9. *Let $\{p_i^t(\mathbf{p}_{-i}) : t \in T\}$ be the optimal solution to problem (3.1) as a*

function of the prices charged by the firms other than firm i . For any two price trajectories $\hat{\mathbf{p}}_{-i} = \{\hat{p}_{-i}^t : t \in T\}$ and $\tilde{\mathbf{p}}_{-i} = \{\tilde{p}_{-i}^t : t \in T\}$ adopted by the firms other than firm i , we have

$$|p_i^t(\hat{\mathbf{p}}_{-i}) - p_i^t(\tilde{\mathbf{p}}_{-i})| \leq \max_{t \in T} \left\{ \frac{\sum_{j \neq i} \gamma_{i,j}^t |\hat{p}_j^t - \tilde{p}_j^t|}{\beta_i^t} \right\}.$$

Proof Outline. Use M_i to denote the right hand side of the inequality in the theorem. Let $\hat{p}_i^t = p_i^t(\hat{\mathbf{p}}_{-i})$ and $\tilde{p}_i^t = p_i^t(\tilde{\mathbf{p}}_{-i})$. Let \hat{v}_i and \tilde{v}_i be such that $G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) = 0$ and $G_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i}) = 0$. Without loss of generality we assume $\hat{v}_i \geq \tilde{v}_i$. Otherwise, we interchange their roles. In the proof, we show that $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2} \max\{M_i, \hat{v}_i - \tilde{v}_i\}$, by considering four cases on whether t is in $\mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$ and $\mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$. In each case, we use Lemma 8 to get expressions for \hat{p}_i^t and \tilde{p}_i^t . Once we have $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2} \max\{M_i, \hat{v}_i - \tilde{v}_i\}$, we only need to show that $\hat{v}_i - \tilde{v}_i \leq M_i$. We use Lemma 9 and the definition of $\mathcal{T}_i(\cdot, \cdot)$ to show that $G_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \geq 0$. Then $\hat{v}_i - M_i \leq \tilde{v}_i$ follows from simple monotonicity properties of $G(\cdot, \tilde{\mathbf{p}}_{-i})$ given in Lemma 24. The details are in Appendix B.1. \square

For the vector $\mathbf{y} = \{y^t : t \in T\}$, define the norm on \mathfrak{R}^T as $\|\mathbf{y}\|_\infty = \max_{t \in T} |y^t|$. Since $\sum_{j \neq i} \gamma_{i,j}^t < \beta_i^t$ for all $i \in N$ and $t \in T$, Theorem 9 implies that firm i 's best response is a contraction under $\|\cdot\|_\infty$, when viewed as a function of the other firms' prices. Therefore, it immediately follows that if the price charged by each firm affects its demand more than the prices charged by the other firms, then there always exists a unique equilibrium without recourse.

The contraction mapping also presents an efficient computation scheme. Let $M = \max_{i \in N, t \in T} \sum_{j \neq i} \gamma_{i,j}^t / \beta_i^t$ and note that since $\sum_{j \neq i} \gamma_{i,j}^t < \beta_i^t$ for all $i \in N$, $t \in T$,

we have $M < 1$. Performing best-response iterations converges linearly to the unique equilibrium at rate M (see [59, Theorem 6.3.3]). In each iteration, one must solve the problems (3.1) for each $i \in N$. Using Lemma 24 in Appendix B.1, each of these n problems can be solved by bisection on v_i , as we can show that $v_i \geq 0$ must lie in a bounded interval. To see this, recall that the firms' prices must satisfy $\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \geq 0$ for all $i \in N$. Rearranging, we get $p_i^t \leq \alpha_i^t / \beta_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t / \beta_i^t \leq \max_{i \in N} \{\alpha_i^t / \beta_i^t\} + M \max_{j \in N} \{p_j^t\}$, which implies that $\max_{i \in N} \{p_i^t\} \leq \max_{i \in N} \{\alpha_i^t / \beta_i^t\} / (1 - M) = P_{\max}$. Then using the definition of $\mathcal{T}(\nu, \mathbf{p}_{-i})$, we obtain that $\mathcal{T}(\nu, \mathbf{p}_{-i})$ is empty if $\nu > P_{\max}$, which implies that $v_i \leq P_{\max}$.

3.2 Equilibrium with Recourse

In this section, we consider *strategies with recourse*, where each firm can change its price at each time period based on its inventory and the inventories of the other firms. In other words, the firms do not commit to a price trajectory at the beginning of the selling horizon. We let x_i^t be the inventory of firm i at the beginning of time period t . Focusing on Markovian strategies without loss of generality, as a function of the inventories $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$ of all of the firms, we use $P_i^t(\mathbf{x}^t)$ to denote the price charged by firm i at time period t . It is useful to view $P_i^t(\cdot)$ as a function that determines the strategy of firm i at time period t as a function of the inventories of all of the firms. We use $\mathbf{P}^t = (P_1^t(\cdot), \dots, P_n^t(\cdot))$ to capture the strategies of all of the

firms at time period t and $\mathbf{P}_{-i}^t = (P_1^t(\cdot), \dots, P_{i-1}^t(\cdot), P_{i+1}^t(\cdot), \dots, P_n^t(\cdot))$ to capture the strategies of the firms other than firm i at time period t . If the firms other than firm i use the strategies $\mathbf{P}_{-i} = \{\mathbf{P}_{-i}^t : t \in T\}$, then we can find the best response strategy of firm i by solving the dynamic program

$$\begin{aligned}
V_i^t(\mathbf{x}^t) = \max & \left\{ \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t P_j^t(\mathbf{x}^t) \right) p_i^t + V_i^{t+1}(\mathbf{x}^{t+1}) : \right. \\
& \alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t P_j^t(\mathbf{x}^t) \geq 0, \\
& x_i^{t+1} = \left[x_i^t - \left(\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t P_j^t(\mathbf{x}^t) \right) \right]^+, \\
& x_\ell^{t+1} = \left[x_\ell^t - \left(\alpha_\ell^t - \beta_\ell^t p_\ell^t + \gamma_{\ell,i}^t p_i^t + \sum_{j \notin \{i,\ell\}} \gamma_{\ell,j}^t P_j^t(\mathbf{x}^t) \right) \right]^+ \\
& \quad \quad \quad \forall \ell \in N \setminus \{i\}, \\
& \left. p_i^t \geq 0, x_\ell^{t+1} \geq 0 \forall \ell \in N \right\},
\end{aligned}$$

with the boundary condition that $V_i^{\tau+1}(\cdot) = 0$. An optimal solution to the problem above characterizes the best response strategy of firm i at time period t .

For the strategies $\{\mathbf{P}^t : t \in T\}$ to form an *equilibrium with recourse*, we require that for each $t \in T$, all inventories \mathbf{x}^t , and each i , the strategy $\{P_i^s(\cdot) : s = t, \dots, \tau\}$ chosen by firm i in the periods subsequent to time t is a best response against other firms' strategies $\{P_{-i}^s(\cdot) : s = t, \dots, \tau\}$ in the subsequent time periods. In other words, we require the strategies $\{\mathbf{P}^t : t \in T\}$ to form a Markov perfect equilibrium [20]. In the previous section, we show that there always exists a unique equilibrium when we focus on strategies without recourse. We give two numerical examples to show

that if we focus on strategies with recourse, then there may not exist an equilibrium or there may be multiple equilibria. Consider the case where there are two firms and the selling horizon has two time periods. For given inventories of the two firms at the second time period, the problem of computing the equilibrium strategy at the second time period is identical to finding an equilibrium without recourse. So, there exists a unique equilibrium strategy for the firms at the second time period for given inventories. Note that the prices charged by the firms in an equilibrium without recourse at the second time period depend on the inventories of the firms at the second time period, which, in turn, depend on the prices charged by the firms at the first time period. To obtain an equilibrium with recourse, we compute the best response strategy of each firm at the first time period as a function of the price of the other firm at the first time period. Recall that if we fix the prices of the firms at the first time period, then we fix the inventories at the second time period, in which case, we can compute the equilibrium strategies at the second time period. We plot the best response of each firm at the first time period as a function of the price of the other firm. An equilibrium with recourse corresponds to the intersection of the two best response curves.

Consider the parameters $\alpha_i^t = 4$, $\beta_i^1 = 4$, $\beta_i^2 = 2$, $\gamma_{i,j}^1 = 16/5$, $\gamma_{i,j}^2 = 1$, $c_i = 3$ for all $i \in \{1, 2\}$, $j \neq i$ and $t \in \{1, 2\}$, which satisfy $\sum_{j \neq i} \gamma_{i,t}^t < \beta_i^t$ for all $i, t \in \{1, 2\}$, so that we know that there exists a unique equilibrium without recourse. In Figure 3.1, the solid line plots the best response of second firm at the first time period on the vertical axis, as a function of the price of the first firm on the horizontal axis, whereas the dashed line plots the best response of the first firm at the first time period on the




Figure 3.1: Best responses when equilibrium does not exist.

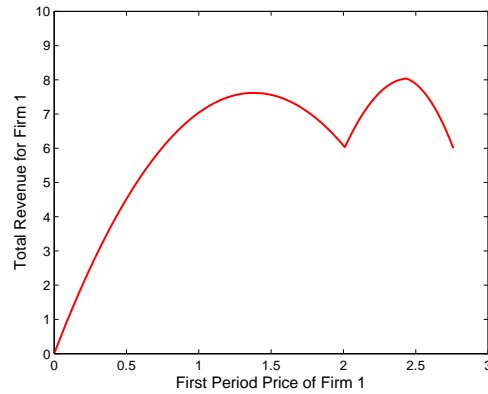


Figure 3.2: Revenue of the first firm as a function of its first period price.

horizontal axis as a function of the price of the second firm on the vertical axis. The two best response functions do not intersect. Therefore, an equilibrium with recourse does not exist. The main driver of the lack of equilibrium is the discontinuity in the best response function, which arises because the revenue of each firm is a multi-modal function of its first time period price. In Figure 3.2, we show first firm's revenue as a function of its price at the first time period, when the second firm's price is fixed at




Figure 3.3: Best responses when there are multiple equilibria.

2.2. So, firm 1 can jump from one mode to another based on the price of the second firm. Considering the parameters $\alpha_i^t = 4$, $\beta_i^1 = 5$, $\beta_i^2 = 2$, $\gamma_{i,j}^1 = 0.1$, $\gamma_{i,j}^2 = 1$ and $c_i = 5$ for all $j \neq i$ and $i, t \in \{1, 2\}$, Figure 3.3 shows the best response of each firm at the first time period as a function of the other firm's price. The best response functions intersect at two points, indicating multiple equilibria with recourse.

3.3 An Approximate Equilibrium

If for each firm i , any deviation from the strategy $\{P_i^t(\cdot) : t \in T\}$ cannot increase the revenue of firm i by more than ϵ given that the other firms use the strategies \mathbf{P}_{-i} , then we say that the price strategies $\{\mathbf{P}^t : t \in T\}$ chosen by the firms is an ϵ -equilibrium with recourse. Since there may not exist an equilibrium with recourse or there may be multiple equilibria with recourse, we focus on ϵ -equilibria with recourse. We consider a low influence regime, where, roughly speaking, the price

charged by a firm affects its demand more than the prices charged by each of the other firms. In particular, we consider the regime where the price charged by a firm affects its demand so much more than the prices charged by each of the other firms such that we have $\gamma_{i,j}^t/\beta_i^t < (1/M) - 1$, where M is as defined at the end of Section 2. When $\sum_{j \neq i} \gamma_{i,j}^t < \beta_i^t$ and the number of firms is large, we expect this assumption to hold. For example, if we have a symmetric setting, where the parameters related to each firm are the same, then under the assumption that $\sum_{j \neq i} \gamma_{i,j}^t < \beta_i^t$, we have $\gamma_{i,j}^t/\beta_i^t < 1/(n-1)$, in which case, the low influence regime naturally holds as the number of firms gets large. In the low influence regime, we show that the equilibrium without recourse studied in the previous section can be used to construct an ϵ -equilibrium with recourse. Intuitively, this result uses the fact that if $\gamma_{i,j}^t/\beta_i^t$ is small, then any deviation of a firm from a given price trajectory has little influence on the prices of the other firms in the subsequent time periods. In the next lemma, we formalize this idea. Throughout the rest of this section, we use $\mu = \max_{i \in N, j \in N \setminus \{i\}, t \in T} \gamma_{i,j}^t/\beta_i^t$ and $\bar{\beta} = \max_{i \in N, t \in T} \beta_i^t / \min_{i \in N, t \in T} \beta_i^t$. Note that the low influence regime is defined as the setting where $\mu < (1/M) - 1$.

Lemma 10. *Fixing the prices $\hat{\mathbf{p}}^1$ charged by the firms at the first time period, let the prices $\{\hat{\mathbf{p}}^t : t \in T \setminus \{1\}\}$ form the equilibrium without recourse in the remaining portion of the selling horizon. Define the prices $\tilde{\mathbf{p}}^1$ at the first time period as $\tilde{p}_i^1 = \hat{p}_i^1 + \delta$ and $\tilde{p}_j^1 = \hat{p}_j^1$ for all $j \in N \setminus \{i\}$ for some $\delta \geq 0$. Fixing the prices $\tilde{\mathbf{p}}^1$ charged by the firms at the first time period, let the prices $\{\tilde{\mathbf{p}}^t : t \in T \setminus \{1\}\}$ form the equilibrium without recourse in the remaining portion of the selling horizon. If we have $\mu < (1/M) - 1$, then $\max_{j \neq i, t \in T \setminus \{1\}} |\hat{p}_j^t - \tilde{p}_j^t| \leq \frac{2\mu\bar{\beta}\delta}{1-M-M\mu}$.*

Consider the problem over the time periods κ, \dots, τ when the inventories of the firms at time period κ are given by $\mathbf{x} = (x_1, \dots, x_n)$. We use $p_i^{N,t}(\kappa, \mathbf{x})$ to denote the price charged by firm i at time period t in the equilibrium without recourse. We consider the following strategy with recourse for firm i . If the inventories of the firms at time period t is given by \mathbf{x} , then firm i charges the price $p_i^{N,t}(t, \mathbf{x})$. In other words, letting $P_i^{R,t}(\cdot)$ be the strategy function of firm i under this strategy with recourse, we have $P_i^{R,t}(\mathbf{x}) = p_i^{N,t}(t, \mathbf{x})$. Using $\mathbf{P}^{R,t} = (P_1^{R,t}(\cdot), \dots, P_n^{R,t}(\cdot))$ to capture the strategies of all of the firms at time period t and $\mathbf{c} = (c_1, \dots, c_n)$ to denote the inventories of the firms at the first time period, note that if all firms use the strategies $\{\mathbf{P}^{R,t} : t \in T\}$ over the selling horizon, then the price charged by each firm i at each time period t is given by $p_i^{N,t}(1, \mathbf{c})$, which is precisely the prices corresponding to the equilibrium without recourse when we consider the problem over the time periods T with the inventories of the firms at the first time period given by \mathbf{c} . However, if one of the firms deviates from the strategies $\{\mathbf{P}^{R,t} : t \in T\}$ at a time period, then the prices charged by the firms will be different from those in the equilibrium without recourse. Therefore, it is not generally true that the strategies $\{\mathbf{P}^{R,t} : t \in T\}$ correspond to an equilibrium with recourse. In the remainder of this section, we show that the strategies $\{\mathbf{P}^{R,t} : t \in T\}$ correspond to an ϵ -equilibrium with recourse in the low influence regime. In the next lemma, we show that if firm i unilaterally deviates from the strategy $\{P_i^{R,t}(\cdot) : t \in T\}$, but the other firms use the strategies $\{\mathbf{P}^{R,t} : t \in T\}$, then firm i does not increase its revenue by more than a simple function of μ .

Lemma 11. *Assume that the strategies of all of the firms are $\{\mathbf{P}^{R,t} : t \in T\}$. Let Π_i^N be the revenue of firm i under these strategies. Also, assume that the strategies*

of the firms other than firm i are $\{\mathbf{P}_{-i}^{R,t} : t \in T\}$, but firm i deviates to charge an arbitrary price at the first time period and uses the strategy $\{P_i^{R,t}(\cdot) : t \in T \setminus \{1\}\}$ at the other time periods. Let Π_i^D be the revenue of firm i under this strategy. Letting $\beta_{\max} = \max_{i \in N, t \in T} \beta_i^t$, we have for $\mu < (1/M) - 1$,

$$\Pi_i^D - \Pi_i^N \leq \frac{2\bar{\beta} M \beta_{\max} P_{\max}^2 (\tau - 1) \mu}{1 - M - M\mu}.$$

In the next theorem, we show that the strategy $\{\mathbf{P}^{R,t} : t \in T\}$ is an ϵ -equilibrium with recourse, when the number of firms is large so that μ is small.

Theorem 10. *Assume that the strategies of all of the firms are $\{\mathbf{P}^{R,t} : t \in T\}$. Let Π_i^N be the revenue of firm i under these strategies. Also, assume that the strategy of the firms other than firm i are $\{\mathbf{P}_{-i}^{R,t} : t \in T\}$, but firm i uses an arbitrary strategy over the whole selling horizon. Let Π_i^A be the revenue of firm i under these strategies. Letting $\Gamma_\mu = \bar{\beta} M \beta_{\max} P_{\max}^2 / (1 - M - M\mu)$, we have, for $\mu < (1/M) - 1$,*

$$\Pi_i^A - \Pi_i^N \leq \Gamma_\mu \tau (\tau - 1) \mu.$$

Thus, $\{\mathbf{P}^{R,t} : t \in T\}$ is an ϵ -equilibrium with recourse, with $\epsilon = \Gamma_\mu \tau (\tau - 1) \mu$.

Proof Outline. We use induction to prove the result. The result trivially holds for $\tau = 1$, as there is no difference between equilibrium with and without recourse for $\tau = 1$. Assume the result is true for $\tau = k$. Let all firms other than i use the strategy $\{\mathbf{P}_{-i}^{R,t} : t \in T\}$. We use $\{Q_i^t : t \in T\}$ to denote the arbitrary strategy of firm i . Let Π_i^N and Π_i^A be firm i 's revenue when it uses strategy $\{P_i^{R,t} : t \in T\}$ and $\{Q_i^t : t \in T\}$, respectively. Let Π_i^D be firm i 's revenue when it uses strategy $\{P_i^{R,1}, Q_i^t : t \in T \setminus \{1\}\}$.

We use Lemma 11 to bound the difference between Π_i^D and Π_i^A . Similarly, we use the induction hypothesis at $\tau = k$ to bound the difference between Π_i^D and Π_i^N . Summing up the two bounds gives us the result for $\tau = k + 1$. The details are in Appendix B.1. \square

We observe that as μ approaches zero, $\Gamma_\mu \tau (\tau - 1) \mu$ approaches zero as well. Therefore, by the theorem above, if we are in the low influence regime, then no firm can improve its revenue significantly by deviating from the policy $\{\mathbf{P}^{R,t} : t \in T\}$, which implies that $\{\mathbf{P}^{R,t} : t \in T\}$ is an ϵ -equilibrium with recourse. As discussed earlier, the price trajectory realized under the strategy $\{\mathbf{P}^{R,t} : t \in T\}$ is same as that in the unique equilibrium without recourse.

3.4 Future Research

A natural research direction is to extend our contraction properties to more general demand models. Also, it would be useful to define an analogue of equilibrium without recourse under stochastic demand and check whether it uniquely exists.

CHAPTER 4
**PERSONALIZED ASSORTMENT RECOMMENDATION AFTER
PURCHASE**

4.1 Introduction

Many online sellers have the opportunity to recommend an assortment of items to a customer after she makes a purchase or engages in an activity. The recommended assortment is often personalized to the customer in one of the following two ways. One type of personalization is based on the customer’s entire profile, which may include her purchase history, demographic or other interaction-specific information such as device type and browser, etc. For example, Walmart.com recommends a list of sponsored products on the first page after a customer logs-in, based on her purchase history. Another type of personalization is only based on the item that is purchased during the current visit. For example, when a customer makes a hotel reservation on Priceline.com, she is given rental car deals in the confirmation email she receives. Similarly, after a customer makes a purchase on Amazon.com, the confirmation email she receives shows two products as “recommended based on the item purchased”. While the second type of personalization can be seen as a special case of the first type, we consider it as a separate type because of the difference in analysis. To be more precise, the product-based nature of the second type of personalization allows us to decompose the analysis by product, while the analysis of the first type of personalization requires decomposition by customer segments.

Both types of personalization are widely used in the web-based industry, and each one has its own advantages. Personalization based on entire profile accounts for more information, thus tend to be more accurate. On the other hand, personalization based on the item purchased is simpler and less costly for the seller because it requires and analyzes much less information. So it is good for sellers who do not have the ability to collect or process complicated customer data. It also tends to be more applicable in industries where customers' interests change fast over time. For example, in the travel industry, a customer is more likely to be interested in a rental car deal in the city which she just booked a flight to, than in a city she has visited in the past.

In this chapter, we focus on the second type of personalization. In Section 4.9, we discuss how our work can be extended to incorporate the first type. We consider a firm selling finite inventories of multiple products over a finite selling horizon. One customer arrives at each period in the selling horizon, and each period has two stages. In stage 1, a customer arrives and chooses to purchase one item, or to leave without a purchase, based on the assortment offered to her. If she makes a purchase in stage 1, she enters stage 2 in which she is offered another assortment, from which she may again choose to purchase an item, or to leave without a purchase. Customers' behavior is modeled using general choice models. The firm chooses an assortment to offer at each stage of each period, based on the remaining inventory and time. The choice of stage 2 assortment is customized to the customer's purchase decision in stage 1 of the same period. The firm's goal is to maximize its expected total revenue.

Although our focus is on the online setting where customer behavior is stochastic,

we begin by analyzing the offline version of the problem. This offline version is obtained via a “fluid approximation”, in which customer demands are deterministic and they can take fractional values. We formulate the offline problem as an LP, using the probability of offering each assortment at each stage as a decision variable. This LP is exponential in the number of products, but it can be solved efficiently using the column generation method (see [16]). We show that the optimal LP revenue is an upper bound on the optimal revenue in the original online problem. We also present an example, under which the expected online revenue under optimal policy is 47% of the optimal offline policy. This suggests that when we do analysis by comparison with the offline LP, the constant performance guarantee cannot exceed 47% for any algorithm.

Using the solution to the offline LP, we propose a novel algorithm, namely the *balancing algorithm*, for the original online problem. The balancing algorithm solves the offline LP only once at the beginning of the selling horizon, and updates the variables in each period to keep them feasible for the offline LP with updated time and inventory information. In each period, the algorithm offers a random assortment suggested by the current feasible solution to the LP, where the support of the random assortment grows as $O(n^2)$ when there are n products.

We show that the balancing algorithm attains an expected revenue of at least $1/3$ of the optimal LP revenue. In a limiting regime in which the number of time periods and initial inventory level increase proportionally and both tend to infinity, we prove that the balancing algorithm is asymptotically optimal. In particular, the revenue

under the balancing algorithm converges in probability to the LP revenue.

The balancing algorithm solves the offline LP only once at the beginning of the selling horizon, thus may not be able to adjust fully to stochastic realization of sales. As a remedy, we also propose a *resolving algorithm* as a natural extension of the balancing algorithm. The resolving algorithm mimics the balancing algorithm, except for that it occasionally resolves the offline LP using updated inventory and time information. We prove that the performance guarantees we have for the balancing algorithm also hold for the resolving algorithm. Namely, it attains an expected revenue of at least $1/3$ of the optimal offline revenue, and when the number of resolves is small, the revenue under the resolving algorithm converges in probability to the optimal offline revenue.

Finally, we numerically test the performance of the balancing and the resolving algorithms under different settings. On average, the balancing algorithm attains 84.9% of the optimal offline revenue. The resolving algorithm with different number of resolves attains from 86.5% to 89.5% of the optimal offline revenue. Both algorithms perform better when the inventory is abundant, when the probability of leaving without a purchase is low, or when the number of time periods in the selling horizon is large. We also observe that resolving once in the middle of the selling horizon may lead to significant improvement in revenue, compared with the balancing algorithm.

4.1.1 Literature Review

The problem studied in this chapter is an online assortment customization problem with inventory constraints, and with recommendation opportunity at checkout.

Some recent papers consider assortment customization. Those papers assume that customers form different segments, and that the seller can observe each customer's segment. Thus, the seller's assortment decision does not only depend on the remaining inventory and time, but is also personalized based on customer characteristics. [31] proposes a family of index-based policies with competitive ratio $1 - 1/e$, which it proves is the best constant ratio. [9] considers assortment customization with identically priced, substitutable products. It proves that it is optimal to follow a threshold policy in some settings. [24] considers an assortment customization problem when each product is composed from a bundle of resources with limited inventories. It proposes an algorithm that achieves a performance guarantee $1/2$, which is the best possible constant ratio.

Similar to this chapter, [13] studies an online assortment customization problem, with opportunity for recommendation at checkout. It minimizes the competitive ratio, assuming an adversary chooses the total number of customers, their arrival sequence, and their type. The algorithm proposed in that paper attains a competitive ratio of $1/4$. On the other hand, we maximize the firm's expected revenue, assuming customer behavior follows a stochastic model. We propose an algorithm with performance guarantee of $1/3$.

We formulate a linear program for the offline problem in which demand is deterministic and can be fractional, and use its solution to develop an algorithm for the online stochastic problem in which customers' purchase behavior is random. In assortment optimization literature, such approach was first used by [23]. It considers settings where there are multiple alternative products serving the same market, and a customer who purchases a "flexible product" in a market is assigned in later dates to one of the products in that market. The paper formulates a choice-based linear program (CDLP) as a deterministic approximation of the original stochastic problem with random demand. The CDLP has exponentially many variables, but it can be solved efficiently using column generation. [42] extends the CDLP formulation to more general settings, and proposes a dynamic control policy based on the CDLP solution. [11] extends the CDLP method to the case when customers belong to overlapping segments. It shows that the associated column generation subproblem is NP-hard, and propose a heuristic to overcome the complexity.

The linear program we formulate for the offline problem has exponentially many variables. But the column generation subproblem can be solved by solving linearly many one-shot static assortment optimization problems. The static assortment optimization problem assumes that there is only one period in the selling horizon, and it chooses an assortment to offer in order to maximize the firm's expected revenue. It is studied in the literature under various settings. [37] provides a literature review in this area. [62] analyzes this problem when customer choice follows the multinomial logit (MNL) model. It assumes that all products have the same price but different attributes, and that there is an inventory cost for each product offered. It shows

that the optimal assortment has a simple structure. [65] shows that under the MNL model and the independent demand model, the optimal policy is to offer products with the largest prices. The optimization problem can thus be solved in linear time by comparing n nested assortments containing the most expensive products when there are n products. [61] studies the problem under the mixed multinomial logit model, under which there are multiple classes of customers, each following a MNL demand model with class-dependent parameters. It shows that the problem is NP-complete even when there are only two customer classes, and proposes a polynomial time approximation scheme. For the same problem, [48] proposes a branch-and-cut algorithm which leads to the nearly-optimal solution very fast. [15] studies the problem under the nested logit model. It shows that the problem is polynomially solvable if the nest dissimilarity parameters are small and the customers always make a purchase. It then proposes an algorithm to deal with the NP-hard cases with worst case performance guarantee. [25] and [18] propose algorithms to solve the problem under nested logit model with various constraints on the offered assortments. [40] studies the problem under a d -level nested logit model in which each product is described by a list of d features, and proposes a polynomial algorithm to solve for the optimal assortment.

4.2 Model

Consider a firm selling n items through a selling horizon with T periods. Each item i has limited inventory denoted by c_i , and we assume c_i is a non-negative integer for all i . The price of item i is p_i , and any inventory left after the selling horizon has zero value. Although we assume each item can be sold at a single exogenous price in this chapter, our analysis can be extended to the case in which the firm can choose from different candidate prices for each item.

Each period is split into two stages. In stage 1, a single customer arrives, and either chooses to purchase one item from the assortment offered to her, or decides to leave without a purchase. If she makes a purchase, she enters stage 2 of this period, in which she is offered another assortment, and she may choose to purchase an item from that assortment, or to leave without a purchase. There is no stage 2 for this period, if the customer does not make purchase in stage 1. Use N and M to denote the set of candidate items for stage 1 and 2, respectively. In this chapter, we assume that N and M are disjoint. This is true when the items recommended in second stage are of different type as the first stage purchase (for example, consider the case when Priceline recommends rental cars in the second stage, to customers who reserve hotels in the first stage). In Section 4.9, we briefly discuss how our results can be extended to the case where N and M may not be disjoint.

We assume that the customers are homogeneous upon arrival, and we model their behavior through choice models. We use $\Pi_i(S)$ to denote the probability of

purchasing item i in stage 1, when assortment $S \subset N$ is offered. We use the item purchased in stage 1 to define the type of the customer, and use $\Phi_j^i(S)$ to denote the probability of a type i customer purchasing item j when she sees assortment $S \subset M$ in stage 2. Note that although we assume each customer only purchases at most one item in each stage, the analysis in this chapter can be extended to allow purchasing of multiple items in each stage. The type of a customer is then defined as the set of items purchased in stage 1.

The only assumption we make on the choice probabilities is as follows. Intuitively, it requires that when one of the items is removed from an assortment, the purchase probability of any other item does not decrease.

Assumption 3 (Substitution Assumption). *For any assortment $T \subset S$, we have*

$$\begin{aligned} \Pi_i(T) &\geq \Pi_i(S) && \forall i \in N \\ \Phi_i^j(T) &\geq \Phi_i^j(S) && \forall i \in N, j \in M. \end{aligned}$$

This is a standard assumption that is satisfied by many widely studied choice models, including multinomial logit model, nested logit mode, Markov chain choice model, and ranking-based choice model. This assumption is true when all items are substitutes. Thus we call it the “substitution assumption”.

The firm maximizes its expected revenue by dynamically offering assortments to customers. More specifically, let c^t be the vector of remaining inventory where there are t time periods left. Note that throughout this chapter, we count time t backwards. We use $S_1^t(c^t)$ to denote the assortment offered in stage 1 of period t ,

when there are c^t inventory left. We use $S_2^t(c^t, i)$ to denote the assortment offered in stage 2 of period t , when item i is sold in stage 1 and the remaining inventory at the beginning of period t is c^t . The firm's strategy at period t can be represented by $S^t = \{S_1^t(\cdot), S_2^t(\cdot, \cdot)\}$. The firm's optimal strategy can be found by solving the following dynamic program:

$$\begin{aligned}
V^t(c^t) = \max_{S_1^t, S_2^t(\cdot)} \sum_{i \in S_1^t} \Pi_i(S_1^t) & \left[p_i + \sum_{j \in S_2^t(i)} \Phi_i^j(S_2^t(i)) (p_j + V^{t-1}(c^t - e_i - e_j)) \right. \\
& \left. + (1 - \sum_{j \in S_2^t(i)} \Phi_i^j) V^{t-1}(c^t - e_i) \right] + (1 - \sum_{i \in S_1^t} \Pi_i(S_1^t)) V^{t-1}(c^t) \quad (4.1)
\end{aligned}$$

subject to $c_i^t \geq 1, \forall i \in S_1^t$,

$$c_j^t \geq 1, \forall j \in \cup_{i \in S_1^t} S_2^t(i),$$

with the boundary condition that $V^0(\cdot) = 0$. Then the firm's expected revenue under optimal policy is $V^T(c)$.

