Essays on Service Operations with Strategic Customers and Innovative Business Models

by

Andrew Edward Frazelle

Business Administration
Duke University

Date: _____

Approved:

_____
Alessandro Arlotto, Co-Supervisor

_____
Yehua Wei, Co-Supervisor

_____
Saša Pekeč

_____
Robert Swinney

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Business Administration
in the Graduate School of Duke University
2018

Abstract

Essays on Service Operations with Strategic Customers and Innovative Business Models

by

Andrew Edward Frazelle

Business Administration
Duke University

Date: _____

Approved:

_____
Alessandro Arlotto, Co-Supervisor

_____
Yehua Wei, Co-Supervisor

_____
Saša Pekeč

_____
Robert Swinney

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Business Administration in the Graduate School of Duke University
2018

Abstract

This dissertation studies operations management problems in service operations and supply chains, analyzing the impact of strategic customers and disruptive technologies. We investigate three problems, and insights from the study of each one illuminate the effect of information and/or strategic customer behavior on decision makers in operations management systems.

The first problem involves the strategic routing behavior of customers in a service network with multiple stations, when customers can choose the order of stations that they visit. We propose a two-station game with all customers present at the start of service and deterministic service times, and we find that strategic customers "herd," i.e., in equilibrium all customers choose the same route. For unobservable systems, we prove that the game is supermodular, and we then identify a broad class of learning rules—which includes both fictitious play and Cournot best-response—that converges to herding in finite time. By combining different theoretical and numerical analyses, we find that the herding behavior is prevalent in many other congested open-routing service networks, including those with arrivals over time, those with stochastic service times, and those with more than two stations. We also find that the system under herding performs very close to the first-best outcome in terms of cumulative system time.

The second problem relates to a disruptive, on-demand delivery platform who provides delivery from an independent sit-down restaurant. Food delivery platforms

maintain a symbiotic relationship with the existing providers in their industry; rather than "stealing" demand from an established player, these platforms work with restaurants to connect customers with the restaurant's product by providing an additional purchase channel. We model the restaurant as a queueing system with customer waiting costs. First, we solve the revenue maximization problem faced by a monopolist who controls both the dine-in and delivery prices and receives all revenues from the system. These results are related to the priority queueing and pricing literature and are of independent interest. We also demonstrate that a coordination problem exists between the restaurant and platform. We then investigate means of coordinating this supply chain via different contracts between the restaurant and the platform. We find that a two-way revenue-sharing contract coordinates the supply chain.

Finally, our third problem is spawned from an industry collaboration aimed at improving the inventory planning decisions of a company that sells high-tech goods with short life cycles. We develop a novel heuristic based on a power approximation in the extant literature. The power approximation computes near-optimal (s, S) policies for infinite-horizon inventory problems. We propose a new form of this approximation, devised using real demand data from our industry partner, and a heuristic based on the approximation that updates the inventory policy as new demand forecasts are generated. We evaluate the performance of our heuristic—also on real demand data, though for a different time period and for different items than were used to fit our model—and find that it performs quite well.

To my Lord and Savior Jesus Christ, and to Gabriela, my life's love.

Contents

Abstract **iv**

List of Tables **x**

List of Figures **xi**

Acknowledgements **xii**

1 Introduction **1**

1.1 Strategic Open Routing 2

1.2 Service Delivery Platforms 3

1.3 Lost-Sales Inventory for Products with Short Life Cycles 4

2 Strategic Open Routing in Service Networks **6**

2.1 Introduction 6

2.2 Literature Review 9

2.3 A Two-Station Service Network 13

2.4 The Open Routing Game 14

2.4.1 Submodularity of Expected System Time 19

2.4.2 Adaptive Dynamics Converge to Herding 25

2.4.3 Equilibrium Refinement 32

2.4.4 The Open Routing Game with Strategic and Non-Strategic Customers 38

2.4.5 The Stochastic Open Routing Game 40

2.5 The Sequential Open Routing Game 46

2.6 Herding vs. Central Optimum 53
2.7 The Prevalence of Herding 65
2.8 Queueing Networks in Steady State 75
2.9 Simulation Study for Systems with Stochastic Arrivals 78
2.9.1 Simulation Setup 78
2.9.2 Results of Simulation 79
2.10 Conclusion 83
3 Service Delivery Platforms: Pricing and Revenue Implications 87
3.1 Introduction 87
3.2 Related Literature 91
3.3 The Model 96
3.4 Customer Equilibrium 98
3.4.1 Customer Equilibrium 98
3.5 Monopoly Pricing 100
3.6 Decentralized System 113
3.6.1 Existence of Equilibrium 113
3.6.2 Coordinating Contracts 115
3.7 Numerical Tests: Coordination Loss and Mutual Gain 117
3.7.1 Coordination and Individual Revenues. 120
3.7.2 Prices and Control Limits. 121
3.7.3 Overall Outcomes. 123
3.8 Numerical Tests: One-Way Revenue Sharing 125
3.9 Conclusion 128
4 Lost-Sales Inventory Problems with Short Lifecycles and Forecast Updates 131
4.1 Introduction 131

4.2 The Problem and the Data 135
4.3 Related Literature 136
4.4 Power Approximations Old and New 139
4.4.1 The Power Approximation for (s, S) Policies 139
4.4.2 A New Power Approximation Using Firm X Data 142
4.4.3 Dynamic Heuristic for Computing (s_t, S_t) Policies 145
4.5 The Optimal Order Quantity: On-Hand vs. On-Order Inventory . . . 149
4.6 GUI-Based Tool and Performance of Heuristic 154
4.6.1 GUI-Based Tool 154
4.6.2 Evaluation of Dynamic Heuristic 154
4.7 Conclusion 161
5 Conclusion **162**
Bibliography **164**
Biography **171**

List of Tables

2.1 Maximum Percent Above Optimal Cumulative System Time Under Herding (i.e., $x = 0$ or $x = N$; $\mu_B = 1$) 64

2.2 Percent Deviation from Submodularity with Dedicated Customers (for μ_B=1, N_S=50). 69

2.3 Summary statistics with $\mu_A = .75$ for number of AB customers (out of 50) in the final round. 81

2.4 Summary statistics with $\mu_A = .5$ for number of AB customers (out of 50) in the final round. 83

3.1 Parameter ranges for numerical tests. In all cases $\mu = 10$ and $c_H = 10$. 117

3.2 Output from numerical tests. Besides frequency, values listed are averages for the given outcome. 125

4.1 Parameter values used to fit new power approximation. 143

4.2 Percentage point change in out of stock rate. 158

4.3 Percentage change in average inventory level. 159

4.4 Percentage point change in out of stock rate with lost-sales correction. 160

4.5 Percentage change in average inventory level with lost-sales correction. 160

List of Figures

2.1 AB customers visit station A first, while BA customers visit station B first. 13

2.2 Cumulative system time as a function of the number of AB customers x. 58

2.3 Frequency charts for number of AB customers in round 250 (100 trials). 80

3.1 Coordination loss versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).118

3.2 Equilibrium platform revenue versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 118

3.3 Restaurant's revenue gain with platform versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 119

3.4 Equilibrium prices and control limits versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 119

3.5 Restaurant's price with and without platform versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 120

3.6 Equilibrium prices versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 126

3.7 Coordination loss versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 126

3.8 Restaurant revenue with and without platform versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$). 127

4.1 Realized demand process for two sample items. 132

4.2 File input screen in GUI-based tool. 155

4.3 Inventory policy for a given date and forecast horizon. 155

4.4 Historical and simulated inventory levels for an example item. 156

Acknowledgements

My heartfelt thanks go to my co-advisors, Alessandro Arlotto and Yehua Wei. For their guidance, patience, brilliance, humor, and above all else, their precious time, I am forever grateful.

Yehua and I began meeting together within a month of my arrival at Fuqua, and since then he has mentored me with unfailing wisdom and kindness. His brilliant example and his perspective on what it means to be a scholar have inspired me to improve myself and my work beyond what I thought possible. I am incredibly thankful to Yehua for taking me on as an advisee, for undertaking a project based on the ramblings of an unseasoned first-semester student, and together with Alessandro, for somehow molding and shaping it and me to produce Chapter 2 of this dissertation.

I met Alessandro even before I joined Fuqua, on my visit as an admitted student. In every encounter since then I have been struck by his intellect, his wit, and his sense of humor. Far more importantly, throughout my time as a PhD student he has trained me—though I have often made the job difficult—to think, speak, and write with clarity and precision. His passion for intellectual rigor, while at home in mathematics, extends to all aspects of life. I have enjoyed many stimulating discussions with Alessandro on wide-ranging topics, and in every case I have benefited from his perspective and wisdom.

Both Alessandro and Yehua have given continually of themselves to support and mentor me these past five years. Each is a first-rate scholar, and each is a first-rate

human being. Through them I have witnessed the path of a world-class researcher, and it is one of my highest ambitions to walk it with anything like the skill and grace that they do.

I am grateful also to Robert Swinney and Pnina Feldman, under whose guidance Chapter 3 of this dissertation was written. Robert's PhD course gave me an introduction to the state of the art for research on supply chain management, and his way of synthesizing and assessing research has had a profound influence on me. Since then, it has been a pleasure to collaborate with him and benefit even more from his perspective. I have similarly benefited from Pnina's knowledge and insight on strategic customer and firm behavior in service systems, which have been incredibly valuable in conducting the research that led to Chapter 3.

Saša Pekeč, who brought me into the program and serves as the PhD coordinator for Decision Sciences, fosters a welcoming and stimulating environment for our PhD students, and I am tremendously thankful to him. Saša also had the wisdom to connect Yehua and me in my first semester, and I cannot thank him enough for that foresight. In addition, Saša has contributed greatly of his time and energy to the research in Chapter 4 of this dissertation, and his insight and ideas have been extremely valuable.

Myriad other members of the Fuqua community have made my time here intellectually rewarding and an absolute pleasure. I am grateful to David Brown for serving on my preliminary committee; to Nilay Argon at UNC–Chapel Hill for her course on design and control of queueing networks; to Peng Sun for his influential course on dynamic programming and for his advice throughout my time at Duke; and to Jim Smith and Bob Winkler for their insight and wisdom. My thanks go to Asa Palley, Levi DeValve, Sajad Modaresi, Mingliu Chen, Yunke Mai, Safak Yucel, and Troy Campbell for their friendship. Asa, Levi, Safak, Troy, and others also shared in my finest basketball moment, as intramural champions in Cameron Indoor Stadium.

Going slightly further back, I thank Jim Dai for inviting me to participate in an undergraduate research project during my time at Georgia Tech. His guidance and my experience conducting research with him were important factors in my decision to seek a PhD.

I would like to thank The Fuqua School of Business for funding my doctoral studies, and also the National Science Foundation, which has supported research in this dissertation under CAREER Award No. 1553274 (Principal Investigator Alessandro Arlotto).

My life has been beautifully enriched by the community at Love Chapel Hill Church led by Reverend Matt LeRoy, and by my wonderful friends: Tyler, Dayton, Christopher, Bennette, Sarah, Laura, Kate, and Glynis. I am incredibly grateful to my family—my parents, Ed and Pat; my sister Kelly; and my brother-in-law Travis—for their support both during and long before my PhD. I love them all. In particular, I am grateful to my father for inspiring me to pursue a PhD from the time when I was a child, so that I could walk in his footsteps.

Finally, words cannot express my love and gratitude for my wife, Gabriela. My love, you make my life a dream every day with your smile, your beautiful mind, your impossibly loving heart, and your sneaky grin. I love you, and I love our life.

1

Introduction

This dissertation studies operations management problems in service operations and supply chains, analyzing the impact of strategic customers and disruptive technologies. Advancing technology continually makes available new and more information to firms and consumers, and as the information structure in a system changes, the strategic interactions of the stakeholders evolve with it. Previously centralized systems may become decentralized, or a disruptive party or technology may provide additional information and options, upending the status quo.

The goal of this work is to provide insights to help managers and consumers adapt to a world that offers more opportunities to be strategic than ever before, and to improve the efficiency of the systems that they inhabit. To study such systems, we employ a variety of analytical techniques from operations management and operations research, including queueing theory, game theory, and inventory management.

We investigate three problems, and insights from the study of each one illuminate the effect of information and/or strategic customer behavior on decision makers in operations management systems. The first problem, which is presented in Chapter 2, involves the strategic routing behavior of customers in a service network with multiple

stations, when customers can choose the order of stations that they visit. The second problem is discussed in Chapter 3, and it relates to a disruptive, on-demand delivery platform who provides delivery from an independent sit-down restaurant; we investigate the effects of congestion on customer purchase decisions and the prices of the restaurant and the platforms. Finally, our third problem is spawned from an industry collaboration aimed at improving the inventory planning decisions of a company that sells high-tech goods with short life cycles. We discuss this problem in Chapter 4, and Chapter 5 concludes.

1.1 Strategic Open Routing

Much research has been devoted to understanding strategic customer behavior in service systems and other congested networks since the seminal works of Naor (1969) and Braess (1968), and there has even been some work looking specifically at the routing of customers in a queueing network, e.g., Cohen and Kelly (1990). My research on strategic routing contributes to these streams of literature with a new twist: we study a service network with *open routing*, i.e., one in which customers must visit multiple stations, but the order of stations that they visit is up to them.

Examples of open-routing service networks include catered receptions (where customers must visit the buffet and the beverage station, in either order) and amusement parks (where customers visit multiple rides, in any order). In these networks, a given customer's system time may be affected by another customer who chose a different route through the system, as even though the two chose different routes, they still may visit the same set of stations, just in a different order. The resulting dynamics reflect a complex interaction between the service and arrival processes of the different stations in the network. Chapter 2 poses two primary questions: in a service network with open routing, what can we discover about the routing decisions of customers whose objective is to minimize the total time that they spend in the network, and

what impact will these routing decisions have on efficiency?

In our base model, we study a two-station service network in which all customers must visit both stations and are present at the start of service availability. Intuitively, it might seem that customers seeking to minimize their expected system times would attempt to avoid each other and visit the less-crowded station first. However, our analysis reveals a surprising "herding" behavior; in any pure-strategy Nash equilibrium, all customers choose the same route through the network. This behavior is driven by the need to achieve a favorable position at the congested station. We also find that the social welfare under herding is within a constant of the optimal social welfare, and that this constant does not depend on the number of players; hence, for large systems herding quickly approaches the socially optimal outcome. By exploiting a key submodularity property of customers' expected system times, we establish the robustness of the herding phenomenon in different settings. Among other results, we show that customers following a broad class of learning rules will converge to herding over time and that any mixed-strategy Nash equilibria are unstable. Finally, we run a simulation study in a more general setting and find that herding continues to emerge in congested systems with stochastic service times and arrivals over time. Research in this chapter was supervised by Alessandro Arlotto and Yehua Wei, and most of the analyses therein have appeared in Arlotto et al. (2018).

1.2 Service Delivery Platforms

Food delivery platforms maintain a symbiotic relationship with the existing providers in their industry; rather than "stealing" demand from an established player, these platforms work with restaurants to connect customers with the restaurant's product by providing an additional purchase channel. While this relationship entails cooperation between the platform and the restaurant, there is substantial opportunity for

misalignment between them. We model the restaurant as an $M/M/1$ queue with customer waiting costs. First, we solve the revenue maximization problem faced by a monopolist who controls both the dine-in and delivery prices and receives all revenues from the system. These results are related to the priority queueing and pricing literature and are of independent interest. We also demonstrate that a coordination problem exists between the restaurant and platform, i.e., the system generates less than the optimal aggregate revenue when the restaurant and platform set their prices independently. We first demonstrate that a pure-strategy equilibrium exists for the pricing game played by the restaurant and platform, and then show that in such an equilibrium the prices can be sub-optimal for the system's aggregate revenue. We then investigate means of coordinating this supply chain via different contracts between the restaurant and the platform. We find that a two-way revenue-sharing contract coordinates the supply chain. This chapter is written under the supervision of Pnina Feldman and Robert Swinney, and results from this chapter have also appeared in Feldman et al. (2018).

1.3 Lost-Sales Inventory for Products with Short Life Cycles

Recently, the pace of innovation in high tech industries has become so rapid that individual products often have extremely short shelf lives. In such a world, some of the prevailing assumptions in research on inventory management have become less applicable to practice. For instance, much of the literature assumes that the distribution for the demand in a period is (i) known and (ii) stationary over time. For a product with a life cycle of little more than a year, it is perhaps unlikely that managers will be able in that time to pin down the true demand distribution, and just as unlikely that the distribution can be considered stationary over the product's lifespan.

We study the inventory management problem faced by a seller of high-tech goods

whose business is growing rapidly, for help with inventory modeling and planning. The demand for this company's products is highly volatile, and an inventory planning process is required that can scale well along with the scope and scale of the business. The key challenges are that the demand is nonstationary and difficult to predict, and also that the company operates in a lost-sales environment. The inventory problem is thus extremely challenging analytically. Because speed of computation and adaptability to changing circumstances are both crucial to our industry partner, we focus on providing a "good," fast solution rather than an optimal one.

In keeping with this philosophy, we develop a novel heuristic based on the power approximation developed in Ehrhardt (1979, 1984). The power approximation computes near-optimal (s, S) policies for infinite-horizon inventory problems. It is a static approximation that assumes stationary demand. We propose a new form of the power approximation, devised using real demand data from our industry partner, and a heuristic based on the approximation that updates the inventory policy as new demand forecasts are generated. We evaluate the performance of our heuristic—also on real demand data, though for a different time period and for different items than were used to fit our model—and find that it performs quite well. We also provide a GUI-based tool to our industry partner which implements our model and allows the recommendations to be easily incorporated into the inventory planning process. This research was supervised by Alessandro Arlotto, Saša Pekeč, and Yehua Wei.

2

Strategic Open Routing in Service Networks

2.1 Introduction

In many large entertainment or commercial service environments, various services are provided at different stations, and individuals are often free to strategically choose their route through the system. For example, consider a catered reception with a buffet line in the main dining room and a beverage cart in the adjacent hall. A guest wants to visit both stations, but she may choose to do so in either order. The guest wishes to select the sequence that minimizes the time that she spends waiting in queues, but given that other guests are probably also attempting to do the same, which sequence should she choose?

This example highlights the strategic behavior of rational customers in service networks with *open routing*, i.e., those in which customers can visit the stations in any order. The dynamics of service networks with open routing are more complex than those of systems in which customers merely choose a server, such as a grocery store where customers choose a checkout lane. In that example, once a customer has

Research in this chapter was supervised by Alessandro Arlotto and Yehua Wei and has also appeared in the joint paper Arlotto et al. (2018), published in *Management Science*.

chosen her preferred lane, she is not impacted by the service process of any lane but her own. However, in a network with open routing, customers who choose different routes can still end up at the same station at the same time, directly impacting each other's waiting times. For instance, a customer who finishes service at station A and joins the queue at station B may encounter another customer there who chose to visit station B first and will visit station A later.