4.3 LP formulation for the offline problem

In this section, we consider a deterministic offline relaxation of the original problem (4.1). The solutions to the offline problem is used later in Section 4.4, in which we propose an algorithm for the original online problem. The offline problem is obtained through a ‘‘fluid approximation’’ where we think of Π and Φ to be deterministic fractional demand. Namely, we let $\Pi_i(S)$ be the demand of item i in stage 1, when assortment S is offered, and let $\Phi_j^i(S)$ be the demand of item j from a type i customer when she sees assortment S in stage 2. The firm solves an offline optimization problem

to maximize its total revenue, subject to inventory constraints. We formulate the firm's problem using linear programming, and show that it can be solved efficiently using column generation.

With its deterministic nature, the offline problem can be reduced to a “single period” problem with initial inventory $\rho \triangleq c/T$. Let α_S be the probability to offer assortment $S \subset N$ in stage 1, and let β_S^i be the probability of seeing a type i customer and offering assortment $S \subset N$ in stage 2. Consider the following linear program:

$$L(\rho) = \max_{\alpha, \beta} \sum_{i \in N} p_i \left[\sum_{S \subset N} \alpha_S \Pi_i(S) \right] + \sum_{j \in M} p_j \left[\sum_{i \in N, S \subset M} \beta_S^i \Phi_j^i(S) \right] \quad (4.2)$$

$$\text{s.t.}, \quad \sum_{S \subset N} \alpha_S \Pi_i(S) \leq \rho_i, \forall i \in N \quad [u_i]$$

$$\sum_{i \in N, S \subset M} \beta_S^i \Phi_j^i(S) \leq \rho_j, \forall j \in M \quad [u_j] \quad (4.3)$$

$$\sum_{S \subset N} \alpha_S = 1, \quad [v] \quad (4.4)$$

$$\sum_{S \subset M} \beta_S^i = \sum_{S \subset N} \alpha_S \Pi_i(S), \forall i \in N \quad [z_i] \quad (4.5)$$

$$\alpha \geq 0, \quad \beta \geq 0.$$

The objective function is the firm's offline revenue in a single period. Thus $L(\rho)T$ is the firm's total offline revenue. Constraint (4.3) is the inventory constraint. Constraint (4.4) requires that the probability of offering one of the assortments in stage 1 sums up to 1. Constraint (4.5) makes sure the probability that the firm shows an assortment to type i customers in stage 2 is equal to the probability that a type i customer arrives to the system. We associate dual variables u, v, z with these constraints.

In the rest of this chapter, we slightly abuse notation and let $L(\rho)$ to denote the LP objective at optimality, as well as the linear program itself with input ρ .

Lemma 12. *The firm's optimal offline revenue $L(\rho)T$ is an upper bound on the firm's expected revenue under optimal strategy, in the original online problem (4.1).*

Proof. This lemma follows immediately from the fact that the optimal strategy can be converted into a feasible solution to this linear program. More specifically, given an optimal strategy to the original online problem (4.1), let $\alpha_S T$ (or $\beta_S^i T$) be the expected number of periods offering assortment S in stage 1 (or stage 2 to type i customers), under that strategy. One can verify that α and β satisfy all constraints in the linear program. \square

The linear program (4.2) has exponentially many random variables. We now present the column generation formulation to this LP. An introduction to the column generation method can be found in [16]. Let $RC(\cdot)$ to denote the reduced cost for each variable. The column generation subproblem seeks a variable with the maximum reduced cost, which can be found by solving $\max_S RC(\alpha_S)$, and $\max_S RC(\beta_S^i)$ for each i . Note that

$$\begin{aligned} \arg \max_S RC(\alpha_S) &= \arg \max_S \sum_{i \in N} \Pi_i(S) [p_i T - u_i + z_i] \\ \arg \max_S RC(\beta_S^i) &= \arg \max_S \sum_{j \in M} \Phi_j^i(S) [p_j T - u_j] \end{aligned}$$

which implies that $\max_S RC(\alpha_S)$ and $\max_S RC(\beta_S^i)$ are equivalent to one-shot assortment problems. Thus, the column generation subproblem can be solved by

solving $n + 1$ one-shot assortment problems. There are many choice models under which the one-shot assortment problem can be solved efficiently. See the literature review section for details.

4.4 The balancing algorithm

In this section, we present an algorithm for the online problem (4.1), which uses the LP solutions obtained in Section 4.3, and guarantees an expected revenue of at least $1/3$ of the upper bound given in Lemma 12.

4.4.1 The algorithm

Let $L(\rho^t)$ be the linear program associated with the offline problem when there are t periods left, and remaining inventory per period is $\rho^t \triangleq c^t/t$. Let α^t, β^t be a feasible solution to $L(\rho^t)$. A randomized policy for period t suggested by this feasible solution is given by the following.

One period randomized policy given α^t, β^t :

In Stage 1, offer assortment S with probability α_S^t .

In stage 2, offer assortment S with probability $\frac{\beta_S^{i,t}}{\sum_{S'} \beta_{S'}^{i,t}}$.

Our algorithm applies the one period randomized policy at each time period, given a feasible solution to the LP at the corresponding period. Namely, the algorithm

solves the offline LP once at the beginning of the selling horizon, and updates α and β in each period so that they are always feasible for the offline LP with updated inventory and time left. Then it uses the one period randomized policy given α^t and β^t , for any time period t . Thus, the key step in the algorithm is to update α and β to keep them feasible for the offline LP. The following algorithm constructs α^{t-1} and β^{t-1} feasible for $L(\rho^{t-1})$, using α^t and β^t feasible for $L(\rho^t)$.

Input: $\alpha^t, \beta^t, \rho^{t-1}$, items sold in period t ;

$\alpha^{t-1} \leftarrow \alpha^t$;

$\beta^{t-1} \leftarrow \beta^t$;

if *item i sold in Stage 1 of period t* **then**

if $\sum_S \alpha_S^{t-1} \Pi_i(S) > \rho_i^{t-1}$ **then**

 Balance item i according to Algorithm 3 so that $\sum_S \alpha_S^{t-1} \Pi_i(S) = \rho_i^{t-1}$;

end

if *item j sold in Stage 2 of period t* **then**

if $\sum_{S,k \in N} \beta_S^{k,t-1} \Phi_j^k(S) > \rho_j^{t-1}$ **then**

 Balance item j according to Algorithm 4 so that

$\sum_{S,k \in N} \beta_S^{k,t-1} \Phi_j^k(S) = \rho_j^{t-1}$;

end

end

end

Output: $\alpha^{t-1}, \beta^{t-1}$;

Algorithm 1: Construct a feasible solution to $L(\rho^{t-1})$

Note that $\sum_S \alpha_S \Pi_i(S)$ is the probability of selling item $i \in N$ when the probability

of offering assortment S is α_S . We call this quantity the “selling probability” of item $i \in N$ under α . The selling probability of item $j \in M$ is similarly defined as $\sum_{S,k} \beta_S^k \Phi_j^k(S)$. In Algorithm 1, an item is balanced when its selling probability exceeds the remaining inventory per period, so that the inventory constraint in the LP (4.2) is violated. In those cases, the balancing step updates α and β to re-satisfy the violated inventory constraints, while keeping the sales probability of all other same-period items the same. The balancing step is the most important piece in this algorithm, and we discuss it in the next subsection.

4.4.2 The balancing step

In this subsection, we present the balancing step. It is used in Algorithm 1 when there exists an item, whose sales probability under the current values of α and β exceeds the remaining inventory per period. In order to keep α and β feasible for the offline LP (4.2), we use the balancing step to decrease the sales probability of this item, while keeping the sales probability of all other same-period items the same. Mathematically, the goal of balancing item $i \in N$ is to construct $\alpha^{t-1}, \beta^{t-1}$ from α^t, β^t and α^T, β^T , such that

$$\begin{aligned} \sum_S \alpha_S^{t-1} \Pi_i(S) &= \rho_i^{t-1} < \sum_S \alpha_S^t \Pi_i(S), \\ \sum_S \alpha_S^{t-1} \Pi_k(S) &= \sum_S \alpha_S^t \Pi_k(S), \forall k \neq i \in N. \end{aligned}$$

The objective for balancing a stage 2 item $j \in M$ can be similarly defined. Due to the structure of the choice probabilities, removing an item from an assortment

may increase the purchasing probability of some other items. Thus, decreasing the sales probability of one item while keeping that of all other items the same is not a trivial task. The following Algorithm presents a method to do this when the purchase probability Π satisfy the substitution assumption. More precisely, given any α , under which the sales probability is x , and any $0 \leq \hat{x} \leq x$, Algorithm 2 constructs $\hat{\alpha}$ under which the sales probability is \hat{x} .

function IncreaseAlphaSingleS($\hat{y}, y, S, \beta, \hat{\alpha}$):

if $|S| = 0$ **then**

| $\hat{\alpha}_\emptyset \leftarrow \hat{\alpha}_\emptyset + \beta;$

else

| Let k and m be minimizer and minimum value of \hat{y}_i/y_i for $i \in S$;

| $\hat{\alpha}_S \leftarrow \hat{\alpha}_S + \beta m;$

| $\hat{\alpha} = \text{IncreaseAlphaSingleS}(\hat{y} - \beta m, \beta(1 - m)\Pi(S/\{k\}), S/\{k\}, \beta(1 - m), \hat{\alpha});$

end

return $\hat{\alpha}$

Input: $\alpha, x, \hat{x}, \Pi, N$;

$\hat{\alpha}_S \leftarrow 0, \forall S \subset N$;

$\gamma_i \leftarrow \frac{\hat{x}_i}{x_i}, \forall i \in N$;

for S such that $\alpha_S > 0$ **do**

| $\hat{\alpha} = \text{IncreaseAlphaSingleS}(\alpha_S \Pi(S) \gamma, \alpha_S \Pi(S), S, \alpha_S, \hat{\alpha});$

end

Output: $\hat{\alpha}$;

Algorithm 2: Decreasing sales probability

Lemma 13. *Given a set of items N with choice probabilities Π that satisfy the substitution assumption, and a vector $\alpha \geq 0$, let $x_i = \sum_S \alpha_S \Pi_i(S)$ for all $i \in N$. Then for any vector $\hat{x} \geq 0$ such that $\hat{x}_i \leq x_i$ for all $i \in N$, there exists a vector $\hat{\alpha}$ given by the output of Algorithm 2, such that*

$$\begin{aligned} \sum_{S \subset N} \hat{\alpha}_S &= \sum_{S \subset N} \alpha_S, \\ \sum_{S \subset N} \hat{\alpha}_S \Pi_i(S) &= \hat{x}_i && \forall i \in N \\ |\{S : \hat{\alpha}_S > 0\}| &\leq \sum_{S: \alpha_S > 0} (|S| + 1) \\ \hat{\alpha}_S &\geq 0 && \forall S \subset N. \end{aligned}$$

Proof. To show the first equation, observe that for any S such that $\alpha_S > 0$, the operations within the *for* loop increases $\sum_{S'} \hat{\alpha}_{S'}$ by α_S . Sum over all S , we conclude that the total increase in $\sum_{S'} \hat{\alpha}_{S'}$ is $\sum_S \alpha_S$. Thus the first equation holds.

The third inequality follows from the observation that each execution of the *IncreaseAlphaSingleS()* function increases $|\{S : \hat{\alpha}_S > 0\}|$ by at most 1, and that the *IncreaseAlphaSingleS()* function is executed for $|S| + 1$ times for each S .

The last inequality holds because the initial value of $\hat{\alpha}$ is zero, and it never decreases in the course of the algorithm.

We now show the second equation. To show this, we first show that starting with any y, \hat{y}, β, S such that $0 \leq \hat{y} \leq y$ and $y_i = \beta \Pi_i(S)$, executing *IncreaseAlphaSingleS*($\hat{y}, y, S, \beta, \hat{\alpha}$) increases $\sum_{S'} \hat{\alpha}_{S'} \Pi_i(S')$ by \hat{y}_i , for any $i \in S$. Without loss of generality, we assume the initial value of $\hat{\alpha}$ is zero. Then we only need to

show that after executing the function, $\sum_{S'} \hat{\alpha}_{S'} \Pi_i(S') = \hat{y}_i$ for any $i \in S$.

We prove this statement by induction on $|S|$.

When $|S| = 1$, without loss of generality assume $S = \{1\}$. The the function is executed only once, and we have

$$\begin{aligned}\hat{\alpha}_S &= \beta \frac{\hat{y}_1}{y_1} \\ \hat{\alpha}_\emptyset &= \beta - \hat{\alpha}_S.\end{aligned}$$

Then the statement follows because

$$\sum_{S'} \hat{\alpha}_{S'} \Pi_1(S') = \hat{\alpha}_S \Pi_1(S) = \beta \frac{\hat{y}_1}{y_1} \Pi_1(S) = \hat{y}_1.$$

Suppose when $|S| = k$, the statement is true. Consider the case when $|S| = k + 1$. Without loss of generality assume $S = \{1, \dots, k + 1\}$ and that $k + 1 = \arg \min_j \frac{\hat{y}_j}{y_j}$. We let $y^{(k)}, \hat{y}^{(k)}, \beta^{(k)}, S^{(k)}$ be the variables when calling the function for the $k + 1$ th time. Before calling the function for the second time, we have

$$\begin{aligned}\hat{\alpha}_S &= \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \\ \beta^{(1)} &= \beta - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \\ \hat{y}_i^{(1)} &= \hat{y}_i - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S) \\ y_i^{(1)} &= \beta^{(1)} \Pi_i(S / \{k + 1\}).\end{aligned}$$

Thus the increase in $\sum_S \hat{\alpha}_S \Pi_i(S)$ before calling the function for the second time is $\beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S)$. We now calculate the increase in $\sum_S \hat{\alpha}_S \Pi_i(S)$ from the second call of the function till its termination, using the induction hypothesis on $y^{(1)}, \hat{y}^{(1)}, \beta^{(1)}$ and

$S/\{k+1\}$.

We have $\hat{y}_i^{(1)} \geq 0$ because

$$\hat{y}_i^{(1)} = \hat{y}_i - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S) = \hat{y}_i - \frac{\hat{y}_{k+1}}{y_{k+1}} y_i \geq 0.$$

We have $y_i^{(1)} \geq \hat{y}_i^{(1)}$ for all i because

$$\hat{y}_i^{(1)} = \hat{y}_i - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S) \leq \beta \Pi_i(S) - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S) \leq \left[\beta \Pi_i(S) - \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \right] \Pi_i(S/\{k+1\}) = y_i^{(1)}.$$

Thus using the induction hypothesis on $y^{(1)}, \hat{y}^{(1)}, \beta^{(1)}$ and $S/\{k+1\}$, the increase in $\sum_S \hat{\alpha}_S \Pi_i(S)$ from the second call of the function till termination is \hat{y}_i^1 . Recall that the increase in $\sum_S \hat{\alpha}_S \Pi_i(S)$ in the first execution is $\beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S)$. The statement follows because $\hat{y}_i^{(1)} + \beta \frac{\hat{y}_{k+1}}{y_{k+1}} \Pi_i(S) = \hat{y}_i$.

Thus, for any S such that $\alpha_S > 0$, $IncreaseAlphaSingleS(\hat{y}, y, S, \beta, \hat{\alpha})$ increases $\sum_{S'} \hat{\alpha}_{S'} \Pi_i(S')$ by $\alpha_S \Pi_i(S) \gamma_i$. Summing up over all S such that $\alpha_S > 0$ and noting that the initial value for $\sum_{S'} \hat{\alpha}_{S'} \Pi_i(S')$ is zero, we get for all i ,

$$\sum_{S'} \hat{\alpha}_{S'} \Pi_i(S') = \sum_S \alpha_S \Pi_i(S) \gamma_i = x_i \frac{\hat{x}_i}{x_i} = \hat{x}_i,$$

as desired. \square

Note that although Lemma 13 is stated in the context of stage 1 choice probability Π and the set of items N , it holds for general choice probability and set of items, as long as all assumptions of the lemma are satisfied. In particular, It holds for stage two items M and choice probability Φ^i , for any i .

Using Algorithm 2, we present the balancing step. Algorithm 3 balances stage 1 item $i \in N$. It uses the initial LP solution α^T to construct α^{t-1} , such that the

sales probability of item i is equal to the remaining inventory per period, and the sales probability of all other stage 1 items are the same as period t . Finally, $\beta^{i,t-1}$ is decreased in proportion to the sales probability of item i , so that constraint (4.5) is satisfied. Algorithm 4 balances stage 2 item $j \in M$. The idea is similar to Algorithm 3. The only difference is that, since different types of customers share stage 2 inventory, we have to proportionally decrease the sales probability of item j to each type $i \in N$.

Input: $\alpha^T, \beta^T, \alpha^t, \beta^t, \rho_i^{t-1}$;

Let x_k be the sales probability of item k under α^T . Namely $x_k \leftarrow \sum_S \alpha_S^T \Pi_k(S)$ for all k ;

Let \hat{x}_i be the remaining inventory per period for item i , and \hat{x}_j be the sales probability of item j in period t for any other item j . Namely $\hat{x}_i \leftarrow \rho_i^{t-1}$ and $\hat{x}_j \leftarrow \sum_S \alpha_S^t \Pi_j(S)$ for all $j \neq i \in N$;

Construct α^{t-1} from $\alpha^T, x, \hat{x}, \Pi, N$, according to Algorithm 2 so that

$\sum_S \alpha_S^{t-1} \Pi_k(S) = \hat{x}_k$ for any k ;

Update $\beta^{i,t-1}$ so that constraint (4.5) is satisfied. Namely for every S ,

$$\beta_S^{i,t-1} \leftarrow \hat{x}_i \frac{\beta_S^{i,t}}{\sum_{S'} \beta_{S'}^{i,t}};$$

Output: $\alpha^{t-1}, \beta^{t-1}$;

Algorithm 3: Balancing step for $i \in N$

Input: $\alpha^T, \beta^T, \alpha^{t-1}, \beta^{t-1}, \rho_j^{t-1}$;

Let y_l^k be the probability of selling item l to type k customer under β^T . Namely $y_l^k \leftarrow \sum_S \beta_S^{k,T} \Phi_l^k(S)$ for all $l \in M$ and $k \in N$;

for $i \in N$ **do**

Let \hat{y}_j^i be the remaining inventory per period for item j , times the probability of selling item j to type i customer under the current values of β . Namely $\hat{y}_j^i \leftarrow \sum_S \beta_S^{i,t-1} \Phi_j^i(S) \frac{\rho_j^{t-1}}{\sum_{s,k \in N} \beta_S^{k,t-1} \Phi_j^k(S)}$;

Let \hat{y}_l^i be the probability of selling item $l \neq j$ to type i customer under the current values of β . Namely, $\hat{y}_l^i \leftarrow \sum_S \beta_S^{i,t-1} \Phi_l^i(S)$ for all $l \neq j \in M$;

end

for $i \in N$ **do**

Construct $\beta^{i,t-1}$ from $\beta^{i,T} \frac{\sum_S \alpha_S^{t-1} \Pi_i(S)}{\sum_S \alpha_S^T \Pi_i(S)}$, $y^i \frac{\sum_S \alpha_S^{t-1} \Pi_i(S)}{\sum_S \alpha_S^T \Pi_i(S)}$, \hat{y}^i, Φ^i, M according to Algorithm 2 so that $\sum_S \beta_S^{i,t-1} \Phi_l^i(S) = \hat{y}_l^i$ for any $l \in M$;

end

Output: β^{t-1} ;

Algorithm 4: Balancing step for $j \in M$

4.4.3 Algorithm complexity

Putting together algorithms 1 - 4, and using the one period randomized policy at each time period, we get the balancing algorithm. The following theorem discusses the complexity of the algorithm.

Theorem 11. *The balancing algorithm randomizes over at most $O(n^2)$ assortments*

at each time period.

Proof. Since the linear program $L(c/T)$ has $2n+1$ constraints, there exists an optimal solution α^T, β^T with $2n+1$ positive elements. By Lemma 13, and note the fact that the inputs to the balance algorithm only has $2n+1$ positive elements, we have for any t

$$|\{S : \alpha_S^t + \beta_S^t > 0\}| \leq \sum_{S: \alpha_S^T + \beta_S^T > 0} (|S| + 1) \leq (2n+1)(n+1).$$

Thus, the algorithm randomizes over at most $O(n^2)$ assortments at each time period.

□

4.5 Feasibility and performance analysis

In this section, we prove that the algorithms proposed in Section 4.4 are feasible. We then proceed to analyzing the algorithm performance. Namely, we show that the balancing algorithm achieves an expected revenue of at least $1/3$ of the optimal offline revenue.

4.5.1 Feasibility of Algorithms

Lemma 14. *At the end of period t , $\alpha^{t-1}, \beta^{t-1}$ produced by Algorithm 1 is feasible for the LP $L(\rho^{t-1})$.*

Proof. We prove the statement by induction. First, observe that α^T, β^T is feasible for $L(\rho^T)$ at the beginning of period T by definition. Now assume that α^t, β^t is feasible for $L(\rho^t)$ at the beginning of period t , we prove that $\alpha^{t-1}, \beta^{t-1}$ is feasible for the LP $L(\rho^{t-1})$ at the end of period t .

Since α^t, β^t is feasible for $L(\rho^t)$, we have

$$\begin{aligned}\sum_S \alpha_S^t &= 1 \\ \sum_S \beta_S^{i,t} &= \sum_S \alpha_S^t \Pi_i(S).\end{aligned}$$

If there is no balancing operation in period t , we have $\alpha^t = \alpha^{t-1}$ and $\beta^t = \beta^{t-1}$. So constraints (4.4) and (4.5) are satisfied. We now show that these two constraints are satisfied after the balancing operations.

Suppose item $i \in N$ is balanced according to Algorithm 3. Since α^{t-1} is constructed from $\alpha(T)$ using Algorithm 2, it follows that $\sum_S \alpha_S^{t-1} = \sum_S \alpha_S^T = 1$. Thus constraint (4.4) is satisfied after performing Algorithm 3. For item $k \neq i \in N$, we have $\sum_S \alpha_S^{t-1} \Pi_k(S) = \sum_S \alpha_S^t \Pi_k(S)$ and $\beta^{k,t-1} = \beta^{k,t}$ after balancing item i . Thus (4.5) is satisfied for $k \neq i \in N$. For item i , (4.5) is satisfied because

$$\sum_S \beta_S^{i,t-1} = \sum_S \hat{x}_i \frac{\beta_S^{i,t}}{\sum_{S'} \beta_{S'}^{i,t}} = \hat{x}_i = \sum_S \alpha_S^{t-1} \Pi_i(S).$$

Therefore, constraints (4.4) and (4.5) are satisfied after balancing item $i \in N$. Similar argument shows that both constraints are satisfied after balancing item $j \in M$.

Now we show that the inventory constraints are satisfied. Note that if item $i \in N$

is not sold in period t , we have

$$\sum_S \alpha_S^{t-1} \Pi_i(S) = \sum_S \alpha_S^t \Pi_i(S) \leq \frac{c_i^t}{t} = \frac{c_i^{t-1}}{t} \leq \frac{c_i^{t-1}}{t-1}.$$

Similar arguments shows that if item $j \in M$ is not sold in period t , its inventory constraint is satisfied. For an item with positive sales in period t , Algorithm 1 tests whether the inventory constraint is satisfied or not, and balances that item to satisfy the constraint if its violated. Thus the inventory constraints (4.3) are satisfied for all items.

Finally, it is straightforward to show that the non-negativity constraints are always satisfied. Thus $\alpha^{t-1}, \beta^{t-1}$ is feasible for $L(\rho^{t-1})$ at the beginning of period $t-1$. The statement of this lemma thus follows by induction. \square

Lemma 15. *When Algorithm 2 is used in Algorithms 3 and 4, the inputs to that algorithm satisfy all assumptions in Lemma 13.*

Proof. We first consider balancing $i \in N$ using Algorithm 3. By definition of x , we have

$$x_k = \sum_S \alpha_S^T \Pi_k(S) \quad \forall k \in N,$$

thus the first assumption of Lemma 13 is satisfied. It is also straightforward to verify that $\hat{x} \geq 0$. We only need to show that $x_k \geq \hat{x}_k$ for all k . Note that when we do balance in period t , we have $\hat{x}_k \leq \sum_S \alpha_S^t \Pi_k(S)$. It is straightforward to show by induction that $\sum_S \alpha_S^t \Pi_k(S) \leq \sum_S \alpha_S^T \Pi_k(S)$. Thus, $x_k \geq \hat{x}_k$ for all k , and the inputs satisfy all assumptions of Lemma 13 when the algorithm balances item i using Algorithm 3.

We now consider balancing $j \in M$ using Algorithm 4. By definition of y , we have for any $k \in N$ and $l \in M$,

$$y_l^k \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^T \Pi_k(S)} = \sum_S \beta_S^{k,T} \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^T \Pi_k(S)} \Phi_l^k(S),$$

thus the first assumption of Lemma 13 is satisfied. It is also straightforward to verify that $\hat{y} \geq 0$. We only need show that $y_l^k \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^T \Pi_k(S)} \geq \hat{y}_l^k$. Note that $\hat{y}_l^k \leq \sum_S \beta_S^{k,t-1} \Phi_l^k(S)$. We thus try to bound $\sum_S \beta_S^{k,t-1} \Phi_l^k(S)$.

First, assume no balance is required in stage 1 of period t , then $\alpha^t = \alpha^{t-1}$, and we have

$$\sum_S \beta_S^{k,t-1} \Phi_l^k(S) = \sum_S \beta_S^k(t) \Phi_l^k(S) \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^t \Pi_k(S)}.$$

Now, assume in stage 1 of period t , item $i \in N$ is sold and balanced. Note that for any item $k \neq i$, Algorithm 3 does not change the value of β^k . And while it changes the values of α_S , it does not change the values of $\sum_S \alpha_S \Pi_k(S)$ for $k \neq i$, by definition of \hat{x}_k . Thus the equation above still holds for item $k \neq i$. For item i which is balanced in stage 1, we have

$$\sum_S \beta_S^{i,t-1} \Phi_l^i(S) = \sum_S \beta_S^{i,t} \frac{\hat{x}_i}{\sum_{S'} \beta_{S'}^{i,t}} \Phi_l^i(S) = \sum_S \beta_S^{i,t} \frac{\sum_S \alpha_S^{t-1} \Pi_i(S)}{\sum_S \alpha_S^t \Pi_i(S)} \Phi_l^i(S),$$

by assignment in Algorithm 3, and by feasibility of β^t which guarantees $\sum_{S'} \beta_{S'}^{i,t} = \sum_S \alpha_S^t \Pi_i(S)$.

Therefore, for any $k \in N, l \in M$ and t , we have

$$\sum_S \beta_S^{k,t-1} \Phi_l^k(S) = \sum_S \beta_S^k(t) \Phi_l^k(S) \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^t \Pi_k(S)},$$

and thus

$$\hat{y}_l^k \leq \sum_S \beta_S^{k,t-1} \Phi_l^k(S) = y_l^k \frac{\sum_S \alpha_S^{t-1} \Pi_k(S)}{\sum_S \alpha_S^t \Pi_k(S)},$$

as desired. Therefore, when called in Algorithm 4, the inputs to Algorithm 2 satisfy

all assumptions of Lemma 13. □

4.5.2 Performance guarantee

In this subsection, we provide a performance guarantee of the balancing algorithm. Let R^b be a random variable representing the revenue from the balancing algorithm. Recall that the optimal offline revenue is $L(c/T)T$. We have the following theorem.

Theorem 12. *The balancing algorithm attains the revenue ratio $\frac{\mathbf{E}[R^b]}{L(c/T)T} \geq \frac{1}{3}$.*

Proof. Let $x_i^t = \sum_S \alpha_S^t \Pi_i(S)$, and let $y_j^t = \sum_{S,i} \beta_S^{i,t} \Phi_j^i(S)$. Let $\mathbf{E}_t(\cdot)$ and $\mathbf{P}_t(\cdot)$ denote the expectation and probability conditioning on all information at the beginning of period t . Let B_i^t denote the event that item i is balanced in period t , and A_i^t denote the event that item i is sold in period t . We use $\neg A$ to denote that event A does not happen.

Observe that

$$L(c/T)T = \sum_{i \in N} T p_i x_i^T + \sum_{j \in M} T y_j^T.$$

Thus by telescoping sum

$$\begin{aligned} L(c/T)T &= \sum_{i \in N} p_i \sum_{t=1}^T \mathbf{E}_T [t x_i^t - (t-1) x_i^{t-1}] + \sum_{j \in M} p_j \sum_{t=1}^T \mathbf{E}_T [t y_j^t - (t-1) y_j^{t-1}] \\ &= \sum_{i \in N} p_i \sum_{t=1}^T \mathbf{E}_T [t x_i^t - (t-1) \mathbf{E}_t x_i^{t-1}] + \sum_{j \in M} p_j \sum_{t=1}^T \mathbf{E}_T [t y_j^t - (t-1) \mathbf{E}_t y_j^{t-1}], \end{aligned}$$

where the second equation holds by tower property.

We first bound $t x_i^t - (t-1) \mathbf{E}_t x_i^{t-1}$. By the randomized policy, item $i \in N$ is sold

in period t with probability x_i^t . If item i is not sold, or if it is sold but no balance is needed, then $x_i^{t-1} = x_i^t$ by Algorithm 1. If item i is balanced in period t , we have

$$(t-1)x_i^{t-1} = c_i^{t-1} = c_i^t - 1 \geq tx_i^t - 1.$$

Thus,

$$\begin{aligned} & tx_i^t - (t-1)\mathbf{E}_t[x_i^{t-1}] \\ & \leq \mathbf{P}_t(B_i^t) [tx_i^t - (tx_i^t - 1)] + (1 - \mathbf{P}_t(B_i^t)) [tx_i^t - (t-1)x_i^t] \\ & \leq \mathbf{P}_t(A_i^t) + x_i^t \\ & = 2\mathbf{P}_t(A_i^t). \end{aligned} \tag{4.6}$$

We now bound $ty_j^t - (t-1)\mathbf{E}_t y_j^{t-1}$. Similarly, if no balance is needed in stage 1 and item j is not balanced in stage 2, we have $y_j^{t-1} = y_j^t$. If item j is balanced, we have $(t-1)y_j^{t-1} \geq ty_j^t - 1$. We now consider the situation that item $i \in N$ is balanced, which changes the value of y_j^{t-1} , but item j itself is not balanced. In this case, we have

$$\begin{aligned} (t-1)y_j^{t-1} &= (t-1) \sum_{S, k \neq i \in N} \beta_S^{k, t-1} \Phi_j^k(S) + (t-1) \sum_S \beta_S^{i, t-1} \Phi_j^i(S) \\ &= (t-1) \sum_{S, k \neq i \in N} \beta_S^{k, t} \Phi_j^k(S) + (t-1) \sum_S \frac{c_i^{t-1}}{t-1} \frac{\beta_S^{i, t}}{\sum_{S'} \beta_{S'}^{i, t}} \Phi_j^i(S) \\ &= (t-1) \sum_{S, k \neq i \in N} \beta_S^{k, t} \Phi_j^k(S) + \sum_S \frac{c_i^t - 1}{\sum_{S'} \beta_{S'}^{i, t}} \beta_S^{i, t} \Phi_j^i(S) \\ &\geq (t-1) \sum_{S, k \neq i \in N} \beta_S^{k, t} \Phi_j^k(S) + \sum_S \frac{tx_i^t - 1}{x_i^t} \beta_S^{i, t} \Phi_j^i(S) \\ &\geq (t-1)y_j^t - \frac{1}{x_i^t} \sum_S \beta_S^{i, t} \Phi_j^i(S), \end{aligned}$$

which implies that $ty_j^t - (t-1)y_j^t \leq y_j^t + \frac{1}{x_i^t} \sum_S \beta_S^{i, t} \Phi_j^i(S)$ when item $i \in N$ is balanced but item j is not balanced.

Thus we have

$$\begin{aligned}
& ty_j^t - (t-1)\mathbf{E}_t[y_j^{t-1}] \\
& \leq \sum_{i \in N} \mathbf{P}_t(B_i^t, \neg B_j^t) \left[y_j^t + \frac{1}{x_i^t} \sum_S \beta_S^{i,t} \Phi_j^i(S) \right] + \mathbf{P}_t(B_j^t) [ty_j^t - (ty_j^t - 1)] \\
& \quad + \sum_{i \in N} \mathbf{P}_t(\neg B_i^t, \neg B_j^t) [ty_j^t - (t-1)y_j^t] \\
& \leq y_j^t + \sum_{i \in N} \mathbf{P}_t(A_i^t) \frac{1}{x_i^t} \sum_S \beta_S^{i,t} \Phi_j^i(S) + \mathbf{P}_t(A_j^t) \\
& = 2\mathbf{P}_t(A_j^t) + \sum_{S,i} \beta_S^{i,t} \Phi_j^i(S) \\
& = 3\mathbf{P}_t(A_j^t). \tag{4.7}
\end{aligned}$$

By inequalities (4.6) and (4.7), we have

$$\begin{aligned}
L(c/T)T &= \sum_{i \in N} p_i \sum_{t=1}^T \mathbf{E}_T [tx_i^t - (t-1)\mathbf{E}_t x_i^{t-1}] + \sum_{j \in M} p_j \sum_{t=1}^T \mathbf{E}_T [ty_j^t - (t-1)\mathbf{E}_t y_j^{t-1}] \\
&\leq 2 \sum_{i \in N} p_i \sum_{t=1}^T \mathbf{E}_T \mathbf{P}_t(A_i^t) + 3 \sum_{j \in M} p_j \sum_{t=1}^T \mathbf{E}_T \mathbf{P}_t(A_j^t) \\
&= 2 \sum_{i \in N} p_i \sum_{t=1}^T \mathbf{P}_T(A_i^t) + 3 \sum_{j \in M} p_j \sum_{t=1}^T \mathbf{P}_T(A_j^t) \\
&\leq 3\mathbf{E}_T[R^b],
\end{aligned}$$

as desired. □

4.6 Asymptotic Optimality

In this section, we study the asymptotic regime in which the time and the initial inventory grow proportionally and both tend to infinity. Namely, we consider the

limiting regime where $T \rightarrow \infty$, with $c = \rho T$, for some fixed vector $\rho \geq 0$. We study the performance of the balancing algorithm in this asymptotic regime, and have the following theorem on asymptotic path-wise optimality of the algorithm.

Theorem 13. *As T goes to infinity with $c = \rho T$, we have*

$$\mathbf{P} \left(\frac{R^b}{L(\rho)T} \geq 1 - \Theta(T^{-1/4}) \right) \geq 1 - n^2 T^{1/2} 2^{-T^{1/2} + T^{1/4} + 1}.$$

We prove Theorem 13 by first showing that the sales of each item in the balancing algorithm converges in probability to its inventory. In particular, let S_i be the total sales of item i in the balancing algorithm. Let x_i be the probability of selling stage 1 item i according to the offline LP. Let y_j^i be the probability of selling stage 2 item j conditioning on seeing a type i customer according to the offline LP. It is straightforward from the LP that x and y only depend on ρ , thus does not change as T grows. We have the following lemma.

Lemma 16. *As T goes to infinity with $c = \rho T$, we have*

$$\mathbf{P} \left(\frac{S_i}{x_i T} > 1 - \Theta(T^{-1/4}) \right) \geq 1 - T^{1/2} 2^{-T^{1/2} + T^{1/4} + 1} \quad \forall i \in N \quad (4.8)$$

$$\mathbf{P} \left(\frac{S_j}{\sum_i x_i y_j^i T} > 1 - \Theta(T^{-1/4}) \right) \geq 1 - n T^{1/2} 2^{-T^{1/2} + T^{1/4} + 1} \quad \forall j \in M. \quad (4.9)$$

To prove this lemma for an item i , we first divide the selling horizon into several long-enough epochs, and show that with large probability, the realized sales of item i in each epoch is not too different from its expectation. Therefore for large enough t , with large probability, the trajectory of the algorithm selling probability is not too different from x_i . Convergence in probability then follows from some algebra. The details is in Appendix C.

We now use Lemma 16 to prove Theorem 13.

Proof of Theorem 13: Recall that $L(\rho) = \sum_{i \in N} p_i x_i + \sum_{j \in M} p_j \left(\sum_{i \in N} x_i y_j^i \right)$.

By Lemma 16, we have

$$\begin{aligned}
& \mathbf{P} \left(R^b > (1 - \Theta(T^{-1/4}))L(\rho)T \right) \\
&= \mathbf{P} \left(\sum_{i \in N \cup M} p_i S_i > (1 - \Theta(T^{-1/4})) \left(\sum_{i \in N} p_i x_i + \sum_{j \in M} p_j \sum_{i \in N} x_i y_j^i \right) T \right) \\
&\geq \mathbf{P} \left(S_i > (1 - \Theta(T^{-1/4}))x_i T, \forall i \in N, \text{ and } S_j > (1 - \Theta(T^{-1/4})) \sum_{i \in N} x_i y_j^i T, \forall j \in M \right) \\
&\geq 1 - \sum_{i \in N} \mathbf{P} \left(S_i < (1 - \Theta(T^{-1/4}))x_i T \right) - \sum_{j \in M} \mathbf{P} \left(S_j < (1 - \Theta(T^{-1/4})) \sum_{i \in N} x_i y_j^i T \right) \\
&\geq 1 - n^2 T^{1/2} 2^{-T^{1/2} + T^{1/4} + 1},
\end{aligned}$$

as desired. □

4.7 The resolving algorithm

As a straightforward extension of the balancing algorithm, we consider occasionally resolving the offline LP using updated inventory and time information. This helps the algorithm to adjust according to stochastic realization of sales. In particular, we divide the selling horizon into k equally spaces subintervals, and solve the LP using updated inventory and time information at the beginning of each subinterval. We then perform the balancing algorithm using the updated LP solutions until the next resolve of LP. We call this the resolving algorithm with parameter k . Note that the resolving algorithm with parameter 1 is equivalent to the balancing algorithm. Let

$R^r(k)$ be a random variable representing the revenue under resolving algorithm with parameter k .

We now show that the revenue ratio guarantee and the asymptotic optimality results of the balancing algorithm continue to hold for the resolving algorithm.

Corollary 1. *The resolving algorithm attains the revenue ratio $\frac{\mathbf{E}[R^r(k)]}{L(\rho)T} \geq \frac{1}{3}$, for any parameters $k = 1, \dots, T$.*

Proof. Let x and y be defined as in the proof of Theorem 12. By a similar telescoping sum argument, we have

$$L(\rho)T = \sum_{t=1}^T \mathbf{E}_T \mathbf{E}_t \left[t \sum_{i \in N} p_i x_i^t + t \sum_{j \in N} p_j y_j^t - (t-1) \sum_{i \in N} p_i x_i^{t-1} - (t-1) \sum_{j \in N} p_j y_j^{t-1} \right].$$

Thus we only need to show that for any t ,

$$\begin{aligned} & \mathbf{E}_t \left[t \sum_{i \in N} p_i x_i^t + t \sum_{j \in N} p_j y_j^t - (t-1) \sum_{i \in N} p_i x_i^{t-1} - (t-1) \sum_{j \in N} p_j y_j^{t-1} \right] \\ & \leq 3 \mathbf{E}_t[\text{Revenue at time } t]. \end{aligned} \tag{4.10}$$

If the LP is not resolved at the end of period t , inequality (4.10) follows directly from inequalities (4.6) and (4.7). We now consider the case when the LP is resolved at the end of period t .

Note that after resolve,

$$(t-1) \sum_{i \in N} p_i x_i^{t-1} - (t-1) \sum_{j \in N} p_j y_j^{t-1} = L(c^{t-1}/(t-1)).$$

Let \hat{x}^{t-1} and \hat{y}^{t-1} be the values of x^{t-1} and y^{t-1} if the LP was not resolved at the end of period t . By Lemma 14, \hat{x}^{t-1} and \hat{y}^{t-1} corresponds to a feasible solution to

$L(c^{t-1}/(t-1))$. Thus, we have

$$L(c^{t-1}/(t-1)) \geq (t-1) \sum_{i \in N} p_i \hat{x}_i^{t-1} - (t-1) \sum_{j \in N} p_j \hat{y}_j^{t-1}.$$

By inequalities (4.6) and (4.7), we have

$$\begin{aligned} & \mathbf{E}_t \left[t \sum_{i \in N} p_i x_i^t + t \sum_{j \in N} p_j y_j^t - (t-1) \sum_{i \in N} p_i \hat{x}_i^{t-1} - (t-1) \sum_{j \in N} p_j \hat{y}_j^{t-1} \right] \\ & \leq 3\mathbf{E}_t[\text{Revenue at time } t]. \end{aligned}$$

Combine the three inequalities above, we get inequality (4.10). □

The following theorem establishes the asymptotic optimality of a resolving algorithm with number of resolves $k = o(T^{-1/4})$.

Theorem 14. *For resolving algorithm with number of resolves k , we have*

$\mathbf{P} \left[R^r(k) > (1 - \Theta(kT^{-1/4}))L(\rho)T \right] > 1 - n^2 T^{1/2} k^{1/2} 2^{-(T/k)^{1/2} + (T/k)^{1/4} + 2} - nk 2^{-T/k}$,
where $\Theta(kT^{-1/4}) \rightarrow 0$ and $n^2 T^{1/2} k^{1/2} 2^{-(T/k)^{1/2} + (T/k)^{1/4} + 2} + nk 2^{-T/k} \rightarrow 0$ as $T \rightarrow \infty$
with $c = \rho T$, when $k = o(T^{-1/4})$.

Proof. By Theorem 13, we have

$$\mathbf{P} \left(R^b \geq (1 - \Theta(T^{-1/4}))L(\rho)T \right) \geq 1 - n^2 T^{1/2} 2^{-T^{1/2} + T^{1/4} + 1}. \quad (4.11)$$

On the other hand, note that the probability of selling item i in the balancing algorithm is at most x_i . Using Hoeffding's inequality, and using the same argument as in the proof of Theorem 13, we have

$$\mathbf{P} \left[R^b < (1 + \Theta(T^{-1/2}))L(c/T)T \right] \geq 1 - n 2^{-T}. \quad (4.12)$$

Consider the resolving algorithm with k resolves of the LP. Let τ be the number of periods between two consecutive resolves. Assume we can show that

$$\mathbf{P} \left[R^r(k) > \left(1 - \left(\sum_{l=1}^{k-1} l^{-1/4}\right) \Theta(\tau^{-1/4})\right) R^b \right] > 1 - n^2 \tau^{1/2} k 2^{-\tau^{1/2} + \tau^{1/4} + 1} - nk 2^{-\tau}. \quad (4.13)$$

Then this theorem follows from inequalities (4.11) and (4.13), as well as the facts that $\tau = T/k$ and $\sum_{l=1}^k l^{-1/4} = \Theta(k^{3/4})$.

It is only left to show inequality (4.13). We show it using induction on k , for fixed τ . When $k = 1$, the balancing algorithm and the resolving algorithm are the same. So $R^r(1) = R^b$, and inequality (4.13) is trivially true.

Now we assume that inequality (4.13) holds for $k - 1$, and prove that it also holds for k . Consider a dummy algorithm which resolves the LP only once at time period $(k - 1)\tau$, and performs balance steps at all other time periods. Let R^d be the revenue under the dummy algorithm. Also, use $R_{t_1}^r(k)$, $R_{t_1}^b$ and $R_{t_1}^d$ to denote the revenue from period t_1 to period 1, under the resolving, balancing, and dummy algorithms, respectively. Use $\mathbf{P}_t(\cdot)$ to denote the probability conditioning on all information before time period t .

By induction hypothesis, and note that after period $(k - 1)\tau$ the dummy algorithm only performs balancing steps, we have

$$\begin{aligned} \mathbf{P}_{(k-1)\tau} \left[R_{(k-1)\tau}^r(k) > \left(1 - \left(\sum_{l=1}^{k-2} l^{-1/4}\right) \Theta(\tau^{-1/4})\right) R_{(k-1)\tau}^d \right] \\ > 1 - n^2 \tau^{1/2} (k - 1) 2^{-\tau^{1/2} + \tau^{1/4} + 1} - n(k - 1) 2^{-\tau}. \end{aligned}$$

On the other hand, let $L^{(k-1)\tau}$ denote the optimal LP objective value at time

period $(k-1)\tau$. We have

$$\begin{aligned}
& \mathbf{P}_{(k-1)\tau} \left[R_{(k-1)\tau}^d > (1 - \Theta((k-1)^{-1/4} \tau^{-1/4})) L^{(k-1)\tau} \right] \\
& \geq 1 - n^2 [(k-1)\tau]^{1/2} 2^{-[(k-1)\tau]^{1/2} + [(k-1)\tau]^{1/4} + 1} \\
& \geq 1 - n^2 \tau^{1/2} 2^{\tau^{1/2} + \tau^{1/4} + 1},
\end{aligned}$$

where the first inequality follows from (4.11), and the second inequality follows from the fact that $n^2 \tau^{1/2} 2^{\tau^{1/2} + \tau^{1/4} + 1}$ is decreasing in τ .

Combine the previous two inequalities, we get

$$\begin{aligned}
& \mathbf{P}_{(k-1)\tau} \left[R_{(k-1)\tau}^r(k) > (1 - \left(\sum_{l=1}^{k-1} l^{-1/4} \right) \Theta(\tau^{-1/4})) L^{(k-1)\tau} \right] \\
& > 1 - n^2 \tau^{1/2} k 2^{-\tau^{1/2} + \tau^{1/4} + 1} - n(k-1) 2^{-\tau}.
\end{aligned}$$

Then combine the inequality above with (4.12), we have

$$\begin{aligned}
& \mathbf{P}_{(k-1)\tau} \left[R_{(k-1)\tau}^r(k) > (1 - \left(\sum_{l=1}^{k-1} l^{-1/4} \right) \Theta(\tau^{-1/4})) R_{(k-1)\tau}^b \right] \\
& > 1 - n^2 \tau^{1/2} k 2^{-\tau^{1/2} + \tau^{1/4} + 1} - nk 2^{-\tau}.
\end{aligned}$$

Finally, note that the resolving algorithm and the balancing algorithm are the same before time period $(k-1)\tau$. Therefore we have

$$\begin{aligned}
& \mathbf{P}_{(k-1)\tau} \left[R^r(k) > (1 - \left(\sum_{l=1}^{k-1} l^{-1/4} \right) \Theta(\tau^{-1/4})) R^b \right] \\
& > 1 - n^2 \tau^{1/2} k 2^{-\tau^{1/2} + \tau^{1/4} + 1} - nk 2^{-\tau}.
\end{aligned}$$

So we proved (4.13) as desired.

□

4.8 Numerical Analysis

In this section, we provide numerical analysis on the balancing and the resolving algorithms. In subsection 4.8.1, we simulate the algorithms under various settings, and calculate the ratio between average algorithm revenue and the optimal offline revenue. In subsection 4.8.2, we provide an example in which the expected revenue under the optimal online policy is 47% of the optimal offline revenue.

4.8.1 Algorithm performance under simulation

In this subsection, we simulate the balancing and the resolving algorithms, and test their performance under different settings. We simulate the algorithm revenue under each setting using the average revenue of 10000 iterations, and calculate the ratio of the simulated algorithm revenue over the optimal offline revenue.

In all of our numerical examples, we assume that the customer choice follows the multinomial logit model for both stages. In particular, we assume that the attractiveness of no purchase option in stage 1 is v_0 , and the attractiveness of item i in stage 1 is v_i . Then the purchase probability of item i in stage 1 when offering assortment $S \subset N$ is

$$\Pi_i(S) = \frac{v_i}{v_0 + \sum_{j \in S} v_j} \mathbb{1}(i \in S).$$

Similarly, we use w_0^i to denote the no purchase attractiveness in stage 2, for a type i customer, and we use w_j^i to denote the attractiveness of item j in stage 2 for such

customer. Then the purchase probability of item j in stage 2 when offering assortment $S \subset M$ to type i customer is

$$\Phi_j^i(S) = \frac{w_j^i}{w_0^i + \sum_{k \in S} w_k^i} \mathbb{1}(j \in S).$$

We generate examples with different number of time periods, inventory levels, no purchase probabilities, and consider the cases when prices and attractiveness are negatively correlated or independent. In particular, we assume that there are 20 products in each stage. We generate their prices independently from uniform $[0, 1]$ distribution, and we generate their attractiveness independently from uniform $[0, 10]$ distribution. We sort the price array from smallest to largest in each stage, and sort the attractiveness array (of stage 1 items, or stage 2 items for each type of customer) from largest to smallest, and pair them to get the parameters when price and attractiveness are negatively correlated. We pair without sorting to get the parameters when they are independent. We consider the cases when the number of periods $T = 100$ or $T = 500$, and when the no purchase attractiveness v_0 and w_0^i are 10, 100 or 1000, which corresponds to high, medium, and low purchase probabilities. To get the initial inventory levels, we solve the linear program (4.2) without the inventory constraint, and let c^* be the total sales under the optimal LP solution. We then consider the case when inventory is a fraction of c^* , namely $c = \lceil \gamma c^* \rceil$, with $\gamma = 0.6, 0.8, 1, 1.2$, corresponding to inventory levels from scarce to abundant.

We use R^b to denote the average revenue under the balancing algorithm, and use $R^r(k)$ to denote the average revenue under the resolving algorithm with parameter k . The results are summarized in the following tables.

T	v_0	γ	correlation of price and attractiveness	$\frac{R^b}{L(\rho)T}$	$\frac{R^r(2)}{L(\rho)T}$	$\frac{R^r(5)}{L(\rho)T}$	$\frac{R^r(10)}{L(\rho)T}$	$\frac{R^r(T)}{L(\rho)T}$
100	10	0.6	independent	0.826	0.873	0.911	0.925	0.935
100	10	0.6	negatively correlated	0.852	0.873	0.911	0.923	0.980
100	10	0.8	independent	0.896	0.895	0.909	0.916	0.940
100	10	0.8	negatively correlated	0.872	0.890	0.909	0.917	0.925
100	10	1	independent	0.939	0.937	0.941	0.947	0.957
100	10	1	negatively correlated	0.933	0.938	0.944	0.944	0.947
100	10	1.2	independent	0.972	0.977	0.979	0.977	0.985
100	10	1.2	negatively correlated	0.974	0.975	0.978	0.978	0.979
100	100	0.6	independent	0.714	0.771	0.804	0.808	0.811
100	100	0.6	negatively correlated	0.735	0.780	0.798	0.809	0.815
100	100	0.8	independent	0.778	0.785	0.795	0.797	0.820
100	100	0.8	negatively correlated	0.763	0.781	0.796	0.797	0.802
100	100	1	independent	0.821	0.829	0.829	0.836	0.856
100	100	1	negatively correlated	0.808	0.825	0.833	0.830	0.831
100	100	1.2	independent	0.888	0.901	0.906	0.908	0.911
100	100	1.2	negatively correlated	0.896	0.897	0.905	0.909	0.911
100	1000	0.6	independent	0.758	0.771	0.791	0.798	0.807
100	1000	0.6	negatively correlated	0.789	0.793	0.795	0.798	0.806
100	1000	0.8	independent	0.768	0.771	0.783	0.786	0.795
100	1000	0.8	negatively correlated	0.794	0.799	0.804	0.809	0.813
100	1000	1	independent	0.771	0.798	0.806	0.805	0.809
100	1000	1	negatively correlated	0.792	0.802	0.803	0.805	0.809
100	1000	1.2	independent	0.803	0.808	0.811	0.822	0.826
100	1000	1.2	negatively correlated	0.809	0.812	0.817	0.821	0.844

T	v_0	γ	correlation of price and attractiveness	$\frac{R^b}{L(\rho)T}$	$\frac{R^r(2)}{L(\rho)T}$	$\frac{R^r(5)}{L(\rho)T}$	$\frac{R^r(10)}{L(\rho)T}$	$\frac{R^r(T)}{L(\rho)T}$
500	10	0.6	independent	0.924	0.932	0.948	0.960	0.986
500	10	0.6	negatively correlated	0.896	0.923	0.948	0.960	0.977
500	10	0.8	independent	0.943	0.942	0.958	0.966	0.983
500	10	0.8	negatively correlated	0.928	0.942	0.958	0.966	0.977
500	10	1	independent	0.957	0.956	0.961	0.963	0.975
500	10	1	negatively correlated	0.958	0.958	0.961	0.964	0.971
500	10	1.2	independent	0.994	0.995	0.996	0.995	0.996
500	10	1.2	negatively correlated	0.993	0.995	0.994	0.997	0.995
500	100	0.6	independent	0.835	0.859	0.901	0.917	0.936
500	100	0.6	negatively correlated	0.824	0.857	0.901	0.917	0.934
500	100	0.8	independent	0.847	0.869	0.897	0.907	0.927
500	100	0.8	negatively correlated	0.828	0.870	0.895	0.907	0.922
500	100	1	independent	0.872	0.880	0.893	0.897	0.913
500	100	1	negatively correlated	0.870	0.879	0.892	0.897	0.906
500	100	1.2	independent	0.955	0.957	0.961	0.966	0.979
500	100	1.2	negatively correlated	0.948	0.958	0.961	0.962	0.969
500	1000	0.6	independent	0.706	0.761	0.788	0.796	0.818
500	1000	0.6	negatively correlated	0.718	0.764	0.792	0.792	0.800
500	1000	0.8	independent	0.732	0.768	0.793	0.793	0.823
500	1000	0.8	negatively correlated	0.739	0.770	0.787	0.796	0.798
500	1000	1	independent	0.802	0.824	0.828	0.837	0.840
500	1000	1	negatively correlated	0.817	0.831	0.831	0.835	0.846
500	1000	1.2	independent	0.871	0.874	0.878	0.878	0.887
500	1000	1.2	negatively correlated	0.858	0.879	0.876	0.877	0.871

For all 48 test cases, the balancing algorithm attains at least 70.6% of the optimal offline revenue, and the resolving algorithm with parameters 2, 5, 10, T attain at least 76.1%, 78.3%, 78.6% and 79.5% of the optimal offline revenue, respectively. The average ratios are 84.9% for the balancing algorithm, and 86.5%, 87.7%, 88.4% and 89.5% for the resolving algorithm with parameters 2, 5, 10, T , respectively.

Comparing the performance of the balancing algorithm and the resolving algorithm with parameter 2, we notice that resolving the LP only once in the middle of the selling horizon may significantly improve revenue. The average improvement is 1.6%. This improvement is more significant when γ is small (i.e., when inventory is scarce). When $\gamma = 0.6$, the average improvement in performance from resolving the LP once is 3.2%, and it can be as large as 5.5%.

From this table, we also observe that both our algorithms perform significantly better when the inventory scarcity factor γ is large (i.e., when there is abundant amount of inventory). For example, the average ratio of the balancing algorithm is 79.8%, 82.4%, 86.2% and 91.4% for $\gamma = 0.6, 0.8, 1$ and 1.2. The average ratio of the resolving algorithm with parameter 2 is 83.0%, 84.0%, 87.2% and 91.9%, respectively.

We also observe that the algorithms perform better as the no purchase attractiveness v_0 decreases, or as the number of total time periods T increases. We do not observe significant difference in performance by comparing different correlation between price and attractiveness.

4.8.2 An upper bound on the performance guarantee

In this subsection, we give an example where the optimal offline revenue is a loose upper bound on the online revenue under optimal policy. This gives an upper bound on the performance guarantee ratio $\frac{R^{alg}}{L(\rho)T}$ under any policy.

We consider an example with T periods, and T items in the first stage, and 1 item in the second stage. Each item has one unit of inventory. The purchase probability in stage 1 given by

$$\Pi_i(S) = \frac{1}{T} \mathbb{1}(i \in S), \forall i \in N.$$

The purchase probability in stage 2 is $1/T$, when the item is shown to a customer, and is independent of the customer's type. We assume that all items in stage 1 have price 0, and the item in stage 2 has price 1.

Since the purchase probabilities for each item only depends on whether it is offered in the assortment, but is independent of all other items, the optimal policy is trivially offering all items that still have inventory left. Let V be the number of items sold in stage 1 under the optimal policy, then the distribution of V is given by

$$\mathbf{P}(V = k) = \binom{T}{k} \sum_{j=0}^k (-1)^j \binom{k}{j} \left(\frac{k-j}{T}\right)^T, \forall k = 0, 1, \dots, T.$$

We obtain this expression by equivalence to the birthday problem. Note that T in our setting corresponds to the number of people in the birthday problem, $1/T$ corresponds to the probability of having one of the birthdays, and V corresponds to the number of different birthdays in this group. See [63] for more information on the birthday problem.

Note that each item sold in stage 1 gives the seller one chance to show the revenue-generating stage 2 item to the buyer, who makes a purchase with probability $1/T$. Thus conditioning on $V = k$, the probability of selling the stage 2 item is $1 - (1 - 1/T)^k$. Therefore, the expected revenue of the seller under this policy is given by

$$R^{opt} = \sum_{k=1}^T \binom{T}{k} \sum_{j=0}^k (-1)^j \binom{k}{j} \left(\frac{k-j}{T}\right)^T \left(1 - \left(1 - \frac{1}{T}\right)^k\right).$$

On the other hand, it is straightforward to see that $\alpha_N = 1, \beta_M^i = 1/T$ for all $i \in N$ is the solution to the linear program (4.2). The corresponding optimal offline revenue is 1. Thus the ratio between revenue from the optimal online policy and optimal offline revenue is given by R^{opt} . When $T = 2$, $R^{opt} = 0.63$. When $T = 3$, $R^{opt} = 0.56$, and when $T = 50$, $R^{opt} = 0.47$. Therefore, any constant performance guarantee may not exceed 47%.

4.9 Discussion and Conclusion

In this chapter, we study an online assortment customization problem with inventory constraints, and with recommendation opportunity at checkout. We propose a balancing algorithm which is based on the solution to the offline problem, and a resolving algorithm as its extension. Both of them attain performance guarantee of $1/3$, and both of them are asymptotically optimal when inventory and time scales proportionally and both tend to infinity.

Our analysis is applicable to more general and complicated settings by easy extensions in various directions. We briefly discuss each possible direction of extension below. Many of the following extensions are based on treating a combination of an item and a feature (for example, price, customer class, or stage) as a “product”, and use the number of that item sold with that feature in the offline optimal solution as the inventory of that product. Then we can reproduce the two algorithms for the products. However, in the optimal offline solution, the number of each item sold may be fractional, which results in fractional inventories. Thus, we need to “protect” one unit of inventory for each product. To do this, we reduce the inventory for each item by the number of features associated with that item when calculating the optimal offline solution. We then round up each fractional inventory to the nearest integer. This incurs a loss in the performance guarantee, and the loss is increasing in the number of features, but decreasing in the initial inventory levels. Thus, these extensions works well when the inventory levels of all items are large compared with the number of features.

Pricing: In this chapter, each item only has one fixed price. However, the algorithms can be extended to incorporate the firm’s pricing problem. We assume that there are h potential price levels for each item. When the firm chooses an assortment, it also chooses a price for each item in the assortment. Our algorithm can be reproduced in this setting, with minor modification. We treat each item and price combination as a product, and define inventory of that product using the number of corresponding item sold at the corresponding price in the optimal offline solution. As discussed above, we need to protect 1 unit of inventory for each price level. The

resulting performance guarantee is thus $\frac{1}{3}(1 - \frac{h}{\min_i c_i})$.

Overlapping sets of candidate items for the two stages: In this chapter, we assume that the sets of candidate items for the two stages are disjoint. This assumption can be removed to allow overlapping sets of candidate items for the two stages. To do this, we define the combination of an item and a stage as a product, and follow the discussion above. Similarly, we need to protect 1 unit of inventory for each stage, and the performance guarantee is $\frac{1}{3}(1 - \frac{2}{\min_i c_i})$.

Different customer classes in the first stage: In the beginning of this chapter, we discussed the two types of personalization, which are personalization based on entire profile, and personalization based on item purchased. In this chapter, we focus on the latter. However, we can also incorporate the former into our model. To do this, we assume that a customer arriving to stage 1 has an observable “class” based on her profile, which the offered assortment is personalized to. Similar to previous extensions, we define the combination of an item and a customer class as a product, and protect 1 unit of inventory for each class. The performance guarantee is $\frac{1}{3}(1 - \frac{h}{\min_i c_i})$, with h classes of customers upon arrival.

Purchasing multiple items at a time: In this chapter, each customer only purchases at most one item in each stage. This model can be extended to allow multiple purchase in each stage, as long as the Assumption 3 on purchase probabilities is satisfied. To do this, define customer type as the set of items they purchase in the first stage, and perform balancing if necessary for all items that are sold in the current stage. None of the other parts of the analysis in this chapter use the assumption

that at most one item is sold in each stage. The performance guarantee is still $1/3$. However, note that by defining a customer's type as the set of items she purchases, the number of types is exponentially large in the number of items. Thus solving the offline LP may become difficult.

CHAPTER 5

CONCLUSION

Strategic behavior is present in many practical settings where dynamic pricing and assortment optimization is required. Modeling strategic behavior in those settings may help firms to better understand their customers and competitors, which is an essential step in making better operational decisions. All three problems studied in this thesis can be extended in various directions, as discussed in the corresponding chapters. We now point out some interesting directions not considered in this thesis.

5.1 Strategic and forward looking customers

In the problems studied in this thesis, we either assume customers are myopic, or assume that they are strategic but live for only one period. However, in practice many customers are strategic and forward looking. This may lead to huge difference in the customers' decision making, thus influences the firm's problem.

For example, a strategic and forward looking customer may predict future prices of products, and may decide to wait for a price cut, even if she would gain a positive utility when making a purchase in the current period. In the context of Chapter 3, with the presence of such customers, demand does not only depend on the prices in the current period, but also depend on the (perceived) future prices. With this change in demand model, the firms' problem is dramatically different.