Open routing characterizes a variety of other service environments in addition to the aforementioned catering example. For instance, Baron et al. (2016) study an outpatient medical clinic in which patients receive a battery of tests over several hours. Tests may be administered in any order, and they numerically study the problem of a central planner who seeks to maximize patient satisfaction. Open-routing service systems also include amusement parks (where customers choose the order in which they visit the attractions), shopping centers (where customers choose the order in which they visit the stores), and college orientations (where students must accomplish multiple tasks, such as a campus tour and a residence hall visit, in any order).

Motivated by these applications, we study how strategic customers choose their routes in a stylized open-routing service environment. In order to achieve a broad understanding of the role of open routing, we characterize analytically the behavior of strategic customers in a two-station network under some restrictive assumptions, and employ numerical simulation to relax some of these assumptions and show that the same insights also hold in more realistic settings.

In the base model, we study a two-station network in which customers require services from both stations. Each station serves its customers on a first-come, first-serve (FCFS) basis. To model strategic open routing, we assume that each customer chooses her own route to acquire service from both stations in the network, with the objective of minimizing her expected total time in the system, and that all customers

are present prior to the start of service. Intuitively, one might expect that a rational customer would attempt to avoid the popular route in order to receive service as quickly as possible. However, closer inspection reveals that a rational customer's decision is driven instead by the need to get into a favorable position at the more congested station. This leads to a surprising *herding* behavior: in any pure-strategy Nash equilibrium, all customers choose exactly the same route. This behavior is related to the notion of "follow the crowd" as discussed in Hassin and Haviv (2003) and Hassin (2016). After proving that any pure-strategy Nash equilibrium must involve herding, we then demonstrate that any mixed-strategy Nash equilibrium must be unstable. Moreover, we show that a large and intuitive class of adaptive learning dynamics—which includes both fictitious play and Cournot best-response—converges to herding at one of the two stations. The herding equilibria therefore emerge as the focal equilibria of the routing game.

While our base model does not directly solve any of the motivating problems discussed above, it uncovers the herding phenomenon in an open-routing service network, which bears further investigation. To verify the robustness of the herding behavior, we analyze strategic open routing in several different settings and its impact on social welfare. Our base two-station model in which all customers are strategic and present prior to the start of service is relaxed in several directions. Specifically, we relax our modeling assumptions by studying games with non-strategic customers, customers who visit only one of the two stations, and service networks with more than two stations. In all of these settings, we find that strategic customers herd. The sequential version of the original game also displays herding, with the minor twist that the last-moving player alone avoids the crowd. Additionally, we run a simulation study to test the robustness of herding. We find that, in the presence of congestion, herding continues to prevail in systems with stochastic service times and arrivals over time, with herding becoming less prominent as the arrival rate decreases

and the systems become less congested. We also find that the social welfare under herding—as measured by the sum of the system times experienced by all customers in the network—is within a constant of the first-best outcome, and such constant is independent of the number of players. For systems with many customers, this result implies that the welfare under herding differs only by a small fraction from the optimal social welfare.

The outline for the rest of the chapter is as follows. We review the related literature in Section 2.2, and we introduce our base model in Section 2.3. Section 2.4 discusses the *open routing game*, in which all customers are present in the system when service becomes available and make routing decisions about which station to visit first before learning their relative priority or the decisions of others. In Section 2.5, we derive the unique subgame perfect equilibrium of the sequential variant of the open routing game; in that setting, customers are aware of their priority and of the routing decisions of those with higher priorities. Section 2.6 studies the system from the perspective of a central planner attempting to optimize the routing assignments of customers. Section 2.7 relaxes some assumptions of the open routing game to allow for customers who wish to visit only one of the two stations, overloaded systems with stochastic arrivals, and systems with more than two stations. Section 2.8 studies systems which do not experience congestion. Section 2.9 analyzes the output from a simulation study designed to test the robustness of herding in networks with stochastic service times and arrivals over time. Finally, Section 2.10 makes concluding remarks.

2.2 Literature Review

Our work is related to the literature on queues with strategic customers as well as the literature on congestion games. The first stream was started by Naor (1969), who studied the use of tolls to induce desired behavior among customers in a queue.

The second was begun by Braess (1968), in his pioneering work demonstrating the now-famous Braess' Paradox.

Hassin and Haviv (2003) provide a comprehensive survey of the existing strategic queueing literature, and Hassin (2016) provides an extremely thorough coverage of the more recent work in the area. Our model is perhaps most closely related to the work of Parlaktürk and Kumar (2004). Theirs is one of only a few papers incorporating a stochastic network with stations in sequence in which customers choose the order of stations that they visit. The authors demonstrate the existence of unstable Nash equilibria for a two-station network in which every self-interested "job" must have two tasks performed on it, where each station can perform either task on a given job, and the first task takes on average a shorter time. Each station has a queue for Task 1 and another queue for Task 2, and the system planner may choose which queue to serve next at each station. By contrast, in our model each station can perform only one task and each station has only one queue; hence stations cannot dynamically distinguish between different classes of customers.

Other work in the vicinity includes that of Adiri and Yechiali (1974), who relay a model of multiple queues for a single server, arranged and priced by priority. Their results mirror ours in that arriving customers choosing a queue must take into account the possibility of being preempted (cut in front of) by an arriving customer who chooses a higher priority queue. Naturally, an important distinction is that in our model there are two queues for two servers, as opposed to two or more queues for one server in their work. Most recently, Honnappa and Jain (2015) study what they call the "network concert queueing game." For feedforward networks of several structures, they use fluid limits to determine symmetric equilibria and the price of anarchy when nonatomic users are allowed to choose both their arrival times and their routes through the network. Cohen and Kelly (1990) present an interesting analog of Braess' Paradox in a stochastic queueing network. They illustrate that when FCFS

nodes are placed in sequence with infinite-server nodes, the mean sojourn time in equilibrium can actually increase if customers are given the ability to switch from one track to the other. The analysis is done in steady state and queue lengths are assumed to be unobservable, simplifying both the queueing and game theoretic portions of the analysis. Other important work in this stream includes Enders et al. (2008) and Glazer and Hassin (1983).

On the empirical side, Pinilla and Prinz (2003) conduct a helpful study of flexible routing schemes (mainly simulation-based) for shape deposition manufacturing. They study the standard sequential model and employ simulation to gain insights on a system which allows flexibility. They propose two examples—shape deposition manufacturing and routing in a coffee shop—which can be cast within the open routing framework, and they find from simulation that performance can be significantly improved by dynamically assigning the sequence of tasks to be performed on each item, instead of following a fixed sequence of tasks.

Su and Zenios (2004) model the U.S. kidney allocation system as an $M/M/1$ queueing system in which potential recipients monitor both their position in the queue and the quality of the organ offered to them. They find that the current benchmark of FCFS service results in socially suboptimal allocation of organs because strategic recipients tend to refuse lower quality organs, knowing that they will likely be offered a better organ later. Modifying the queueing discipline to last-come, first-serve leads to the socially optimal outcome, although the authors note that such a discipline will likely be considered too unfair to implement in practice. Schummer (2017) and Leshno (2017) also study allocation of objects to strategic customers on queue-like waiting lists.

Finally, Veeraraghavan and Debo (2009, 2011) investigate competition between two service providers where queues build up and customers have private information regarding the quality of each provider, and they find herding behavior similar to what

we discover here. However, the driving force behind their result is quite different from what we observe. In Veeraraghavan and Debo (2009, 2011), customers are motivated by service quality, while in our model the herding behavior occurs because customers require services at both stations and attempt to minimize their expected total time in the system. In our setting, although by starting with the less-crowded station a customer may glean a shorter wait time before beginning her first service, she will afterwards face a severe penalty upon arriving to the congested station and being near the end of the queue there. Additional related work includes Afèche and Mendelson (2004), Debo and Veeraraghavan (2014), and Cui and Veeraraghavan (2016).

Our work is also related to the broad literature on strategic consumers in operations management. Notably, Cachon and Swinney (2009) introduce a model of fast-fashion retailing in which some consumers are strategic while others are not. They find that the benefit of quick-response, in which inventory can be replenished during the selling season, is greater in the presence of strategic consumers than otherwise. Other related work includes Swinney (2011) and Cachon and Swinney (2011).

Several papers in the congestion literature also warrant discussion. Feldman and Tamir (2012) also consider "jobs" (customers) to be strategic. Their paper focuses mainly on "conflicting congestion effects," both positive and negative, associated with the level of participation or traffic on a network. In their work, customers are allowed to choose from a set of machines working in parallel. They show that best-response dynamics do not always converge to a Nash equilibrium, but that the schedule generated by the longest processing time heuristic is indeed a Nash equilibrium if the number of machines is "right." An important difference from our analysis is that they do not incorporate the ordering of customers on a machine, instead modeling each player's cost function as merely the overall load of the machine chosen by that player. In the standard congestion model of which Braess' Paradox

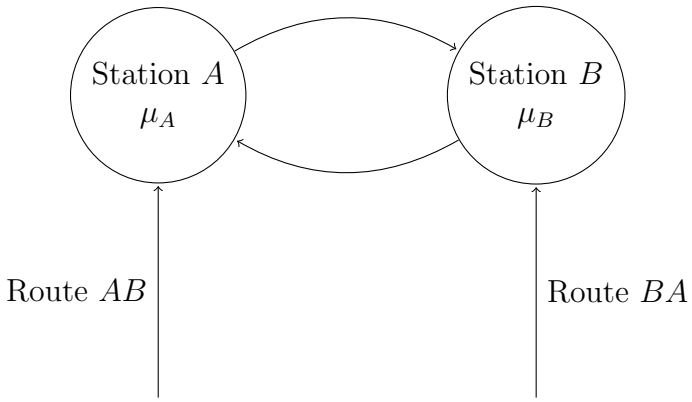

FIGURE 2.1: AB customers visit station A first, while BA customers visit station B first.

comprises a special case, Roughgarden and Tardos (2002) show that the price of anarchy is at most 4/3 when the performance measure is the total latency of the system and the latency functions are linear.

2.3 A Two-Station Service Network

Our setting is that of a service network with two stations, station A and station B, each with a single server, and with service rates μ_A and μ_B, respectively. The network serves N customers (or "players") who are all present in the system when service becomes available. We focus on the case in which the service rates are nonidentical, and thus without loss of generality we assume that the expected service time at station A is greater than that at station B (i.e., we have the service rate relation $\mu_A < \mu_B$). Similar to many service environments, each station operates on a FCFS basis. Figure 2.1 gives a visual depiction of the network.

Every customer must visit each station exactly once, but the order in which to visit the stations is unrestricted. Customers seek to minimize their expected total time in the system; their action space is the set $\{AB, BA\}$, where the first letter denotes the station visited first. We will at times refer to AB customers to identify

those who visit station A first. Similarly, BA customers are those who visit station B first.

We note that the centralized (non-strategic) version of the problem can be viewed as an *open-shop* scheduling problem with jobs that need to be processed by two servers, where either sequence is permissible for each job. A concise summary of the available results on open-shop scheduling can be found in Pinedo (2012, Chapter 8 and Section 13.4) and Baron et al. (2016), which we paraphrase here. For deterministic systems, there are polynomial-time algorithms available for makespan, but the problem of minimizing total completion time is NP-hard. For stochastic systems, the optimal policy for makespan is the longest expected remaining processing time first rule, and when the metric is expected total completion time, the preemptive shortest expected remaining processing time first policy is optimal among preemptive dynamic policies.

2.4 The Open Routing Game

In this section, we study the *open routing game*, in which all customers are present in the system when service becomes available and positions in the priority order are drawn uniformly at random. Players are aware that this randomization takes place, but they must make their routing decisions before learning their realized priorities or the routing decisions made by others. The open routing game is therefore equivalent to a symmetric one-shot game in which all players make routing decisions simultaneously. Players seek to minimize their expected total time spent in the system. Service times at stations A and B are first assumed to be deterministic with rates μ_A and μ_B, respectively, such that $\mu_A < \mu_B$. In Section 2.4.5, we relax the deterministic service assumption to incorporate stochastic service times. These dynamics resemble service environments in which a large number of customers are present before the service starts.

We first make two observations that help us understand the system time experienced by customers choosing each of the two routes through the network.

Property 1. *If station B ever becomes idle, then it will never build up a queue again.*

To understand Property 1, we note that, because $\mu_A < \mu_B$, the service time at station B is shorter than that at station A. Hence, after the first service begins, the arrivals to station B occur deterministically with an arrival rate that is smaller than station B's service rate. Therefore, once station B becomes idle, arrivals will never occur close enough together to form a queue. So, the system time for a customer who chooses route AB, and who has priority j at station A, depends on whether or not station B becomes idle before the customer departs station A. If station B idles, then the customer's system time is the sum of j service times at station A plus her own service time at station B. If station B does not idle, then the customer's system time is the sum of all of the service times at station B that must be completed up to and including herself. This will include all of the BA players, plus the $j-1$ players in front of her at station A, as well as herself.

Property 2. *Station A never idles from the time it begins its first service until it finishes serving all N customers.*

Similarly to Property 1, Property 2 follows from the fact that, because $\mu_A < \mu_B$, the service time at station B is shorter than that at station A. As soon as station A begins service, it will complete a service every $1/\mu_A$ units of time, but it will receive an arrival every $1/\mu_B < 1/\mu_A$ units of time until the last BA player departs station B. Station A will then never become idle until it finishes its workload. The one way in which station A can idle is at the beginning, if all N players choose route BA: in that case, station A will idle during the first service time at station B, after which it will work continually until it finishes with all N customers.

One might expect that players minimizing their system times would attempt to avoid each other and seek a less congested route. Instead, when the number of players N is large enough, we find that in equilibrium players "herd"—that is, all players take the same route through the network.

Proposition 1 (Herding Equilibria in the Open Routing Game)**.** *For $N \geqslant 2\mu_A/\mu_B + 1$, the open routing game has a Nash equilibrium in which all players "herd" at station A, that is, take route AB. Furthermore, if $\mu_B < 2\mu_A$ and $N \geqslant \max\{\mu_B/\mu_A + 1, (2\mu_A + \mu_B)/(2\mu_A - \mu_B)\}$, then the game also has a Nash equilibrium in which all players "herd" at station B.*

Proof. The fact that priorities are drawn uniformly at random implies that the game is symmetric, and we thus consider player (or customer) i, where i is an arbitrary player index. Suppose that all customers take route AB and that $N \geqslant 2\mu_A/\mu_B + 1$. The system time experienced by player i depends on which priority she is assigned. Because all customers are taking route AB, by Property 1, customer i will always find station B idle when she finishes service at station A. When assigned priority j, customer i will wait for $j-1$ players to be served at station A, be served herself, and then immediately be served at station B. Thus, as a function of her priority j, customer i's total system time $S^A(j)$ is given by

$$S^A(j) = \frac{j}{\mu_A} + \frac{1}{\mu_B}, \quad j = 1, \ldots, N.$$

Let $T(1, m)$ denote customer i's expected system time if she chooses route AB and m other customers also choose route AB, and let $T(0, m)$ denote customer i's expected system time if she chooses route BA and m other customers choose route AB. As priorities are drawn uniformly at random, customer i's expected system time when

following the candidate equilibrium strategy is given by

$$T(1, N-1) = \sum_{j=1}^{N} \frac{1}{N} S^A(j) = \sum_{j=1}^{N} \frac{1}{N} (\frac{j}{\mu_A} + \frac{1}{\mu_B}) = \frac{1}{\mu_B} + \frac{N+1}{2\mu_A}.$$

If she takes route BA, then player i will be behind all $N-1$ other players when she gets to station A. So, her total system time $T(0, N-1)$ is deterministic and is given by

$$T(0, N-1) = \frac{1}{\mu_B} + (\frac{N}{\mu_A} - \frac{1}{\mu_B}) = \frac{N}{\mu_A}.$$

Now, our assumption that $N \geqslant 2\mu_A/\mu_B + 1$ implies that

$$\frac{1}{\mu_B} + \frac{1}{2\mu_A} \leqslant \frac{N}{2\mu_A} \implies \frac{1}{\mu_B} + \frac{N+1}{2\mu_A} \leqslant \frac{N}{\mu_A}.$$

Therefore, customer i has no incentive to deviate as $T(1, N-1) \leqslant T(0, N-1)$, and we have a Nash equilibrium.

With the herding equilibrium at station A established, assume now that $\mu_B < 2\mu_A$ and $N \geqslant \max\{\mu_B/\mu_A + 1, 2\mu_A + \mu_B/(2\mu_A - \mu_B)\}$, and suppose that all customers take route BA. We will evaluate whether any customer has incentive to deviate. The condition $N \geqslant (\mu_B/\mu_A) + 1$ ensures that, if player i deviates and takes route AB, then station B will not finish serving all $N-1$ other players before player i finishes at station A. Applying the same notation as the previous case, customer i's total system time $T(1, 0)$ from taking route AB is deterministic and is given by

$$T(1, 0) = \frac{N}{\mu_B}.$$

If customer i takes route BA, then her priority at station B is drawn uniformly at random, and with probability $1/N$ she will be in position j, for $j = 1, \ldots, N$. Suppose that she draws priority j; she will wait for $j-1$ customers to be served at

station B, be served herself there, and then wait in a queue at station A. Because all customers take route BA, station A will idle for the first $1/\mu_B$ units of time, and then will work continuously until it has processed all N customers. Player i's system time corresponds to the time when station A finishes with the j-th job. Thus, we have that customer i's system time $S^B(j)$ is

$$S^B(j) = \frac{j}{\mu_A} + \frac{1}{\mu_B}.$$

We can now calculate her expected total system time $T(0,0)$ as

$$T(0,0) = \sum_{j=1}^{N} \frac{1}{N} S^B(j) = \sum_{j=1}^{N} \frac{1}{N} (\frac{j}{\mu_A} + \frac{1}{\mu_B}) = \frac{1}{\mu_B} + \frac{N+1}{2\mu_A}.$$

By our assumption that $\mu_B < 2\mu_A$ and $N \geqslant (2\mu_A + \mu_B)/(2\mu_A - \mu_B)$, we have

$$2\mu_A + \mu_B \leqslant (2\mu_A - \mu_B)N \implies \frac{1}{\mu_B} + \frac{N+1}{2\mu_A} \leqslant \frac{N}{\mu_B}.$$

Therefore we have $T(0,0) \leqslant T(1,0)$, implying that no customer has incentive to deviate, and we have a Nash equilibrium. □

We next give an intuitive explanation of Proposition 1. In the first candidate profile, all players visit station A first, so a player will have on average half of the other customers behind her if she visits station A first. However, if she visits station B first while everyone else visits station A first, then she will be the last customer to receive service at station A. Therefore, if N is large enough—in fact, $N \geqslant 3$ is sufficient here—then it is in her best interest not to deviate from the candidate profile. The herding equilibrium at station B can be similarly explained; intuitively, if N is large and the rest of the customers are slated to visit station B first, then a customer is better off being in front of an average of half of the other customers

at station B. Otherwise, after visiting station A first she will certainly have to wait behind all of the other players at station B.