Another example appears when customers are recurring, and are recognized and differentiated by the firm upon their return. A strategic and forward looking customer in this setting may choose to act suboptimally to “disguise” their type, for potential future benefits. In the context of Chapter 4, such customer may choose a stage 1 item that is different from the one that maximizes her stage 1 utility, in the hope that she will be shown a potentially more rewarding stage 2 assortment. As a result, with the presence of such customers, the choice probabilities in stage 1 is a function of the firm’s strategy to recommend stage 2 assortment. This results in a game between the firm and the customers.

5.2 The role of information

In this thesis, we do not consider what information the firm should make available to the public. However in practice, choosing the correct amount of information to reveal to public and the correct time to do so, may make a difference. The firm’s problem can also change significantly based on the amount of information revealed.

In the context of Chapter 2, the bid-based firm may consider whether to reveal information about current queue length and the bids submitted by customers in the queue. The current model does not reveal any of such information. In the setting where such information is available, the analysis is more difficult due to a much larger state space, and whether or not equilibrium exists (or is unique) is an open question.

In the context of Chapter 3, in the game with recourse and assuming stochastic demand, a firm may consider whether or not to reveal its real-time inventory levels. When such information is not available, a firm has to maintain Bayesian beliefs on other firms' current inventory levels, and update these beliefs based on the prices they set. The analysis of the game is thus very different from our Chapter 3, and it is an interesting direction for future research.

APPENDIX A
APPENDIX FOR CHAPTER 2

A.1 Extensions to general distribution F

All our results in Sections 2.3–2.4, except uniqueness, can be extended to the case where the distribution of the unit waiting cost is any bounded continuously differentiable distribution F . Since we use the intermediate value theorem and monotonicity with respect to c for the proof of existence of equilibrium, our proof can be extended in a straightforward manner to general distribution F , on replacing c_1, c_2 with $F(c_1), F(c_2)$ respectively, and replacing α with $F(c_2) - F(c_1)$. This argument would show the existence of equilibrium strategies with thresholds in the values of $F(c)$, which translates to the thresholds in values of c . For uniqueness, we need the additional requirement that the density of F is strictly positive on its support. This is a reasonable requirement, since otherwise there may exist multiple values of thresholds for which effectively the same set of customers obtain service from each firm.

A.2 Proofs of the results in Section 2.3

Proof of Lemma 1. Fix a symmetric equilibrium (x, b) . Consider two customers, with unit waiting costs c and $c' > c$, such that both choose to obtain service. The total expected cost of the customer with unit waiting cost c is given by $b(c) + cw(c|x, b)$.

(Here, we adopt the convention that if $x(c) = \text{FIX}$, then $b(c) \triangleq P$.) If this customer deviates, and chooses the service decision $x(c')$ with bid $b(c')$, her total expected cost would be $b(c') + cw(c'|x, b)$. As (x, b) is an equilibrium, we must have that the customer's total expected cost on following her equilibrium action should be at most as that from deviating to the action $(x(c'), b(c'))$. This implies that

$$b(c) + cw(c|x, b) \leq b(c') + cw(c'|x, b). \quad (\text{A.1})$$

Similarly for customer with unit waiting cost c' cannot become better off from deviating to the action $(x(c), b(c))$. This implies that

$$b(c') + c'w(c'|x, b) \leq b(c) + c'w(c|x, b). \quad (\text{A.2})$$

Adding the two inequalities above, we obtain that $(c' - c)(w(c'|x, b) - w(c|x, b)) \leq 0$. Since $c' > c$, we have $w(c|x, b) \geq w(c'|x, b)$, thereby proving that the expected waiting cost in equilibrium is non-increasing in the unit waiting cost.

Using the fact that $w(c|x, b) \geq w(c'|x, b)$ in (A.1), we obtain that $b(c) \leq b(c')$, thus proving that the expected payment in equilibrium is non-decreasing in the unit waiting cost.

Moreover, using $c' > c$ and $w(c'|x, b) > 0$, we get $b(c') + cw(c'|x, b) < b(c') + c'w(c'|x, b)$. Combining this with equation (A.1), we obtain that $b(c) + cw(c|x, b) < b(c') + c'w(c'|x, b)$. Thus, the total expected cost is strictly increasing in the unit waiting cost.

Now we only need to show that the bids in bid-based firm $b(c)$ is uniquely determined by and strictly increasing in the fraction of customers in the bid-based

firm with unit waiting cost less than c , which we denote by $B(c) = \frac{\int_0^c \mathbf{I}_{\{x(\hat{c})=\text{BID}\}} d\hat{c}}{\int_0^1 \mathbf{I}_{\{x(\hat{c})=\text{BID}\}} d\hat{c}}$. We first show that for any $c' > c$, $B(c') = B(c)$ implies $b(c') = b(c)$. Earlier in the proof, we show that the bids are non-decreasing in unit waiting cost, so $b(c') \geq b(c)$. Suppose $b(c') > b(c)$, then if the customer with unit waiting cost c' decreases her bid from $b(c')$ to $b(c)$, her priority decreases. And the set of unit waiting costs with higher priority than the customer when she bids $b(c)$ while having lower priority when she bids $b(c')$ is a subset of $\{\hat{c} : B(\hat{c}) = B(c')\}$. As a result, the expected waiting time increases by at most the expected time waiting for customers with unit waiting cost in the set $\{\hat{c} : B(\hat{c}) = B(c')\}$. Since the set has measure zero, the increase in expected waiting time is zero. So the customer with unit waiting cost c' is strictly better off submitting a bid $b(c)$, which contradicts the fact that (x, b) is a symmetric equilibrium. Thus, $B(c') = B(c)$ implies $b(c') = b(c)$, and the bidding function is uniquely determined by $B(\cdot)$.

Now we want to show that for any $c' > c$, $B(c') > B(c)$ implies $b(c') > b(c)$. Since bids are non-decreasing in unit waiting cost, we have $b(c') \geq b(c)$. Suppose $b(c') = b(c)$, then every customer with unit waiting cost between c and c' has the same bid. If the customer with unit waiting cost c' increases her bid to $b(c') + \epsilon$, then she can get service before anyone in the set $\{\hat{c} : B(c) \leq B(\hat{c}) < B(c'), x(\hat{c}) = \text{BID}\}$, and her expected waiting time decreases by at least the expected time she had to spend waiting for customers in the set $\{\hat{c} : B(c) \leq B(\hat{c}) < B(c'), x(\hat{c}) = \text{BID}\}$ when her bid was $b(c')$. Since the set has measure $B(c') - B(c) > 0$ and we break the ties among the customers in this set uniformly at random, the expected time one has to spend waiting for these customers to complete their service is positive (and bounded

below). Thus for small enough ϵ , the customer with unit waiting cost c' is better off using bid $b(c') + \epsilon$, which again contradicts the fact that (x, b) is an equilibrium. So in equilibrium, $B(c') > B(c)$ implies $b(c') > b(c)$, i.e., the bidding function is strictly increasing in $B(\cdot)$. \square

Proof of Lemma 2. Fix a symmetric equilibrium (x, b) . For those customers that choose not to obtain service, the expected payment, the expected waiting time, and the bid are all zero by definition. Hence, we focus only on those customers that choose to obtain service.

For a customer choosing to obtain service in the bid-based firm, by Lemma 1, the bid is strictly increasing in $B(c)$, the fraction of customers in the bid-based firm with unit waiting cost less than c . Since the priority of service is determined in the descending order of the bids (with ties broken uniformly at random), service is provided in the decreasing order of $B(\cdot)$ and thus, the expected waiting time in the bid-based firm depends only on $B(\cdot)$. Similarly, since the service discipline in the fixed-price firm is first-in-first-out, the expected waiting time in the fixed-price firm depends only on the arrival rate of the customers into the fixed-price firm, which is again determined by the service decision $x(\cdot)$. This argument proves that the expected waiting cost in equilibrium is completely specified by the service decision $x(\cdot)$. Henceforth, we denote the expected waiting time of a customer with unit waiting cost c in the symmetric equilibrium (x, c) by $w(c|x)$.

Next, we show that the bidding function is uniquely determined by service decision

in equilibrium. Fix a symmetric equilibrium (x, b) , and consider a customer with unit waiting cost c with $x(c) \neq \text{LEAVE}$. Suppose, for some \hat{c} with $x(\hat{c}) \neq \text{LEAVE}$, the customer deviates to the service decision $x(\hat{c})$, with the corresponding bid $b(\hat{c})$. Her total expected cost is given by $\pi(\hat{c}, c) = cw(\hat{c}|x) + b(\hat{c})$. (Here, recall our convention that $b(\hat{c}) = P$ if $x(\hat{c}) = \text{FIX}$.) By the fact that (x, b) is a symmetric equilibrium, we obtain

$$\pi(c, c) = \max_{\hat{c}: x(\hat{c}) \neq \text{LEAVE}} \pi(\hat{c}, c),$$

for all $c \in [0, 1]$ with $x(c) \neq \text{LEAVE}$, with the maximum being attained at $\hat{c} = c$. We use the Mirrlees trick [51] to compute $\pi(c, c)$ in equilibrium. Towards that goal, observe that by the envelope theorem, we obtain

$$\frac{d\pi(c, c)}{dc} = \left. \frac{\partial \pi(\hat{c}, c)}{\partial c} \right|_{\hat{c}=c} = w(c|x),$$

for all c such that $x(c) \neq \text{LEAVE}$. Now, note that if $x(c) \neq \text{LEAVE}$, then from Lemma 1, we obtain that $x(c') \neq \text{LEAVE}$ for all $c' \leq c$. This, combined with the fact that $\pi(0, 0) = 0$, we obtain by integrating

$$\pi(c, c) = \int_0^c w(t|x) dt.$$

Finally, since $\pi(c, c) = cw(c|x) + b(c)$, we get

$$b(c) = \int_0^c w(t|x) dt - cw(c|x),$$

for all c such that $x(c) \neq \text{LEAVE}$. Thus, the bidding function is completely determined from the expected waiting time $w(\cdot|x)$, which in turn depends only on the service decision x . □

A.3 Proofs of auxiliary results in Section 2.4

Proof of Lemma 3. The first two statements follow directly from a straightforward stochastic coupling argument. We provide a brief sketch of the continuity and monotonicity of Γ for completeness. To show that $\Gamma(a) < \Gamma(b)$ for $0 \leq a < b < k$, consider two coupled copies Q_a, Q_b of the preemptive bid-based priority queue with arrival rate $\lambda \in (b, k)$. Suppose two customers, one per queue, arrive at time 0 with identical service requirement, with the only difference being that the customer to the copy Q_i has a priority level such that the arrival rate of customers to Q_i with higher priority is exactly equal to i , for each $i \in \{a, b\}$. Through this coupling, it follows directly that the waiting time of the customer to the queue Q_a is almost surely less than the waiting time of the customer to the queue Q_b . Taking expectations, we obtain from the definition of Γ that $\Gamma(a) \leq \Gamma(b)$. The strict inequality and continuity follow from the fact that the arrival rate of customers with intermediate priority between a and b is positive, and decreases to zero as a approaches b .

The third statement follows directly from the fact that the expected service time is one.

Since $\int_0^y \Gamma(t)dt/y$ is the average waiting time in the system when there are k servers and the arrival rate is y , it is no less than the average waiting time in a work conserving system with the same number of servers and the same arrival rate. Since the latter tends to infinity as y approaches k , we have

$$\int_0^y \Gamma(t)dt \rightarrow \infty, \quad \text{as } y \rightarrow k.$$

□

Proof of Proposition 1. We prove a more general statement, which will be useful later when we prove Lemma 4 and Lemma 5. The statement we prove here is that suppose for some $c_\ell \leq 1$, service is mandatory for anyone with cost $c \leq c_\ell$, and service is forbidden for anyone with cost $c > c_\ell$, then $\bar{c} = (c_1, c_2, c_\ell)$, with $0 < c_1 < c_2 < c_\ell \leq 1$ is a symmetric equilibrium if and only if (ContT-P) and (ContW-P) are satisfied.

We split the proof into two steps showing first the necessity and then the sufficiency of the conditions for equilibrium.

Necessity of (ContT-P) and (ContW-P): Suppose $\bar{c} = (c_1, c_2, c_\ell)$, with $0 < c_1 < c_2 < c_\ell \leq 1$ is a symmetric equilibrium. We begin by showing that the condition (ContT-P) holds.

Consider a customer with unit waiting cost $c_1 + \epsilon$ for some $\epsilon \in (0, c_2 - c_1)$, who obtains service from the fixed-price firm, and incurs a total expected cost equal to $(c_1 + \epsilon)w_F(\bar{c}) + P$. If such a customer decides instead to obtain service from the bid-based firm and submit the same bid as a customer with unit waiting cost c_1 , then her total expected cost is given by

$$(c_1 + \epsilon)w(c_1|\bar{c}) + b(c_1|\bar{c}) = \int_0^{c_1} w(t|\bar{c})dt + \epsilon w(c_1|\bar{c}),$$

where the right hand side follows from (2.2). In equilibrium such a unilateral deviation must be non-preferable. Since this is true for any $\epsilon \in (0, c_2 - c_1)$, we obtain $\int_0^{c_1} w(t|\bar{c})dt \geq c_1 w_F(\bar{c}) + P$. On the other hand, since a customer with unit waiting cost c_1 prefers to obtain service from the bid-based firm instead of the fixed-price firm,

we have $\int_0^{c_1} w(t|\bar{c})dt \leq c_1 w_F(\bar{c}) + P$. Together, these inequalities yield (ContT-P).

We now show that condition (ContW-P) holds. As the expected waiting time is non-increasing in c , we have $w_B(c_1|\bar{c}) \geq w_F(\bar{c}) \geq w_B(c_2|\bar{c})$. Since customers in the bid-based firm are served in the decreasing order of their waiting costs, by the fact that F is continuous and hence has no atom, it follows that the expected waiting time for a customer with unit waiting cost c_1 and c_2 must be equal. From this, we obtain (ContW-P).

Sufficiency of (ContT-P) and (ContW-P): Suppose all customers adopt a strategy $\bar{c} = (c_1, c_2, c_\ell)$ with $0 < c_1 < c_2 < c_\ell \leq 1$ satisfying the conditions (ContT-P) and (ContW-P). We begin by obtaining the expression for the expected waiting time and the expected total cost under this strategy profile.

Observe that since the expected waiting times in the bid-based firm and the fixed-price firm satisfy the condition (ContW-P), the expected waiting time function defined by

$$w(t|\bar{c}) = \begin{cases} w_B(t|\bar{c}), & \text{for } t \in [0, c_1] \cup [c_2, c_\ell]; \\ w_F(\bar{c}), & \text{for } t \in (c_1, c_2) \end{cases} \quad (\text{A.3})$$

is continuous and non-increasing over $t \in [0, c_\ell]$.

Next, note that for a customer with unit waiting cost c obtaining service in the bid-based firm, by (2.2), her total expected cost is given by $\int_0^c w(t|\bar{c})dt$. We show that this expression holds for all customers, even those obtaining service from the fixed-price firm. To see this, note that for such a customer with unit waiting cost

$c \in (c_1, c_2)$, the total expected cost is given by

$$\begin{aligned} cw_F(\bar{c}) + P &= c_1 w_F(\bar{c}) + P + (c - c_1) w_F(\bar{c}) \\ &= \int_0^{c_1} w_B(t|\bar{c}) dt + (c - c_1) w_F(\bar{c}) = \int_0^c w(t|\bar{c}) dt, \end{aligned}$$

where the second equality follows (ContT-P), and the third from (A.3).

Next, consider a customer with unit waiting cost $c \leq 1$. If she adopts the actions of a customer with unit waiting cost $c' \neq c$, then her total expected cost, using (2.2), is given by

$$\begin{aligned} cw(c'|\bar{c}) + b(c'|\bar{c}) &= \int_0^{c'} w(t|\bar{c}) dt + (c - c') w(c'|\bar{c}) \\ &= \int_0^c w(t|\bar{c}) dt + \int_c^{c'} w(t|\bar{c}) dt + (c - c') w(c'|\bar{c}) \\ &\geq \int_0^c w(t|\bar{c}) dt + (c' - c) w(c'|\bar{c}) + (c - c') w(c'|\bar{c}) \\ &= \int_0^c w(t|\bar{c}) dt. \end{aligned}$$

Here, the inequality follows from the fact that the expected waiting time $w(\cdot|\bar{c})$ is non-increasing. Since the right hand side denotes the total expected cost for the customer under strategy \bar{c} , this implies that a best response of a customer with unit waiting cost c is the action suggested by the strategy \bar{c} .

Taken together, this implies that the strategy \bar{c} is a best response, assuming all others follow \bar{c} , and hence it is a symmetric equilibrium. \square

Proof of Lemma 4. We already proved the necessity of condition (IND) in the beginning of Subsection 2.4.2. It is left to show the sufficiency of condition (IND). Suppose condition (IND) holds, we show that $\bar{c}(u)$ is an equilibrium for $\text{SYS}_{\text{op}}(\frac{\lambda}{u}, Pu, V)$.

First consider the case when $0 < c_1 < c_2 < 1$. By Proposition 1, since $\bar{c} = (c_1, c_2, 1)$ is a symmetric equilibrium for $\text{SYS}_{\text{man}}(\lambda, P)$, it satisfies (ContT-P) and (ContW-P). First, we prove that $\bar{c}(u)$ satisfies conditions (ContT-P) and (ContW-P). To prove this observe that

$$\begin{aligned} \int_0^{c_1 u} w_B(t|\bar{c}(u))dt &= \int_0^{c_1 u} \Gamma\left(\frac{\lambda}{u}u - \frac{\lambda}{u}(c_2 u - c_1 u) - \frac{\lambda}{u}t\right) dt \\ &= u \int_0^{c_1} \Gamma(\lambda - \lambda(c_2 - c_1) - \lambda t)dt = u \int_0^{c_1} w_B(t|\bar{c})dt. \end{aligned}$$

Here, the first equality follows from the definition of $w_B(t|\bar{c}(u))$, the second equality follows from a change of variable and the third follows from the definition of $w_B(t|\bar{c})$.

Now, since \bar{c} satisfies (ContT-P), we obtain

$$\begin{aligned} u \int_0^{c_1} w_B(t|\bar{c})dt &= u(c_1 w_F(\bar{c}) + P) = u(c_1 \Phi(\lambda(c_2 - c_1)) + P) \\ &= c_1 u \Phi\left(\frac{\lambda}{u}(c_2 u - c_1 u)\right) + Pu = (c_1 u)w_F(\bar{c}(u)) + Pu. \end{aligned}$$

Here, the second and the fourth equalities follow from the definition of $w_F(\bar{c})$ and $w_F(\bar{c}(u))$ respectively. From this, we obtain that condition (ContT-P) holds for $\bar{c}(u)$. Through a similar argument, we can show that condition (ContW-P) holds for $\bar{c}(u)$.

As we discussed in the proof of Proposition 1, since $\bar{c}(u)$ satisfies conditions (ContT-P) and (ContW-P), the action suggested by $\bar{c}(u)$ is optimal for customers with cost $c \leq u$ when the choice of **LEAVE** is not available. It is only left to show that **LEAVE** is not optimal for anyone with $c \leq u$, and it is optimal for anyone with $c > u$.

Now we show **LEAVE** is not optimal for anyone with $c \leq u$. Note that the expected total cost for a customer with unit cost $c < u$ is $\int_0^c w(t|\bar{c}(u))dt$, which is less than $\int_0^u w(t|\bar{c}(u))dt$, the expected total cost for a customer with unit cost u . By condition (IND), $\int_0^u w(t|\bar{c}(u))dt \leq V$, thus the expected total cost for customer with unit cost

$c \leq u$ is also no more than V , and LEAVE is not optimal for such a customer.

Finally we show LEAVE is optimal for anyone with $c > u$, when $u < 1$. Suppose there exists $c > u$ such that customers with unit cost c have an action that strictly dominates LEAVE. Such an action has expected total cost strictly less than V when the unit cost is c . Consider a customer with unit cost u using this action, and her expected total cost is also strictly less than V . This contradicts with (IND), which suggests that when $u < 1$, customer with a unit cost u has expected total cost of V , and we have a contradiction. Thus, LEAVE is optimal for any customer with $c > u$.

Thus we proved when $0 < c_1 < c_2 < 1$, all customers do not want to deviate from the action suggested by $\bar{c}(u)$. Thus $\bar{c}(u)$ is an equilibrium for $\text{SYS}_{\text{op}}(\frac{\lambda}{u}, Pu, V)$.

On the other hand, suppose $c_1 = c_2 = 1$. By the same argument used in the proof of Theorem 2, the action suggested by $\bar{c}(u)$ is optimal for customers with cost $c \leq u$ when the choice of LEAVE is not available. We can also use the same argument as for the case with $0 < c_1 < c_2 < 1$, to show that LEAVE is not optimal for anyone with $c \leq u$, and it is optimal for anyone with $c > u$. Thus it follows that $\bar{c}(u)$ is also an equilibrium for $\text{SYS}_{\text{op}}(\frac{\lambda}{u}, Pu, V)$ in this case. \square

Proof of Lemma 5. If $0 < c_1 < c_2 < u \leq 1$, by Proposition 1, (c_1, c_2, u) satisfies conditions (ContT-P) and (ContW-P). Use the same argument as in the proof of Lemma 4, we can show that $(\frac{c_1}{u}, \frac{c_2}{u}, 1)$ also satisfies conditions (ContT-P) and (ContW-P), thus $(\frac{c_1}{u}, \frac{c_2}{u}, 1)$ is an equilibrium for $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$.

If $c_1 = c_2 = u$, by the same argument as used in the proof of Theorem 2, (c_1, c_2, u) satisfies condition (Pref-BID). We can show that $(\frac{c_1}{u}, \frac{c_2}{u}, 1)$ also satisfies the same condition and thus is an equilibrium for $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$. \square

A.4 Proofs of Theorem 2 and Theorem 3

Define $\nu = \min\{n, \lambda\}$ and $\kappa = \min\{k, \lambda\}$. Let ξ be defined as

$$\xi \triangleq \sup\{\lambda - \kappa < x < \nu : \Gamma(\lambda - x) > \Phi(x)\}.$$

As $x \downarrow \lambda - \kappa$, we have $\Gamma(\lambda - x) \rightarrow \Gamma(\kappa) > 1$ if $\kappa < k$ and $\Gamma(\lambda - x) \rightarrow \infty$ if $\kappa = k$. On the other hand, $\Phi(\lambda - \kappa) = \Phi(0) = 1$ if $\kappa < k$, and $\Phi(\lambda - \kappa) < \Phi(n) = \infty$ if $\kappa = k$. Hence, $\xi \in (\lambda - \kappa, \nu]$. By continuity and strict monotonicity of Γ and Φ , we obtain that for all $x \in [\lambda - \kappa, \xi)$, $\Gamma(\lambda - x) > \Phi(x)$, and $\Gamma(\xi) \geq \Phi(\xi)$, with equality if $\xi < \nu$.

Define the function s as follows: for $x \in (\lambda - \kappa, \nu)$, and $z \in [0, \lambda - x]$,

$$s(z, x) \triangleq \int_x^{z+x} \Gamma(\lambda - t) dt - z\Phi(x).$$

First observe that from Assumption 1, we obtain that $s(z, x)$ is twice-differentiable in z and x (almost everywhere). Note that we have $\frac{\partial^2 s(z, x)}{\partial z^2} = -\Gamma'(\lambda - z - x) < 0$. Thus $s(z, x)$ is strictly concave in z for each $x \in (\lambda - \kappa, \nu)$. Thus, there exists a unique maximizer of $s(z, x)$ over $z \in [0, \lambda - x]$ for all $x \in (\lambda - \kappa, \nu)$. Define $z_{\max}(x)$ to be this unique maximizer for $x \in (\lambda - \kappa, \nu)$:

$$z_{\max}(x) = \arg \max_{z \in [0, \lambda - x]} s(z, x).$$

We have the following properties of $z_{\max}(x)$.

Lemma 17. *The function $z_{\max}(x)$ is continuous over $x \in (\lambda - \kappa, \nu)$. Further, $z_{\max}(\cdot)$ is strictly decreasing over $(\lambda - \kappa, \xi)$ and $z_{\max}(x) = 0$ for all $x \in [\xi, \nu)$. Moreover, $z_{\max}(x) = \lambda - x - \Gamma^{-1}(\Phi(x)) \in (0, \lambda - x)$ for $x \in (\lambda - \kappa, \xi)$, and if $\kappa = \lambda < k$, then $z_{\max}(0) \triangleq \lim_{x \downarrow 0} z_{\max}(x) = \lambda$.*

Proof. The continuity of $z_{\max}(x)$ follows from the application of Berge's maximum theorem [5] to the function $s(z, x)$, which is strictly concave in z for each $x \in (\lambda - \kappa, \nu)$.

Next, observe that for all $x \in (\lambda - \kappa, \nu)$, we have

$$\left. \frac{\partial s(z, x)}{\partial z} \right|_{z=0} = \Gamma(\lambda - x) - \Phi(x),$$

which is positive for all $x \in (\lambda - \kappa, \xi)$, equals zero for $x = \xi$ if $\xi < \nu$, and is negative for $x \in (\xi, \nu)$. Moreover, we have

$$\left. \frac{\partial s(z, x)}{\partial z} \right|_{z=\lambda-x} = \Gamma(0) - \Phi(x),$$

which is negative for all $x \in (\lambda - \kappa, \nu)$. From this, it follows that the unique maximizer $z_{\max}(x)$ of $s(z, x)$ over $z \in [0, \lambda - x]$ must lie in $(0, \lambda - x)$ for all $x \in (\lambda - \kappa, \xi)$, and equals zero for $x \in [\xi, \nu)$.

Since $z_{\max}(x) \in (0, \lambda - x)$ for all $x \in (\lambda - \kappa, \xi)$, we have by first order necessary conditions,

$$\left. \frac{\partial s(z, x)}{\partial z} \right|_{z=z_{\max}(x)} = \Gamma(\lambda - z_{\max}(x) - x) - \Phi(x) = 0, \quad \text{for all } x \in (\lambda - \kappa, \xi).$$

Thus, we obtain $z_{\max}(x) = \lambda - x - \Gamma^{-1}(\Phi(x))$ for all $x \in (\lambda - \kappa, \xi)$. Since Γ and Φ are strictly increasing, this implies that $z_{\max}(x)$ is strictly decreasing over $x \in (\lambda - \kappa, \xi)$. Finally, suppose $\kappa = \lambda < k$. Then, for all small enough $\epsilon > 0$, we have $z_{\max}(\epsilon) = \lambda - \epsilon - \Gamma^{-1}(\Phi(\epsilon))$. Taking limits as $\epsilon \rightarrow 0$, and observing that

$\Gamma(0) = \Phi(0) = 1$, we obtain that $z_{\max}(0) = \lambda$. \square

Lemma 18. *The function $s(z_{\max}(x), x)$ is continuous over $(\lambda - \kappa, \nu)$ and strictly decreasing over $(\lambda - \kappa, \xi)$. Further, we have $s(z_{\max}(x), x) = 0$ for all $x \in [\xi, \nu)$.*

Proof. The continuity follows trivially from the continuity of $s(z, x)$ and Lemma 17.

For $x \in [\xi, \nu)$, the result follows directly from $z_{\max}(x) = 0$ and the definition of $s(z, x)$.

For $x \in (\lambda - k, \xi)$, we have, by the envelope theorem,

$$\frac{ds(z_{\max}(x), x)}{dx} = \Gamma(\lambda - z_{\max}(x) - x) - \Gamma(\lambda - x) - z_{\max}(x)\Phi'(x) < 0,$$

where the last inequality follows from the fact that Γ is strictly increasing, Φ is strictly increasing, and $z_{\max}(x) > 0$. \square

Lemma 19. *We have*

$$\lim_{x \downarrow \lambda - \kappa} s(z_{\max}(x), x) = \lambda P_{\max}(\lambda) = \begin{cases} \int_0^\lambda \Gamma(t) dt - \lambda & \text{if } \kappa = \lambda < k; \\ \infty & \text{if } \kappa = k. \end{cases}$$

Proof. If $\kappa = \lambda < k$, then $\lambda - \kappa = 0$. Then, by continuity of $s(z, x)$ and $z_{\max}(x)$, we obtain

$$\begin{aligned} \lim_{x \downarrow \lambda - \kappa} s(z_{\max}(x), x) &= s(z_{\max}(0), 0) \\ &= s(\lambda, 0) \\ &= \int_0^\lambda \Gamma(\lambda - t) dt - \lambda \Phi(0) \\ &= \int_0^\lambda \Gamma(t) dt - \lambda, \end{aligned}$$

where we use that fact that $z_{\max}(0) = \lambda$ if $\kappa = \lambda < k$, and that $\Phi(0) = 1$.

Next, for $\kappa = k \leq \lambda$, we have for $\lambda - \kappa < x < \xi$,

$$\begin{aligned}
s(z_{\max}(x), x) &\geq s(\lambda - x, x) \\
&= \int_x^\lambda \Gamma(\lambda - t) dt - (\lambda - x)\Phi(x) \\
&= \int_0^{\lambda-x} \Gamma(t) dt - (\lambda - x)\Phi(x) \\
&\geq \int_0^{\lambda-x} \Gamma(t) dt - \kappa\Phi(\xi).
\end{aligned}$$

For $\kappa = k$, by Assumption 1, $\lim_{x \downarrow \lambda - \kappa} s(z_{\max}(x), x) = \infty$.

□

Lemma 20. *Suppose $P \geq P_{\max}(\lambda)$. Then, for all $x \in (\lambda - \kappa, \nu)$ and for all $z \in [0, \lambda - x]$, we have $s(z, x) < \lambda P$.*

Proof. The result follows directly from Lemma 18 and Lemma 19.

□

For any $0 < P < P_{\max}(\lambda)$, define $F(P)$ as follows:

$$F(P) = \sup\{\lambda - \kappa < x < \nu : s(z_{\max}(x), x) > \lambda P\}$$

From Lemma 18 and Lemma 19, we have $F(P) \in (\lambda - \kappa, \xi)$ for all $P \in (0, P_{\max}(\lambda))$.

Note that for $P \in (0, P_{\max}(\lambda))$, and for each $x \in (\lambda - \kappa, F(P)]$, we have $s(0, x) = 0 < \lambda P$. Also, by continuity of $s(z_{\max}(x), x)$, we have $s(z_{\max}(x), x) > \lambda P$ for all $x \in (\lambda - k, F(P))$. Thus, there exists a unique solution $z = v_1(x, P) \in (0, z_{\max}(x))$ to the equation $s(z, x) = \lambda P$ for each $x \in (\lambda - k, F(P))$ and $0 < P < P_{\max}(\lambda)$. (Here uniqueness follows from the strict concavity of $s(z, x)$ in z). For $x \in [F(P), \nu)$, define $v_1(x, P) = z_{\max}(F(P))$. (It is straightforward to show that $v_1(x, P)$ is continuous at

$x = F(P).$)

Define

$$v_2(x, P) \triangleq \begin{cases} v_1(x, P) + x & \text{if } x \in (\lambda - \kappa, F(P)]; \\ v_1(F(P), P) + F(P) & \text{if } x \in (F(P), \nu). \end{cases}$$

Note that for all $0 < P < P_{\max}(\lambda)$, and $x \in (\lambda - \kappa, F(P)]$, we have $v_2(x, P) = v_1(x, P) + x < z_{\max}(x) + x \leq \lambda$. Hence, $v_2(x, P) < \lambda$. Define $\Psi(x, P) \triangleq \Phi^{-1}(\Gamma(\lambda - v_2(x, P)))$ for all $x \in (\lambda - \kappa, \nu)$ and $0 < P < P_{\max}(\lambda)$.

Lemma 21. *For all $P \in (0, P_{\max}(\lambda))$, $\Psi(x, P)$ is strictly decreasing in x over $(\lambda - \kappa, F(P)]$.*

Proof. It suffices to show that $v_1(x, P)$ is strictly increasing in x over $(\lambda - \kappa, F(P)]$. This is because, then so is $v_2(x, P) = v_1(x, P) + x$, and the proof follows from observing that both Γ and Φ are strictly increasing.

Note that for $P \in (0, P_{\max}(\lambda))$, we have by definition, $s(v_1(x, P), x) = \lambda P$ for all $x \in (\lambda - \kappa, F(P)]$. This implies, on differentiating with respect to x ,

$$\frac{\partial v_1(x, P)}{\partial x} \frac{\partial s(z, x)}{\partial z} \Big|_{z=v_1(x, P)} + \frac{\partial s(z, x)}{\partial x} \Big|_{z=v_1(x, P)} = 0.$$

Observe that, for $x \in (\lambda - \kappa, F(P))$,

$$\frac{\partial s(z, x)}{\partial z} \Big|_{z=v_1(x, P)} = \Gamma(\lambda - v_1(x, P) - x) - \Phi(x) > 0,$$

and

$$\frac{\partial s(z, x)}{\partial x} \Big|_{z=v_1(x, P)} = \Gamma(\lambda - v_1(x, P) - x) - \Gamma(\lambda - x) - v_1(x, P)\Phi'(x) < 0.$$

This implies that $\frac{\partial v_1(x, P)}{\partial x} > 0$, and hence we are done. \square

Finally, we have the following characterization of the fixed point of $\Psi(\cdot, P)$ for all $P \in [0, P_{\max}(\lambda))$.

Lemma 22. *For each $P \in (0, P_{\max}(\lambda))$, the equation $\Psi(x, P) = x$ with $x \in (\lambda - \kappa, \nu)$ has a unique solution given by $F(P) \in (\lambda - \kappa, \xi)$.*

Proof. Since $F(P) \in (\lambda - \kappa, \xi)$, from Lemma 17 we obtain that $z_{\max}(F(P)) = \lambda - F(P) - \Gamma^{-1}(\Phi(F(P)))$. By continuity of $s(z_{\max}(x), x)$, we have $s(z_{\max}(F(P)), F(P)) = \lambda P$. This implies that $v_1(F(P), P) = z_{\max}(F(P)) = \lambda - F(P) - \Gamma^{-1}(\Phi(F(P)))$. Substituting the expression for $v_1(F(P), P)$, we obtain $\Psi(F(P), P) = \Phi^{-1}(\Gamma(\lambda - v_2(F(P), P))) = \Phi^{-1}(\Gamma(\lambda - v_1(F(P), P) - F(P))) = F(P)$. Finally, from Lemma 21, we obtain that $\Psi(x)$ is strictly decreasing in $(\lambda - \kappa, F(P)]$. Since $\Psi(x, P) = F(P) < x$ for $x \in (F(P), \nu)$, we obtain that $F(P)$ is the only solution to $\Psi(x, P) = x$ in $(\lambda - \kappa, \nu)$ for all $P \in (0, P_{\max}(\lambda))$. \square

The following theorem relates the fixed points of $\Psi(\cdot, P)$ to symmetric equilibria.

Theorem 15. *1. Suppose for some $P > 0$, there exists a symmetric equilibrium $\bar{c} = (c_1, c_2, 1)$ with $0 < c_1 < c_2 < 1$. Then, $P < P_{\max}(\lambda)$, and $x \triangleq \lambda(c_2 - c_1) \in (\lambda - \kappa, \nu)$ satisfies $\Psi(x, P) = x$, with $\lambda c_1 = v_1(x, P)$.*

2. Conversely, suppose for some $x \in (\lambda - \kappa, \nu)$ and $P \in (0, P_{\max}(\lambda))$, we have $\Psi(x, P) = x$. Then, $\bar{c} = (c_1, c_2, 1)$ with $c_1 = \frac{v_1(x, P)}{\lambda} > 0$ and $c_2 = \frac{v_2(x, P)}{\lambda} \in (c_1, 1)$, constitutes a symmetric equilibrium.

Proof. We provide the proof in two steps corresponding to the two statements.

Step 1. Suppose $\bar{c} = (c_1, c_2, 1)$ is a symmetric equilibrium with $0 < c_1 < c_2 < 1$. Let $x = \lambda(c_2 - c_1)$. By stability of the fixed-price firm in equilibrium, we have $x < \min\{n, \lambda\} = \nu$. By stability of the bid-based firm in equilibrium, we have $\lambda - x < \min\{k, \lambda\}$, implying $x > \lambda - \kappa$. Thus, $x \in (\lambda - \kappa, \nu)$.

Now, in the symmetric equilibrium $\bar{c} = (c_1, c_2, 1)$, we obtain, from (2.3) and (2.4), that

$$w_F(\bar{c}) = \Phi(x)$$

$$w_B(t|\bar{c}) = \begin{cases} \Gamma(\lambda - \lambda t) & t \in [c_2, 1]; \\ \Gamma(\lambda - x - \lambda t) & t \in [0, c_1], \end{cases}$$

where we have used the fact that $x = \lambda(c_2 - c_1)$.

Recall, from Theorem 1, that the necessary conditions for \bar{c} to be an equilibrium are

$$\int_0^{c_1} w_B(t|\bar{c}) dt = c_1 w_F(\bar{c}) + P,$$

$$w_B(c_1|\bar{c}) = w_B(c_2|\bar{c}) = w_F(\bar{c}).$$

Using the expressions for $w_F(\bar{c})$ and $w_B(t|\bar{c})$, we obtain

$$\int_0^{c_1} \Gamma(\lambda - x - \lambda t) dt = c_1 \Phi(x) + P$$

$$\Gamma(\lambda - x - \lambda c_1) = \Phi(x).$$

On substituting $u = x + \lambda t$ in the integral in the first equation and rearranging, we obtain

$$\int_x^{x+\lambda c_1} \Gamma(\lambda - u) du - \lambda c_1 \Phi(x) = \lambda P$$

Note that, by definition, the left hand side is equal to $s(\lambda c_1, x)$. Thus, we obtain

$s(\lambda c_1, x) = \lambda P$. Now, $x \in (\lambda - \kappa, \nu)$ and $\lambda c_1 = \lambda c_2 - x < \lambda - x$. Hence, the equation $s(\lambda c_1, x) = \lambda P$, together with Lemma 20, implies that $P < P_{\max}(\lambda)$.

Next, from the second necessary condition, $\Gamma(\lambda - \lambda c_1 - x) = \Phi(x)$, we obtain that $\Gamma(\lambda - x) > \Phi(x)$, which yields $x \in (\lambda - \kappa, \xi)$. This, along with Lemma 17, yields $z_{\max}(x) = \lambda c_1$. Hence, $s(z_{\max}(x), x) = \lambda P$. By Lemma 18, we know that $s(z_{\max}(t), t)$ is strictly decreasing in t over $(\lambda - \kappa, \xi)$. Further, by definition of $F(P)$, we obtain $s(z_{\max}(F(P)), F(P)) = \lambda P$. Taken together, we obtain $x = F(P)$, and hence $\Psi(x, P) = x$. Moreover, we have $\lambda c_1 = z_{\max}(x) = z_{\max}(F(P)) = v_1(F(P), P) = v_1(x, P)$.

Step 2. Suppose for some $x \in (\lambda - \kappa, \nu)$ and $P \in (0, P_{\max}(\lambda))$, we have $\Psi(x, P) = P$. By Lemma 22, we obtain that $x = F(P) \in (\lambda - \kappa, \xi)$. Let $c_1 = v_1(x, P)/\lambda > 0$, and $c_2 = v_2(x, P)/\lambda > c_1$. Note that $\lambda c_2 = v_2(x, P) < \lambda$, and hence $c_2 < 1$.

Since $x = F(P)$, we obtain $z_{\max}(x) = v_1(x, P) = \lambda c_1$ and $s(\lambda c_1, x) = s(z_{\max}(x), x) = \lambda P$. Thus, we obtain

$$\int_x^{x+\lambda c_1} \Gamma(\lambda - t) dt - \lambda c_1 \Phi(x) = \lambda P.$$

Now, observe that for $\bar{c} = (c_1, c_2, 1)$, we have

$$w_F(\bar{c}) = \Phi(x)$$

$$w_B(t|\bar{c}) = \begin{cases} \Gamma(\lambda - \lambda t) & t \in [c_2, 1]; \\ \Gamma(\lambda - x - \lambda t) & t \in [0, c_1]. \end{cases}$$

Substituting these expressions, and making a change of variables, yields,

$$P = \frac{1}{\lambda} \left(\int_x^{x+\lambda c_1} \Gamma(\lambda - t) dt - \lambda c_1 \Phi(x) \right)$$

$$\begin{aligned}
&= \frac{1}{\lambda} \left(\int_0^{\lambda c_1} \Gamma(\lambda - x - t) dt - \lambda c_1 \Phi(x) \right) \\
&= \int_0^{c_1} \Gamma(\lambda - x - \lambda t) dt - c_1 \Phi(x) \\
&= \int_0^{c_1} w_B(t|\bar{c}) dt - c_1 w_F(\bar{c}).
\end{aligned}$$

Thus, we obtain,

$$\int_0^{c_1} w_B(t|\bar{c}) dt = c_1 w_F(\bar{c}) + P.$$

Finally, since $\Psi(x, P) = x$, we obtain

$$\Phi(x) = \Gamma(\lambda - v_2(x, P)) = \Gamma(\lambda - x - v_1(x, P)) = \Gamma(\lambda - x - \lambda c_1).$$

This implies that $w_F(\bar{c}) = w_B(c_1|\bar{c}) = w_B(c_2|\bar{c})$. Taken together, this implies that $(c_1, c_2, 1)$ satisfies the sufficient conditions in Theorem 1 for being a symmetric equilibrium. \square

Proof of Theorem 2. For $\bar{c} = (1, 1, 1)$ to be a symmetric equilibrium, a necessary condition is that the customer with unit waiting cost $c = 1$ must prefer to obtain service from the bid-based firm over the fixed-price firm. Using the expression for the total expected cost from (2.2), this yields

$$\int_0^1 w_B(t|\bar{c}) dt \leq w_F(\bar{c}) + P.$$

Note that $w_F(\bar{c}) = 1$, whereas, from (2.4), we have $w_B(t|\bar{c}) = \Gamma(\lambda - \lambda t)$ for all $t \in [0, 1]$.

Thus, a necessary condition for $\bar{c} = (1, 1, 1)$ to be a symmetric equilibrium is

$$P \geq \int_0^1 \Gamma(\lambda - \lambda t) dt - 1 = \frac{1}{\lambda} \left(\int_0^\lambda \Gamma(t) dt - \lambda \right) = P_{\max}(\lambda).$$

Next, suppose $P \geq P_{\max}(\lambda)$ with $\kappa = \lambda < k$. From Lemma 20, we obtain $s(z, x) < \lambda P$

for all $x \in (\lambda - \kappa, \nu)$ and $z \in [0, \lambda - x]$. (Note that $\lambda - \kappa = 0$.) This implies

$$z\Phi(x) + \lambda P > \int_x^{z+x} \Gamma(\lambda - t) dt, \quad \text{for all } x \in (0, \nu), \text{ and } z \in [0, \lambda - x].$$

Taking limits as $x \downarrow 0$, and letting $c = z/\lambda$, we obtain

$$c\Phi(0) + \lambda P \geq \int_0^c \Gamma(\lambda - \lambda t) dt, \quad \text{for all } c \in [0, 1].$$

For $\bar{c} = (1, 1, 1)$, we have $\Phi(0) = w_F(\bar{c})$, and $\Gamma(\lambda - \lambda t) = w_B(t|\bar{c})$ for all $t \in [0, 1]$.

Thus, the preceding equation implies

$$cw_F(\bar{c}) + \lambda P \geq \int_0^c w_B(t|\bar{c}) dt, \quad \text{for all } c \in [0, 1].$$

This implies that if all other customers follow the strategy \bar{c} , then it is preferable for a customer with unit waiting cost $c \in [0, 1]$ to obtain service from the bid-based firm over the fixed-price firm. Thus, $\bar{c} = (1, 1, 1)$ is a symmetric equilibrium. To obtain uniqueness, observe that if there exists another equilibrium $(c_1, c_2, 1)$ with $0 < c_1 < c_2 < 1$, then from the first statement of Theorem 15, we obtain that $P < P_{\max}(\lambda)$, which contradicts our assumption on P . \square

Proof of Theorem 3. From Lemma 22, we know that for each $P \in (0, P_{\max}(\lambda))$, there exists a unique solution x^* to the equation $\Psi(x, P) = x$ with $x \in (\lambda - \kappa, \nu)$. From the second statement of Theorem 15, we obtain that there exists a symmetric equilibrium $\bar{c} = (c_1, c_2, 1)$ with $0 < c_1 < c_2 < 1$, and $\lambda c_1 = v_1(x^*, P)$.

To obtain uniqueness, observe that for an equilibrium $\bar{C} = (C_1, C_2, 1)$ with $0 < C_1 < C_2 < 1$, by the first part of Theorem 15, we obtain that $X = \lambda(C_2 - C_1) \in (\lambda - \kappa, \nu)$ satisfies $\Psi(X, P) = X$ with $\lambda C_1 = v_1(X, P)$. But since x^* is the unique solution, we have $X = x^*$, and hence $\lambda C_1 = v_1(X, P) = v_1(x^*, P) = c_1$. Finally, note that since $P < P_{\max}(\lambda)$, from Theorem 2, there cannot be an equilibrium of the form $\bar{c} = (1, 1, 1)$. This proves the uniqueness of the symmetric equilibrium. \square

A.5 Proof of Theorem 4

Let $U_\lambda = (0, 1] \cap (0, \frac{n+k}{\lambda})$. Lemma 4 and Lemma 5, together with Theorem 2 and Theorem 3, imply that there exist functions $\mathcal{C}_i : U_\lambda \rightarrow [0, 1]$ for $i = 1, 2$, such that for each $u \in U_\lambda$, we have $\mathcal{C}_1(u) \leq \mathcal{C}_2(u) \leq u$, and the strategy $(\frac{\mathcal{C}_1(u)}{u}, \frac{\mathcal{C}_2(u)}{u}, 1)$ is the unique symmetric equilibrium of the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$. Define $\mathcal{A}(u) = \mathcal{C}_2(u) - \mathcal{C}_1(u)$. We have the following lemma:

Lemma 23. *The function $\mathcal{A}(u)$ is non-decreasing and continuous in u over U_λ . Further, $u - \mathcal{A}(u)$ is non-decreasing in u over U_λ .*

Proof. First, we consider the set of values of $u \in U_\lambda$ for which $\frac{P}{u} \geq P_{\text{max}}(\lambda u)$. Note that since $P_{\text{max}}(\lambda)$ is non-decreasing in λ , we obtain that this set is an interval $(0, u_0]$ for some $u_0 \in U_\lambda$. Since $(\frac{\mathcal{C}_1(u)}{u}, \frac{\mathcal{C}_2(u)}{u}, 1)$ is the unique symmetric equilibrium of the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$, we obtain from Theorem 2, that $\mathcal{C}_1(u) = \mathcal{C}_2(u) = u$ for all $u \in (0, u_0]$, and hence $\mathcal{A}(u) = 0$ for $u \in (0, u_0]$.

Now consider $u \in U_\lambda$ with $u > u_0$. Again, applying Theorem 3 to the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$, we obtain that $0 < \mathcal{C}_1(u) < \mathcal{C}_2(u) < u$, and hence $\mathcal{A}(u) > 0$, for all $u \in U_\lambda$ with $u > u_0$. From the necessary condition (ContT-P) for equilibrium for this system, we obtain for all $u \in U_\lambda$ with $u > u_0$,

$$\begin{aligned}
 P &= \int_0^{\mathcal{C}_1(u)} \Gamma(\lambda u - \lambda \mathcal{C}_2(u) + \lambda \mathcal{C}_1(u) - \lambda t) dt - \mathcal{C}_1(u) \Phi(\lambda (\mathcal{C}_2(u) - \mathcal{C}_1(u))) \\
 &= \int_0^{\mathcal{C}_1(u)} \Gamma(\lambda u - \lambda \mathcal{A}(u) - \lambda t) dt - \mathcal{C}_1(u) \Phi(\lambda \mathcal{A}(u)) \\
 &= \int_{u - \mathcal{A}(u) - \mathcal{C}_1(u)}^{u - \mathcal{A}(u)} \Gamma(\lambda t) dt - \mathcal{C}_1(u) \Phi(\lambda \mathcal{A}(u)). \tag{A.4}
 \end{aligned}$$

Similarly, from the necessary condition (ContW-P), we obtain for all $u \in U_\lambda$ with $u > u_0$,

$$\Phi(\lambda\mathcal{A}(u)) = \Gamma(\lambda(u - \mathcal{C}_2(u))) = \Gamma(\lambda(u - \mathcal{A}(u) - \mathcal{C}_1(u))). \quad (\text{A.5})$$

We begin with the proof of the first statement in lemma. Suppose, for the sake of arriving at a contradiction, we have $\mathcal{A}(u_1) > \mathcal{A}(u_2)$ for $u_1, u_2 \in U_\lambda$ with $u_2 > u_1 > u_0$. Since Γ and Φ are strictly increasing, from (A.5), we obtain $u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1) > u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)$. Since $u_2 - \mathcal{A}(u_2) > u_1 - \mathcal{A}(u_1)$, we obtain $\mathcal{C}_1(u_2) > \mathcal{C}_1(u_1)$. From this, we have

$$\begin{aligned} & \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt - \mathcal{C}_1(u_1) \Phi(\lambda\mathcal{A}(u_2)) \\ & > \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma(\lambda t) dt - \mathcal{C}_1(u_1) \Phi(\lambda\mathcal{A}(u_1)) = P, \end{aligned}$$

which yields,

$$\int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt > \mathcal{C}_1(u_1) \Phi(\lambda\mathcal{A}(u_2)) + P. \quad (\text{A.6})$$

Now, note that since $\mathcal{C}_1(u_2) > \mathcal{C}_1(u_1)$, under the symmetric equilibrium strategy $(\frac{\mathcal{C}_1(u_2)}{u_2}, \frac{\mathcal{C}_2(u_2)}{u_2}, 1)$ for the system $\text{SYS}_{\text{man}}(\lambda u_2, \frac{P}{u_2})$, the customer with unit waiting cost $\frac{\mathcal{C}_1(u_1)}{u_2}$ prefers to obtain service from the bid-based firm as opposed to the fixed-price firm. Using (2.2) and (2.4), the expected total cost of this customer in equilibrium is given by

$$\int_0^{\frac{\mathcal{C}_1(u_1)}{u_2}} \Gamma(\lambda u_2 - \lambda\mathcal{A}(u_2) - \lambda u_2 t) dt = \frac{1}{u_2} \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt.$$

On the other hand, the expected total cost of this customer if she obtains service from the fixed-price firm is given by $\frac{\mathcal{C}_1(u_1)}{u_2} \Phi(\lambda\mathcal{A}(u_2)) + \frac{P}{u_2}$. Thus, in equilibrium for the system $\text{SYS}_{\text{man}}(\lambda u_2, \frac{P}{u_2})$, we obtain

$$\frac{1}{u_2} \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt \leq \frac{\mathcal{C}_1(u_1)}{u_2} \Phi(\lambda\mathcal{A}(u_2)) + \frac{P}{u_2}.$$

This contradicts (A.6), and hence, we must have $\mathcal{A}(u_2) \geq \mathcal{A}(u_1) > 0$ for all $u_2 > u_1 > u_0$. Since $\mathcal{A}(u) = 0$ for $u \in U_\lambda$ with $u \leq u_0$, this completes the proof of the statement $\mathcal{A}(u)$ is non-decreasing over U_λ .

Next, we show that $u - \mathcal{A}(u)$ is non-decreasing over U_λ . Since $\mathcal{A}(u) = 0$ for $u \in (0, u_0]$, the statement holds trivially over $(0, u_0]$. If $u_0 = 1$, we are done. Hence, suppose $u_0 < 1$. Then, by continuity of $P_{\max}(\cdot)$, we obtain that $u_0 P_{\max}(\lambda u_0) = P$. Using the expression for $P_{\max}(\cdot)$ from (2.5), we obtain

$$\begin{aligned} P &= \frac{1}{\lambda} \int_0^{\lambda u_0} \Gamma(t) dt - u_0 \\ &= \int_0^{u_0} \Gamma(\lambda t) dt - u_0 \\ &= \int_{u_0 - \mathcal{A}(u_0) - \mathcal{C}_1(u_0)}^{u_0 - \mathcal{A}(u_0)} \Gamma(\lambda t) dt - \mathcal{C}_1(u_0) \Phi(\lambda \mathcal{A}(u_0)), \end{aligned}$$

where, in the last equality we use the fact that $\mathcal{A}(u_0) = 0$, $\mathcal{C}_1(u_0) = u_0$, and $\Phi(0) = 1$.

This implies that (A.4) holds for all $u \in U_\lambda$ with $u \geq u_0$, when $u_0 < 1$. Similarly, it is straightforward to verify that (A.5) also holds for all $u \in U_\lambda$ with $u \geq u_0$. Thus,

for all $u_1, u_2 \in U_\lambda$ with $u_2 > u_1 \geq u_0$, we have

$$\begin{aligned} 0 &= \left(\int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt - \mathcal{C}_1(u_2) \Phi(\lambda \mathcal{A}(u_2)) \right) \\ &\quad - \left(\int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma(\lambda t) dt - \mathcal{C}_1(u_1) \Phi(\lambda \mathcal{A}(u_1)) \right) \\ &= \int_{u_1 - \mathcal{A}(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt - \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)} \Gamma(\lambda t) dt \\ &\quad - \mathcal{C}_1(u_2) \Phi(\lambda \mathcal{A}(u_2)) + \mathcal{C}_1(u_1) \Phi(\lambda \mathcal{A}(u_1)) \\ &= \int_{u_1 - \mathcal{A}(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt - \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)} \Gamma(\lambda t) dt - \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)} \Gamma(\lambda t) dt \\ &\quad - (\mathcal{C}_1(u_2) - \mathcal{C}_1(u_1)) \Phi(\lambda \mathcal{A}(u_2)) - \mathcal{C}_1(u_1) (\Phi(\lambda \mathcal{A}(u_2)) - \Phi(\lambda \mathcal{A}(u_1))). \end{aligned}$$

This implies,

$$\begin{aligned} \int_{u_1 - \mathcal{A}(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma(\lambda t) dt - \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)} \Gamma(\lambda t) dt &= \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)} \Gamma(\lambda t) dt \\ &+ (\mathcal{C}_1(u_2) - \mathcal{C}_1(u_1)) \Phi(\lambda \mathcal{A}(u_2)) \\ &+ \mathcal{C}_1(u_1) (\Phi(\lambda \mathcal{A}(u_2)) - \Phi(\lambda \mathcal{A}(u_1))). \end{aligned}$$

After some algebra and rearranging, we obtain

$$\begin{aligned} \int_{u_1 - \mathcal{A}(u_1)}^{u_2 - \mathcal{A}(u_2)} (\Gamma(\lambda t) - \Gamma(\lambda t - \mathcal{C}_1(u_1))) dt &= \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)} \Gamma(\lambda t) - \Phi(\lambda \mathcal{A}(u_2)) dt \\ &+ \mathcal{C}_1(u_2) (\Phi(\lambda \mathcal{A}(u_2)) - \Phi(\lambda \mathcal{A}(u_1))). \end{aligned}$$

Now note that since $\Gamma(\cdot)$ is increasing, from (A.5), we obtain that $\Gamma(\lambda t) > \Phi(\lambda \mathcal{A}(u_2))$ for $t > u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)$, and $\Gamma(\lambda t) < \Phi(\lambda \mathcal{A}(u_2))$ for $t < u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)$. Thus the integral on the right hand side is non-negative. Further, since $\mathcal{A}(u)$ is non-decreasing and Φ is strictly increasing, the second term on the right hand side is also non-negative. This implies that the left hand side is non-negative. Since $\Gamma(\lambda t) - \Gamma(\lambda t - \mathcal{C}_1(u_1)) \geq 0$, this implies that $u_2 - \mathcal{A}(u_2) \geq u_1 - \mathcal{A}(u_1)$ for $u_2 > u_1 \geq u_0$. Hence $u - \mathcal{A}(u)$ is non-decreasing over all $u \in U_\lambda$ with $u \geq u_0$. Since $\mathcal{A}(u) = 0$ for $u \in U_\lambda$ with $u \leq u_0$, the statement extends to all $u \in (0, 1]$.

Finally, we show that $\mathcal{A}(u)$ is continuous over $u \in U_\lambda$. For each $u \in U_\lambda$ with $u \geq u_0$, we obtain that $\mathcal{A}(u)$ and $\mathcal{C}_1(u)$ satisfy (A.4) and (A.5). By continuity of both sides of these equations in u , we obtain that for a sequence $u_n \rightarrow u_\infty \in U_\lambda$ with $u_\infty \geq u_0$, the limits $\lim_{n \rightarrow \infty} \mathcal{A}(u_n)$ and $\lim_{n \rightarrow \infty} \mathcal{C}_1(u_n)$ (along a subsequence if necessary for existence of the limits) also satisfy the same equations for $u = u_\infty$. However, since (A.4) and (A.5) also constitute the sufficient conditions for equilibrium (from Theorem 1), we obtain that $\mathcal{A}(u_\infty) = \lim_{n \rightarrow \infty} \mathcal{A}(u_n)$ and $\mathcal{C}_1(u_\infty) = \lim_{n \rightarrow \infty} \mathcal{C}_1(u_n)$. This implies that

$\mathcal{A}(u)$ is continuous over $u \in U_\lambda$ with $u \geq u_0$. Observe that $\mathcal{A}(u) = 0$ for $u \in (0, u_0]$. Taken together, this implies that $\mathcal{A}(u)$ is continuous over $u \in U_\lambda$. \square

Proof of Theorem 4. From Lemma 4, we obtain that to show the existence of a symmetric equilibrium for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$, it suffices to show that there exists a $u \in U_\lambda$ such that the condition (IND) holds for the strategy $(\mathcal{C}_1(u), \mathcal{C}_2(u), u)$.

For $u \in U_\lambda$, observe that the total expected cost of a customer with unit waiting cost u under the strategy $\bar{\mathcal{C}}(u) = (\mathcal{C}_1(u), \mathcal{C}_2(u), u)$ is given by

$$\begin{aligned} TC(u) &= \int_0^{\mathcal{C}_1(u)} \Gamma(\lambda u - \lambda \mathcal{A}(u) - \lambda t) dt + \Phi(\lambda \mathcal{A}(u)) (\mathcal{C}_2(u) - \mathcal{C}_1(u)) + \int_{\mathcal{C}_2(u)}^u \Gamma(\lambda u - \lambda t) dt \\ &= \int_0^{u - \mathcal{A}(u)} \Gamma(\lambda t) dt + \mathcal{A}(u) \Phi(\lambda \mathcal{A}(u)). \end{aligned}$$

Now, from Lemma 23, we obtain that both $\mathcal{A}(u)$ and $u - \mathcal{A}(u)$ are non-decreasing and continuous over U_λ . Since one of these two functions must strictly increase at any u , we obtain that $TC(u)$ is strictly increasing and continuous over $u \in U_\lambda$. Thus, for all $V \geq 0$, there exists a unique $u = u(\lambda, P, V) \in U_\lambda$ such that either $TC(u) = V$ with $u \leq 1$, or $TC(u) \leq V$ and $u = 1$. Note that this is exactly the condition (IND) for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$. Thus, we obtain that the strategy $\bar{\mathcal{C}}(u)$ for $u = u(\lambda, P, V)$ satisfies the (necessary and) sufficient equilibrium conditions for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$, and hence constitutes the unique symmetric equilibrium for the system.

Define $\Delta_\lambda(u) = \int_0^u \Gamma(\lambda t) dt$. Observe that Δ_λ is a strictly increasing function in u over $(0, 1] \cap (0, k/\lambda)$. Let $u_\lambda = u_\lambda(V) \triangleq \min\{\Delta_\lambda^{-1}(V), 1\} > 0$ for $V > 0$ and $P(\lambda, V) = \Delta_\lambda(u_\lambda(V)) - u_\lambda(V) = u_\lambda(V) P_{\max}(\lambda u_\lambda(V))$. Now, if $P \geq P(\lambda, V) =$

$u_\lambda P_{\max}(\lambda u_\lambda)$, then from Theorem 2, we obtain that $(1, 1, 1)$ is the unique symmetric equilibrium for the system $\text{SYS}_{\text{man}}(\lambda u_\lambda, \frac{P}{u_\lambda})$. Observe that $u_\lambda \leq \Delta^{-1}(V) \leq k/\lambda$, and hence $\lambda u_\lambda < k < n + k$. Furthermore, using the definition of u_λ , it is straightforward to verify that for this strategy, we have $\Delta_\lambda(u_\lambda) = TC(u_\lambda)$ and that the condition (IND) also holds. Hence, from Lemma 4, we obtain that $(u_\lambda, u_\lambda, u_\lambda)$ constitutes the unique symmetric equilibrium for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$ for $P \geq P(\lambda, V)$. This implies that the arrival rate of the customers to the fixed-price firm is zero.

Conversely, if (u, u, u) is a symmetric equilibrium for the system $\text{SYS}_{\text{op}}(\lambda, P, V)$, then we have $\Delta_\lambda(u) = TC(u)$ for this strategy, and condition (IND) implies $u = u_\lambda$. Also, from Lemma 5, we obtain that $(1, 1, 1)$ is a symmetric equilibrium for the system $\text{SYS}_{\text{man}}(\lambda u, \frac{P}{u})$, and hence from Theorem 2, we obtain $P \geq u P_{\max}(\lambda u) = P(\lambda, V)$. \square

A.6 Proofs in Section 2.5

Proof of Theorem 6. Since the strategy space is compact, we only need to show that the revenue of each operator is continuous in both P and r . Then the existence of a mixed strategy equilibrium follows from Glicksberg's theorem [30].

To show the continuity of firms' revenue, observe that both firms' revenue is continuous in the equilibrium thresholds $\bar{c} = (c_1, c_2, c_\ell)$. Thus it suffices to show that \bar{c} is continuous in P and r . We show this for $P > r$, and the case with $P \leq r$ can be shown using the same argument.