Proposition 1 establishes the existence of Nash equilibria which exhibit herding behavior. However, we must answer several questions to verify that these herding equilibria are indeed plausible: (i) can the herding equilibria be implemented via simple, decentralized learning dynamics? (ii) are there other, non-herding Nash equilibria? and (iii) what happens when some of the players are not strategic? To address these questions, we establish a key submodularity property for the players' expected system times in the open routing game. This property then allows us to pinpoint the herding profiles as the focal equilibria of the game.

2.4.1 Submodularity of Expected System Time

Before deriving the submodularity property of the expected system time, we require some additional notation. Let i be an arbitrary player (or customer) index, s_i the strategy of player i, and $\mathbf{s}_{-i}$ the vector of strategies for all of the other players. We say that $s_i = 1$ if player i chooses route AB, and $s_i = 0$ if she chooses route BA. Similarly, a value of 1 for a given entry in $\mathbf{s}_{-i}$ means that the corresponding player has chosen route AB, while a value of 0 for a given entry in $\mathbf{s}_{-i}$ means that the corresponding player has chosen route BA. Denote by $T(s_i, \mathbf{s}_{-i})$ the expected system time for a player who employs strategy s_i when her opponents play the profile $\mathbf{s}_{-i}$. Note that the uniform priority randomization means that the game is symmetric, and thus we require no player index on T. Following the definition of Topkis (1998) and letting $\leqslant$ denote the usual partial order, we will show that the function $T(s_i, \mathbf{s}_{-i})$ is *submodular*, i.e., that it has *decreasing differences*. Specifically, we will find that

$$T(s_i, \mathbf{s}_{-i}) - T(\widetilde{s}_i, \mathbf{s}_{-i}) \leqslant T(s_i, \widetilde{\mathbf{s}}_{-i}) - T(\widetilde{s}_i, \widetilde{\mathbf{s}}_{-i}) \quad \text{for all } \widetilde{s}_i \leqslant s_i \text{ and } \widetilde{\mathbf{s}}_{-i} \leqslant \mathbf{s}_{-i}. \tag{2.1}$$

The decreasing differences condition (2.1) trivially holds if $\widetilde{s}_i = s_i$, so we can focus

on the case in which $\widetilde{s}_i = 0$ and $s_i = 1$. Moreover, because all customers are present when service starts and priorities are drawn uniformly at random, we can replace $\mathbf{s}_{-i}$ with the sum of its entries, m, and $T(s_i, \mathbf{s}_{-i})$ with $T(s_i, m)$. The variable m then simply represents the number of players—excluding player i—who have chosen route AB. With a slight abuse of notation, we can replace the decreasing differences condition (2.1) with

$$T(1, m) - T(0, m) \leqslant T(1, \widetilde{m}) - T(0, \widetilde{m}) \quad \text{for all } 0 \leqslant \widetilde{m} \leqslant m \leqslant N - 1. \tag{2.2}$$

We now introduce the shorthand

$$d_m := T(1, m) - T(0, m)$$

and rewrite condition (2.2) as

$$d_m \leqslant d_{\widetilde{m}} \quad \text{for all } 0 \leqslant \widetilde{m} \leqslant m \leqslant N - 1. \tag{2.3}$$

The difference d_m represents the relative preference of player i between route AB and route BA, given that m other players chose route AB; a negative value indicates that route AB will result in a shorter expected system time, and a positive value means that route BA will yield a shorter expected system time. Similarly, a relatively smaller value of d_m indicates a relatively greater preference for route AB. Therefore, if equation (2.3) holds—and we will show that it does—then the greater the number of other customers who have chosen route AB, the greater relative preference each customer will have for route AB.

Proposition 2 (Submodularity of Expected System Time)**.** *If*

$$N > N_{\text{sub}} := \frac{2\mu_B}{\mu_A} + 1,$$

then each player's expected system time in the open routing game is submodular. Moreover, we have

$$d_m < d_{m-1} \quad \text{for all } 1 \leqslant m < N - \frac{\mu_B}{\mu_A}. \tag{2.4}$$

Proof. As discussed, submodularity is equivalent to the decreasing differences condition (2.3). For our proof, we divide $T(1,m)$, the expected system time for a fixed player choosing route AB when m other players are also choosing the same route, into two components. We first define $\widetilde{T}(1,m)$ to be the expected system time for the player choosing route AB, not counting any waiting time that she may experience in the queue at station B. Because priorities are drawn uniformly at random, we have

$$\widetilde{T}(1,m) := \sum_{j=1}^{m+1} \frac{1}{m+1}\left(\frac{j}{\mu_A} + \frac{1}{\mu_B}\right) = \frac{1}{\mu_B} + \frac{m+2}{2\mu_A}. \tag{2.5}$$

Next, we define δ_m as the difference between the total expected system time and the expression in equation (2.5), i.e., $\delta_m := T(1,m) - \widetilde{T}(1,m)$. By definition, then, δ_m is equal to the expected time spent waiting in the queue at station B.

Station A never idles until it finishes as long as at least one player chooses route AB. Therefore, for $m \geqslant 1$, the expected system time for a player choosing route BA when m other players choose route AB is given by

$$T(0,m) = \sum_{\ell=1}^{N-m} \frac{1}{N-m}\left(\frac{m+\ell}{\mu_A}\right) = \frac{N+m+1}{2\mu_A}. \tag{2.6}$$

By equations (2.5) and (2.6), we have that for all $m \geqslant 1$,

$$d_m := T(1,m) - T(0,m) = \delta_m + \widetilde{T}(1,m) - T(0,m) = \delta_m + \frac{1}{\mu_B} - \frac{N-1}{2\mu_A}, \tag{2.7}$$

which implies that for all $1 \leqslant \widetilde{m} \leqslant m \leqslant N-1$,

$$d_m \leqslant d_{\widetilde{m}} \text{ if and only if } \delta_m \leqslant \delta_{\widetilde{m}}. \tag{2.8}$$

For $m \geqslant 1$, define $\delta_m^{(j)}$ to be the wait time in queue at station B experienced by the fixed player when m other players choose route AB, given that she chooses route AB and receives priority j at station A. Then we have the representations

$$\delta_m^{(1)} = \left(\frac{N-m-1}{\mu_B} - \frac{1}{\mu_A}\right)^+, \text{ and } \delta_m^{(j+1)} = \left(\delta_m^{(j)} + \frac{1}{\mu_B} - \frac{1}{\mu_A}\right)^+ \text{ for all } j = 1, \dots m.$$

Because $1/\mu_B - 1/\mu_A < 0$, we have that

$$\delta_m^{(j+1)} \leqslant \delta_m^{(j)} \text{ for all } j = 1, \ldots, m. \tag{2.9}$$

Moreover, because the function $f(x) := (x - 1/\mu_A)^+$ is non-decreasing, we also have

$$\delta_{m+1}^{(1)} \leqslant \delta_m^{(1)},$$

and therefore, for $j = 2, \ldots, m+1$,

$$\delta_{m+1}^{(j)} = \left(\delta_{m+1}^{(j-1)} + \frac{1}{\mu_B} - \frac{1}{\mu_A}\right)^+ \leqslant \left(\delta_m^{(j-1)} + \frac{1}{\mu_B} - \frac{1}{\mu_A}\right)^+ = \delta_m^{(j)}. \tag{2.10}$$

We have thus established that $\delta_m^{(j)}$ is monotonically decreasing both in j for any given m and in m for any given j. Next, because priorities are drawn uniformly at random, we have that

$$\delta_m = \frac{1}{m+1} \sum_{j=1}^{m+1} \delta_m^{(j)} \quad \text{for all } m \geqslant 1. \tag{2.11}$$

Equation (2.11) expresses δ_m as the average of the terms $\delta_m^{(1)}, \ldots, \delta_m^{(m+1)}$. By equation (2.9), $\delta_m^{(j)}$ is non-increasing in j, and therefore the average of $\delta_m^{(1)}, \ldots, \delta_m^{(m+1)}$ is bounded above by the average of $\delta_m^{(1)}, \ldots, \delta_m^{(m)}$. We then have

$$\delta_m = \frac{1}{m+1} \sum_{j=1}^{m+1} \delta_m^{(j)} \leqslant \frac{1}{m} \sum_{j=1}^{m} \delta_m^{(j)} \leqslant \frac{1}{m} \sum_{j=1}^{m} \delta_{m-1}^{(j)} = \delta_{m-1} \text{ for all } 2 \leqslant m \leqslant N-1, \tag{2.12}$$

where the second inequality comes from the monotonicity in m in equation (2.10). Equations (2.8) and (2.12) then imply that

$$d_m \leqslant d_{m-1} \quad \text{for all } 2 \leqslant m \leqslant N-1, \tag{2.13}$$

which satisfies the decreasing differences condition. Moreover, when $2 \leqslant m < N - \mu_B/\mu_A$, we have that

$$\delta_m^{(1)} = \left(\frac{N-m-1}{\mu_B} - \frac{1}{\mu_A}\right)^+ < \left(\frac{N-m}{\mu_B} - \frac{1}{\mu_A}\right)^+ = \delta_{m-1}^{(1)}, \tag{2.14}$$

so we conclude that the inequalities in equations (2.12) and (2.13) hold strictly in this range.

Lastly, we directly evaluate and compare d_1 and d_0. Observe that $N > N_{\text{sub}}$ implies that $2/\mu_A < (N-1)/\mu_B$, so when $m = 0$ or $m = 1$, if the fixed player chooses route AB, then she always faces a queue at station B. Therefore,

$$T(1,1) = \frac{1}{2}\left(\frac{N-1}{\mu_B} + \frac{N}{\mu_B}\right) = \frac{2N-1}{2\mu_B} \quad \text{and} \quad T(1,0) = \frac{N}{\mu_B}.$$

When the player chooses route BA, we note that station A will not idle if $m = 1$, but it idles for $1/\mu_B$ units of time if $m = 0$. Therefore,

$$T(0,1) = \frac{1}{N-1}\sum_{j=1}^{N-1} \frac{j+1}{\mu_A} = \frac{N+2}{2\mu_A}$$

and

$$T(0,0) = \sum_{j=1}^{N} \frac{1}{N}\left(\frac{j}{\mu_A} + \frac{1}{\mu_B}\right) = \frac{1}{\mu_B} + \frac{N+1}{2\mu_A}.$$

Applying the equations above, we have the relation

$$d_1 = \frac{N-1}{\mu_B} - \frac{N+1}{2\mu_A} + \left(\frac{1}{2\mu_B} - \frac{1}{2\mu_A}\right) < \frac{N-1}{\mu_B} - \frac{N+1}{2\mu_A} = T(1,0) - T(0,0) = d_0. \tag{2.15}$$

Equations (2.13) and (2.15) imply that the expected system time has decreasing differences and is therefore submodular, and equations (2.14) and (2.15) imply equation (2.4). □

Observe that $d_{\widetilde{m}} < 0$ implies that, if a player were aware that exactly $\widetilde{m}$ others were choosing route AB, then her expected system time for route AB would be less than that for route BA, and thus she would prefer route AB. Furthermore, the decreasing differences property gives us that if $d_{\widetilde{m}} < 0$, then $d_m < 0$ also for $\widetilde{m} \leqslant m$, and therefore for any $\widetilde{m} \leqslant m$, a player who knew that m others had chosen route AB would also want to choose route AB. Intuitively, a critical mass of players choosing a given route tends to attract the remaining players to the same route; if many players go to station A first, then the others should join them, and similarly for station B. We will revisit this idea in the next subsection when we discuss convergence of adaptive dynamics to the herding equilibria. The strict inequality in equation (2.4) of Proposition 2 plays a key role in establishing this convergence.

Proposition 2 implies that if, as we henceforth assume, players' utility functions decrease linearly with their expected system times, then their utility functions are supermodular, and the open routing game is a *supermodular game* in the sense of Topkis (1998, Section 4.1). Supermodular games have received much attention in the literature. For example, it is well documented that if a supermodular game has a unique Nash equilibrium, then a wide range of learning rules will converge to it (see Milgrom and Roberts, 1990). However, because the game that we study has multiple equilibria, we cannot directly apply the classical convergence result. Nevertheless, in the subsequent subsection, we apply supermodularity and the strict inequalities in equation (2.4) to prove that in our game, a large class of learning rules converges to one of the herding equilibria. Next, we state a corollary that follows directly from the proof of Proposition 2, which will also be used in the derivation of the subsequent convergence result.

Corollary 1. *If $N > N_{\text{sub}}$, then for any $m = 1, \ldots, N-1$, either $d_m < 0$ or $d_m < d_{m-1}$.*

Proof. If $m < N - \mu_B/\mu_A$, then the statement is immediately verified by Proposition 2. If $m \geqslant N - \mu_B/\mu_A$, we then must have

$$\frac{N-m}{\mu_B} \leqslant \frac{1}{\mu_A},$$

which implies that the queue at station B clears before the first AB customer arrives. Recall that $\delta_m = T(1,m) - \widetilde{T}(1,m)$ is equal to the expected time spent waiting in the queue at station B. Because the queue at station B clears before the first AB customer arrives, we have $\delta_m = 0$, and equation (2.7) implies that

$$d_m = \frac{1}{\mu_B} - \frac{N-1}{2\mu_A} < \frac{1}{\mu_B} - \frac{1}{\mu_A} < 0,$$

where the first inequality follows because $N > N_{\text{sub}} > 3$. □

2.4.2 Adaptive Dynamics Converge to Herding

As we have seen, when the service rates are close together and N is large, the herding strategy profiles are Nash equilibria. We next show that in addition, a general class of decentralized learning processes will converge to one of these herding equilibria.

We propose a model of learning which allows customers to update their beliefs in each period based on the play observed and also incorporates a "memory" of past actions. First, we assume that in each period players choose their routes to minimize their expected system times based on their current beliefs about other customers' strategies. When a player faces a tie, we assume that the player always chooses route AB. Let $\boldsymbol{\beta}_i^{(t)} = (\beta_{i,0}^{(t)}, \ldots, \beta_{i,N-1}^{(t)})$ be the vector of player i's beliefs, i.e., $\beta_{i,j}^{(t)}$ is player i's probability assessment, at the beginning of period t, that exactly j players (not including herself) will take route AB in period t, for $j = 0, 1, \ldots, N-1$. We allow the initial beliefs, $\boldsymbol{\beta}_i^{(1)}$, to be an arbitrary vector of probabilities summing to one. Next, let $x^{(t)}$ be the realized total number of players who take route AB in period

t, and let $x_i^{(t)}$ be the decision of player i (with $x_i^{(t)} = 1$ for choosing route AB and 0 otherwise). We also let $x_{-i}^{(t)}$ denote the realized number of players who take route AB in period t, excluding player i; that is, $x_{-i}^{(t)} = x^{(t)} - x_i^{(t)}$. Player i is making a decision in period t based on $\boldsymbol{\beta}_i^{(t)}$, her belief at the beginning of period t, which incorporates her experience up to and including period $t-1$.

For $m \in \{0, 1, \ldots, N-1\}$, let $\boldsymbol{e}(m) \in \mathbb{R}^N$ be the vector with a one in the $(m+1)$-st entry, and zeros in all of the remaining entries. Given a sequence of real numbers $\{\alpha_t\}$ with $0 \leqslant \alpha_t \leqslant 1$ for all $t = 1, 2, \ldots$, the beliefs in our model satisfy the recursion

$$\boldsymbol{\beta}_i^{(t+1)} = (1 - \alpha_t)\boldsymbol{\beta}_i^{(t)} + \alpha_t \boldsymbol{e}(x_{-i}^{(t)}) \quad t = 1, 2, \ldots. \tag{2.16}$$

We give the name $\{\alpha_t\}$-*learning* to the process in which beliefs evolve according to equation (2.16). Intuitively, a larger α_t implies that the players are giving more weight to their experience in period t, and less weight to earlier periods. We also note that choosing $\alpha_t = 1/t$ results in the familiar learning rule known as fictitious play introduced in Brown (1951). Within fictitious play learning, players best-respond to the empirical frequency of past moves. Similarly, by letting $\alpha_t = 1$ we recover the Cournot best-response model, in which players best-respond to the path of play realized in the prior period. The interested reader is referred to Fudenberg and Levine (1998) for detailed discussion of the Cournot model (Chapter 1), fictitious play (Chapter 2), and other learning models.

Next we show that under a mild regularity condition on the sequence $\{\alpha_t\}$ and on the initial beliefs $\boldsymbol{\beta}_i^{(1)}$, if all players update their beliefs according to the learning model (2.16), then play will converge to one of the herding equilibria.

Proposition 3 (Convergence of $\{\alpha_t\}$-Learning to Herding Equilibria)**.** *Define* m^* *by*

$$m^* = \min\{m \in \{1, \ldots, N\} : d_m \leqslant 0\},$$

and consider the $\{\alpha_t\}$*-learning process in equation* (2.16)*. If* (i) $N > N_{\text{sub}}$*,* (ii) *the sequence* $\{\alpha_t\}$ *is such that*

$$\lim_{\bar{t}\to\infty} \prod_{t=t'}^{\bar{t}} (1-\alpha_t) = 0 \quad \textit{for all } t' \geqslant 1, \tag{2.17}$$

and (iii) *there exists some* $\ell \geqslant 1$ *such that* $x^{(\ell)} \neq m^*$*, then players will converge to one of the herding equilibria in finitely many periods. That is, there exists* $t_0 < \infty$ *such that either*

$$x^{(\bar{t}+1)} = N \quad \textit{for all } \bar{t} \geqslant t_0,$$

$$\textit{or} \quad x^{(\bar{t}+1)} = 0 \quad \textit{for all } \bar{t} \geqslant t_0.$$

Proof. Let $\langle \cdot, \cdot \rangle$ denote the standard vector product, and let $\boldsymbol{d}$ be the vector of differences $d_m = T(1,m) - T(0,m)$ of waiting times for choosing routes AB and BA, i.e., $\boldsymbol{d} := (d_0, d_1, \ldots, d_{N-1})$. Define $\pi_i^{(t)}$ as the expected difference in system times for player i, given her belief at time t, which is given by

$$\pi_i^{(t)} := \langle \boldsymbol{\beta}_i^{(t)}, \boldsymbol{d} \rangle.$$

We proceed by cases.

Case 1: suppose that $x^{(\ell)} \geqslant m^* + 1$. Then for any player i who chose route AB in period ℓ, we must have $x_{-i}^{(\ell)} \geqslant m^*$, and also $\pi_i^{(\ell)} \leqslant 0$ because otherwise player i would have chosen route BA in period ℓ. This implies that

$$\pi_i^{(\ell+1)} = (1-\alpha_\ell)\pi_i^{(\ell)} + \alpha_\ell \langle \boldsymbol{e}(x_{-i}^{(\ell)}), \boldsymbol{d} \rangle \leqslant (1-\alpha_\ell)\pi_i^{(\ell)} + \alpha_\ell d_{m^*} \leqslant 0, \tag{2.18}$$

where the first inequality follows from the submodularity of the expected system time, and the second inequality follows by the definition of m^*. Equation (2.18) implies that all of the customers who chose route AB in period ℓ will do so again in periods $\ell+1$, $\ell+2$, ... Therefore, for any $t \geqslant \ell$, we have that $x_i^{(t)} = 1$ for each

player i that chose route AB in period ℓ, and this implies that

$$m^* + 1 \leqslant x^{(\ell)} \leqslant x^{(t)} \quad \text{for all } t \geqslant \ell.$$

For $\bar{t} \geqslant \ell$ and for any customer i who chose route BA in period ℓ, we have that

$$\begin{aligned}
\pi_i^{(\bar{t}+1)} &= \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \sum_{t=\ell}^{\bar{t}} \alpha_t \langle \boldsymbol{e}(x_{-i}^{(t)}), \boldsymbol{d}\rangle \prod_{s=t+1}^{\bar{t}}(1-\alpha_s) \\
&\leqslant \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \sum_{t=\ell}^{\bar{t}} \alpha_t d_{m^*+1} \prod_{s=t+1}^{\bar{t}}(1-\alpha_s) \\
&= \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \left(1 - \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\right) d_{m^*+1},
\end{aligned}$$

where the inequality follows by the submodularity of the expected system time. Note that we always have $d_{m^*+1} < 0$, because if $d_{m^*} < 0$, then $d_{m^*+1} \leqslant d_{m^*} < 0$; and if $d_{m^*} = 0$, then by Corollary 1, we have $d_{m^*+1} < 0$. Combining this with the assumption in equation (2.17), there must exist t_0 such that for all $\bar{t} \geqslant t_0$,

$$\pi_i^{(\bar{t}+1)} \leqslant \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \left(1 - \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\right) d_{m^*+1} < 0.$$

This implies that for all $\bar{t} \geqslant t_0$, we have $x_i^{(\bar{t}+1)} = 1$ for every customer i that chose route BA in period ℓ.

Case 2: suppose that $x^{(\ell)} \leqslant m^* - 1$. Then for any customer i who chose route BA in period ℓ, we must have $\pi_i^{(\ell)} > 0$ and $x_{-i}^{(\ell)} \leqslant m^* - 1$. This implies that

$$\pi_i^{(\ell+1)} = (1-\alpha_\ell)\pi_i^{(\ell)} + \alpha_\ell \langle \boldsymbol{e}(x_{-i}^{(\ell)}), \boldsymbol{d}\rangle \geqslant (1-\alpha_\ell)\pi_i^{(\ell)} + \alpha_\ell d_{m^*-1} > 0,$$

where the first inequality follows from the submodularity of the expected system time, and the second inequality follows from the definition of m^*. Therefore, for any $t \geqslant \ell$, we have that $x_i^{(t)} = 0$ for every customer i that chose route BA in period ℓ.

Finally, for $\bar{t} \geqslant \ell$ and for any customer i that chose route AB in period ℓ, we have

$$\pi_i^{(\bar{t}+1)} \geqslant \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \left(1 - \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\right) d_{m^*-2}.$$

The assumption in equation (2.17) and the fact that $d_{m^*-2} > 0$ together imply that there must exist t_0 such that for all $\bar{t} \geqslant t_0$,

$$\pi_i^{(\bar{t}+1)} \geqslant \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\pi_i^{(\ell)} + \left(1 - \prod_{t=\ell}^{\bar{t}}(1-\alpha_t)\right) d_{m^*-2} > 0.$$

This implies that for all $\bar{t} \geqslant t_0$, we have $x_i^{(\bar{t}+1)} = 0$ for each customer i that chose route AB in period ℓ. □

In Proposition 3, condition (2.17) essentially enforces that players' earlier beliefs and actions must eventually fade so that play does not get stuck on strategy profiles that are anchored to initial beliefs. For example, natural learning rules such as fictitious play and Cournot best-response satisfy this condition. This result reinforces the intuition that many players choosing a given route exerts a pull on others to do the same. The only circumstance which can possibly avoid herding is that in which $x^{(t)} = m^*$ for all $t = 1, 2, \ldots$. For this stagnation to occur, we must have $x^{(1)} = m^*$, an unlikely event if N is large and, say, initial beliefs are drawn as independent random vectors uniformly distributed on the N-dimensional probability simplex. Even if this first-period event is realized, route switching in later periods is inevitable because the beliefs of AB customers will be moving to favor route BA, and those of BA customers will be moving to favor route AB. When this route switching occurs, it must always be perfectly symmetric to maintain $x^{(t)} = m^*$, which is also unlikely.

The proof of Proposition 3 leverages the submodularity of the expected system time established by Proposition 2. Intuitively, if the number of players taking route AB is strictly greater than m^*, then the players currently taking route AB have no incentive to switch, in this period or in any later period, because the beliefs which led

them to select route AB in the current period will be further reinforced. Then, given that there will always be more than m^* players on route AB in the future periods, the players who chose route BA will eventually switch to route AB, provided that they sufficiently update their beliefs, as in condition (2.17). Similarly, if the number of players taking route AB is strictly less than m^* in some period, one can deduce that eventually all players will converge to route BA. This result reiterates the herding phenomenon: many players on a single route attract the others to join them.

Next, we state a corollary which shows that for Cournot best-response, the requirement of $x^{(\ell)} \neq m^*$ is not required when N is large enough.

Corollary 2 (Convergence of Cournot best-response)**.** *Under Cournot best-response, if* $N > \max\{N_{\text{sub}}, 2\mu_A/(\mu_B - \mu_A)\}$*, then players will converge to one of the herding equilibria in finitely many periods.*

Proof. By Proposition 3, we know that if there exists some $\ell \geqslant 0$ such that $x^{(\ell)} \neq m^*$, then Cournot best-response will converge to herding because it is a special case of $\{\alpha_t\}$-learning. We proceed to show that with Cournot best-response, there always exists some such ℓ. Consider an arbitrary path of play in which customers play Cournot best-response, that is, $\{\alpha_t\}$-learning with $\alpha_t = 1$ for all $t \geqslant 1$. In period t, $x^{(t)}$ players choose route AB. If $x^{(t)} \neq m^*$, then Proposition 3 implies that play will converge to herding in finitely many periods. If instead $x^{(t)} = m^*$, then we have $d_{x^{(t)}} \leqslant 0$ and $d_{x^{(t)}-1} > 0$ by the definition of m^*. Clearly, if N is odd, then $m^* \neq N/2$, and next, we show that $m^* \neq N/2$ when N is even. Define $Q_{\frac{N}{2}}$ by

$$Q_{\frac{N}{2}} := \min\{\frac{N}{2}, \lfloor \frac{\mu_A(\frac{N}{2} - 1)}{\mu_B - \mu_A} \rfloor\}.$$

The quantity $Q_{\frac{N}{2}}$ represents the number of AB customers who face a queue at station B when they depart station A, given that a total of $N/2$ customers chose route AB.

We now have

$$T(1, \frac{N}{2} - 1) = \sum_{k=1}^{Q_{\frac{N}{2}}} \frac{1}{(\frac{N}{2})} \left(\frac{k + \frac{N}{2}}{\mu_B} \right) + \sum_{k=Q_{\frac{N}{2}}+1}^{\frac{N}{2}} \frac{1}{(\frac{N}{2})} \left(\frac{k}{\mu_A} + \frac{1}{\mu_B} \right)$$

and

$$T(0, \frac{N}{2} - 1) = \sum_{j=1}^{\frac{N}{2}+1} \left(\frac{1}{\frac{N}{2} + 1} \right) \left(\frac{j + \frac{N}{2} - 1}{\mu_A} \right)$$

$$\geqslant \sum_{k=1}^{\frac{N}{2}} \frac{1}{(\frac{N}{2})} \left(\frac{j + \frac{N}{2} - 1}{\mu_A} \right), \tag{2.19}$$

where the inequality in equation (2.19) comes from the fact that the average of a set of $(N/2) + 1$ real numbers is larger than the average of the smallest $N/2$ numbers in the set. We can then write

$$d_{\frac{N}{2}-1} = T(1, \frac{N}{2} - 1) - T(0, \frac{N}{2} - 1)$$

$$\leqslant \sum_{k=1}^{Q_{\frac{N}{2}}} \frac{1}{(\frac{N}{2})} \left(\frac{k + \frac{N}{2}}{\mu_B} - \frac{k + \frac{N}{2} - 1}{\mu_A} \right) + \sum_{k=Q_{\frac{N}{2}}+1}^{\frac{N}{2}} \frac{1}{(\frac{N}{2})} \left(\frac{k}{\mu_A} + \frac{1}{\mu_B} - \frac{k + \frac{N}{2} - 1}{\mu_A} \right).$$

From the assumption that $N > 2\mu_A/(\mu_B - \mu_A)$, we get

$$\frac{k + \frac{N}{2}}{\mu_B} - \frac{k + \frac{N}{2} - 1}{\mu_A} = -\frac{N(\mu_B - \mu_A)}{2\mu_A\mu_B} + \frac{k}{\mu_B} - \frac{k-1}{\mu_A} < \frac{k-1}{\mu_B} - \frac{k-1}{\mu_A} \leqslant 0;$$

which implies that every term in the first summation above is strictly negative. Also, by the assumption that $N \geqslant N_{\text{sub}} + 1 = 2\mu_B/\mu_A + 2$, we have

$$\frac{1}{\mu_B} - \frac{N}{2\mu_A} + \frac{1}{\mu_A} \leqslant \frac{1}{\mu_B} - \frac{\mu_B}{\mu_A^2} < \frac{1}{\mu_B} - \frac{1}{\mu_A} < 0;$$

which implies that every term in the second summation is strictly negative. By this reasoning and the decreasing differences condition, we conclude that

$$d_{\frac{N}{2}} \leqslant d_{\frac{N}{2}-1} < 0,$$

contradicting the definition of m^*, which requires that $d_{m^*-1} > 0$. Therefore $m^* \neq N/2$. If $x^{(t)} = m^*$, then in period $t+1$ the customers who played route AB in period t will switch to route BA, and the customers who played route BA will switch to route AB. The fact that $m^* \neq N/2$ implies that $x^{(t+1)} \neq m^*$, and therefore play will converge to herding by Proposition 3. □

2.4.3 Equilibrium Refinement

We have established that in the open routing game, players following intuitive learning rules such as Cournot best-response will converge to one of the herding equilibria. We now again invoke the submodularity of the expected system time to show that the herding profiles are the only pure-strategy Nash equilibria for this game.

Corollary 3 (Equilibrium Refinement)**.** *If $N > N_{\text{sub}}$, then the only pure-strategy Nash equilibria of the open routing game are the symmetric herding equilibria of Proposition 1.*

Proof. Assume by contradiction that there exists a pure-strategy Nash equilibrium in which $0 < N_{AB} < N$ players choose route AB, and $N_{BA} = N - N_{AB}$ players choose route BA. These assumptions imply that $d_{N_{AB}} \geqslant 0$ and $d_{N_{AB}-1} \leqslant 0$. But by Proposition 2, we must also have $d_{N_{AB}} \leqslant d_{N_{AB}-1}$. This implies that $d_{N_{AB}} = d_{N_{AB}-1} = 0$, which contradicts Corollary 1. Therefore, there exist no pure-strategy Nash equilibria besides the herding equilibria of Proposition 1. □

However, Corollary 3 does not completely rule out the existence of Nash equilibria in which players adopt mixed strategies. Our next result rules out any Nash equilibria in which some but not all players use properly mixed strategies, and shows that any mixed-strategy Nash equilibrium is not stable enough to survive even the slightest perturbation in another player's strategy.

Proposition 4 (Elimination and Instability of Mixed Equilibria)**.** *If $N > N_{\text{sub}}$, then the open routing game has no Nash equilibria in which some players mix and others use pure strategies. Moreover, any Nash equilibrium in which all N players use mixed strategies is unstable; that is, a small perturbation in the strategy of any player will cause other players to strictly prefer a pure strategy.*

Proof. For players $i = 1, 2, \ldots, N$, let $s_i \in [0, 1]$ denote player i's strategy, specifically, the probability that player i chooses route AB. Assume by way of contradiction that there exists a Nash equilibrium with some players adopting mixed strategies and other players adopting pure strategies. Let N_{AB} be the number of customers playing the pure strategy of choosing route AB, N_{BA} be the number of players playing the pure strategy of choosing route BA, and $N_M := N - N_{AB} - N_{BA}$ be the number of customers playing "properly" mixed strategies, i.e., placing strictly positive probability on both routes. Let the index i be defined such that $s_i = 1$ for $i = 1, \ldots, N_{AB}$; $s_i = 0$ for $i = N_{AB} + 1, \ldots, N_{AB} + N_{BA}$; and $0 < s_i < 1$ for $i = N_{AB} + N_{BA} + 1, \ldots, N$. By our assumption, $N_{AB} + N_{BA} < N$ and either $N_{AB} \geqslant 1$ or $N_{BA} \geqslant 1$.

Let Γ_i denote the difference between player i's expected system time from choosing route AB and the expected system time from choosing route BA. Let η_k be the probability, given their strategies, that exactly k players choose route AB among players in the set $\{N_{AB} + N_{BA} + 2, \ldots, N\}$. Player $N_{AB} + N_{BA} + 1$'s difference between her expected system times from taking routes AB and BA can be expressed as

$$0 = \Gamma_{N_{AB}+N_{BA}+1} = \sum_{k=0}^{N_M - 1} \eta_k d_{N_{AB}+k}, \tag{2.20}$$

where the fact that $\Gamma_{N_{AB}+N_{BA}+1} = 0$ is true by assumption; if player $N_{AB} + N_{BA} + 1$ is employing a mixed strategy in this Nash equilibrium, then she must be indifferent between route AB and route BA.

Assume first that $N_{AB} \geqslant 1$. Then player 1 is choosing route AB. In this case, the difference between her expected system times from taking routes AB and BA can be expressed as

$$\Gamma_1 = s_{N_{AB}+N_{BA}+1} \sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k} + (1 - s_{N_{AB}+N_{BA}+1}) \sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k-1}$$

$$= (1 - s_{N_{AB}+N_{BA}+1}) \sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k-1},$$

where the last equality holds by equation (2.20). Moreover, because player 1 is choosing route AB, that route must be weakly better for her, implying that

$$\sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k-1} \leqslant 0. \tag{2.21}$$

Now, combining equations (2.20), (2.21), and the decreasing differences property, we get

$$\sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k-1} = \sum_{k=0}^{N_M-1} \eta_k d_{N_{AB}+k} = 0 \tag{2.22}$$

If the decreasing differences condition holds strictly (that is, if $d_{N_{AB}+k} < d_{N_{AB}+k-1}$) for some $k \in \{0, 1, \ldots N_M - 1\}$, then equation (2.22) cannot hold. Therefore we must have

$$d_{N_{AB}-1} = d_{N_{AB}} = \ldots = d_{N_{AB}+N_M-1} = 0. \tag{2.23}$$

Equation (2.23) contradicts Corollary 1, and thus N_{AB} must be equal to zero and therefore $N_{BA} \geqslant 1$. In that case, player 1 is choosing route BA, and the difference between her expected times from choosing routes AB and BA can be expressed as

$$\Gamma_1 = s_{N_{AB}+N_{BA}+1} \sum_{k=0}^{N_M-1} \eta_k d_{1+k} + (1 - s_{N_{AB}+N_{BA}+1}) \sum_{k=0}^{N_M-1} \eta_k d_k$$

$$= s_{N_{AB}+N_{BA}+1} \sum_{k=0}^{N_M-1} \eta_k d_{1+k},$$

where the last equality holds by equation (2.20). Moreover, because player 1 is choosing route BA, that route must be weakly better for her, implying that

$$\sum_{k=0}^{N_M-1} \eta_k d_{1+k} \geqslant 0. \tag{2.24}$$

Now, applying a similar argument as in the case with $N_{AB} \geqslant 1$, combining equations (2.20), (2.24), and the decreasing differences property gives us

$$d_0 = d_1 = \ldots = d_{N_M} = 0. \tag{2.25}$$

Equation (2.25) also contradicts Corollary 1. Therefore, there are no Nash equilibria of the open routing game in which some players adopt mixed strategies and other players adopt pure strategies.

Consider now an equilibrium in which all N players mix, that is, $N_M = N$ and $N_{AB} = N_{BA} = 0$. Let ν_k be the probability, given their strategies, that exactly k players choose route AB among players in the set $\{3, \ldots, N\}$. Player 1's expected savings from taking route BA instead of route AB is given by

$$\Gamma_1 = s_2 \sum_{k=0}^{N-2} \nu_k d_{1+k} + (1 - s_2) \sum_{k=0}^{N-2} \nu_k d_k = 0,$$

by assumption since player 1 must be indifferent if she plays a mixed strategy in equilibrium. Denote by $\widetilde{\Gamma}_1^{\epsilon}$ the perturbed expected savings for player 1 if player 2 increases her probability of route AB to $s_2 + \epsilon$. We have

$$\begin{aligned}
\widetilde{\Gamma}_1^{\epsilon} &= (s_2 + \epsilon) \sum_{k=0}^{N-2} d_{1+k} \nu_k + (1 - s_2 - \epsilon) \sum_{k=0}^{N-2} d_k \nu_k \\
&= \epsilon \left(\sum_{k=0}^{N-2} \nu_k (d_{1+k} - d_k) \right) \\
&< 0,
\end{aligned}$$

where the strict inequality follows from decreasing differences and the fact that $d_1 < d_0$, from the proof of Proposition 2. Therefore, for any $\epsilon > 0$, if player 2 perturbs her strategy by increasing her probability of route AB to $s_2 + \epsilon$, then s_1 is no longer a best response for player 1, and thus this mixed-strategy Nash equilibrium is not "ϵ-stable." The same argument applies to the perturbation of any one player's strategy. It should be observed that this notion of ϵ-stability is related to but much stronger than the notion of an "evolutionarily stable strategy" as defined in Hassin and Haviv (2003). □

Therefore, the herding equilibria are the only pure-strategy Nash equilibria of the open routing game, and in any mixed-strategy Nash equilibrium it must be that all players employ properly mixed strategies (i.e, no one plays a pure strategy). Such mixed-strategy equilibria, however, are quite unstable and unlikely to be implemented.