By the discussion in subsection 2.4.3, a system with positive r can be mapped to a system with $r = 0$, fixed price $P - r$, and value $V - r$. Thus, it suffices to show that in any system with zero reserve price, \bar{c} is continuous in P and V .

First, consider $\text{SYS}_{\text{man}}(\lambda, P)$. Note that by Lemma 18, $s(z_{\text{max}}(x), x)$ is continuous and strictly decreasing, thus $F(P) = \sup\{\lambda - \kappa < x < \nu : s(z_{\text{max}}(x), x) > \lambda P\}$ is continuous in P . Then by Lemma 15, α in $\text{SYS}_{\text{man}}(\lambda, P)$ is given by $F(P)$ for $P < P_{\text{max}}(\lambda)$, and zero otherwise. Thus α is continuous in P in $\text{SYS}_{\text{man}}(\lambda, P)$, for fixed λ . It then follows from condition (ContW-P) that c_1 is continuous in P in $\text{SYS}_{\text{man}}(\lambda, P)$.

Now consider $\text{SYS}_{\text{op}}(\lambda, P, V)$. Recall from the proof of Theorem 4, $TC(u)$, which is the total expected cost of a customer with unit waiting cost u under the strategy $\bar{\mathcal{C}}(u) = (\mathcal{C}_1(u), \mathcal{C}_2(u), u)$, is continuous and strictly increasing in u . Thus, $u(\lambda, P, V)$ that satisfies condition (IND) ($TC(u) = V$ with $u \leq 1$, or $TC(u) \leq V$ with $u = 1$) is continuous in V for fixed λ and P . On the other hand, since c_1 and α are continuous in P in $\text{SYS}_{\text{man}}(\lambda, P)$, it follows that $\mathcal{A}(u)$ is continuous in P for fixed u . Thus, $TC(u)$ is continuous in P . It then follows from monotonicity of $TC(u)$ in u that $u(\lambda, P, V)$ that satisfies condition (IND) is continuous in both P and V . The result then follows immediately from the observation that $\mathcal{C}_1(u), \mathcal{C}_2(u)$ are continuous in both u and P . \square

The rest of the results in this section use the assumption on the expected waiting time expressions, under which we can simplify the conditions (ContW-P), (ContT-P),

and (IND). In particular, with the expected waiting time expressions and under the assumption that $r = 0$, the conditions (ContW-P), (ContT-P) become

$$\frac{1}{1 - \rho_B} - \frac{1}{1 - \rho_B + \rho_B \frac{c_1}{1-\alpha}} - \frac{1}{q_B} \frac{c_1}{(1 - \rho_B + \rho_B \frac{c_1}{1-\alpha})^2} = \lambda P, \quad (\text{ContT-P}')$$

$$\frac{1}{q_B(1 - \rho_B + \rho_B \frac{c_1}{1-\alpha})^2} - \frac{1}{q_B} = \frac{\rho_F}{1 - \rho_F} \frac{1}{q_F}, \quad (\text{ContW-P}')$$

where $\rho_s = (1 - \alpha)/(q_B)$, and $\rho_F = \alpha/(q_F)$. Condition (Pref-BID) can be obtained similarly. And the expected total cost of the customer with unit waiting cost c_ℓ is given by

$$\int_0^{c_\ell} w(t; \bar{c}) dt = \frac{1}{\lambda - \frac{c_\ell - \alpha}{q_B} \lambda} - \frac{1}{\lambda} + c_\ell \left(1 - \frac{1}{k}\right) + \frac{\alpha}{k} + \left(\frac{1}{1 - \rho_F} - 1\right) \frac{\alpha}{n}, \quad (\text{A.7})$$

which is used in condition (IND).

Proof of Lemma 6. We prove the first case here, and the other cases can be shown using the same analysis.

We start by computing the limiting thresholds in $\text{SYS}_{\text{man}}(\lambda, P)$ as λ approaches infinity, when $r = 0$. Note that as λ approaches infinity, the right hand side of condition (ContT-P') tends to infinity, thus the left hand side also tends to infinity. As a result, $\frac{1}{1 - \rho_B}$ tends to infinity and thus α^λ tends to $1 - q_B$, as λ approaches infinity. Substituting the limiting value of α^λ into (ContW-P'), we obtain that c_1^λ tends to $\sqrt{q_B / \left(\frac{1}{q_F + q_B - 1} - \frac{1}{q_F} + \frac{1}{q_B}\right)}$ as λ approaches infinity. Thus, using equation (A.7) with $c_\ell = 1$, the expected total cost of the customer with unit waiting cost 1 is given by $P + 1$. Using the mapping considered in Section 2.4.2, we can show that in the system $\text{SYS}_{\text{op}}(\lambda, P, V)$ (with $r = 0$), the expected total cost of the customer with unit cost u when all customers follows the strategy

given by $\bar{\mathcal{C}}(u) = (\mathcal{C}_1(u), \mathcal{C}_2(u), u)$ is $P + c_\ell$. Thus, we obtain from condition (IND) that $c_\ell^\infty(P, r) = \min\{V - P, 1\}$. It immediately follows from the mapping that c_1 tends to $c_1^\infty(P, r) = c_\ell^\infty(P, r) \sqrt{q_B / \left(\frac{1}{q_F + q_B - c_\ell^\infty(P, r)} - \frac{1}{q_F} + \frac{1}{q_B} \right)}$, and α tends to $\alpha^\infty(P, r) = c_\ell^\infty(P, r) - q_B$, as λ approaches infinity.

The limiting revenue of the fixed-price firm per arrival rate can be obtained immediately by multiplying P and the proportion of the customers obtaining service in the fixed-price firm. Thus $R_F^\infty(P, r) = (c_\ell^\infty(P, r) - q_B) P$. To compute the limiting revenue of the bid-based firm, we consider the expected total revenue of the two firms in the $\text{SYS}_{\text{man}}(\lambda, P)$, which is

$$\begin{aligned} & \int_0^1 \left[\int_0^c w(t) dt \right] dc - \int_0^1 cw(c) dc \\ &= -2\alpha^\lambda P + \frac{1}{\lambda} \left[\frac{\alpha^\lambda(1-\alpha^\lambda)}{q_B(1-\frac{1-\alpha^\lambda}{q_B} + \frac{c_1^\lambda}{q_B})^2} + 1 + \frac{1+2\alpha^\lambda}{1-\frac{1-\alpha^\lambda}{q_B}} + 2q_B \log\left(1 - \frac{1-\alpha^\lambda}{q_B}\right) \right] \quad (\text{A.8}) \end{aligned}$$

Note that by the limit of α and c_1 , $\frac{\alpha^\lambda(1-\alpha^\lambda)}{q_B(1-\frac{1-\alpha^\lambda}{q_B})} + 1$ is finite in the limit, and thus the limit of $\frac{1}{\lambda} \left[\frac{\alpha^\lambda(1-\alpha^\lambda)}{q_B(1-\frac{1-\alpha^\lambda}{q_B} + \frac{c_1^\lambda}{q_B})^2} + 1 \right]$ is zero. On the other hand, by condition (ContT-P'), we have

$$\frac{1}{\lambda} \frac{1}{1-\frac{1-\alpha^\lambda}{q_B}} = P + \frac{1}{\lambda} \left[\frac{1}{1-\rho_B + \rho_B \frac{c_1^\lambda}{1-\alpha^\lambda}} + \frac{1}{q_B} \frac{c_1^\lambda}{(1-\rho_B + \rho_B \frac{c_1^\lambda}{1-\alpha^\lambda})^2} \right],$$

which tends to P in the limit, as the term in the square bracket has a finite limit.

And since $\rho_B \rightarrow 1$, the limit of $[\log(1 - \rho_B)] / \frac{1}{1-\rho_B}$ is zero.

Therefore, the limit of (A.8) is P , and in $\text{SYS}_{\text{man}}(\lambda, P)$, the limiting revenue of the bid-based firm is $q_B P$. By the mapping considered in Section 2.4.2, its limiting

revenue in $\text{SYS}_{\text{op}}(\lambda, P, V)$ (with $r = 0$) is also $q_B P$.

Thus we proved the first case when $r = 0$. The first case when $r > 0$ then immediately follows from the mapping discussed in subsection 2.4.3. The other three cases can be shown using the same techniques. \square

Proof of Theorem 7. We only need to show that the firms' revenue is continuous in λ . This proof is analogous to the proof of Lemma 6, so we only briefly outline the argument here.

As shown in the proof of Lemma 6, the thresholds c_1, c_2 that solves $\text{SYS}_{\text{man}}(\lambda, P)$ is continuous in P . Also observe that conditions (ContT-P') and (ContW-P') only depend on λ and P through λP , thus the thresholds that solves $\text{SYS}_{\text{man}}(\lambda, P)$ is also continuous in λ . Moreover, from equation (A.7) we observe that the expected total cost of the customer with unit waiting cost u when all customers follows the strategy $\bar{C}(u)$ is continuous and strictly decreasing in λ . Recall that it is also continuous and strictly monotone in u . Thus, $u(\lambda, P, V)$ that satisfies condition (IND) is continuous in λ . So far we have shown that the equilibrium thresholds are continuous in λ , and thus the revenues are also continuous in λ . \square

Proof of Lemma 7. Lemma 6 gives the payoffs in the limiting game R_F^∞ and R_B^∞ as functions of (P, r) , from which we can calculate the best response functions $P^{BR}(r)$ and $r^{BR}(P)$. In particular, if $V \geq 2 - q_B$, the best response of the fixed-price operator

to the reserve price set by the bid-based firm is given by

$$P^{BR}(r) = \begin{cases} V - 1 & \text{if } r \leq (V - 1) \frac{1 - q_B}{q_F} \\ r & \text{if } r > (V - 1) \frac{1 - q_B}{q_F}. \end{cases}$$

And when $V \geq 2 - q_B$ the best response of the bid-based firm to the fixed-price firm's fixed price is given by

$$r^{BR}(P) \in \begin{cases} V - 1 & \text{if } P \leq (V - 1) \frac{1 - q_F}{q_B} \\ [0, P) & \text{if } P > (V - 1) \frac{1 - q_F}{q_B}. \end{cases}$$

Finding the fixed point of the best response functions, we prove the first case of this lemma. The second case can be shown using the same approach. \square

Proof of Theorem 8. We first consider the optimization problem under collusion. The firms choose (P, r) to maximize the total payoff, which is given by

$$R_F^\infty(P, r) + R_B^\infty(P, r) = \begin{cases} P & \text{if } P \leq V - 1, r < P \\ (V - P)P & \text{if } V > P \geq V - 1, r < P \\ Pq_F + r(1 - q_F) & \text{if } P \leq r \leq V - 1 \\ Pq_F + r(V - r - q_F) & \text{if } r > V - 1, r \geq P. \end{cases}$$

Maximizing the quantity above, we get the policy under collusion given by

$$(P, r) \in \begin{cases} \{(P, r) : P = V - 1, r \leq P\} & \text{if } V \geq 2 \\ \{(P, r) : P = V/2, r \leq P\} & \text{if } V < 2. \end{cases}$$

And the total payoff under collusion is $V - 1$ if $V \geq 2$, and $V^2/4$ if $V < 2$.

By Lemmas 6 and 7, the total payoff in Nash equilibrium of the limiting game is

$V - 1$ if $V \geq 2 - q_B$, and $\frac{V^2 - q_B^2}{4}$ if $V < 2 - q_B$. Thus, the price of stability is given by

$$\begin{cases} \frac{V-1}{V-1} = 1 & \text{if } V \geq 2 \\ \frac{V^2/4}{V-1} & \text{if } 2 - q_B \leq V < 2 \\ \frac{V^2/4}{(V^2 - q_B^2)/4} = \frac{1}{1 - q_B^2/V^2} & \text{if } V < 2 - q_B. \end{cases}$$

It is easy to verify that the price of stability is non-increasing in V . Moreover, when $q_B \geq 0.5$, the price of stability at $V = 1.5$ is given by $(V^2)/4(V - 1) = 9/8$. When $q_B < 0.5$, the price of stability at $V = 1.5$ is $1/(1 - q_B^2/V^2) < 1/(1 - 1.5^2/V^2) = 9/8$. The theorem follows from this.

□

A.7 Extension to non-preemptive queues for the bid-based firm

To extend our analysis in Sections 2.3 and 2.4 to the case where the bid-based firm operates a non-preemptive queue, we need to change our assumptions of $\Gamma(\cdot)$, so that it not only depends on the arrival rate of the customers with higher bids, but also depends on the arrival rate of all customers to the bid-based queue. More specifically, we let $\Gamma_y(x)$ to be the expected waiting time in the bid-based firm when the arrival rate of the customers with higher bids is x , and the arrival rate of all customers to the bid-based firm is y . Similar to our previous assumptions, we assume that $\Gamma_y(x)$ is continuous and strictly increasing in both x and y for $x \leq y < k$, and that $\Gamma_y(0) > 1$.

It is straightforward to verify that any $G/M/k$ non-preemptive queues satisfy these assumptions.

In Section 2.3: Almost all analysis in Section 2.3 follow in this non-preemptive case. The only change is that in Theorem 1, the strict inequality in $0 < c_1 < c_2 < c_\ell \leq 1$ has to be changed to a less than or equal to as in $0 < c_1 < c_2 \leq c_\ell \leq 1$. In other words, it is possible that in equilibrium, the fixed-price firm is non-empty yet there is no customer bidding higher than P in the bid-based firm. This is because when queues are non-preemptive, $\Gamma_y(0) > 1$, and thus there exists some ϵ_0 such that $\Gamma_y(0) = \Phi(\epsilon_0)$. Thus there may exist an equilibrium in which the fixed-price firm has arrival rate less than ϵ_0 and $c_2 = c_\ell$.

In Section 2.4: Multiple changes need to be made to Section 2.4. The idea is to add appropriate subscripts to all occurrence of $\Gamma(\cdot)$ to represent the total arrival rate to the bid-based firm. Some proofs need to be carefully modified, especially when there are inequalities comparing the Γ functions under different total arrival rates. Moreover, the possibility of an equilibrium in which $0 < c_1 < c_2 = c_\ell \leq 1$ as discussed above leads to modification of condition (ContW-P'), as in such equilibrium the expected waiting time of the highest priority customer in the bid-based firm can be strictly larger than the expected waiting time in fixed-price firm. We discuss the changes in the main text and the appendix separately, as follows.

In the main text of Section 2.4: In equation (2.4) which connects $w_B(\cdot)$ with $\Gamma(\cdot)$, the subscript $\lambda - \alpha\lambda$ needs to be added to both occurrences. In equation (2.5) where we define $P_{\max}(\lambda)$, the subscript λ needs to be added to Γ . In Theorem 3, we

need to change $0 < c_1 < c_2 < 1$ to $0 < c_1 < c_2 \leq 1$. In equation (2.6), the subscript $\lambda - \alpha\lambda$ needs to be added to Γ . And in all occurrences of Γ in subsection 2.4.2, the subscript $\lambda u - \mathcal{A}(u)$ needs to be added.

Moreover, condition (ContW-P') needs to be changed to

$$c_2 < c_\ell, \quad w_B(c_2|\bar{c}) = w_B(c_1|\bar{c}) = w_F(\bar{c}), \quad \text{OR} \quad c_2 = c_\ell, \quad w_B(c_1|\bar{c}) \geq w_F(\bar{c}).$$

In Appendix A.3: Almost all analysis in Appendix A.3 follow without modification. We only need to add appropriate subscripts to the equation in the proof of Lemma 4 involving Γ , and it becomes

$$\begin{aligned} \int_0^{c_1 u} w_B(t|\bar{c}(u)) dt &= \int_0^{c_1 u} \Gamma_{\frac{\lambda}{u}u - \frac{\lambda}{u}(c_2 u - c_1 u)} \left(\frac{\lambda}{u}u - \frac{\lambda}{u}(c_2 u - c_1 u) - \frac{\lambda}{u}t \right) dt \\ &= u \int_0^{c_1} \Gamma_{\lambda - \lambda(c_2 - c_1)}(\lambda - \lambda(c_2 - c_1) - \lambda t) dt = u \int_0^{c_1} w_B(t|\bar{c}) dt. \end{aligned}$$

In Appendix A.4: First, define η as

$$\eta \triangleq \sup\{\lambda - \kappa < x < \nu : \Gamma_{\lambda-x}(0) > \Phi(x)\}.$$

Add subscript $\lambda - x$ to Γ in the definition of ξ and $s(z, x)$.

The statement in Lemma 17 becomes “*The function $z_{\max}(x)$ is continuous over $x \in (\lambda - \kappa, \nu)$. Further, $z_{\max}(\cdot) = \lambda - x$ for all $x \in (\lambda - \kappa, \eta)$, $z_{\max}(\cdot)$ is strictly decreasing over (η, ξ) and $z_{\max}(x) = 0$ for all $x \in [\xi, \nu)$. Moreover, $z_{\max}(x) = \lambda - x - \Gamma_{\lambda-x}^{-1}(\Phi(x)) \in (0, \lambda - x)$ for $x \in (\eta, \xi)$, and $z_{\max}(\eta) = \lim_{x \downarrow \eta} z_{\max}(x) = \lambda - \eta$.”*

The proof can be done using the same argument as the original proof.

In the proof of Lemma 18, the derivative needs to be changed to

$$\frac{ds(z_{\max}(x), x)}{dx} = \Gamma_{\lambda-x}(\lambda - z_{\max}(x) - x) - \Gamma_{\lambda-x}(\lambda - x) - z_{\max}(x)\Phi'(x)$$

$$+ \int_x^{z_{\max}(x)+x} \frac{\partial \Gamma_{\lambda-x}(\lambda-t)}{\partial x} dt.$$

By monotonicity of Γ , $\frac{\partial \Gamma_{\lambda-x}(\lambda-t)}{\partial x} < 0$. Thus the original result still follows.

Lemma 19 and its proof hold after adding subscript λ to Γ in the first two displayed equations, and adding subscript $\lambda - x$ to Γ in the last displayed equation.

Lemmas 20, 21, Theorem 15 and their proofs hold after adding subscript $\lambda - x$ to all occurrences of Γ .

In the proof of Lemma 22, add subscript $\lambda - F(P)$. More importantly, note that $\Gamma_{\lambda-F(P)}^{-1}(\Phi(F(P)))$ may not be well defined when $\Phi(F(P)) < \Gamma_{\lambda-F(P)}(0)$. In those cases, we let $v_1(F(P), P) = \lambda - F(P)$, and it corresponds to an equilibrium in which $c_1 < c_2 = c_\ell$. We can verify that all other statements in the proof hold after we make this change.

The proof of Theorem 2 holds after adding subscript λ .

In Appendix A.5: The proof of Theorem 4 holds after adding the subscript $\lambda u - \lambda \mathcal{A}(u)$ to all occurrences of Γ , and changing the definition of $\Delta_\lambda(u)$ to $\Delta_\lambda(u) = \int_0^u \Gamma_{u\lambda}(\lambda t) dt$.

The first part of the proof of Lemma 23 holds after adding the subscripts $\lambda u - \lambda \mathcal{A}(u)$, $\lambda u_1 - \lambda \mathcal{A}(u_1)$, $\lambda u_2 - \lambda \mathcal{A}(u_2)$ or λu_0 (which one to add should be clear given context). However, to show that $u - \mathcal{A}(u)$ is non-decreasing in u , we need a new proof as follows. For all $u_1, u_2 \in U_\lambda$ with $u_2 > u_1 \geq u_0$, we have

$$0 = \left(\int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)}^{u_2 - \mathcal{A}(u_2)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t) dt - \mathcal{C}_1(u_2) \Phi(\lambda \mathcal{A}(u_2)) \right)$$

$$\begin{aligned}
& - \left(\int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t) dt - \mathcal{C}_1(u_1) \Phi(\lambda \mathcal{A}(u_1)) \right) \\
= & \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)}^{u_2 - \mathcal{A}(u_2)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t) dt - \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t) dt \\
& - (\mathcal{C}_1(u_2) - \mathcal{C}_1(u_1)) \Phi(\lambda \mathcal{A}(u_2)) - \mathcal{C}_1(u_1) (\Phi(\lambda \mathcal{A}(u_2)) - \Phi(\lambda \mathcal{A}(u_1))).
\end{aligned}$$

This implies,

$$\begin{aligned}
& \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t) dt - \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t) dt \\
= & \int_{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_1)}^{u_2 - \mathcal{A}(u_2) - \mathcal{C}_1(u_2)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t) dt + (\mathcal{C}_1(u_2) - \mathcal{C}_1(u_1)) \Phi(\lambda \mathcal{A}(u_2)) \\
& + \mathcal{C}_1(u_1) (\Phi(\lambda \mathcal{A}(u_2)) - \Phi(\lambda \mathcal{A}(u_1))).
\end{aligned}$$

Using the same argument as in the original proof, we can show that the right hand side of the equation above is non-negative. Thus, the left hand side is also non-negative.

After rearranging the left hand side, we have

$$\int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t + (u_2 - \mathcal{A}(u_2)) - (u_1 - \mathcal{A}(u_1))) - \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t) dt \geq 0.$$

Suppose $u_2 - \mathcal{A}(u_2) < u_1 - \mathcal{A}(u_1)$, we have by monotonicity of $\Gamma_y(x)$ in both x and y , that

$$\begin{aligned}
& \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t + (u_2 - \mathcal{A}(u_2)) - (u_1 - \mathcal{A}(u_1))) < \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t) < \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t). \\
\text{Thus, } & \int_{u_1 - \mathcal{A}(u_1) - \mathcal{C}_1(u_1)}^{u_1 - \mathcal{A}(u_1)} \Gamma_{\lambda u_2 - \lambda \mathcal{A}(u_2)}(\lambda t + (u_2 - \mathcal{A}(u_2)) - (u_1 - \mathcal{A}(u_1))) - \Gamma_{\lambda u_1 - \lambda \mathcal{A}(u_1)}(\lambda t) dt < \\
& 0, \text{ which contradicts with the previous inequality.}
\end{aligned}$$

Therefore $u_2 - \mathcal{A}(u_2) \geq u_1 - \mathcal{A}(u_1)$, as desired.

APPENDIX B

APPENDIX FOR CHAPTER 3

B.1 Omitted Results

In the following lemma, we show elementary properties for the function $G_i(\cdot, \mathbf{p}_{-i})$. We use these properties throughout the paper.

Lemma 24. *For fixed \mathbf{p}_{-i} , let $\nu^* = \inf\{\nu \in \mathfrak{R}_+ : \mathcal{T}_i(\nu, \mathbf{p}_{-i}) = \emptyset\}$. The function $G_i(\cdot, \mathbf{p}_{-i})$ satisfies the following properties.*

- (a) *The function $G_i(\nu, \mathbf{p}_{-i})$ is continuous in $\nu \in (0, \infty)$.*
- (b) *The function $G_i(\nu, \mathbf{p}_{-i})$ is strictly decreasing in $\nu \in [0, \nu^*)$ and constant in $\nu \in [\nu^*, \infty)$ satisfying $G_i(\nu, \mathbf{p}_{-i}) = -2c_i$ for all $\nu \in [\nu^*, \infty)$.*
- (c) *There exists a unique $\hat{\nu} \in [0, \infty)$ satisfying $G_i(\hat{\nu}, \mathbf{p}_{-i}) = 0$.*

Proof. First, we show Part a. Fix $\nu > 0$ and $\epsilon > 0$ small enough that $\nu - \epsilon > 0$. The definition of $\mathcal{T}_i(\nu, \mathbf{p}_{-i})$ implies that $\mathcal{T}_i(\nu - \epsilon, \mathbf{p}_{-i}) \supseteq \mathcal{T}_i(\nu, \mathbf{p}_{-i})$. Also, if $t \in \mathcal{T}_i(\nu - \epsilon, \mathbf{p}_{-i}) \setminus \mathcal{T}_i(\nu, \mathbf{p}_{-i})$, then we have $\nu \geq \frac{\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t}{\beta_i^t} > \nu - \epsilon$, which implies that we have $0 \geq \alpha_i^t - \beta_i^t \nu + \sum_{j \neq i} \gamma_{i,j}^t p_j^t > -\beta_i^t \epsilon$. For notational brevity, we let $\mathcal{T}_i^+ = \mathcal{T}_i(\nu, \mathbf{p}_{-i})$, $\mathcal{T}_i^- = \mathcal{T}_i(\nu - \epsilon, \mathbf{p}_{-i})$ and $\mathcal{U}_i = \mathcal{T}_i^- \setminus \mathcal{T}_i^+$ so that $\mathcal{T}_i^- = \mathcal{T}_i^+ \cup \mathcal{U}_i$. Noting the definition of $G_i(\nu, \mathbf{p}_{-i})$, we have

$$G_i(\nu, \mathbf{p}_{-i}) - G_i(\nu - \epsilon, \mathbf{p}_{-i})$$

$$\begin{aligned}
&= \left\{ \sum_{t \in \mathcal{T}_i^+} \left(\alpha_i^t - \beta_i^t \nu + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) - 2c_i \right\} - \left\{ \sum_{t \in \mathcal{T}_i^-} \left(\alpha_i^t - \beta_i^t (\nu - \epsilon) + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) - 2c_i \right\} \\
&= - \sum_{t \in \mathcal{U}_i} \left(\alpha_i^t - \beta_i^t \nu + \sum_{j \neq i} \gamma_{i,j}^t p_j^t \right) - \sum_{t \in \mathcal{T}_i^+ \cup \mathcal{U}_i} \beta_i^t \epsilon,
\end{aligned}$$

Since $0 \geq \alpha_i^t - \beta_i^t \nu + \sum_{j \neq i} \gamma_{i,j}^t p_j^t > -\beta_i^t \epsilon$ for all $t \in \mathcal{U}_i$, the equality above yields $-\sum_{t \in \mathcal{T}_i^+ \cup \mathcal{U}_i} \beta_i^t \epsilon \leq G_i(\nu, \mathbf{p}_{-i}) - G_i(\nu - \epsilon, \mathbf{p}_{-i}) \leq -\sum_{t \in \mathcal{T}_i^+} \beta_i^t \epsilon$, so that $G_i(\nu, \mathbf{p}_{-i})$ is continuous in $\nu \in (0, \infty)$.

Second, we show Part b. Fix $\nu \in (0, \nu^*)$, in which case, by the definition of ν^* , we have $\mathcal{T}_i(\nu, \mathbf{p}_{-i}) \neq \emptyset$. In the proof of Part a, we show that $G_i(\nu, \mathbf{p}_{-i}) - G_i(\nu - \epsilon, \mathbf{p}_{-i}) \leq -\sum_{t \in \mathcal{T}_i(\nu, \mathbf{p}_{-i})} \beta_i^t \epsilon$ for all $\epsilon > 0$ small enough that $\nu - \epsilon > 0$. Since $\beta_i^t > 0$ for all $t \in T$ and $\mathcal{T}_i(\nu, \mathbf{p}_{-i}) \neq \emptyset$, the last inequality implies that $G_i(\nu - \epsilon, \mathbf{p}_{-i}) > G_i(\nu, \mathbf{p}_{-i})$ for all $\nu \in (0, \nu^*)$ and $\epsilon > 0$ small enough that $\nu - \epsilon > 0$. Also, noting that $\alpha_i^t > 0$ and $\beta_i^t > 0$, by the definition of $\mathcal{T}_i(\nu, \mathbf{p}_{-i})$, we have $\mathcal{T}_i(\epsilon, \mathbf{p}_{-i}) = T$ for small enough $\epsilon > 0$. In this case, by the definition of $G_i(\nu, \mathbf{p}_{-i})$, we obtain $G_i(0, \mathbf{p}_{-i}) \geq \sum_{t \in T} (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i > \sum_{t \in T} (\alpha_i^t - \beta_i^t \epsilon + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i = G_i(\epsilon, \mathbf{p}_{-i})$, which implies that $G_i(0, \mathbf{p}_{-i}) > G_i(\epsilon, \mathbf{p}_{-i})$ for small enough $\epsilon > 0$. Therefore, we have $G_i(\nu - \epsilon, \mathbf{p}_{-i}) > G_i(\nu, \mathbf{p}_{-i})$ for all $\nu \in (0, \nu^*)$ and $\epsilon > 0$ small enough that $\nu - \epsilon > 0$. Also, we have $G_i(0, \mathbf{p}_{-i}) > G_i(\epsilon, \mathbf{p}_{-i})$ for small enough $\epsilon > 0$. The last two statements establish that $G_i(\nu, \mathbf{p}_{-i})$ is strictly decreasing in $\nu \in [0, \nu^*)$. Lastly, fix $\nu \in (\nu^*, \infty)$. By the definition of ν^* , we have $\mathcal{T}_i(\nu, \mathbf{p}_{-i}) = \emptyset$, in which case, by the definition of $G_i(\nu, \mathbf{p}_{-i})$, we obtain $G_i(\nu, \mathbf{p}_{-i}) = -2c_i$. Since $G_i(\nu, \mathbf{p}_{-i}) = -2c_i$ for all $\nu \in (\nu^*, \infty)$ and $G_i(\nu, \mathbf{p}_{-i})$ is continuous in $\nu \in (0, \infty)$, it must be the case that $G_i(\nu^*, \mathbf{p}_{-i}) = -2c_i$ as well. Therefore, we have $G_i(\nu, \mathbf{p}_{-i}) = -2c_i$ for all $\nu \in [\nu^*, \infty)$.

Third, we show Part c. Assume that $G_i(0, \mathbf{p}_{-i}) > 0$. Since $\alpha_i^t > 0$ and $\beta_i^t > 0$, we have $\mathcal{T}_i(0, \mathbf{p}_{-i}) = T$ by the definition of $\mathcal{T}_i(\nu, \mathbf{p}_{-i})$. In this case, by the definition of $G_i(\nu, \mathbf{p}_{-i})$, we get $G_i(0, \mathbf{p}_{-i}) = \sum_{t \in T} (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i > 0$. Similarly, since $\alpha_i^t > 0$ and $\beta_i^t > 0$, we have $\mathcal{T}_i(\epsilon, \mathbf{p}_{-i}) = T$ for small enough $\epsilon > 0$. In this case, by the definition of $G_i(\nu, \mathbf{p}_{-i})$, we have $G_i(\epsilon, \mathbf{p}_{-i}) = \sum_{t \in T} (\alpha_i^t - \beta_i^t \epsilon + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i$. Therefore, we have $\lim_{\epsilon \rightarrow 0} G_i(\epsilon, \mathbf{p}_{-i}) = G_i(0, \mathbf{p}_{-i})$, indicating that $G_i(\nu, \mathbf{p}_{-i})$ is continuous at $\nu = 0$. Noting Part a, it follows that $G_i(\nu, \mathbf{p}_{-i})$ is continuous in $\nu \in [0, \infty)$. Since $G_i(\nu, \mathbf{p}_{-i})$ is strictly decreasing in $\nu \in [0, \nu^*)$ and $G_i(\nu, \mathbf{p}_{-i}) < 0$ for all $\nu \in [\nu^*, \infty)$ by Part b and $G_i(\nu, \mathbf{p}_{-i})$ is continuous in $\nu \in [0, \infty)$ with $G_i(0, \mathbf{p}_{-i}) > 0$, there exists a unique $\hat{\nu}$ such that $G_i(\hat{\nu}, \mathbf{p}_{-i}) = 0$. Next, assume that $G_i(0, \mathbf{p}_{-i}) = 0$. Clearly $\hat{\nu} = 0$ satisfies $G_i(\hat{\nu}, \mathbf{p}_{-i}) = 0$. Also, since $G_i(\nu, \mathbf{p}_{-i})$ is strictly decreasing in $\nu \in [0, \nu^*)$ and constant at a negative value for $\nu \in [\nu^*, \infty)$ by Part b, there cannot be another $\hat{\nu}$ such that $G_i(\hat{\nu}, \mathbf{p}_{-i}) = 0$. \square

In the next lemma, we show that if we fix the price trajectories of the firms other than firm i , then the best response of firm i , when viewed as a function of its initial inventory, is Lipschitz.