Exploiting the decreasing differences property of the expected system time, we also find that if there is significant service rate disparity—specifically, if service at station B is more than twice as fast as service at station A—then route AB is a strictly dominant strategy for all players.

Corollary 4 (Dominant Strategy)**.** *If $2\mu_A \leqslant \mu_B$ and $N > N_{\text{sub}}$, then route AB is a strictly dominant strategy for all players.*

Proof. Proposition 2 tells us that if $N > N_{\text{sub}}$, then the game has decreasing differences, i.e., $d_{N-1} \leqslant d_{N-2} \leqslant \cdots \leqslant d_1 \leqslant d_0$. The relation $2\mu_A \leqslant \mu_B$ then gives us that

$$\frac{N-1}{\mu_B} < \frac{N+1}{2\mu_A},$$

which in turn implies that

$$d_0 = T(1,0) - T(0,0) = \frac{N-1}{\mu_B} - \frac{N+1}{2\mu_A} < 0.$$

Decreasing differences and the fact that $d_0 < 0$ gives us that $d_m = T(1,m) - T(0,m) < 0$ for all $m \in \{0, 1, \ldots, N-1\}$. Therefore a player's expected system time is always smaller for route AB than route BA, no matter how many other players choose route AB. □

The dominant strategy result can be understood via the following rough argument. Consider a customer contemplating her move, and suppose that the other $N-1$ customers are playing route BA. If she follows the crowd on route BA, then, on average, she will be behind half of the other customers, first at station B and then again at station A. Since station A is the bottleneck, her expected system time will be on the order of $N/(2\mu_A)$. If instead she chooses route AB, then her system time will be a deterministic N/μ_B. Because $2\mu_A \leqslant \mu_B$, it is better for her to choose route AB over route BA. Then, because $N > N_{\text{sub}}$, if any customer shifts to route AB, then the decreasing differences property implies that her preference for route AB will only increase, and therefore route AB is her optimal strategy regardless of the moves of the other customers. So, if the service rates are close together, then both herding profiles are Nash equilibria, while if they are far apart, then it is a dominant strategy to visit the slower station A first.

With these results in favor of the herding equilibria and the fact that an intuitive class of learning rules converges to them, we see that the herding effect exerts a strong influence on the behavior of rational customers in the open routing game. We note that this herding effect coincides with several examples of open routing. At catered events, anecdotal evidence suggests that guests often flock to the buffet line (the slower station) even when others are doing the same, presumably to avoid waiting in an even longer queue there if they instead got their drink first; also in theme parks, although customers do not all start at the same station, many customers rush to the most popular rides first to experience them before the queue grows long.

2.4.4 The Open Routing Game with Strategic and Non-Strategic Customers

In practice some (perhaps many) customers may not attempt to minimize their overall expected system times. For example, at a catered reception a guest may have a preference which dictates a certain route irrespective of its effect on her total system time. Accordingly, we now examine the effect on the system of customers who are not strategic, that is, customers who must visit both stations but who have a pre-determined route which they will follow without contemplating any alternative.

To this end, consider a system as described in Section 2.3 with N customers, and again assume that priorities are drawn uniformly at random. Suppose now that $N = N_{AB} + N_{BA} + N_S$, where N_{AB} is the number of non-strategic customers who will take route AB no matter what, N_{BA} is the number of non-strategic customers who will take route BA no matter what, and N_S is the number of strategic customers. In this system, a strategic customer could feasibly encounter $N_{AB}, \ldots, N_{AB} + N_S - 1$ other customers choosing route AB. Therefore, we can focus on

$$d_m = T(1, m) - T(0, m), \quad \text{for } m = N_{AB}, \ldots, N_{AB} + N_S - 1.$$

Recall that d_m represents the relative preference between route AB and route BA for a strategic customer, when there are m other customers choosing route AB.

By Proposition 2, we immediately have that d_m is decreasing in m in the range of $N_{AB}, \ldots, N_{AB} + N_S - 1$ as long as $N > N_{\text{sub}}$. Therefore, as in Section 2.4.1, the expected system times for the customers that are strategic are submodular. The submodularity in turn implies that the herding equilibria (among the strategic customers) prevail as the only pure-strategy Nash equilibria. This is summarized in the next corollary.

Corollary 5 (Equilibria with Non-Strategic Customers)**.** *If $N > N_{\text{sub}}$, then exactly one of the following holds:*

(i) *route AB is a dominant strategy for all strategic customers;*

(ii) *route BA is a dominant strategy for all strategic customers;*

(iii) *both herding profiles are Nash equilibria, and there are no other pure-strategy Nash equilibria.*

Proof. Similar to the argument of Corollary 3, suppose that there is a pure-strategy Nash equilibrium besides the herding equilibria. Then there must exist some $N_{AB} < m < N_{AB} + N_S - 1$ such that $d_m \geqslant 0$ and $d_{m-1} \leqslant 0$. Combining this with Proposition 2, we must have $d_m = d_{m-1} = 0$, but this contradicts Corollary 1. Therefore, there cannot be non-herding pure-strategy Nash equilibria. Next, we identify three regimes corresponding to the three conclusions in the statement of the proposition.

Case 1: $d_{N_{AB}} < 0$.

If $d_{N_{AB}} < 0$, then even if only non-strategic customers take route AB, it is in the interest of each strategic customer to deviate to route AB. Decreasing differences implies that it is a dominant strategy for all of the strategic players to choose route AB.

Case 2: $d_{N_{AB}+N_S-1} > 0$.

If $d_{N_{AB}+N_S-1} > 0$, then even if all of the other strategic customers are taking route AB, a strategic customer would rather take route BA. Decreasing differences gives that it is a dominant strategy for all of the strategic players to choose route BA.

Case 3: $d_{N_{AB}} \geqslant 0$ and $d_{N_{AB}+N_S-1} \leqslant 0$.

In this case, both herding profiles are Nash equilibria. □

We note that Proposition 3 also holds in this setting with non-strategic customers, in the sense that strategic customers will converge to one of the herding equilibria under a general class of learning dynamics. These results suggest that customers who are aware that some others are not rational can still implement an equilibrium profile which involves herding. Moreover, depending on the system parameters, herding at

one station can be a dominant strategy.

2.4.5 The Stochastic Open Routing Game

We now discuss how the herding behavior that we observe in the open routing game continues to emerge when the service times are stochastic. We refer to this new setting as the *stochastic open routing game*. Specifically, for $\mu_A < \mu_B$ we now suppose that the service times at stations A and B are independent and identically distributed random variables with means $1/\mu_A$ and $1/\mu_B$ and variances σ_A^2 and σ_B^2, respectively.

Properties 1 and 2 do not extend to this environment because the uncertainty in service times may induce idle and busy periods that did not arise in the deterministic setting. As a result, we can no longer express customers' expected system times in closed form. Nevertheless, we can bound the expected system times when all customers herd by applying existing results from queueing theory. More specifically, when all customers choose route AB, the waiting times at station B behave as in an underloaded $GI/GI/1$ queue (single-server queue with general independent arrival and service time distributions) that started empty, with arrival rate μ_A and service rate μ_B. Then, a stochastic dominance argument from Müller and Stoyan (2002, Theorem 6.2.1) tells us that we can bound the expected waiting time at station B for any customer in this $GI/GI/1$ queue by the expected steady-state waiting time. Combining this argument with the classical steady-state result of Kingman (1962), we then bound the expected system time for all of the customers who take route AB, and this bound holds for any value of N. Under suitable conditions, we have that our upper bound for the expected system time on route AB is smaller than the expected system time on route BA, and therefore it is a Nash equilibrium for all customers to choose route AB.

Proposition 5 (Nash Equilibrium with Route AB—Stochastic Service)**.** *If*

$$N > 1 + \frac{2\mu_A}{\mu_B} + \frac{\mu_A^2(\sigma_A^2 + \sigma_B^2)}{1 - \mu_A/\mu_B}, \tag{2.26}$$

then it is a symmetric Nash equilibrium for all customers to choose route AB.

Proof. Assume that all N players visit station A first, so station B is initially empty until the first departure from station A. Observe that if all N customers visit station A first, then station B behaves like a $GI/GI/1$ queueing system with arrival rate μ_A and service rate μ_B (recall that $\mu_A < \mu_B$, so such a system would be stable). Let $F_{W_0^B}$ be the distribution function for a random variable which is independent of the arrival and service processes and which may modify the initial state of the queueing system. Let W_k^B, $k \geqslant 1$, be the waiting time that the k-th departure from station A experiences at station B, and let $F_{W_k^B}$ be the distribution function of W_k^B. Note that with probability 1 we have $W_0^B = 0$ and $W_1^B = 0$ because station B is initially empty. Therefore, $F_{W_1^B}$ stochastically dominates $F_{W_0^B}$, which we denote as

$$F_{W_0^B} \leqslant_{st} F_{W_1^B}.$$

Define $F_{W_\infty^B}$ as the stationary waiting-time distribution function for a $GI/GI/1$ queueing system with arrival rate μ_A and service rate μ_B. By Theorem 6.2.1 in Müller and Stoyan (2002), we then have that

$$F_{W_k^B} \leqslant_{st} F_{W_\infty^B} \quad \text{for all } k = 1, 2, \ldots. \tag{2.27}$$

Armed with the stochastic dominance relation (2.27), we consider the candidate equilibrium profile in which all customers visit station A first and evaluate the prospect of deviating to route BA. If the customer follows her current strategy, she will receive priority k at station A, for $k = 1, \ldots N$, with probability $1/N$. Conditional on her priority k, her expected total system time $\mathbb{E}[S^A|k]$ is given by

$$\mathbb{E}[S^A|k] = \frac{k}{\mu_A} + \mathbb{E}[W_k^B] + \frac{1}{\mu_B}.$$

Equation (2.27) implies that $\mathbb{E}[W_k^B] \leqslant \mathbb{E}[W_\infty^B]$, where $\mathbb{E}[W_\infty^B]$ is the steady-state expected waiting time from the distribution function $F_{W_\infty^B}$. We then have the bound

$$\mathbb{E}[S^A|k] \leqslant \frac{k}{\mu_A} + \mathbb{E}[W_\infty^B] + \frac{1}{\mu_B} \leqslant \frac{k}{\mu_A} + \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B}, \tag{2.28}$$

where the last inequality follows from bounds for the steady-state expected waiting time in queue for a $GI/GI/1$ queue found in Kingman (1962). Taking expectation over the priority in equation (2.28), we have

$$\begin{aligned}\mathbb{E}[S^A] &\leqslant \frac{1}{N}\sum_{k=1}^{N}(\frac{k}{\mu_A} + \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B}) \\ &= \frac{N+1}{2\mu_A} + \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B}.\end{aligned}$$

If a customer deviates and visits station B first, then she will be the last to be served at station A, so her expected system time $\mathbb{E}[S^B]$ satisfies $\mathbb{E}[S^B] \geqslant N/\mu_A$. Finally, by equation (2.26), we have

$$\begin{aligned}\frac{2\mu_A}{\mu_B} + \frac{\mu_A^2(\sigma_A^2 + \sigma_B^2)}{1-\mu_A/\mu_B} &< N-1 \\ \Rightarrow \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B} &< \frac{N-1}{2\mu_A} \\ \implies \mathbb{E}[S^A] \leqslant \frac{N+1}{2\mu_A} + \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B} &< \frac{N}{\mu_A} \leqslant \mathbb{E}[S^B].\end{aligned}$$

As a customer's expected total system time is shorter if she follows the candidate profile and visits station A first, she has no incentive to deviate. We conclude that it is a Nash equilibrium for all customers to visit station A first. □

Next, we consider the case in which all customers are choosing route BA. In this case, the waiting times at station A behave as in an overloaded $GI/GI/1$ queue with arrival rate μ_B and service rate μ_A, which does not have a steady-state distribution.

To bound the expected waiting time in an overloaded $GI/GI/1$ queue, we apply Lindley's equation to show that the total idling time is represented by the "dual" process of the overloaded system, and that this "dual" process has the same distribution as the waiting time in an underloaded $GI/GI/1$ queue (see Grimmett and Stirzaker, 2001, Section 11.5). This allows us to bound the total expected idling time and to show that under suitable conditions it is a Nash equilibrium for all customers to choose route BA. The result is summarized in the next proposition.

Proposition 6 (Nash Equilibrium with Route BA—Stochastic Service)**.** *If we have* $\mu_A < \mu_B < 2\mu_A$ *and*

$$N \geqslant \frac{\mu_B + 2\mu_A}{2\mu_A - \mu_B} + \frac{\mu_A^2 \mu_B^2 (\sigma_A^2 + \sigma_B^2)}{(2\mu_A - \mu_B)(\mu_B - \mu_A)}, \tag{2.29}$$

then it is a symmetric Nash equilibrium for all customers to choose route BA*.*

Proof. Let X_k^A and X_k^B be the service time experienced by the k-th customer to be served at stations A and B, respectively. Suppose that all players visit station B first, and let W_k^A be the waiting time at station A experienced by the k-th departure from station B. Define U_k by

$$U_k := X_k^A - X_{k+1}^B.$$

By Lindley's equation, then, we have

$$\begin{aligned} W_1^A &= 0 \\ W_2^A &= \max\{0, W_1^A + U_1\} = \max\{0, U_1\} \\ &\cdots \\ W_{k+1}^A &= \max\{0, U_k, U_k + U_{k-1}, \ldots, U_k + U_{k-1} + \cdots + U_1\}. \end{aligned}$$

Now, define I_k^A to be the cumulative idle time experienced by station A before the arrival of the k-th customer, when all customers visit station A first. We can relate

the waiting time W_k^A of the k-th customer to the idle time and the excess workload by

$$\begin{aligned} I_k^A &= W_k^A - \left(\sum_{i=1}^{k-1} X_i^A - \sum_{j=1}^{k} X_j^B \right) \\ &= W_k^A - \sum_{i=1}^{k-1} U_i + X_1^B \\ &= X_1^B + \max \left\{ -\sum_{i=1}^{k-1} U_i, -\sum_{i=1}^{k-2} U_i, \ldots, -U_1, 0 \right\}. \end{aligned}$$

Note then that $-U_i$ has the same distribution as $X_{k-i}^B - X_{k-i+1}^A$, and therefore, $I_k^A - X_1^B$ has the same distribution as

$$W_k^B := \max \left\{ \sum_{i=1}^{k-1} (X_i^B - X_{i+1}^A), \sum_{i=2}^{k-1} (X_i^B - X_{i+1}^A), \ldots, (X_{k-1}^B - X_k^A), 0 \right\}, \tag{2.30}$$

which is the "dual process" of W_k^A. Moreover, by Lindley's equation, we can view equation (2.30) as the wait time of a single server queue with interarrival time distribution corresponding to station A's service time distribution, and service time distribution corresponding to station A's interarrival time distribution. From equation (2.27) in the proof of Proposition 5 and Kingman (1962), we have that

$$\mathbb{E}[W_k^B] \leqslant \mathbb{E}[W_\infty^B] \leqslant \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1 - \mu_A/\mu_B)}. \tag{2.31}$$

Let S^A be the system time associated with choosing route AB, and S^B be the system time associated with choosing route BA. Now, given a priority k at station

B, we have

$$\begin{aligned}\mathbb{E}[S^B|k] &= \frac{k}{\mu_A} + \mathbb{E}[I_k^A] \\ &= \frac{k}{\mu_A} + \frac{1}{\mu_B} + \mathbb{E}[W_k^B] \\ &\leqslant \frac{k}{\mu_A} + \frac{1}{\mu_B} + \mathbb{E}[W_\infty^B],\end{aligned}$$

where the equality follows from the fact that $I_k^A - X_1^B$ has the same distribution as W_k^A, while the inequality follows from equation (2.31). Applying this inequality, we get that

$$\begin{aligned}\mathbb{E}[S^B] &\leqslant \mathbb{E}[W_\infty^B] + \frac{1}{\mu_B} + \sum_{k=1}^{N} \frac{1}{N}\left(\frac{k}{\mu_A}\right) \\ &= \mathbb{E}[W_\infty^B] + \frac{1}{\mu_B} + \frac{N+1}{2\mu_A} \\ &\leqslant \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} + \frac{1}{\mu_B} + \frac{N+1}{2\mu_A},\end{aligned}$$

with the last inequality following from equation (2.31). Finally, because a player deviating to route AB will be the last customer served at station B, we clearly have

$$\mathbb{E}[S^A] \geqslant \frac{N}{\mu_B}.$$

Equation (2.29) then gives

$$\begin{aligned}&\frac{1}{2\mu_A} + \frac{1}{\mu_B} + \frac{\mu_A\mu_B(\sigma_A^2 + \sigma_B^2)}{2(\mu_B - \mu_A)} \leqslant N\left(\frac{2\mu_A - \mu_B}{2\mu_A\mu_B}\right) \\ \implies \mathbb{E}[S^B] &\leqslant \frac{N+1}{2\mu_A} + \frac{1}{\mu_B} + \frac{\mu_A(\sigma_A^2 + \sigma_B^2)}{2(1-\mu_A/\mu_B)} \leqslant \frac{N}{\mu_B} \leqslant \mathbb{E}[S^A].\end{aligned}$$

As it is, the expected system time from choosing route BA is less than that from choosing route AB. Therefore, no customer has incentive to deviate, and we have that it is a Nash equilibrium for every player to choose route BA. □

2.5 The Sequential Open Routing Game

We have thus far assumed that customers had no visibility into the state of the network when making routing decisions, i.e., that they were not aware of their relative priorities or of the decisions made by others. In many applications, however, customers may have some knowledge about their position in the queue and about the routes that others have chosen. For example, in service systems such as amusement parks, at times customers may pre-queue, allowing them at the preset opening time to choose their route conditioned on the actions of those preceding them. To investigate a system where customers have knowledge about their position, we consider the *sequential open routing game*, in which customers make decisions according to their relative position among the N customers and can observe the routes chosen by those who move before them. This setting facilitates a natural representation as an extensive-form game, and we will find its unique subgame perfect equilibrium using backward induction.

We remark that the open routing game of Section 2.4 represents one extreme, in which customers have no information apart from the system parameters (μ_A, μ_B, and N), while the sequential open routing game that we study here corresponds to the other extreme, in which customers have perfect information regarding their position and the state of the network prior to their decision. Interestingly, in this sequential environment we also observe a phenomenon that is similar to the herding behavior observed in the open routing game: in equilibrium all customers but one visit the slower station A first.

Here we again consider a two-station service network with deterministic service times and N customers present at the beginning of service availability. We note, however, that unlike the open routing game of Section 2.4, the sequential open routing game is not symmetric, and accordingly the customer index i will no longer be

arbitrary. We index customers $i = 1, 2, \ldots, N$, by the order in which they make routing decisions, so customer 1 is the first to move, customer 2 is the second, etc. We also assume that a customer's position in the order corresponds to her priority, so customer 1 will be served first at whichever station she chooses, and customer $i \geqslant 2$ will wait behind any customers in the set $\{1, \ldots, i-1\}$ who have chosen the same route as her. Next, for $i = 1, 2, \ldots, N$, we define $y_i^A \in \{0, 1, \ldots, i\}$ to be the number of players that have chosen route AB, up to and including customer i. To derive the subgame perfect equilibrium via backward induction, we need to analyze the total time that player N, the last customer to make her routing decision, spends in the system under different strategy profiles. The state observed by player N is fully described by y_{N-1}^A, the number of the first $N-1$ customers that chose route AB.

System Time for Customer N.

Given a value of y_{N-1}^A, we use S_N^A to denote the system time that customer N, the last customer to move, experiences if she takes route AB, and S_N^B to denote the system time that she experiences if she takes route BA.

Because $\mu_A < \mu_B$ and the service times are deterministic, we see that if customer N first visits the faster station B, followed by the slower station A (i.e., she chooses route BA), then she experiences exactly the same system time as if she merely joined a queue of $N-1$ customers all waiting at station A, as long as at least one customer before her chose route AB (i.e., as long as $y_{N-1}^A \geqslant 1$). If no one else chose route AB, that is, if $y_{N-1}^A = 0$, then the above is still true except that station A sits idle until the first departure from station B, so customer N's system time is increased by $1/\mu_B$. Therefore, we can express S_N^B, the system time for the case in which customer

N chooses route BA, as

$$S_N^B = \begin{cases} \dfrac{N}{\mu_A} + \dfrac{1}{\mu_B} & \text{if } y_{N-1}^A = 0, \\ \dfrac{N}{\mu_A} & \text{if } y_{N-1}^A \geqslant 1. \end{cases} \tag{2.32}$$

Now, suppose that customer N chooses route AB, visiting the slower station A first, followed by the faster station B. If player N arrives at station B and finds it idle, then S_N^A, the system time that she experiences if she visits station A first, is simply equal to $(y_{N-1}^A + 1)/\mu_A$, the time that she spends at station A, plus the time it takes for her to be served at station B, which is equal to $1/\mu_B$. If player N arrives at station B and finds it busy, then by Property 1, station B must never have been idle since it started service. As player N will be the last customer served at station B, we have then that S_N^A is equal to N/μ_B. To summarize, the system time that customer N experiences from visiting station A first, S_N^A, is given by

$$S_N^A = \begin{cases} \dfrac{y_{N-1}^A + 1}{\mu_A} + \dfrac{1}{\mu_B} & \text{if } \dfrac{y_{N-1}^A + 1}{\mu_A} \geqslant \dfrac{N-1}{\mu_B}, \\ \dfrac{N}{\mu_B} & \text{if } \dfrac{y_{N-1}^A + 1}{\mu_A} < \dfrac{N-1}{\mu_B}, \end{cases} \tag{2.33}$$

where $(y_{N-1}^A + 1)/\mu_A \geqslant (N-1)/\mu_B$ is the condition that ensures that station B is idle upon customer N's departure from station A.

Equilibrium Strategy Profile.

With the evaluation of customer N's system times we deduce her strategy in the set of subgame perfect equilibria. Because $\mu_A < \mu_B$, whenever the second condition in equation (2.33) is met we have that $S_N^A < S_N^B$, so customer N's system time from choosing route AB is the shorter.

If instead the first condition in equation (2.33) holds, then we have two cases. First, if the number of preceding customers choosing route AB satisfies the inequality $y^A_{N-1} \leqslant N-2$, then S^A_N is bounded above by $(N-1)/\mu_A + 1/\mu_B$, so we again have that $S^A_N < S^B_N$.

Second, if $y^A_{N-1} = N-1$ and customer N chooses route AB, then she will be served last at station A and finish there after N/μ_A units of time, but she will still require a service time of $1/\mu_B$ at station B. On the other hand, if she takes route BA, then she will immediately be served at station B and leave the system exactly when station A completes its workload, after N/μ_A time units. The extra $1/\mu_B$ from route AB causes S^A_N to be greater than S^B_N for $y^A_{N-1} = N-1$.

Summarizing, we have that $S^B_N < S^A_N$ if and only if $y^A_{N-1} = N-1$, and otherwise $S^A_N < S^B_N$. Therefore, customer N's optimal strategy is to choose route BA if and only if $y^A_{N-1} = N-1$, and route AB otherwise.

With the optimal strategy profile for customer N, we now inductively determine the subgame perfect equilibrium strategy profile for customers $N-n$, for every n in the set $\{1, ..., N-1\}$.

Proposition 7 (Unique Subgame Perfect Equilibrium)**.** *The following strategies form the unique subgame perfect equilibrium of the sequential open routing game:*

1. *Customers* $1, \ldots, N-1$ *visit station A first in every subhistory.*
2. *Customer N visits station B first if and only if she observes* $y^A_{N-1} = N-1$*; otherwise, she visits station A first.*

Moreover, along the equilibrium path the first $N-1$ *customers to move visit station A first, and the final customer visits station B first.*

Proof. We establish the subgame perfect equilibrium using backward induction. First, equations (2.32) and (2.33) tell us that the unique optimal strategy for customer N is to take route BA if $y^A_{N-1} = N-1$ and route AB otherwise. Note that

this strategy ensures that in equilibrium always at least one customer will take route AB. Consequently, being the slower station, in equilibrium station A will never idle until it has processed all N customers. Next, we use induction to identify the optimal strategy for all other players. For the induction hypothesis, assume for some integer $2 \leqslant n \leqslant N-1$ that the strategy of subsequent customers $N-n'$, where $1 \leqslant n' < n$, is to always take route AB, and that the final customer to move, customer N, follows the optimal strategy derived above.

To serve as the base case, we first verify the induction hypothesis for $n = 2$ by deriving the equilibrium strategy for customer $N-1$. Player N, the last to move, is the only customer that follows player $N-1$, and in equilibrium customer N will follow the strategy of visiting station B first if $y^A_{N-1} = N-1$, and station A otherwise. So, if customer $N-1$ visits station B first, then customer N will visit station A first, and customer $N-1$ will be the last person served at station A. She will then experience system time given by

$$S^B_{N-1} = \frac{N}{\mu_A} \tag{2.34}$$

because, as noted, station A never idles given player N's equilibrium strategy. If customer $N-1$ takes route AB, then her system time will depend on how many players before her made the same choice. If a small enough number of them chose route AB that when player $N-1$ departs from station A she will find station B busy, then her system time S^A_{N-1} is given by

$$S^A_{N-1} = \frac{N-1}{\mu_B} < \frac{N-1}{\mu_A} < \frac{N-1}{\mu_A} + \frac{1}{\mu_A} = S^B_{N-1}. \tag{2.35}$$

On the other hand, if enough customers before customer $N-1$ chose route AB that she will find station B idle when she departs from station A, then her system

time S^A_{N-1} is given by

$$S^A_{N-1} = \frac{y^A_{N-2}+1}{\mu_A} + \frac{1}{\mu_B} < \frac{N-1}{\mu_A} + \frac{1}{\mu_A} = S^B_{N-1} \tag{2.36}$$

Thus, the equilibrium strategy for customer $N-1$ is to visit station A first in every subhistory. Combined with the equilibrium strategy for customer N, equations (2.34)-(2.36) verify the induction hypothesis for $n = 2$.

Now, assume that the induction hypothesis holds for some integer $2 \leqslant n \leqslant N-1$. Let S^B_{N-n} be the system time that customer $N-n$ experiences if she chooses to join station B first. Similarly, let S^A_{N-n} be the system time that customer $N-n$ experiences if she chooses to join station A first. If customer $N-n$ chooses route AB, then everyone after her will join station A first, and customer $N-n$ will be the last to be served at station A. Because station A never idles, we then must have $S^B_{N-n} = N/\mu_A$. To study S^A_{N-n}, we need to consider the following two cases.

Case 1: The number of AB customers before customer $N-n$, denoted by y^A_{N-n-1}, is small enough that customer $N-n$ finds station B busy when she finishes at station A. Then by Property 2, station B has never idled since starting service and we must have $S^A_{N-n} = (N-n)/\mu_B$. Moreover,

$$S^A_{N-n} = \frac{N-n}{\mu_B} < \frac{N-n}{\mu_A} < \frac{N-1}{\mu_A} + \frac{1}{\mu_A} = S^B_{N-n}. \tag{2.37}$$

Case 2: The number of AB customers before customer $N-n$, denoted by y^A_{N-n-1}, is big enough that customer $N-n$ finds station B idle when she finishes at station A. In this case, we have $S^A_{N-n} = (y^A_{N-n-1}+1)/\mu_A + (1/\mu_B)$. Therefore, we have that

$$S^A_{N-n} = \frac{y^A_{N-n-1}+1}{\mu_A} + \frac{1}{\mu_B} < \frac{N-1}{\mu_A} + \frac{1}{\mu_A} = S^B_{N-n}.$$

In either case, choosing route AB results in a strictly shorter system time for customer $N-n$. Thus, if the induction hypothesis holds for some $n \geqslant 2$, we now

have that for any $n \leqslant \widetilde{n} \leqslant N-1$, the unique optimal strategy for customer $N-\widetilde{n}$ is to take route AB regardless of the subhistory, given that all subsequent customers act optimally. Having verified the induction hypothesis for $n = 2$, we obtain the unique subgame perfect equilibrium of our game, comprised of the strategies stated in the proposition. Furthermore, inspection of the strategies reveals that the resulting equilibrium path entails the first $N-1$ players taking route AB, and the final player N taking route BA. □

Intuitively, one might expect that customers should join the queue with the shortest wait time because that minimizes their waiting time before getting served at one of the two stations. However, Proposition 7 states that if customers are rational, then all except the last will first visit station A, which is slower and which each customer except the first one observes to have more customers in its queue than station B. This behavior is best explained by what we call delayed overtaking. If customers later in the order will visit the slower station A first, then it is in the best interest of a customer who moves earlier to immediately join the queue at station A because otherwise, during her time at the faster station B, others will overtake her at station A.

Interestingly, we observe that the optimal actions of the first $N-1$ customers are completely independent of the state that they observe and are essentially driven by the strategies of those later in the order. Note also that the equilibrium actions—i.e., the herding profiles—from the open routing game of Section 2.4 cannot be supported as equilibria (Nash or subgame perfect) for the sequential game. Namely, if the last customer to move is aware of her position and the decisions of the other customers, then her best response to all of the earlier customers visiting the same station (either A or B) is always to visit the other station, breaking both of the equilibria of Proposition 1. However, in equilibrium only one customer visits station B first here,

so the subgame perfect equilibrium of this section is quite similar to the herding equilibrium at station A discussed in Section 2.4. Therefore, in an environment in which decisions are made sequentially, the pull to the slower station has the strongest impact on behavior, and the delayed overtaking effect dominates.

On the whole, we find that behavior in the sequential open routing game parallels that of the one-shot version. The most important departures are (i) that the delayed overtaking effect always dominates in the sequential environment, supporting a form of herding only at the slower station and (ii) that in equilibrium the last customer, recognizing that she is the last to move, takes the opposite route of all of the other players.

2.6 Herding vs. Central Optimum

We consider now a central planner who wishes to optimize social welfare in the open-routing service network of Sections 2.4 and 2.5. We take as our measure of social welfare the sum of the system times of all customers. The central planner's problem is to minimize this quantity, which we call *cumulative system time*. We compare the cumulative system time under herding with the optimal cumulative system time, first by deriving bounds on the sub-optimality of herding and then by numerically solving for the optimal routing. We end by discussing the implications of these results on the price of anarchy.

Cumulative System Time Under Herding.

If $0 \leqslant x \leqslant N$ customers are assigned to route AB, then denote by $D(x)$ the resulting cumulative system time. Let Q be the number of AB customers who, after their service at station A, find station B busy. We use $D_A^i(x), i \in \{1, \ldots, x\}$, to denote the total time spent in the system by the i-th AB customer; similarly, we use $D_B^j(x), j \in \{1, \ldots, N-x\}$, to denote the total time spent in the system by the j-th BA customer.

Finally, define $D(x)$ to be the cumulative system time, i.e.,

$$D(x) := \sum_{i=1}^{x} D_A^i(x) + \sum_{j=1}^{N-x} D_B^j(x). \tag{2.38}$$

If the i-th AB customer finds station B busy upon completing her service at station A, then her total system time is given by

$$D_A^i(x) = \frac{N-x}{\mu_B} + \frac{i}{\mu_B}, \quad i = 1, \ldots, Q. \tag{2.39}$$

We may understand this by recalling from Property 1 that if the i-th customer finds station B busy, then station B must necessarily have never been idle since the start of service availability. Thus, the i-th AB customer will enter service at station B exactly when station B has finished processing all $N-x$ of the BA customers plus the $i-1$ customers who visited station A first and are in front of the i-th AB customer. Adding her own service time, she will experience the total system time related in equation (2.39).

On the other hand, if the i-th AB customer finishes service at station A and finds station B idle, then her total system time is given by

$$D_A^i(x) = \frac{i}{\mu_A} + \frac{1}{\mu_B}, \qquad \text{for } i = Q+1, \ldots, x. \tag{2.40}$$

Finding station B idle is equivalent to not facing a wait at station B, so a customer's total system time is the sum of the time that she spends at station A and her service time at station B. She must wait at station A for the service of the $i-1$ customers in front of her there, and then she must be served. She then moves to station B, where she immediately enters service, resulting in the system time given by equation (2.40).

Finally, we have that the j-th BA customer has total system time given by

$$D_B^j(x) = \frac{x+j}{\mu_A}, \qquad \text{for } j = 1, \ldots, N-x, \tag{2.41}$$

if $x > 0$. To build some intuition around this last quantity, we note that the j-th BA customer will enter service at station A exactly when station A finishes processing all x of the AB customers plus all $j-1$ of the BA customers in front of her at station B. We recall from Property 2 that station A never idles as long as $x > 0$, so the time until she enters service at station A is given by the sum of the service times of $x+j-1$ customers at station A. Adding her own service time at station A, we see that her system time is as expressed in equation (2.41). Note that if $x = 0$, then we have $D_B^j(x) = j/\mu_A + 1/\mu_B$, i.e., the quantity in equation (2.41) plus $1/\mu_B$, where $1/\mu_B$ is the time that station A is idle before its first customer departs from station B.

Therefore, the cumulative system time, $D(x)$, is simply the sum of all of the terms in equations (2.39)-(2.41), and we present this quantity in the following lemma as a function of the number x of customers that visit station A first.

Lemma 1 (Cumulative System Time)**.** *Take $x \in \{0, \ldots, N\}$, and define $\tilde{Q}$ as*

$$\tilde{Q} := \max\left\{k \in \mathbb{N} : \frac{k}{\mu_A} \leqslant \frac{N-x+k-1}{\mu_B}\right\} = \left\lfloor \frac{\mu_A(N-x-1)}{\mu_B - \mu_A} \right\rfloor.$$

The number Q of AB customers who find station B busy upon service completion at station A is given by

$$Q = \min\{x, \tilde{Q}\},$$

and the cumulative system time $D(x)$ is given by

$$D(x) = \begin{cases} \dfrac{N}{\mu_B} + \dfrac{N^2+N}{2\mu_A} & x = 0, \\ \dfrac{x(2N-x+1)}{2\mu_B} + \dfrac{(N-x)(N+x+1)}{2\mu_A} & 1 \leqslant x \leqslant \dfrac{\mu_A(N-1)}{\mu_B}, \\ \dfrac{\tilde{Q}(2N-2x+\tilde{Q}-1)+2x}{2\mu_B} + & \\ \qquad \dfrac{N+N^2-\tilde{Q}-\tilde{Q}^2}{2\mu_A} & \dfrac{\mu_A(N-1)}{\mu_B} < x \leqslant N - \dfrac{\mu_B}{\mu_A}, \\ \dfrac{x}{\mu_B} + \dfrac{N^2+N}{2\mu_A} & N - \dfrac{\mu_B}{\mu_A} < x \leqslant N. \end{cases} \tag{2.42}$$

Proof. First, observe that

$$\begin{aligned} x &> \frac{\mu_A(N-1)}{\mu_B} \\ \mu_B x - \mu_A x &> \mu_A(N-1) - \mu_A x \\ x &> \frac{\mu_A(N-1-x)}{\mu_B - \mu_A} \\ x &> \left\lfloor \frac{\mu_A(N-1-x)}{\mu_B - \mu_A} \right\rfloor = \tilde{Q}, \end{aligned}$$

where the last equivalence condition holds because x is an integer. Therefore we have

$$Q = \tilde{Q} < x \text{ if } x > \frac{\mu_A(N-1)}{\mu_B} \quad \text{and} \quad Q = x \text{ if } x \leqslant \frac{\mu_A(N-1)}{\mu_B}. \tag{2.43}$$

Next, we will study $D(x)$ under the four cases listed in equation (2.42).