Lemma 25. *Fix the prices \mathbf{p}_{-i} charged by the firms other than firm i and let $\{p_i^t(c_i) : t \in T\}$ be the optimal solution to problem (3.1) as a function of the initial inventory of firm i . Letting $\beta_{\min} = \min_{i \in N, t \in T} \beta_i^t$, for any two initial inventory levels \hat{c}_i and \tilde{c}_i , we have*

$$\max_{t \in T} \left\{ |p_i^t(\hat{c}_i) - p_i^t(\tilde{c}_i)| \right\} \leq \frac{1}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i|.$$

Proof. Since the prices \mathbf{p}_{-i} charged by the firms other than firm i are fixed and

we work with two different initial inventory levels, we drop the argument \mathbf{p}_{-i} from $G_i(\nu, \mathbf{p}_{-i})$ and make the dependence of $G_i(\nu, \mathbf{p}_{-i})$ on c_i explicit. Thus, we use $G_i(\nu, c_i)$ to denote $G_i(\nu, \mathbf{p}_{-i})$ throughout the proof. For notational brevity, we let $\hat{p}_i^t = p_i^t(\hat{c}_i)$ and $\tilde{p}_i^t = p_i^t(\tilde{c}_i)$. Noting Lemma 9, we let \hat{v}_i and \tilde{v}_i be such that $G_i(\hat{v}_i, \hat{c}_i) = 0$ and $G_i(\tilde{v}_i, \tilde{c}_i) = 0$. Without loss of generality, we assume that $\hat{v}_i \geq \tilde{v}_i$. Since $G_i(\nu, c_i)$ is non-increasing in $\nu \in [0, \infty)$ by Lemma 24 and $G_i(\tilde{v}_i, \tilde{c}_i) = 0$, we have $G_i(\hat{v}_i, \tilde{c}_i) \leq 0$. Repeating the same argument in the first paragraph of the proof of Theorem 9, we also get $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}(\hat{v}_i - \tilde{v}_i)$ for all $t \in T$. The only difference is that we have $M_i = 0$ in this context since the prices charged by the firms other than firm i are fixed. If $\hat{v}_i \leq 2|\hat{c}_i - \tilde{c}_i|/\beta_{\min}$, then the last inequality implies that $|\hat{p}_i^t - \tilde{p}_i^t| \leq |\hat{c}_i - \tilde{c}_i|/\beta_{\min}$, which is the result that we want to show! In the rest of the proof, we proceed under the assumption that $\hat{v}_i > 2|\hat{c}_i - \tilde{c}_i|/\beta_{\min}$.

Noting that $G_i(\hat{v}_i, \hat{c}_i) = 0$ and $G_i(\nu, \hat{c}_i) < 0$ for all $\nu \in [\nu^*, \infty)$ by Lemma 24, the definition of ν^* implies that $\mathcal{T}_i(\hat{v}_i, \mathbf{p}_{-i}) \neq \emptyset$. If, otherwise, $\mathcal{T}_i(\hat{v}_i, \mathbf{p}_{-i}) = \emptyset$, then we obtain $\nu^* \leq \hat{v}_i$ by the definition of ν^* , which contradicts the fact that $G_i(\hat{v}_i, \hat{c}_i) = 0$, $G_i(\nu^*, \hat{c}_i) < 0$ and $G_i(\cdot, \hat{c}_i)$ is decreasing. Also, since $\hat{v}_i > 2|\hat{c}_i - \tilde{c}_i|/\beta_{\min} \geq 0$, by the definition of $G_i(\nu, c_i)$, we obtain $G_i(\hat{v}_i, \tilde{c}_i) - G_i(\hat{v}_i, \hat{c}_i) = -2(\tilde{c}_i - \hat{c}_i)$. Noting that $G_i(\hat{v}_i, \hat{c}_i) = 0$, the last equality yields $G_i(\hat{v}_i, \tilde{c}_i) = -2(\tilde{c}_i - \hat{c}_i)$. In the proof of Part a of Lemma 24, we show that $G_i(\nu, c_i) - G_i(\nu - \epsilon, c_i) \leq -\sum_{t \in \mathcal{T}_i(\nu, \mathbf{p}_{-i})} \beta_i^t \epsilon$ for all $\nu \in (0, \infty)$ and $\epsilon > 0$ small enough that $\nu - \epsilon > 0$. Using this inequality with $\nu = \hat{v}_i$ and $c_i = \tilde{c}_i$, we obtain $G_i(\hat{v}_i - \epsilon, \tilde{c}_i) \geq G_i(\hat{v}_i, \tilde{c}_i) + \sum_{t \in \mathcal{T}_i(\hat{v}_i, \mathbf{p}_{-i})} \beta_i^t \epsilon$. Using the last inequality with $\epsilon = 2|\hat{c}_i - \tilde{c}_i|/\beta_{\min}$, since $\mathcal{T}_i(\hat{v}_i, \mathbf{p}_{-i}) \neq \emptyset$ and $G_i(\hat{v}_i, \tilde{c}_i) = -2(\tilde{c}_i - \hat{c}_i)$, we get

$$G_i\left(\hat{v}_i - \frac{2}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i|, \tilde{c}_i\right) \geq -2(\tilde{c}_i - \hat{c}_i) + \beta_{\min} \frac{2}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i| \geq 0 = G_i(\tilde{v}_i, \tilde{c}_i).$$

Noting that $G_i(\nu, c_i)$ is strictly decreasing in $\nu \in [0, \nu^*)$ and constant at a negative value for $\nu \in [\nu^*, \infty)$ by Part b of Lemma 24, having $G_i(\hat{v}_i - \frac{2}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i|, \tilde{c}_i) \geq 0 = G_i(\tilde{v}_i, \tilde{c}_i)$ implies that $\hat{v}_i - \frac{2}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i| \leq \tilde{v}_i$. So, we obtain $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}(\hat{v}_i - \tilde{v}_i) \leq \frac{1}{\beta_{\min}} |\hat{c}_i - \tilde{c}_i|$ for all $t \in T$. \square

B.2 Omitted Proofs

Proof of Lemma 8. Since the solution $\{p_i^t : t \in T\}$, along with the dual multipliers v_i and $\{u_i^t : t \in T\}$, satisfies the KKT conditions in (3.2), solving for u_i^t in the third KKT condition, we have

$$u_i^t = \frac{\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t}{\beta_i^t} - 2p_i^t + v_i. \quad (\text{B.1})$$

For notational brevity, we let $\Delta_i^t = (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) / \beta_i^t$. Therefore, we can write (B.1) as $u_i^t = \Delta_i^t - 2p_i^t + v_i$. Furthermore, noting that $\beta_i^t > 0$ and dividing the second KKT condition in (3.2) by β_i^t , we observe that $(\Delta_i^t - p_i^t) u_i^t = 0$ for all $t \in T$. Consider any $t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})$. By the definition of $\mathcal{T}_i(v_i, \mathbf{p}_{-i})$, we have $\Delta_i^t > v_i$. In this case, by (B.1), it follows that $u_i^t = \Delta_i^t - 2p_i^t + v_i < 2(\Delta_i^t - p_i^t)$. Multiplying the last chain of inequalities by u_i^t and noting that $(\Delta_i^t - p_i^t) u_i^t = 0$, we get $(u_i^t)^2 \leq 0$, which implies that $u_i^t = 0$. Using this value of u_i^t in (B.1) and solving for p_i^t , we have

$p_i^t = \Delta_i^t/2 + v_i/2$. Therefore, the desired result holds for any $t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})$. Consider any $t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i})$. By the definition of $\mathcal{T}_i(v_i, \mathbf{p}_{-i})$, we have $\Delta_i^t \leq v_i$. In this case, using (B.1), it follows that $u_i^t = \Delta_i^t - 2p_i^t + v_i \geq 2(\Delta_i^t - p_i^t)$. Multiplying the last chain of inequalities by $\Delta_i^t - p_i^t$ and noting that $(\Delta_i^t - p_i^t)u_i^t = 0$, we have $(\Delta_i^t - p_i^t)^2 \leq 0$, which implies that $p_i^t = \Delta_i^t$. Using this value of p_i^t in (B.1) and noting the definition of Δ_i^t , we get $u_i^t = v_i - \Delta_i^t$. Therefore, the desired result holds for any $t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i})$. \square

Proof of Lemma 9. As in the proof of Lemma 8, we let $\Delta_i^t = (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) / \beta_i^t$ for notational brevity. By Lemma 8, we have $p_i^t = (\Delta_i^t + v_i)/2$ for all $t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})$ and $p_i^t = \Delta_i^t$ for all $t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i})$. First, we assume that $v_i = 0$. Since $\alpha_i^t > 0$, we have $\mathcal{T}_i(v_i, \mathbf{p}_{-i}) = T$ by the definition of $\mathcal{T}_i(v_i, \mathbf{p}_{-i})$, which implies that $p_i^t = (\Delta_i^t + v_i)/2 = \Delta_i^t/2$ for all $t \in T$. In this case, we obtain $\frac{1}{2} \sum_{t \in T} \beta_i^t \Delta_i^t = \sum_{t \in T} \beta_i^t (\Delta_i^t - p_i^t) = \sum_{t \in T} (\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) \leq c_i$, where the second equality uses the definition of Δ_i^t and the inequality follows from the fact that $\{p_i^t : t \in T\}$ is a feasible solution to problem (3.1). The last chain of inequalities imply that $\sum_{t \in T} \beta_i^t \Delta_i^t - 2c_i \leq 0$. Noting the definition of Δ_i^t and the fact that $\mathcal{T}_i(v_i, \mathbf{p}_{-i}) = T$, we obtain $\sum_{t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})} (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i \leq 0$, which implies that $G_i(v_i, \mathbf{p}_{-i}) = G_i(0, \mathbf{p}_{-i}) = 0$. Therefore, the desired result holds when $v_i = 0$. Second, we assume that $v_i > 0$. Using the fact that $p_i^t = (\Delta_i^t + v_i)/2$ for all $t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})$ and $p_i^t = \Delta_i^t$ for all $t \notin \mathcal{T}_i(v_i, \mathbf{p}_{-i})$, we have $\sum_{t \in T} (\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) = \sum_{t \in T} \beta_i^t (\Delta_i^t - p_i^t) = \sum_{t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})} \beta_i^t (\Delta_i^t - v_i)/2$. Since $v_i > 0$, by the first KKT condition in (3.2), we also have $c_i = \sum_{t \in T} (\alpha_i^t - \beta_i^t p_i^t + \sum_{j \neq i} \gamma_{i,j}^t p_j^t)$. In this

case, by the last chain of equalities, we get

$$c_i = \frac{1}{2} \sum_{t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})} \beta_i^t (\Delta_i^t - v_i).$$

By the definition of Δ_i^t , the equality above is equivalent to $\sum_{t \in \mathcal{T}_i(v_i, \mathbf{p}_{-i})} (\alpha_i^t - \beta_i v_i + \sum_{j \neq i} \gamma_{i,j}^t p_j^t) - 2c_i = 0$, which implies that $G_i(v_i, \mathbf{p}_{-i}) = 0$. Therefore, the desired result holds when $v_i > 0$. \square

Proof of Theorem 9. For notational brevity, we let $\hat{p}_i^t = p_i^t(\hat{\mathbf{p}}_{-i})$ and $\tilde{p}_i^t = p_i^t(\tilde{\mathbf{p}}_{-i})$. In other words, $\{\hat{p}_i^t : t \in T\}$ is the optimal solution to problem (3.1) when we solve this problem after replacing \mathbf{p}_{-i} with $\hat{\mathbf{p}}_{-i}$. Similarly, $\{\tilde{p}_i^t : t \in T\}$ is the optimal solution to problem (3.1) when we solve this problem after replacing \mathbf{p}_{-i} with $\tilde{\mathbf{p}}_{-i}$. Also, we let \hat{v}_i and \tilde{v}_i be such that $G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) = 0$ and $G_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i}) = 0$. Without loss of generality, we assume that $\hat{v}_i \geq \tilde{v}_i$. Otherwise, we interchange the roles of \hat{v}_i and \tilde{v}_i . Finally, we let $\hat{\Delta}_i^t = (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t \hat{p}_j^t) / \beta_i^t$ and $\tilde{\Delta}_i^t = (\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t) / \beta_i^t$ for notational brevity. Note that $|\hat{\Delta}_i^t - \tilde{\Delta}_i^t| \leq \sum_{j \neq i} \gamma_{i,j}^t |\hat{p}_j^t - \tilde{p}_j^t| / \beta_i^t$. In this case, using $M_i = \max_{t \in T} \{\sum_{j \neq i} \gamma_{i,j}^t |\hat{p}_j^t - \tilde{p}_j^t| / \beta_i^t\}$, we have $|\hat{\Delta}_i^t - \tilde{\Delta}_i^t| \leq M_i$ for all $t \in T$. We proceed to examining four cases to show that $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2} \max\{M_i, \hat{v}_i - \tilde{v}_i\}$ for all $t \in T$. First, we assume that $t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$ and $t \in \mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$. Using Lemma 8, we have $|\hat{p}_i^t - \tilde{p}_i^t| = \frac{1}{2} |\hat{\Delta}_i^t + \hat{v}_i - \tilde{\Delta}_i^t - \tilde{v}_i| \leq \frac{1}{2} |\hat{\Delta}_i^t - \tilde{\Delta}_i^t| + \frac{1}{2} (\hat{v}_i - \tilde{v}_i) \leq \frac{1}{2}M_i + \frac{1}{2}(\hat{v}_i - \tilde{v}_i) \leq \frac{1}{2}M_i + \frac{1}{2} \max\{M_i, \hat{v}_i - \tilde{v}_i\}$, as desired. Second, we assume that $t \notin \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$ and $t \notin \mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$. Using Lemma 8 once more, we have $|\hat{p}_i^t - \tilde{p}_i^t| = |\hat{\Delta}_i^t - \tilde{\Delta}_i^t| \leq M_i \leq \frac{1}{2}M_i + \frac{1}{2} \max\{M_i, \hat{v}_i - \tilde{v}_i\}$, as desired. Third, we assume that $t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$ and $t \notin \mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$. Since $t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$, we have $\hat{\Delta}_i^t > \hat{v}_i$, which implies $\hat{v}_i - \tilde{\Delta}_i^t < \hat{\Delta}_i^t - \tilde{\Delta}_i^t \leq M_i$. Also, since $t \notin \mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$, we have $\tilde{\Delta}_i^t \leq \tilde{v}_i$,

which implies $\hat{v}_i - \tilde{\Delta}_i^t \geq \hat{v}_i - \tilde{v}_i$. Noting the last two inequalities, it follows that $|\hat{v}_i - \tilde{\Delta}_i^t| \leq \max\{M_i, \hat{v}_i - \tilde{v}_i\}$. In this case, using Lemma 8 one last time and using the fact that $|\hat{v}_i - \tilde{\Delta}_i^t| \leq \max\{M_i, \hat{v}_i - \tilde{v}_i\}$, we obtain $|\hat{p}_i^t - \tilde{p}_i^t| = |\frac{1}{2}\hat{\Delta}_i^t + \frac{1}{2}\hat{v}_i - \tilde{\Delta}_i^t| \leq \frac{1}{2}|\hat{\Delta}_i^t - \tilde{\Delta}_i^t| + \frac{1}{2}|\hat{v}_i - \tilde{\Delta}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2}|\hat{v}_i - \tilde{\Delta}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2}\max\{M_i, \hat{v}_i - \tilde{v}_i\}$, as desired. Fourth, we assume that $t \notin \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})$ and $t \in \mathcal{T}_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i})$, in which case, we can follow the same argument in the third case to obtain $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2}\max\{M_i, \hat{v}_i - \tilde{v}_i\}$. The preceding discussion shows that $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2}\max\{M_i, \hat{v}_i - \tilde{v}_i\}$. If $\hat{v}_i \leq M_i$, then noting that $\tilde{v}_i \geq 0$, the last inequality implies that $|\hat{p}_i^t - \tilde{p}_i^t| \leq M_i$, which is the result we want to show! In the rest of the proof, we proceed under the assumption that $\hat{v}_i > M_i$.

Consider the function $G_i(\cdot, \tilde{\mathbf{p}}_{-i})$. By Lemma 24 in the Appendix A, the function $G_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is strictly decreasing over the interval $[0, \nu^*)$ for some ν^* and constant over the interval $[\nu^*, \infty)$. By the same lemma, we also have $G_i(\nu^*, \tilde{\mathbf{p}}_{-i}) = -2c_i < 0$. In the rest of the proof, we show that $G_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \geq 0$. Also, we have $G_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i}) = 0$ by Lemma 9. In this case, since $G_i(\nu^*, \tilde{\mathbf{p}}_{-i}) < 0$ and $G_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is strictly decreasing over the interval $[0, \nu^*)$ and constant over the interval $[\nu^*, \infty)$, having $G_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \geq 0$ and $G_i(\tilde{v}_i, \tilde{\mathbf{p}}_{-i}) = 0$ implies that $\hat{v}_i - M_i \leq \tilde{v}_i$. Therefore, we have $\hat{v}_i - \tilde{v}_i \leq M_i$, so that we get $|\hat{p}_i^t - \tilde{p}_i^t| \leq \frac{1}{2}M_i + \frac{1}{2}\max\{M_i, \hat{v}_i - \tilde{v}_i\} = M_i$, which is the result we want to show. It remains to show that $G_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \geq 0$. Using $\mathbf{1}(\cdot)$ to denote the indicator function, since $\hat{v}_i > M_i$, by the definition of $G_i(\cdot, \tilde{\mathbf{p}}_{-i})$, we have

$$\begin{aligned} G_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) &= \sum_{t \in \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})} \left(\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \right) - 2c_i \\ &= \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \left(\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \right) - 2c_i \end{aligned}$$

$$\begin{aligned}
& + \sum_{t \in T} \mathbf{1}(t \in \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \setminus \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})) \left(\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \right) \\
& - \sum_{t \in T} \mathbf{1}(t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) \setminus \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})) \left(\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \right).
\end{aligned} \tag{B.2}$$

We consider each one of the three terms on the right side above one by one. For the first term, by Lemma 9, we have $G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) = 0$. By the definition of M_i , we also have $\sum_{j \neq i} \gamma_{i,j}^t |\hat{p}_j^t - \tilde{p}_j^t| / \beta_i^t \leq M_i$ for all $t \in T$, so that $\sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \sum_{j \neq i} \gamma_{i,j}^t |\hat{p}_j^t - \tilde{p}_j^t| \leq M_i \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \beta_i^t$. Thus, we get

$$\begin{aligned}
& \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \left(\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \right) - 2c_i \\
& = \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \left(\alpha_i^t - \beta_i^t \hat{v}_i + \sum_{j \neq i} \gamma_{i,j}^t \hat{p}_j^t \right) - 2c_i + \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \sum_{j \neq i} \gamma_{i,j}^t (\tilde{p}_j^t - \hat{p}_j^t) + M_i \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \beta_i^t \\
& = G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) + \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \sum_{j \neq i} \gamma_{i,j}^t (\tilde{p}_j^t - \hat{p}_j^t) + M_i \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \beta_i^t \\
& \geq G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) - \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \sum_{j \neq i} \gamma_{i,j}^t |\tilde{p}_j^t - \hat{p}_j^t| + M_i \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} \beta_i^t \geq 0,
\end{aligned}$$

where the second equality uses the fact that $\hat{v}_i > M_i \geq 0$ so that we have $G_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) = \sum_{t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})} (\alpha_i^t - \beta_i^t \hat{v}_i + \sum_{j \neq i} \gamma_{i,j}^t \hat{p}_j^t) - 2c_i$. Therefore, the first term on the right side of (B.2) is non-negative. For the second term, by the definition of $\mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})$, we have $\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t > \beta_i^t (\hat{v}_i - M_i)$ for all $t \in \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})$. Therefore, we have $\mathbf{1}(t \in \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i}) \setminus \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i})) (\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t) \geq 0$, which implies that the second term on the right side of (B.2) is non-negative. For the third term, we have $\alpha_i^t + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t \leq \beta_i^t (\hat{v}_i - M_i)$ for all $t \notin \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})$. Therefore, we have $\mathbf{1}(t \in \mathcal{T}_i(\hat{v}_i, \hat{\mathbf{p}}_{-i}) \setminus \mathcal{T}_i(\hat{v}_i - M_i, \tilde{\mathbf{p}}_{-i})) \times (\alpha_i^t - \beta_i^t (\hat{v}_i - M_i) + \sum_{j \neq i} \gamma_{i,j}^t \tilde{p}_j^t) \leq 0$, indicating that the third term on the right side of (B.2) is non-positive. So, the first and second terms on the right side of (B.2) are non-negative, whereas the third term

is non-positive, in which case, we have $G_i(\hat{v}_i - M_i, \tilde{p}_{-i}) \geq 0$. \square

Proof of Lemma 10. We consider the problem over the time periods $T \setminus \{1\}$ where the inventory for each firm ℓ at the second time period is given by $[c_\ell - D_\ell^1(\tilde{\mathbf{p}}^1)]^+$. By the discussion that follows Theorem 9, if we start with any set of prices for the firms at the initial iteration and iteratively compute the best response of each firm to prices at the previous iteration, then we reach the equilibrium without recourse. Therefore, to compute $\{\tilde{\mathbf{p}}^t : t \in T \setminus \{1\}\}$, we consider the problem over the time periods $T \setminus \{1\}$ with the inventory of each firm ℓ at the second time period given by $[c_\ell - D_\ell^1(\tilde{\mathbf{p}}^1)]^+$ and starting with the prices $\{\hat{\mathbf{p}}^t : t \in T \setminus \{1\}\}$ at the initial iteration, we iteratively compute the best response of each firm to the prices at the previous iteration. Letting $\{\tilde{\mathbf{p}}^{t,k} : t \in T \setminus \{1\}\}$ be the price trajectories for the firms at iteration k , we know that $\lim_{k \rightarrow \infty} \tilde{\mathbf{p}}_\ell^{t,k} = \tilde{\mathbf{p}}_\ell^t$ for all $\ell \in N$, $t \in T \setminus \{1\}$. Therefore, for all $t \in T \setminus \{1\}$ and $\ell \in N$, we have $|\hat{p}_\ell^t - \tilde{p}_\ell^t| = |\hat{p}_\ell^t - \tilde{p}_\ell^{t,1} + \sum_{k=1}^{\infty} (\tilde{p}_\ell^{t,k} - \tilde{p}_\ell^{t,k+1})| \leq \sum_{k=1}^{\infty} |\tilde{p}_\ell^{t,k} - \tilde{p}_\ell^{t,k+1}|$, where the inequality uses the fact that $\tilde{p}_\ell^{t,1} = \hat{p}_\ell^t$. In this case, to bound $|\hat{p}_\ell^t - \tilde{p}_\ell^t|$, we can bound $|\tilde{p}_\ell^{t,k} - \tilde{p}_\ell^{t,k+1}|$ for all $k = 1, 2, \dots$ and add up the bounds on the latter quantity. We proceed to bounding $|\tilde{p}_\ell^{t,k} - \tilde{p}_\ell^{t,k+1}|$. By definition, the price trajectory $\{\tilde{p}_\ell^{t,k+1} : t \in T \setminus \{1\}\}$ of firm ℓ at iteration $k+1$ is the best response of firm ℓ to the price trajectories $\{\tilde{\mathbf{p}}_{-\ell}^{t,k} : t \in T \setminus \{1\}\}$ of the other firms at iteration k . In this case, Theorem 9 implies that $|\tilde{p}_\ell^{t,k+1} - \tilde{p}_\ell^{t,k}| \leq \max_{t \in T \setminus \{1\}} \{\sum_{j \neq \ell} \gamma_{\ell,j}^t |\tilde{p}_j^{t,k} - \tilde{p}_j^{t,k-1}| / \beta_\ell^t\}$ for all $\ell \in N$, $t \in T \setminus \{1\}$, $k = 2, 3, \dots$. For notational brevity, we let $\Phi_\ell^k = \max_{t \in T \setminus \{1\}} \{|\tilde{p}_\ell^{t,k+1} - \tilde{p}_\ell^{t,k}|\}$ so that the last inequality yields $\Phi_\ell^k \leq \max_{t \in T \setminus \{1\}} \sum_{j \neq \ell} \gamma_{\ell,j}^t \Phi_j^{k-1} / \beta_\ell^t$ for all $\ell \in N$, $k = 2, 3, \dots$. Using the inequality $\Phi_\ell^k \leq \max_{t \in T \setminus \{1\}} \sum_{j \neq \ell} \gamma_{\ell,j}^t \Phi_j^{k-1} / \beta_\ell^t$ for firm $\ell = j$ with $j \neq i$ and

noting the definitions of M and μ , it follows that

$$\Phi_j^k \leq \max_{t \in T \setminus \{1\}} \left\{ \frac{\sum_{\ell \in N \setminus \{j,i\}} \gamma_{j,\ell}^t \Phi_\ell^{k-1}}{\beta_j^t} + \frac{\gamma_{j,i}^t \Phi_i^{k-1}}{\beta_j^t} \right\} \leq M \max_{\ell \in N \setminus \{j,i\}} \left\{ \Phi_\ell^{k-1} \right\} + \mu \Phi_i^{k-1} \quad (\text{B.3})$$

for all $j \neq i$ and $k = 2, 3, \dots$. Using the inequality $\Phi_\ell^k \leq \max_{t \in T \setminus \{1\}} \sum_{j \neq \ell} \gamma_{\ell,j}^t \Phi_j^{k-1} / \beta_\ell^t$ again for firm $\ell = i$, we get $\Phi_i^k \leq \max_{t \in T \setminus \{1\}} \sum_{j \neq i} \gamma_{i,j}^t \Phi_j^{k-1} / \beta_i^t \leq M \max_{j \neq i} \Phi_j^{k-1}$ for all $k = 2, 3, \dots$. If we use the last inequality in (B.3), then for all $j \neq i$ and $k = 3, 4, \dots$, we have $\Phi_j^k \leq M \max_{\ell \in N \setminus \{j,i\}} \{\Phi_\ell^{k-1}\} + M\mu \max_{j \neq i} \{\Phi_j^{k-2}\}$. So, letting $\Theta^k = \max_{j \neq i} \Phi_j^k$, the last inequality yields

$$\Theta^k \leq M\Theta^{k-1} + M\mu\Theta^{k-2}$$

for all $k = 3, 4, \dots$. Adding the inequality above over all $k = 3, 4, \dots$, we obtain $\sum_{k=3}^{\infty} \Theta^k \leq M \sum_{k=2}^{\infty} \Theta^k + M\mu \sum_{k=1}^{\infty} \Theta^k$, which is equivalent to $\sum_{k=1}^{\infty} \Theta^k \leq M \sum_{k=2}^{\infty} \Theta^k + M\mu \sum_{k=1}^{\infty} \Theta^k + \Theta^1 + \Theta^2 = (M + M\mu) \sum_{k=1}^{\infty} \Theta^k + (1 - M) \Theta^1 + \Theta^2$. Rearranging the terms in the last chain of inequalities, we get $\sum_{k=1}^{\infty} \Theta^k \leq ((1 - M) \Theta^1 + \Theta^2) / (1 - M - M\mu)$.