Case 1: $x = 0$. In this case, station A will be idle for the first $1/\mu_B$ units of time. Since all customers visit station B first, each customer must wait for the customers in front of her to finish at station A before she is served there. Hence the customer with priority j at station B will finish at station A exactly when station A completes

its j-th service, which because of the idling will occur after $(1/\mu_B) + (j/\mu_A)$ units of time. Thus, the cumulative system time is equal to

$$D(x) = \sum_{j=1}^{N} D_B^j(0) = \sum_{j=1}^{N} \left(\frac{1}{\mu_B} + \frac{j}{\mu_A} \right) = \frac{N}{\mu_B} + \frac{N^2 + N}{2\mu_A}.$$

Case 2: $1 \leqslant x \leqslant \mu_A(N-1)/\mu_B$. By equation (2.43), here we have $Q = x$. Equations (2.38) and (2.39) give us that

$$D(x) = \sum_{i=1}^{x} D_A^i(x) + \sum_{j=1}^{N-x} D_B^j(x) = \frac{x(2N - x + 1)}{2\mu_B} + \frac{(N-x)(N+x+1)}{2\mu_A}.$$

Case 3: $\mu_A(N-1)/\mu_B < x \leqslant N - \mu_B/\mu_A$. By equation (2.43), in this case $Q = \tilde{Q}$. By equations (2.38), (2.39), and (2.40), we get

$$\begin{aligned} D(x) &= \sum_{i=1}^{\tilde{Q}} D_A^i(x) + \sum_{i=\tilde{Q}+1}^{x} D_A^i(x) + \sum_{j=1}^{N-x} D_B^j(x) \\ &= \frac{\tilde{Q}(2N - 2x + \tilde{Q} - 1) + 2x}{2\mu_B} + \frac{N + N^2 - \tilde{Q} - \tilde{Q}^2}{2\mu_A}. \end{aligned}$$

Case 4: $N - \mu_B/\mu_A < x \leqslant N$. Recall that if station B idles then it will never build up a queue again. In this case, even the first AB customer does not face a queue at station B, and therefore the cumulative system time is equal to

$$D(x) = \sum_{i=1}^{x} D_A^i(x) + \sum_{j=1}^{N-x} D_B^j(x) = \frac{x}{\mu_B} + \frac{N^2 + N}{2\mu_A}.$$

□

The four ranges for x that appear in equation (2.42) correspond to four regimes for the number of AB customers: if $x = 0$, then there are no AB customers and station A will be idle during the first $1/\mu_B$ time periods; if $1 \leqslant x \leqslant \mu_A(N-1)/\mu_B$, then all of the AB customers will face a queue at station B; if $\mu_A(N-1)/\mu_B < x \leqslant N - \mu_B/\mu_A$, then

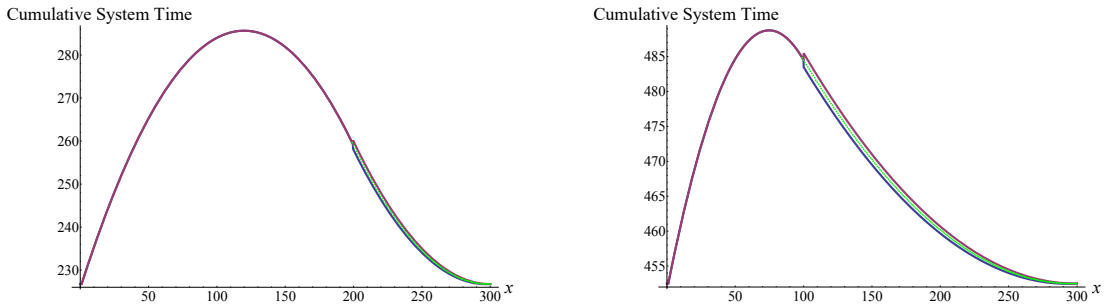

(a) $N = 300$, $\mu_A = 2/3$, $\mu_B = 1$. (b) $N = 300$, $\mu_A = 1/3$, $\mu_B = 1$.

FIGURE 2.2: Cumulative system time as a function of the number of AB customers x.

some of the earlier AB customers will face a queue at station B, while the later AB customers will be served at station B immediately; and finally, if $N - \mu_B/\mu_A < x \leqslant N$, then no AB customers will face a queue at station B.

Lemma 1 provides an explicit formula for the cumulative system time $D(x)$, when x customers visit station A first. Because $\tilde{Q}$ is a discrete function of x, we now develop tight upper and lower bounds for $D(x)$ when $x \in \{\lceil \mu_A(N-1)/\mu_B \rceil, \ldots, \lfloor N - (\mu_B/\mu_A) \rfloor\}$. Both the upper and lower bounds are outputs of continuous functions, which are much more amenable to analysis, and the lower bound will be used in the proof of Proposition 8.

Lemma 2 (Bounds on Cumulative System Time)**.** *Defining the functions* $S_1(\cdot)$ *and* $S_2(\cdot)$ *and the quantity* S_3 *by*

$$S_1(t) = \frac{t}{2}\left(\frac{N-x+1}{\mu_B} + \frac{N-x+t}{\mu_B}\right),$$

$$S_2(t) = \frac{x-t}{2}\left(\frac{t+1}{\mu_A} + \frac{1}{\mu_B} + \frac{x}{\mu_A} + \frac{1}{\mu_B}\right),$$

$$\text{and } S_3 = \frac{N-x}{2}\left(\frac{x+1}{\mu_A} + \frac{N}{\mu_A}\right),$$

we have that if $\mu_A(N-1)/\mu_B < x \leqslant N - (\mu_B/\mu_A)$*, then* $D(x)$ *is bounded by*

$$D(x) \geqslant S_1\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}-1\right) + S_2\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}\right) + S_3, \tag{2.44}$$

$$D(x) \leqslant S_1\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}\right) + S_2\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}-1\right) + S_3.$$

Proof. Recall from the proof of Lemma 1 that if $(\mu_A(N-1)/\mu_B) < x \leqslant N-(\mu_B/\mu_A)$, then

$$D(x) = \sum_{i=1}^{\tilde{Q}} \frac{N-x+i}{\mu_B} + \sum_{i=\tilde{Q}+1}^{x} \left(\frac{i}{\mu_A}+\frac{1}{\mu_B}\right) + \sum_{j=1}^{N-x} \frac{x+j}{\mu_A} = S_1(\tilde{Q}) + S_2(\tilde{Q}) + S_3.$$

The function $S_1(t)$ is non-decreasing on $\mathbb{R}_+$, and the function $S_2(t)$ is non-increasing on $\mathbb{R}_+$. Here we have that $x \leqslant N - \mu_B/\mu_A$, so

$$1 \leqslant \tilde{Q} \leqslant \frac{\mu_A(N-x-1)}{\mu_B-\mu_A},$$

and the monotonicity properties of $S_1(\cdot)$ and $S_2(\cdot)$ give us that

$$D(x) \geqslant S_1\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}-1\right) + S_2\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}\right) + S_3$$

$$\text{and } D(x) \leqslant S_1\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}\right) + S_2\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}-1\right) + S_3.$$

□

The lemma follows from the upper and lower bounds on $\tilde{Q}$ given by

$$\frac{\mu_A(N-x-1)}{\mu_B-\mu_A} - 1 < \tilde{Q} = \left\lfloor \frac{\mu_A(N-x-1)}{\mu_B-\mu_A} \right\rfloor \leqslant \frac{\mu_A(N-x-1)}{\mu_B-\mu_A}.$$

In turn, the bounds that we obtain are extremely tight when N is large, as they are essentially the bounds for the discretization error. Figure 2.2 plots the function

$D(x)$, along with these upper and lower bounds, for two problem instances. We note that

$$D(0) = D(N) = \frac{N}{\mu_B} + \frac{N^2 + N}{2\mu_A}, \tag{2.45}$$

i.e., the cumulative system time is the same under either herding profile. We let D_N^H denote the quantity in equation (2.45). Let D_N^* denote the minimal cumulative system time, that is,

$$D_N^* = \min_{x \in \{0,\ldots,N\}} D(x).$$

Next, we show that the gap, $D_N^* - D_N^H$, is uniformly bounded by a constant.

Proposition 8 (Optimality Gap of Herding)**.** *The optimality gap $D_N^H - D_N^*$ satisfies the bounds*

$$\frac{1}{\mu_A} \leqslant D_N^H - D_N^* \leqslant \frac{1}{\mu_A} + \frac{17\mu_B^2 - 6\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)} \quad \textit{for all } N \geqslant 1 + \frac{2\mu_B}{\mu_A}. \tag{2.46}$$

Proof. We will establish that $D(1)$ is a lower bound on the value of $D(x)$ everywhere except the region where $\tilde{Q}$ appears in the expression for $D(x)$, and we will bound the difference between $D(1)$ and the minimum possible value of $D(x)$ in that region.

Case 1: $x = 0$. Substituting 1 for x in equation (2.42), we get

$$\begin{aligned} D(1) &= \frac{N}{\mu_B} + \frac{(N-1)(N+2)}{2\mu_A} \\ &= \frac{N}{\mu_B} + \frac{N^2 + N}{2\mu_A} - \frac{1}{\mu_A} \\ &= D(0) - \frac{1}{\mu_A}, \end{aligned}$$

and therefore $1/\mu_A \leqslant D_N^H - D_N^*$.

Case 2: $1 \leqslant x \leqslant \mu_A(N-1)/\mu_B$. In this interval, the continuous extension of $D(x)$ is a concave quadratic function. Thus, the minimum of $D(x)$ in this region

must be at one of the endpoints. We evaluate $D(\mu_A(N-1)/\mu_B)$ and get

$$D\left(\frac{\mu_A(N-1)}{\mu_B}\right) = \frac{N^2+N}{2\mu_A} + \frac{N^2(\mu_A^2\mu_B - \mu_A^3)}{2\mu_A\mu_B^3} + \frac{N(2\mu_A^3 + \mu_A^2\mu_B - \mu_A\mu_B^2)}{2\mu_A\mu_B^3} + \frac{\mu_B^2 - 2\mu_A\mu_B - \mu_A^2}{2\mu_B^3}.$$

The difference $D(\mu_A(N-1)/\mu_B) - D(1)$ is then given by

$$D\left(\frac{\mu_A(N-1)}{\mu_B}\right) - D(1) = \frac{N^2\mu_A(\mu_B-\mu_A)}{2\mu_B^3} - \frac{N(\mu_B-\mu_A)(2\mu_A+3\mu_B)}{2\mu_B^3} + \frac{(\mu_B-\mu_A)(\mu_A^2+3\mu_A\mu_B+2\mu_B^2)}{2\mu_A\mu_B^3}.$$

The larger root of this convex quadratic function of N occurs at $N = 1 + 2\mu_B/\mu_A$, and thus

$$N \geqslant 1 + \frac{2\mu_B}{\mu_A} \implies D\left(\frac{\mu_A(N-1)}{\mu_B}\right) - D(1) \geqslant 0.$$

We deduce therefore that $D(1)$ is the minimum value of $D(x)$ for $1 \leqslant x \leqslant \mu_A(N-1)/\mu_B$.

Case 3: $\mu_A(N-1)/\mu_B < x \leqslant N - \mu_B/\mu_A$. In this region we will work relative to the lower bound in equation (2.44). Evaluating this bound, we have

$$\begin{aligned} D(x) &\geqslant S_1\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A} - 1\right) + S_2\left(\frac{\mu_A(N-x-1)}{\mu_B-\mu_A}\right) + S_3 \\ &= x^2\left(\frac{\mu_A}{2\mu_B(\mu_B-\mu_A)}\right) + x\left(\frac{5}{2(\mu_B-\mu_A)} - \frac{\mu_A(2N+1)}{2\mu_B(\mu_B-\mu_A)}\right) \\ &\quad + N^2\left(\frac{\mu_B-\mu_A}{2\mu_A\mu_B} + \frac{1}{2(\mu_B-\mu_A)}\right) + N\left(\frac{\mu_B+\mu_A}{2\mu_A\mu_B} - \frac{2}{\mu_B-\mu_A}\right) \\ &\quad + \frac{2\mu_A+\mu_B}{2\mu_B(\mu_B-\mu_A)} \\ &=: \underline{D}(x). \end{aligned}$$

We note that $\underline{D}(x)$ is a convex quadratic function in x. Differentiating with respect to x, we get

$$\frac{\partial \underline{D}}{\partial x} = \frac{5\mu_B - \mu_A(1 + 2N - 2x)}{2\mu_B(\mu_B - \mu_A)}.$$

Because the first-order condition is sufficient for a minimum, we set the derivative $\partial \underline{D}/\partial x$ equal to zero and solve for the root x^*, obtaining

$$\begin{aligned} 5\mu_B - \mu_A(1 + 2N - 2x^*) &= 0 \\ 2\mu_A x^* &= \mu_A(2N + 1) - 5\mu_B \\ x^* &= N + \frac{1}{2} - \frac{5\mu_B}{2\mu_A}. \end{aligned}$$

We then have that $\underline{D}(x^*)$ is a lower bound on the value of $D(x)$ for any $\mu_A(N - 1)/\mu_B < x \leqslant N - \mu_B/\mu_A$. The bound holds whether or not x^* falls within the relevant interval, as $\underline{D}(x^*)$ is the global minimum of the quadratic function. Evaluating $\underline{D}(x^*)$, we get

$$\begin{aligned} \underline{D}(x^*) &= \frac{N^2 + N}{2\mu_A} + \frac{N}{\mu_B} - \frac{25\mu_B^2 - 14\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)} \\ &= D(0) - \frac{25\mu_B^2 - 14\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)} \\ &= D(1) + \frac{1}{\mu_A} - \frac{25\mu_B^2 - 14\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)}. \end{aligned}$$

Because $\mu_A < \mu_B$, we have

$$\begin{aligned} 0 &< 17\mu_B^2 - 6\mu_A\mu_B - 7\mu_A^2 \\ 8\mu_B^2 - 8\mu_A\mu_B &< 25\mu_B^2 - 14\mu_A\mu_B - 7\mu_A^2 \\ \frac{1}{\mu_A} &< \frac{25\mu_B^2 - 14\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)}, \end{aligned}$$

which implies that $\underline{D}(x^*) < D(1)$.

Case 4: $N - \mu_B/\mu_A < x \leqslant N$. The cumulative system time $D(x)$ is an increasing linear function on the half-open interval $(N - \mu_B/\mu_A, N]$. Therefore, the infimum of $D(x)$ on this interval is given by the output of the linear function evaluated at $N - \mu_B/\mu_A$, i.e.,

$$\begin{aligned}\inf_{x \in (N-\mu_B/\mu_A, N]} D(x) &= \frac{N - \mu_B/\mu_A}{\mu_B} + \frac{N^2 + N}{2\mu_A} \\ &= \frac{N}{\mu_B} + \frac{N^2 + N}{2\mu_A} - \frac{1}{\mu_A} \\ &= D(1).\end{aligned}$$

Collecting the results from the four cases, we conclude that

$$\frac{1}{\mu_A} \leqslant D_N^H - D_N^* \leqslant \frac{1}{\mu_A} + \frac{17\mu_B^2 - 6\mu_A\mu_B - 7\mu_A^2}{8\mu_A\mu_B(\mu_B - \mu_A)}.$$

□

The bounds (2.46) are independent of the population size N, and the ratio of this gap to the optimal cumulative system time rapidly approaches 0 as N grows. As we discuss below, numerical evidence suggests that in most cases assigning exactly one customer to route AB minimizes the cumulative system time. In those cases the optimality gap is usually equal to one service time at station A.

Exact Central Optimal Solution.

With Proposition 8 bounding the optimality gap of herding, we now numerically solve the central planner's problem for a variety of parameter values. The main observation worth noting is that, for all of the parameter combinations that we have studied, the minimal cumulative system time D_N^* is attained by setting the number of AB customers to one. The reason for doing so, instead of assigning all of the customers to herd on route BA (i.e., setting $x = 0$), is to avoid inducing idle time

Table 2.1: Maximum Percent Above Optimal Cumulative System Time Under Herding (i.e., $x = 0$ or $x = N$; $\mu_B = 1$)

μ_A	10^{-5}	$(10^{-5}, .2]$	(.2,.4]	(.4,.6]	(.6,.8]	(.8,1]
Max % Gap	.00004	.078	.078	.077	.077	.078

at station A during the first service time at station B. The opposite extreme of assigning all customers to herd on route AB induces a similar idle time problem at station B, resulting in equivalent cumulative system time to the case with $x = 0$. Still, the difference in cumulative system time when setting $x = 0$ instead of $x = 1$ is usually negligible, as we will observe. We compute and study the cumulative system time for a range of parameter values. Fixing $\mu_B = 1$ and taking $N = 50$, we compute the cumulative system time vector (letting $x = 0, 1, \ldots, N$) for all values of $\mu_A \in \{1 \times 10^{-5}, 2 \times 10^{-5}, \ldots, 99,999 \times 10^{-5}\}$. Table 2.1 shows the maximum percentage difference between the cumulative system time under herding and the minimal cumulative system time, over ranges of μ_A. In all cases the cumulative system time under herding is less than a tenth of a percent above optimal.

Price of Anarchy.

In Sections 2.4 and 2.5, we found that self-interested customers followed the crowd to the same route. The optimal outcome under a social planner is quite similar. Rather than balancing the workload across the two queues, a social planner should send almost all of the customers to the same station first. The intuition has to do with the fact that if the queues are of similar length at both stations, then the first customer to exit the system will not do so until she has waited a substantial amount of time for her second service. The performance of such a load-balancing routing assignment is close to that of the worst possible routing scheme (see also Figure 2.2). Conversely, if all customers are assigned to the same route, then the early customers have a short (if any) wait at both stations. Whoever the final customers are to leave

the system, they cannot possibly leave until station A processes everyone before them, and that time serves as a lower bound on their system times. But for the customers early in the order, their system times are significantly shorter if customers are concentrated on one route.

These results lead to an interesting observation regarding the price of anarchy. The price of anarchy is defined as the ratio between the social welfare in the worst equilibrium outcome and the best outcome achievable by a social planner (Koutsoupias and Papadimitriou, 2009). In the open routing game of Section 2.4, the price of anarchy for pure strategies is close to 1 because the equilibrium outcome sends everyone to the same station, and the difference in social welfare between this and the optimum is usually only a single service time at station A. Moreover, as N grows the price of anarchy converges to 1 by Proposition 8. As rational customers will herd in equilibrium, the good overall performance of herding with respect to the cumulative system time is encouraging for managers. In the event that customers are not fully rational, there is a strong argument for managers to incentivize herding among their customers, as it is a simple strategy with near-optimal performance.

2.7 The Prevalence of Herding

We have seen that the herding behavior predominates in the open routing game. With enough players, the herding equilibria are always the only pure-strategy Nash equilibria of the game. Furthermore, when the service rates are far apart, customers have a strictly dominant strategy which implements the herding equilibrium at station A. The herding equilibria continue to hold even when some customers are not strategic and have their routes fixed in advance, or when we allow for stochastic service times. In the sequential variant of the open routing game, we similarly observe a form of herding, albeit only at the slower station A. In this section, we discuss additional settings in which the herding behavior continues to prevail.

Customers Who Visit Only One Station

First, we discuss the impact of "dedicated" customers, who require service only at one of the two stations. In the context of the open routing game of Section 2.4, we now relax the assumption that all customers must visit both stations. Specifically, we decompose N as $N = N_A + N_B + N_S$, where N_A customers visit only station A, N_B customers visit only station B, and the remaining N_S customers are strategic, i.e., they visit both stations but may choose the order. When there are enough strategic players, both the herding profile at station A and that at station B are Nash equilibria of this game. We first show the result for the equilibrium on route AB.