Therefore, if we can bound Θ^1 and Θ^2 , then we can bound $\sum_{k=1}^{\infty} \Theta^k$. When we increase the price of firm i at the first time period by δ , the inventory of firm i at the second time period changes by at most $\beta_i^1 \delta$ and the inventory of firm $j \neq i$ at the second time period changes by at most $\gamma_{j,i}^1 \delta$. In Appendix A, Lemma 25 shows that if we fix the price trajectories of the firms other than firm i , then the best response of firm i , when viewed as a function of its initial inventory, is Lipschitz with constant $1/\beta_{\min}$, where we let $\beta_{\min} = \min_{i \in N, t \in T} \beta_i^t$. Note that the best response of firm i does not depend on the inventories of the other firms, since

the price trajectories of the other firms is fixed. By definition, if we consider the problem over the time periods $T \setminus \{1\}$ with the inventory of each firm ℓ at the second time period given by $[c_\ell - D_\ell^1(\hat{\mathbf{p}}^1)]^+$, by definition, $\{\hat{\mathbf{p}}_\ell^t : t \in T \setminus \{1\}\}$ is the best response to the price trajectories $\{\hat{\mathbf{p}}_{-\ell}^t : t \in T \setminus \{1\}\}$. Also, if we consider the problem over the time periods $T \setminus \{1\}$ with the inventory of each firm ℓ at the second time period given by $[c_\ell - D_\ell^1(\tilde{\mathbf{p}}^1)]^+$, by definition, $\{\tilde{\mathbf{p}}_\ell^{t,2} : t \in T \setminus \{1\}\}$ is the best response to the price trajectories $\{\tilde{\mathbf{p}}_{-\ell}^{t,1} : t \in T \setminus \{1\}\}$. Since the price trajectories $\{\hat{\mathbf{p}}_{-\ell}^t : t \in T \setminus \{1\}\}$ and $\{\tilde{\mathbf{p}}_{-\ell}^{t,1} : t \in T \setminus \{1\}\}$ are the same, Lemma 25 in Appendix A implies that $|\tilde{p}_\ell^{t,2} - \tilde{p}_\ell^{t,1}| = |\tilde{p}_\ell^{t,2} - \hat{p}_\ell^t| \leq |(c_\ell - D_\ell^1(\tilde{\mathbf{p}}^1))^+ - (c_\ell - D_\ell^1(\hat{\mathbf{p}}^1))^+|/\beta_{\min}$ for all $t \in T \setminus \{1\}$. As discussed at the beginning of this paragraph, the expression on the right side of the last inequality is bounded by $\beta_i^1 \delta$ when $\ell = i$ and bounded by $\gamma_{j,i}^1 \delta$ when $\ell = j$ with $j \neq i$. Therefore, we obtain $|\hat{p}_i^{t,2} - \tilde{p}_i^{t,1}| \leq \beta_i^1 \delta / \beta_{\min} \leq \bar{\beta} \delta$ and $|\tilde{p}_j^{t,2} - \tilde{p}_j^{t,1}| \leq \gamma_{j,i}^1 \delta / \beta_{\min} \leq \bar{\beta} \mu \delta$ for all $j \neq i$. The second one of the last two inequalities yields $\Theta^1 = \max_{j \neq i, t \in T \setminus \{1\}} \{|\tilde{p}_j^{t,2} - \tilde{p}_j^{t,1}|\} \leq \mu \bar{\beta} \delta$. The first one of the last two inequalities yields $\Phi_i^1 = \max_{t \in T \setminus \{1\}} \{|\hat{p}_i^{t,2} - \tilde{p}_i^{t,1}|\} \leq \bar{\beta} \delta$, in which case, noting (B.3), we get $\Theta^2 = \max_{j \neq i} \{\Phi_j^2\} \leq M \Theta^1 + \mu \Phi_i^1 \leq M \mu \bar{\beta} \delta + \mu \bar{\beta} \delta$. Thus, Θ^1 and Θ^2 are respectively bounded by $\mu \bar{\beta} \delta$ and $(1 + M) \mu \bar{\beta} \delta$. In this case, for all $j \neq i$ and $t \in T \setminus \{1\}$, we have

$$\begin{aligned} |\hat{p}_j^t - \tilde{p}_j^t| &\leq \sum_{k=1}^{\infty} |\tilde{p}_j^{t,k+1} - \tilde{p}_j^{t,k}| \leq \sum_{k=1}^{\infty} \max_{j \neq i, t \in T \setminus \{1\}} \left\{ |\tilde{p}_j^{t,k+1} - \tilde{p}_j^{t,k}| \right\} = \sum_{k=1}^{\infty} \Theta^k \\ &\leq \frac{(1 - M) \Theta^1 + \Theta^2}{1 - M - M\mu} \leq \frac{(1 - M) \mu \bar{\beta} \delta + (1 + M) \mu \bar{\beta} \delta}{1 - M - M\mu} \leq \frac{2 \mu \bar{\beta} \delta}{1 - M - M\mu}, \end{aligned}$$

where the first inequality follows from the discussion at the beginning of the proof

and the equality is by the definition of Θ^k and Φ_ℓ^k . \square

Proof of Lemma 11. We let \hat{p}_i^t be the price charged by firm i at time period t in the equilibrium without recourse. As discussed right before the lemma, given that all of the firms use the strategy $\{\mathbf{P}^{R,t} : t \in T\}$, the realized prices are $\{\hat{p}_i^t : i \in N, t \in T\}$. We use \hat{q}_i^1 to denote the arbitrary price charged by firm i at the first time period. Given that firm i uses the strategy $\{P_i^{R,t}(\cdot) : t \in T \setminus \{1\}\}$ at the other time periods and the other firms use the strategy $\{\mathbf{P}_{-i}^{R,t} : t \in T\}$, we let $\{\hat{q}_i^t : i \in N, t \in T\}$ be the realized prices. For each firm $j \neq i$, note that $\hat{q}_j^1 = P_j^{R,1}(\mathbf{c}) = p_j^{N,1}(1, \mathbf{c}) = \hat{p}_j^1$. Also, by Lemma 10, for all $j \neq i$ and $t \in T \setminus \{1\}$, we have $|\hat{p}_j^t - \hat{q}_j^t| \leq 2\mu\bar{\beta}|\hat{p}_i^1 - \hat{q}_i^1|/(1 - M - M\mu) \leq 2\mu\bar{\beta}P_{\max}/(1 - M - M\mu)$. We use $\pi_i^t(p_i^t, \mathbf{p}_{-i}^t)$ to denote the revenue of firm i at time period t when firm i charges the price p_i^t and the other firms charge the price $\mathbf{p}_{-i}^t = (p_1^t, \dots, p_{i-1}^t, p_{i+1}^t, \dots, p_n^t)$. We have $\Pi_i^N = \sum_{t \in T} \pi_i^t(\hat{p}_i^t, \hat{\mathbf{p}}_{-i}^t)$ and $\Pi_i^D = \sum_{t \in T} \pi_i^t(\hat{q}_i^t, \hat{\mathbf{q}}_{-i}^t)$. In this case, we get

$$\Pi_i^N = \sum_{t \in T} \pi_i^t(\hat{p}_i^t, \hat{\mathbf{p}}_{-i}^t) \geq \sum_{t \in T} \pi_i^t(\hat{q}_i^t, \hat{\mathbf{p}}_{-i}^t) = \Pi_i^D - \sum_{t \in T \setminus \{1\}} \pi_i^t(\hat{q}_i^t, \hat{\mathbf{q}}_{-i}^t) + \sum_{t \in T \setminus \{1\}} \pi_i^t(\hat{q}_i^t, \hat{\mathbf{p}}_{-i}^t),$$

where the inequality uses the fact that $\{\hat{p}_i^t : t \in T\}$ is the best response of firm i to the prices $\{\hat{\mathbf{p}}_{-i}^t : t \in T\}$ and the second equality uses the fact that $\hat{q}_{-i}^1 = \hat{\mathbf{p}}_{-i}^1$. Using $D_i^t(p_i^t, \mathbf{p}_{-i}^t)$ to denote the demand of firm i at time period t when firm i charges the price p_i^t and the other firms charge the prices \mathbf{p}_{-i}^t , by the inequality above, we get $\Pi_i^D - \Pi_i^N \leq \sum_{t \in T \setminus \{1\}} |\pi_i^t(\hat{q}_i^t, \hat{\mathbf{q}}_{-i}^t) - \pi_i^t(\hat{q}_i^t, \hat{\mathbf{p}}_{-i}^t)| = \sum_{t \in T \setminus \{1\}} |\hat{q}_i^t D_i^t(\hat{q}_i^t, \hat{\mathbf{q}}_{-i}^t) - \hat{q}_i^t D_i^t(\hat{q}_i^t, \hat{\mathbf{p}}_{-i}^t)|$. Since $D_i^t(\hat{q}_i^t, \hat{\mathbf{q}}_{-i}^t) - D_i^t(\hat{q}_i^t, \hat{\mathbf{p}}_{-i}^t) = (\alpha_i^t + \beta_i^t \hat{q}_i^t - \sum_{j \neq i} \gamma_{i,j}^t \hat{q}_j^t) - (\alpha_i^t + \beta_i^t \hat{q}_i^t - \sum_{j \neq i} \gamma_{i,j}^t \hat{p}_j^t) = \sum_{j \neq i} \gamma_{i,j}^t (\hat{p}_j^t - \hat{q}_j^t)$, the last chain of inequalities yields $\Pi_i^D - \Pi_i^N \leq \sum_{t \in T \setminus \{1\}} \hat{q}_i^t \sum_{j \neq i} \gamma_{i,j}^t |\hat{q}_j^t - \hat{p}_j^t| \leq \sum_{t \in T \setminus \{1\}} P_{\max} M \beta_i^t \max_{j \neq i} |\hat{q}_j^t - \hat{p}_j^t|$,

where we use the fact that $\hat{q}_i^t \leq P_{\max}$ and $M \geq \sum_{j \neq i} \gamma_{i,j}^t / \beta_i^t$. At the beginning of the proof, we show that $|\hat{p}_j^t - \hat{q}_j^t| \leq 2\mu\bar{\beta}P_{\max}/(1-M-M\mu)$ for all $j \neq i$ and $t \in T \setminus \{1\}$. In this case, we obtain $\Pi_i^D - \Pi_i^N \leq \sum_{t \in T \setminus \{1\}} P_{\max} M \beta_i^t \max_{j \neq i} \{|\hat{q}_j^t - \hat{p}_j^t|\} = 2\bar{\beta}M\beta_{\max}P_{\max}^2(\tau-1)\mu/(1-M-M\mu)$. \square

Proof of Theorem 10. Consider the problem over the time periods κ, \dots, τ . We use $\text{Rev}_i^\kappa(P_i^\kappa(\cdot), \dots, P_i^\tau(\cdot), \mathbf{P}_{-i}, \mathbf{x})$ to denote the revenue of firm i over the time periods κ, \dots, τ , when the firm uses the strategy $\{P_i^\kappa(\cdot), \dots, P_i^\tau(\cdot)\}$, the other firms use the strategy \mathbf{P}_{-i} and the inventories at time period κ are given by \mathbf{x} . We let $\{Q_i^t(\cdot) : t \in T\}$ be an arbitrary strategy used by firm i . We use induction over the time periods to show that $\text{Rev}_i^\kappa(Q_i^\kappa(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^\kappa(P_i^{R,\kappa}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) \leq \Gamma_\mu(\tau - \kappa + 1)(\tau - \kappa)\mu$. In this case, the result follows by noting that $\Pi_i^A = \text{Rev}_i^1(Q_i^1(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{c})$, $\Pi_i^N = \text{Rev}_i^1(P_i^{R,1}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{c})$ and using the last inequality with $\kappa = 1$. Consider the case $\kappa = \tau$. We have $\text{Rev}_i^\tau(Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^\tau(P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) \leq 0$, where the inequality follows from the fact that $P_i^{R,\tau}(\mathbf{x})$ is the best response of firm i to the prices $\mathbf{P}_{-i}^{R,t}(\mathbf{x})$. Therefore the result holds for $\kappa = \tau$. Assuming that the result holds for $\kappa = t + 1$, we show that the result holds for $\kappa = t$. Using $\mathbf{D}^t(p_i^t, \mathbf{p}_{-i}^t)$ to denote the vector of demands for the firms when the prices are $(p_i^t, \mathbf{p}_{-i}^t)$ and letting $\mathbf{x}' = [\mathbf{x} - \mathbf{D}^t(Q_i^t(\mathbf{x}), \mathbf{P}_{-i}^t(\mathbf{x}))]^+$, observe that $\text{Rev}_i^t(Q_i^t(\cdot), Q_i^{t+1}(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^t(Q_i^t(\cdot), P_i^{R,t+1}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) = \text{Rev}_i^{t+1}(Q_i^{t+1}(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}') - \text{Rev}_i^{t+1}(P_i^{R,t+1}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}')$, since all firms make the same pricing decisions at time period t in the two revenue expressions on the left side of the equality. By the induction assumption, the right

side of the last equality is bounded by $\Gamma_\mu(\tau - t)(\tau - t - 1)\mu$. Also, considering $\text{Rev}_i^t(Q_i^t(\cdot), P_i^{R,t+1}(\cdot), \dots, R_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^t(P_i^{R,t}(\cdot), P_i^{R,t+1}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x})$, this expression is the change in the revenue of firm i when firm i deviates from the strategy $\{P_i^{R,t}(\cdot) : t \in T\}$ only at the initial period and there are $\tau - t + 1$ time periods in the problem. By Lemma 11, this expression is bounded by $2\Gamma_\mu(\tau - t)\mu$. In this case, noting that

$$\begin{aligned} & \text{Rev}_i^t(Q_i^t(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^t(P_i^{R,t}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) \\ &= \text{Rev}_i^t(Q_i^t(\cdot), Q_i^{t+1}(\cdot), \dots, Q_i^\tau(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^t(Q_i^t(\cdot), P_i^{R,t+1}(\cdot), \dots, R_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) \\ & \quad + \text{Rev}_i^t(Q_i^t(\cdot), P_i^{R,t+1}(\cdot), \dots, R_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}) - \text{Rev}_i^t(P_i^{R,t}(\cdot), \dots, P_i^{R,\tau}(\cdot), \mathbf{P}_{-i}^R, \mathbf{x}), \end{aligned}$$

the two differences on the right side above are bounded by $\Gamma_\mu(\tau - t)(\tau - t + 1)\mu$ and $2\Gamma_\mu(\tau - t)\mu$. Since $\Gamma_\mu(\tau - t)(\tau - t - 1)\mu + 2\Gamma_\mu(\tau - t)\mu = \Gamma_\mu(\tau - t + 1)(\tau - t)\mu$, the result holds for $\kappa = t$. \square

APPENDIX C

APPENDIX FOR CHAPTER 4

Proof to Lemma 16:

In the balancing algorithm, let x_i^t be the probability of selling stage 1 item i in time t . Let $y_j^{t,i}$ be the probability of selling stage 2 item j to type i customer in time t , conditioning on that a type i customer arrives to the system. Note that $x^T = x$ and $y^T = y$. For any stage 1 item i , x_i^t changes in time period t if only if item i is balanced in period t , and it is nonincreasing in t . If item i is balanced at time period t , then

$$x_i^t = \frac{c_i^t}{t}.$$

Similarly, $y_j^{t,i}$ changes in time period t if only if item j is balanced in period t , and it is nonincreasing in t . If item i is balanced at time period t , then

$$\sum_i x_i^t y_j^{t,i} = \frac{c_j^t}{t}.$$

Proving inequality (4.8): We first divide the selling horizon into several long-enough epochs, and show that with large probability, the realized sales of item i in each epoch is not too different from its expectation. It follows that for large enough t , with large probability, x_i^t is bounded from below by $x_i - \epsilon(T)$, for some $\epsilon(T)$ that goes to 0 as T goes to infinity. Convergence in probability then follows from some algebra.

We divide the selling horizon into epochs, each containing τ periods. We let D_k be a random variable denoting the realized sales of item i in the k th epoch of time

periods, under the balancing algorithm. Namely, D_k is the realized sales of item i from period $T - (k - 1)\tau$ to period $T - k\tau + 1$, where $k < T/\tau$. Let $B(x_i, \tau)$ denote a binomial random variable when there are τ trials, each with success rate x_i . Note that in each period, item i is sold with probability at most x_i , thus D_k is stochastically dominated by $B(x_i, \tau)$. Then by Hoeffding's inequality, we have for any $a > 0$,

$$\mathbf{P}(D_k - x_i\tau \geq a) \leq \mathbf{P}(B(x_i, \tau) - x_i\tau \geq a) \leq 2^{-\tau a^2}.$$

Then by the union bound, for any $l < T/\tau$,

$$\mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l) \geq 1 - l2^{-\tau a^2}. \quad (\text{C.1})$$

If item i is balanced in the k th epoch of time periods, a lower bound of the sales probability x_i^t after balancing is obtained when all D_k sales are realized at the beginning D_k periods of the k th epoch. Namely for any period t in the k th epoch,

$$x_i^t \geq \frac{c_i^T - \sum_{k'=1}^{k-1} D_{k'} - D_k}{T - (k-1)\tau - D_k}.$$

Now assume that $D_{k'} < \mu_i\tau + a, \forall k' = 1, \dots, k$. We have

$$\begin{aligned} x_i^t &\geq \frac{x_i T - (k-1)(x_i\tau + a) - D_k}{T - (k-1)\tau - D_k} \\ &\geq x_i - \frac{(k-1)a + (1-x_i)D_k}{T - (k-1)\tau - D_k} \\ &\geq x_i - \frac{(k-1)a + (1-x_i)(x_i\tau + a)}{T - (k-1)\tau - (x_i\tau + a)}. \end{aligned}$$

Note that x_i^t only decreases when balance step is performed, and we can use the inequality to bound the updated x_i^t after each balance. Thus, assuming that $D_k < x_i\tau + a, \forall k = 1, \dots, l$, we have for any period t in the first l epochs of time

periods,

$$x_i^t \geq x_i - \frac{(l-1)a + (1-x_i)(x_i\tau + a)}{T - (l-1)\tau - (x_i\tau + a)} =: \hat{x}_i. \quad (\text{C.2})$$

Using inequalities (C.1) and (C.2), we conclude that with probability at least $1 - l2^{-a^2\tau}$, the probability of selling item i in the balancing algorithm is at least \hat{x}_i in any period t in the first l epochs of time periods.

We now provide a lower bound on the probability of D_k greater than a given threshold, for any k . To be more precise, we construct a lower bound on $P(D_k - \hat{x}_i\tau > -b, \forall k = 1, \dots, l)$. Recall that given the conditioned event $\{D_k < x_i\tau + a, \forall k = 1, \dots, l\}$, the sales probability of item i in each period of the first $l\tau$ time periods is at least \hat{x}_i . Thus, by a coupling argument, conditioning on this event, D_k stochastically dominates $B(\hat{x}_i, \tau)$. Therefore, again by Hoeffding's inequality, we have for any k and any $b > 0$,

$$\begin{aligned} & \mathbf{P}(D_k - \hat{x}_i\tau < -b | D_k < x_i\tau + a, \forall k = 1, \dots, l) \\ & \leq \mathbf{P}(B(\hat{x}_i, \tau) - \hat{x}_i\tau < -b | D_k < x_i\tau + a, \forall k = 1, \dots, l) \\ & \leq \frac{\mathbf{P}(B(\hat{x}_i, \tau) - \hat{x}_i\tau < -b)}{\mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l)} \\ & \leq \frac{2^{-b^2\tau}}{\mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l)}. \end{aligned}$$

Thus by union bound,

$$\begin{aligned} & \mathbf{P}(D_k - \hat{x}_i\tau > -b, \forall k = 1, \dots, l | D_k < x_i\tau + a, \forall k = 1, \dots, l) \\ & \geq 1 - \frac{l2^{-b^2\tau}}{\mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l)}. \end{aligned}$$

Combine the previous inequalities, we get

$$\begin{aligned}
& \mathbf{P}(D_k - \hat{x}_i\tau > -b, \forall k = 1, \dots, l) \\
& \geq \mathbf{P}(D_k - \hat{x}_i\tau > -b, \forall k = 1, \dots, l | D_k < x_i\tau + a, \forall k = 1, \dots, l) \mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l) \\
& \geq \mathbf{P}(D_k < x_i\tau + a, \forall k = 1, \dots, l) - l2^{-b^2\tau} \\
& \geq 1 - l2^{-a^2\tau} - l2^{-b^2\tau}.
\end{aligned}$$

Note that the total sales of item i is at least $\sum_k D_k$, we have

$$\begin{aligned}
& \mathbf{P}(S_i > (\hat{x}_i\tau - b)l) \\
& \geq \mathbf{P}(D_k - \hat{x}_i\tau > -b, \forall k = 1, \dots, l) \\
& \geq 1 - l2^{-a^2\tau} - l2^{-b^2\tau}.
\end{aligned}$$

Let $\tau = T^{1/2} - T^{1/4}$, $l = T^{1/2}$, and let $a = b = 1$, then

$$1 - l2^{-\tau a^2} - l2^{-\tau b^2} = 1 - T^{1/2}2^{-T^{1/2} + T^{1/4} + 1},$$

and

$$\begin{aligned}
\hat{x}_i &= x_i - \frac{(l-1)a + (1-x_i)(x_i\tau + a)}{T - (l-1)\tau - (x_i\tau + a)} \\
&\geq x_i - \frac{la + \tau}{T - l\tau - \tau - a} \\
&= x_i - \frac{T^{1/2} + (T^{1/2} - T^{1/4})}{T - (T^{1/2} - T^{1/4})(T^{1/2} + 1) - 1} \\
&= x_i - \frac{2T^{1/2} - T^{1/4}}{T^{3/4} - T^{1/2} + T^{1/4} - 1} \\
&= x_i - \Theta(T^{-1/4}).
\end{aligned}$$

Thus, for given values of a, b, τ, l , we have

$$1 - T^{1/2}2^{-T^{1/2} + T^{1/4} + 1}$$

$$\begin{aligned}
&\leq \mathbf{P} \left(\frac{S_i}{x_i T} > \frac{(\hat{x}_i \tau - b)l}{x_i T} \right) \\
&\leq \mathbf{P} \left(\frac{S_i}{x_i T} > \frac{(x_i - \Theta(T^{-1/4}))(T - T^{3/4}) - T^{1/2}}{x_i T} \right) \\
&= \mathbf{P} \left(\frac{S_i}{x_i T} > 1 - \Theta(T^{-1/4}) \right).
\end{aligned}$$

Thus we obtain inequality (4.8).

Proving inequality (4.9):

Recall that in the proof of inequality (4.8), we show that with probability at least $1 - T^{1/2}2^{-T^{1/2}+T^{1/4}}$, for any period t in the first $T^{1/2}$ epochs of time periods, and any stage 1 item i ,

$$x_i^t \geq x_i - \Theta(T^{-1/4}).$$

We now consider stage item j . Similarly, we can show that with probability at least $1 - T^{1/2}2^{-T^{1/2}+T^{1/4}}$, for any period t in the first $T^{1/2}$ epochs of time periods, and any type i ,

$$y_j^{t,i} \geq y_j^i - \Theta(T^{-1/4}).$$

Thus, by union bound, with probability at least $1 - nT^{1/2}2^{-T^{1/2}+T^{1/4}}$, for any period t in the first $T^{1/2}$ epochs of time periods, the probability of selling j in period t satisfies

$$\begin{aligned}
\sum_i y_j^{t,i} x_i^t &\geq \sum_i (y_j^i - \Theta(T^{-1/4}))(x_i - \Theta(T^{-1/4})) \\
&= \sum_i y_j^i x_i - \Theta(T^{-1/4}).
\end{aligned}$$

Same argument as before then shows that

$$\mathbf{P} \left(\frac{S_j}{\sum_{i \in N} y_j^i x_i T} > 1 - \Theta(T^{-1/4}) \right) \geq 1 - nT^{1/2} 2^{-T^{1/2} + T^{1/4} + 1},$$

as desired. □

BIBLIOGRAPHY

- [1] Vineet Abhishek, Ian A Kash, and Peter Key. Fixed and market pricing for cloud services. *arXiv preprint arXiv:1201.5621*, 2012.
- [2] Philipp Afèche. Incentive-compatible revenue management in queueing systems: optimal strategic delay. *Manufacturing & Service Operations Management*, 15(3):423–443, 2013.
- [3] Philipp Afèche and Haim Mendelson. Pricing and priority auctions in queueing systems with a generalized delay cost structure. *Management Science*, 50(7):869–882, 2004.
- [4] Philipp Afèche and J Michael Pavlin. Optimal price/lead-time menus for queues with customer choice: Segmentation, pooling, and strategic delay. *Management Science*, 62(8):2412–2436, 2016.
- [5] Charalambos D Aliprantis and Kim Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer Science & Business Media, 2006.
- [6] Gad Allon and Awi Federgruen. Competition in service industries. *Operations Research*, 55(1):37–55, 2007.
- [7] Amazon. Amazon EC2 instance purchasing options. <http://aws.amazon.com/ec2/purchasing-options>, 2015. [Online; accessed 26-June-2015].
- [8] K.R. Balachandran. Purchasing priorities in queues. *Management Science*, 18(5-part-1):319–326, 1972.
- [9] Fernando Bernstein, A Gürhan Kök, and Lei Xie. Dynamic assortment customization with limited inventories. *Manufacturing & Service Operations Management*, 17(4):538–553, 2015.
- [10] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

- [11] Juan José Miranda Bront, Isabel Méndez-Díaz, and Gustavo Vulcano. A column generation algorithm for choice-based network revenue management. *Operations Research*, 57(3):769–784, 2009.
- [12] René Caldentey and Gustavo Vulcano. Online auction and list price revenue management. *Management Science*, 53(5):795–813, 2007.
- [13] Xi Chen, Will Ma, David Simchi-Levi, and Linwei Xin. Dynamic recommendation at checkout under inventory constraint. 2016.
- [14] Usman W. Chohan. Dididache - hailing the future, one cab at a time. <https://www.mcgill.ca/channels/news/dididache-%E5%98%80%E5%98%80%E6%89%93%E8%BD%A6-hailing-future-one-cab-time-235312>, 2014. [Online; accessed 4-Apr-2017].
- [15] James M Davis, Guillermo Gallego, and Huseyin Topaloglu. Assortment optimization under variants of the nested logit model. *Operations Research*, 62(2):250–273, 2014.
- [16] Guy Desaulniers, Jacques Desrosiers, and Marius M Solomon. *Column generation*, volume 5. Springer Science & Business Media, 2006.
- [17] Hila Etzion, Edieal Pinker, and Abraham Seidmann. Analyzing the simultaneous use of auctions and posted prices for online selling. *Manufacturing & Service Operations Management*, 8(1):68–91, 2006.
- [18] Jacob B Feldman and Huseyin Topaloglu. Capacity constraints across nests in assortment optimization under the nested logit model. *Operations Research*, 63(4):812–822, 2015.
- [19] Drew Fudenberg and David K Levine. Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory*, 44(1):1–18, 1988.
- [20] Drew Fudenberg and Jean Tirole. *Game theory*. The MIT Press, Cambridge, Massachusetts, 1991.

- [21] Guillermo Gallego and Ming Hu. Dynamic pricing of perishable assets under competition. *Management Science*, 60(5):1241–1259, 2014.
- [22] Guillermo Gallego, Woonghee Tim Huh, Wanmo Kang, and Robert Phillips. Price competition with the attraction demand model: Existence of unique equilibrium and its stability. *Manufacturing & Service Operations Management*, 8(4):359–375, 2006.
- [23] Guillermo Gallego, Garud Iyengar, R Phillips, and Abha Dubey. Managing flexible products on a network. 2004.
- [24] Guillermo Gallego, Anran Li, Van-Anh Truong, and Xinshang Wang. Online personalized resource allocation with customer choice. Technical report, Working Paper. <http://arxiv.org/abs/1511.01837> v1, 2016.
- [25] Guillermo Gallego and Huseyin Topaloglu. Constrained assortment optimization for the nested logit model. *Management Science*, 60(10):2583–2601, 2014.
- [26] Guillermo Gallego and Ruxian Wang. Multiproduct price optimization and competition under the nested logit model with product-differentiated price sensitivities. *Operations Research*, 62(2):450–461, 2014.
- [27] Natarajan Gautam. *Analysis of Queues: Methods and Applications*. CRC Press, 2012.
- [28] Srinagesh Gavirneni and Vidyadhar G Kulkarni. Self-selecting priority queues with burr distributed waiting costs. *Production and Operations Management*, 2016.
- [29] Amihai Glazer and Refael Hassin. Stable priority purchasing in queues. *Operations Research Letters*, 4(6):285–288, 1986.
- [30] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [31] Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-time

- optimization of personalized assortments. *Management Science*, 60(6):1532–1551, 2014.
- [32] Refael Hassin. Decentralized regulation of a queue. *Management Science*, 41(1):163–173, 1995.
- [33] Refael Hassin and Moshe Haviv. *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59. Springer, 2003.
- [34] Cinar Kilcioglu and Justin M Rao. Competition on price and quality in cloud computing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1123–1132. International World Wide Web Conferences Steering Committee, 2016.
- [35] Thomas Kittsteiner and Benny Moldovanu. Priority auctions and queue disciplines that depend on processing time. *Management Science*, 51(2):236–248, 2005.
- [36] Leonard Kleinrock. Optimum bribing for queue position. *Operations Research*, 15(2):304–318, 1967.
- [37] A Gürhan Kök, Marshall L Fisher, and Ramnath Vaidyanathan. Assortment planning: Review of literature and industry practice. In *Retail supply chain management*, pages 99–153. Springer, 2008.
- [38] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [39] Yuri Levin, Jeff McGill, and Mikhail Nediak. Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management science*, 55(1):32–46, 2009.
- [40] Guang Li, Paat Rusmevichientong, and Huseyin Topaloglu. The d-level nested logit model: Assortment and price optimization problems. *Operations Research*, 63(2):325–342, 2015.
- [41] Kyle Y Lin and Soheil Y Sibdari. Dynamic price competition with discrete

- customer choices. *European Journal of Operational Research*, 197(3):969–980, 2009.
- [42] Qian Liu and Garrett Van Ryzin. On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management*, 10(2):288–310, 2008.
- [43] Qian Liu and Dan Zhang. Dynamic pricing competition with strategic customers under vertical product differentiation. *Management Science*, 59(1):84–101, 2013.
- [44] Francis T Lui. An equilibrium queuing model of bribery. *Journal of Political Economy*, pages 760–781, 1985.
- [45] Costis Maglaras, John Yao, and Assaf Zeevi. Optimal price and delay differentiation in queueing systems. *Available at SSRN 2297042*, 2013.
- [46] Victor Martinez-de Albeniz and Kalyan Talluri. Dynamic price competition with fixed capacities. *Management Science*, 57(6):1078–1093, 2011.
- [47] Haim Mendelson and Seungjin Whang. Optimal incentive-compatible priority pricing for the $M/M/1$ queue. *Operations Research*, 38(5):870–883, 1990.
- [48] Isabel Méndez-Díaz, Juan José Miranda-Bront, Gustavo Vulcano, and Paula Zabala. A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Applied Mathematics*, 164:246–263, 2014.
- [49] Microsoft. Azure pricing. <https://azure.microsoft.com/en-us/pricing/>, 2017. [Online; accessed 4-Apr-2017].
- [50] Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, pages 1255–1277, 1990.
- [51] James A Mirrlees. An exploration in the theory of optimum income taxation. *The review of economic studies*, 38(2):175–208, 1971.

- [52] Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317, 1978.
- [53] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [54] Hamid Nazerzadeh and Georgia Perakis. Non-linear pricing competition with private capacity information. *Social Science Electronic Publishing*, 2015.
- [55] Hamid Nazerzadeh and Ramandeep S Randhawa. Asymptotic optimality of two service grades for customer differentiation in queueing systems. *Available at SSRN 2438300*, 2014.
- [56] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.
- [57] XZ Palmer. [how to]: Call a cab with didi chuxing and kuaidi dache. <http://www.smartshanghai.com/articles/sms/ how-to-call-a-taxi-with-didi-chuxing-and-kuaidi-dache>, 2015. [Online; accessed 4-Apr-2017].
- [58] Margaret P Pierson, Gad Allon, and Awi Federgruen. Price competition under mixed multinomial logit demand functions. *Management Science*, 59:8, 2013.
- [59] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [60] J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534, 1965.
- [61] Paat Rusmevichientong, David Shmoys, and Huseyin Topaloglu. Assortment optimization with mixtures of logits. Technical report, Tech. rep., School of IEOR, Cornell University, 2010.
- [62] Garrett van Ryzin and Siddharth Mahajan. On the relationship between inventory costs and variety benefits in retail assortments. *Management Science*, 45(11):1496–1509, 1999.

- [63] Kyle Siegrist. The birthday problem. <http://www.math.uah.edu/stat/urn/Birthday.html>, 2017. [Online; accessed 15-Aug-2017].
- [64] Springwise. In china, cab passengers bid tips to secure a float home. <https://www.springwise.com/china-taxi-passengers-bid-tips-secure-ride-home/>, 2013. [Online; accessed 4-Apr-2017].
- [65] Kalyan Talluri and Garrett Van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004.
- [66] Major Tian. Didi dache-kuaidi dache wants to do more than just drive you home. <http://knowledge.ckgsb.edu.cn/2015/05/12/china-business-strategy/didi-dache-and-kuaidi-dache-to-do-more-than-just-drive-you-home/>, 2015. [Online; accessed 4-Apr-2017].
- [67] Jan A Van Mieghem. Price and service discrimination in queuing systems: Incentive compatibility of $g\mu$ scheduling. *Management Science*, 46(9):1249–1267, 2000.
- [68] Xavier Vives. *Oligopoly pricing: old ideas and new tools*. MIT press, 2001.
- [69] Wei Wang, Baochun Li, and Ben Liang. Towards optimal capacity segmentation with hybrid cloud pricing. In *2012 IEEE 32nd International Conference on Distributed Computing Systems (ICDCS)*, pages 425–434. IEEE, 2012.
- [70] Tomer Yahalom, J Michael Harrison, and Sunil Kumar. Designing and pricing incentive compatible grades of service in queueing systems. Stanford University Working paper, 2006.
- [71] Luyi Yang, Laurens Debo, and Varun Gupta. Trading time in a congested environment. *Chicago Booth Research Paper*, (15-07), 2015.