Proposition 9 (Nash Equilibrium with Route AB—Dedicated Customers)**.** *If*

$$N_S \geqslant (N_B + 1)\Big(\frac{2\mu_A}{\mu_B}\Big) + 1 - N_A, \tag{2.47}$$

then it is a Nash equilibrium for all N_S strategic players to visit station A first.

Proof. Suppose that all of the strategic players choose route AB. If so, then a player who deviates to take route BA will be served at station A after all N_A of the "A-only" players and after the other $N_S - 1$ strategic players. We therefore have the following lower bound on the expected system time $\mathbb{E}[S^B]$ from deviation given by

$$\mathbb{E}[S^B] \geqslant \frac{N_A + N_S}{\mu_A}. \tag{2.48}$$

Next we calculate an upper bound on the expected system time from following the profile and taking route AB. Here we note that Property 1 still applies to this system. Namely, if station B ever idles, then it will never build up a queue again. First, consider the case with no B-only players (that is, $N_B = 0$). In this case, there will never be a queue at station B, and strategic customers departing from station

A will immediately enter service at station B. Therefore, with no B-only customers we have an exact expression for $\mathbb{E}[S^A]$ given by

$$\mathbb{E}[S^A] = \sum_{j=1}^{N_A+N_S} \frac{1}{N_A+N_S}\left(\frac{j}{\mu_A}+\frac{1}{\mu_B}\right).$$

Now, for a given priority j, if N_B is greater than 0, then the maximum increase in the system time S^A from the case with no B-only players is equal to N_B/μ_B, the increased initial workload at station B. Thus, for any value of N_B we have the bound

$$\mathbb{E}[S^A] \leqslant \sum_{j=1}^{N_A+N_S} \frac{1}{N_A+N_S}\left(\frac{j}{\mu_A}+\frac{1}{\mu_B}+\frac{N_B}{\mu_B}\right) = \frac{N_A+N_S+1}{2\mu_A}+\frac{N_B+1}{\mu_B}. \quad (2.49)$$

Equations (2.47), (2.48), and (2.49) give

$$(N_B+1)(\frac{2\mu_A}{\mu_B})+1 \leqslant N_A+N_S$$

$$\implies \mathbb{E}[S^A] \leqslant \frac{N_B+1}{\mu_B}+\frac{N_A+N_S+1}{2\mu_A} \leqslant \frac{N_A+N_S}{\mu_A} \leqslant \mathbb{E}[S^B].$$

We conclude that strategic players have no incentive to deviate, and therefore it is a symmetric Nash equilibrium for all strategic players to visit station A first. □

The intuition behind the proof is similar to that of the open routing game. If all of the strategic customers herd at station A, then a customer who deviates will again be the last customer to be served at station A, and therefore she prefers to follow the crowd. We next prove a similar result for route BA, with similar intuition.

Proposition 10 (Nash Equilibrium with Route BA—Dedicated Customers)**.** *If* $\mu_B < 2\mu_A$ *and*

$$N_S \geqslant \frac{\mu_B(2N_A+1)-\mu_A(N_B-2)}{2\mu_A-\mu_B}, \quad (2.50)$$

then it is a Nash equilibrium for all N_S *strategic players to visit station* B *first.*

Proof. Suppose that all of the strategic customers are following the profile of visiting station B first, and consider a player who contemplates deviating and visiting station A first. If she deviates, then she will certainly not enter service at station B until after station B processes all N_B B-only players as well as the other $N_S - 1$ strategic customers. Therefore, we can bound her expected system time $\mathbb{E}[S^A]$ from deviating by

$$\mathbb{E}[S^A] \geqslant \frac{N_B + N_S}{\mu_B}. \tag{2.51}$$

Next, suppose that she follows the profile and visits station B first, and further suppose that she receives priority k at station B. Let Z_k be the random variable representing the number of strategic customers among the first $k-1$ players at station B. Given that there are $N_B + N_S - 1$ other players at station B, $N_S - 1$ of which are strategic, the random variable Z_k has a hypergeometric distribution with $k-1$ trials from a population of size $N_B + N_S - 1$ containing $N_S - 1$ successes. Therefore, its mean is

$$\mathbb{E}[Z_k] = (k-1)\frac{N_S - 1}{N_B + N_S - 1}. \tag{2.52}$$

With a large number of B-only players it is possible that station A will become idle while some strategic customers have not yet been served at station B. Consider a strategic customer who is assigned priority k at station B. The greatest amount of idle time that could possibly be introduced at station A is the sum of the service times at station B of the B-only customers with a higher priority. There are exactly $k - 1 - Z_k$ such players. Define I_k as the amount of time that station A spends idle before the customer with priority k at station B finishes her service at station B. If $N_A = 0$, then station A would also idle for $1/\mu_B$ before it receives its first customer. An upper bound on the mean of $\mathbb{E}[I_k]$ is then

$$\mathbb{E}[I_k] \leqslant \frac{1}{\mu_B} + \mathbb{E}\left[\frac{k - 1 - Z_k}{\mu_B}\right] = \frac{1}{\mu_B} + \frac{k - 1 - \mathbb{E}[Z_k]}{\mu_B}. \tag{2.53}$$

Table 2.2: Percent Deviation from Submodularity with Dedicated Customers (for μ_B=1, N_S=50).

	μ_A				
N_A, N_B	1.1	1.3	1.5	1.7	1.9
0,5	2	0	0	0	0
0,10	0	2	0	0	0
5,0	0	2	2	4	4
5,10	0	0	0	0	0
10,0	2	0	4	4	6
10,5	0	0	0	2	4
10,10	0	0	0	0	0

We can now express the expected system time from choosing route BA by

$$\begin{aligned}
\mathbb{E}[S^B] &= \sum_{k=1}^{N_B+N_S} \frac{1}{N_B+N_S}\left(\frac{N_A+1+\mathbb{E}[Z_k]}{\mu_A} + \mathbb{E}[I_k]\right) \\
&\leqslant \sum_{k=1}^{N_B+N_S} \frac{1}{N_B+N_S}\left(\frac{N_A+1+\mathbb{E}[Z_k]}{\mu_A} + \frac{k-1-\mathbb{E}[Z_k]}{\mu_B} + \frac{1}{\mu_B}\right) \\
&= \frac{1}{\mu_B} + \frac{N_A+1}{\mu_A} + \frac{N_S-1}{2\mu_A} + \frac{N_B}{2\mu_B}, \qquad (2.54)
\end{aligned}$$

where the inequality follows from equation (2.53) and the last equality follows from substituting the expression in equation (2.52) for $\mathbb{E}[Z_k]$ and then evaluating the summation. Because $\mu_B < 2\mu_A$, equations (2.50), (2.51), and (2.54) then give

$$\begin{aligned}
&\frac{\mu_B(2N_A+1)-\mu_A(N_B-2)}{2\mu_A-\mu_B} \leqslant N_S \\
&\implies \frac{2N_A+1}{2\mu_A} - \frac{N_B-2}{2\mu_B} \leqslant N_S\left(\frac{1}{\mu_B}-\frac{1}{2\mu_A}\right) \\
\implies \mathbb{E}[S^B] \leqslant \frac{1}{\mu_B} + \frac{N_A+1}{\mu_A} + \frac{N_S-1}{2\mu_A} + \frac{N_B}{2\mu_B} &\leqslant \frac{N_B+N_S}{\mu_B} \leqslant \mathbb{E}[S^A].
\end{aligned}$$

Because $\mathbb{E}[S^B] \leqslant \mathbb{E}[S^A]$, we conclude that no player has incentive to deviate, and thus it is a Nash equilibrium for all N_S strategic players to visit station B first. □

We also observe in a series of numerical experiments that, if there is not too much imbalance in the number of dedicated customers at each station (i.e., only

station A vs. only station B), then the expected system time for the strategic customers is often submodular (or close to submodular) even in the presence of dedicated customers. The results of the numerical study are summarized in Table 2.2. We fix $N_S = 50$ and $\mu_B = 1$. The number of customers who are dedicated to each station varies; there are always 50 strategic customers, and any non-strategic customers increase the total population size. For each routing profile of the strategic customers (i.e., for $x = 0, 1, \ldots, 50$ customers on route AB, with the remainder on route BA), we simulate the system 1,000 times. Service times are deterministic, so the only randomness in the system is the arrangement of customers of different types in the queues. Using the sample average system time for AB and BA customers for each routing profile, we compute the sample differences vector. Table 2.2 depicts the percentage of difference comparisons (i.e., d_0 vs. d_1, d_1 vs. d_2, etc.) which do *not* satisfy the decreasing differences condition. We note that all instances satisfy the condition within a .5% tolerance. These numerical experiments provide evidence that the expected system time continues to be submodular or almost submodular in the presence of dedicated customers, if most of the customers are strategic and there is not too much imbalance in the number of customers dedicated to each station. We conclude that the herding behavior is robust to the presence of "dedicated" customers who require service only at one of the two stations.

Overloaded System with Stochastic Arrivals

Next, we consider a setting in which exogenous arrivals occur according to a Poisson process with rate λ, while service times are again deterministic with rates $0 < \mu_A < \mu_B$. We also assume that the system is overloaded, that is, $\lambda \geqslant \mu_A$. As usual, we assume that each customer attempts to minimize her expected total system time.

On arrival, we assume that a customer has complete visibility into queue lengths; that is, she can observe the number of customers queued at each station as well as

the amount of time that the current customers have so far been in service, which is relevant because service times are deterministic. We then let n_A and n_B denote the queue lengths at stations A and B, respectively. For instance, if a customer arrives and sees 6 customers in line at station A, and a customer in service with 45% of her service time remaining, then the queue length at station A is given by $n_A = 6.45$.

We now show that this game has a Nash equilibrium where all players herd on route AB.

Proposition 11 (Nash Equilibrium with Route AB—Stochastic Arrivals)**.** *If $\lambda \geqslant \mu_A$, then the strategy "visit station A first for any n_A and n_B" forms a symmetric Nash equilibrium.*

Proof. Suppose that a customer arrives and observes queue lengths n_A and n_B, and that all customers before and after her are following the strategy "visit station A first in every state." Recall that n_A and n_B may be fractional, as each represents the number of customers in the queue at the respective station, plus the remaining fraction of service time for the customer currently being served. Because service is deterministic and all players are following route AB, the arrival process to station B is also deterministic, at a rate of at most μ_A. Thus, with this strategy profile station B never builds up a queue, so if the arriving customer visits station A first, then her system time $S^A_{n_A,n_B}$ is exactly

$$S^A_{n_A,n_B} = \frac{n_A + 1}{\mu_A} + \frac{1}{\mu_B}. \tag{2.55}$$

To determine the exact expected system time for an arriving customer who visits station B first is less straightforward, but it suffices to establish a lower bound. On arrival, the time T^B until a customer who visits station B first finishes service at station B is exactly

$$T^B = \frac{n_B + 1}{\mu_B}.$$

The number of customers Z_A to arrive during this time window is a Poisson random variable with parameter λT^B, and by assumption, these later customers follow the equilibrium strategy of visiting station A first.

Let $S^B_{n_A,n_B}(Z_A)$ be the system time experienced by an arriving customer who observes queue lengths n_A and n_B and chooses route BA, as a function of the number Z_A of customers who arrive during her time at station B. For a given realization k of Z_A, the system time for our arriving customer depends on the arrival times of these k subsequent customers because of possible idle time at station A. However, given k arrivals during T^B, our customer cannot possibly begin service at station A until that server has finished processing the $n_A + k$ customers that must be served before her. Incorporating her own service time at station A, we have that her total system time $S^B_{n_A,n_B}(k)$ from visiting station B first must satisfy

$$S^B_{n_A,n_B}(k) \;\geqslant\; \frac{n_A + k}{\mu_A} + \frac{1}{\mu_A}.$$

We thus have the relation

$$\mathbb{E}[S^B_{n_A,n_B}(Z_A)] \geqslant \sum_{k=0}^{\infty} \Big(\frac{n_A+1}{\mu_A} + \frac{k}{\mu_A}\Big)\mathbb{P}[Z_A = k] = \frac{n_A+1}{\mu_A} + \frac{1}{\mu_A}\mathbb{E}[Z_A]. \tag{2.56}$$

As mentioned, the random variable Z_A has a Poisson distribution, and its expected value is $\mathbb{E}[Z_A] = \lambda T^B$. Equation (2.56) can therefore be expressed as

$$\mathbb{E}[S^B_{n_A,n_B}(Z_A)] \;\geqslant\; \frac{n_A+1}{\mu_A} + \frac{1}{\mu_A}\Big[\lambda\Big(\frac{n_B+1}{\mu_B}\Big)\Big]. \tag{2.57}$$

We have $\lambda/\mu_A \geqslant 1$, which implies that

$$\begin{aligned} &\frac{1}{\mu_B} \leqslant \frac{\lambda}{\mu_A}\Big(\frac{n_B+1}{\mu_B}\Big) \\ \iff\; &\frac{n_A+1}{\mu_A} + \frac{1}{\mu_B} \leqslant \frac{n_A+1}{\mu_A} + \frac{1}{\mu_A}\Big[\lambda\Big(\frac{n_B+1}{\mu_B}\Big)\Big]. \end{aligned} \tag{2.58}$$

Applying equations (2.55), (2.57), and (2.58), we now have

$$S^A_{n_A,n_B} = \frac{n_A+1}{\mu_A} + \frac{1}{\mu_B} \leqslant \frac{n_A+1}{\mu_A} + \frac{1}{\mu_A}\Big[\lambda\Big(\frac{n_B+1}{\mu_B}\Big)\Big] \leqslant \mathbb{E}[S^B_{n_A,n_B}(Z_A)].$$

Each customer's expected system time is weakly smaller if she follows the candidate equilibrium strategy profile. Thus, when the system is overloaded and all customers are following the strategy "visit station A first in every state," no player has incentive to deviate, and we have a Nash equilibrium. □

Note that the system is overloaded, so over an infinite horizon, the queue lengths tend to infinity. However, the strategy is specified for all values of n_A and n_B, and any arriving customer finds a queue with finite-length, so the instability of the queue does not preclude the Nash equilibrium. Thus, in an overloaded (i.e., heavily congested) system with stochastic arrivals, we again observe the herding phenomenon.

$\mathcal{S}$-Station Open Routing Game

We now discuss systems with more than two stations. Consider a generalized version of the open routing game of Section 2.4 for a system with $\mathcal{S}$ stations and N players. For stations $\xi = 1, 2, \ldots, \mathcal{S}$, let μ_ξ be the service rate at station ξ, and assume that $\mu_1 < \mu_2 < \cdots < \mu_\mathcal{S}$.

If players make routing decisions simultaneously and if priorities are drawn uniformly at random, then players must choose from among the $\mathcal{S}!$ possible routing vectors. In this case, we have a Nash equilibrium which is analogous to that in the open routing game with two stations. Specifically, it is a Nash equilibrium for players to choose a route that visits stations in order of increasing service rate, from the slowest to the fastest.

Proposition 12 (Nash Equilibrium for Unobservable $\mathcal{S}$-Station System)**.** *If we have*

$$N \geqslant 1 + 2\mu_1 \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi}, \tag{2.59}$$

then the unobservable $\mathcal{S}$-station open routing game has a Nash equilibrium in which all players choose the routing vector $(1, 2, \ldots, \mathcal{S})$.

Proof. Suppose that all players are following the routing vector $(1, 2, \ldots, \mathcal{S})$, and consider a player who contemplates deviation. On path (that is, if everyone follows the prescribed profile), the only time any customer will face a queue is at station 1. All of the other stations start empty and then receive arrivals only when the station immediately before them completes a service. Because the routing vector is in order of increasing service rate, a queue will never build up at any station other than station 1. Thus, if the customer follows the prescribed profile, then her expected system time $\mathbb{E}[S^{\mathrm{EQ}}]$ is given by

$$\mathbb{E}[S^{\mathrm{EQ}}] = \sum_{j=1}^{N} \frac{1}{N}(\frac{j}{\mu_1}) + \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi} = \frac{N+1}{2\mu_1} + \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi}.$$

Because the player will not face a queue at any station besides station 1, she cannot possibly improve her system time by changing the order of stations that she visits after station 1. Thus the only deviations that we must consider are those which involve a vector that starts at a station other than station 1. If the player starts at station $\xi \geqslant 2$, then she will necessarily be the last customer to be served at station 1. Thus, we easily have a lower bound on her expected system time $\mathbb{E}[S^{\mathrm{D}}]$ from deviating given by

$$\mathbb{E}[S^{\mathrm{D}}] \geqslant \frac{N}{\mu_1}.$$

Equation (2.59) then gives

$$N \geqslant 1 + 2\mu_1 \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi}$$

$$\implies \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi} \leqslant \frac{N-1}{2\mu_1}$$

$$\implies \mathbb{E}[S^{\mathrm{EQ}}] = \frac{N+1}{2\mu_1} + \sum_{\xi=2}^{\mathcal{S}} \frac{1}{\mu_\xi} \leqslant \frac{N}{\mu_1} \leqslant \mathbb{E}[S^{\mathrm{D}}].$$

The player's expected system time is less if she follows the prescribed profile, and we therefore have a symmetric Nash equilibrium where all customers herd at station 1 and follow the routing vector $(1, 2, \ldots, \mathcal{S})$. □

Thus once again we have a herding effect, where players follow the crowd and congregate at a single station.

2.8 Queueing Networks in Steady State

In all of the systems that we have studied thus far, the open-routing service networks were *congested*, i.e., until all customers had arrived, the service rates were lower than the customer arrival rates. In all of those systems, the herding phenomenon has prevailed. We now turn to the question of whether herding behavior continues to emerge in a service network which does not face congestion. Specifically, we investigate equilibrium behavior when customers arrive according to a Poisson process with rate λ, rather than all being present in the system when service becomes available. Again every customer must be processed at both station A and station B. Service times at stations A and B are independent and exponentially distributed with rates μ_A and μ_B, respectively. Here we take $\lambda < \mu_A < \mu_B$, and we assume that arriving customers find the queueing network in steady state and are served in order of arrival at each station. Note that $\lambda < \mu_A$ implies that, in contrast to earlier settings, with

positive probability station A will have idle time in between services. The state of the system is not observable to the customers, i.e., customers may not condition their actions on the system state which includes the queue lengths and types of customers at both stations.

The distribution of the queue length encountered by a customer at the first station in her route is the steady-state distribution of the queue length at that station. In addition, classical results from Kelly (1979) give us that the distribution for the queue length observed by a customer at the second station in her route also matches the steady-state distribution for the queue length at that station. Next, we state a lemma which is a version of Corollary 3.5 from Kelly (1979) modified to our queueing network, which has customer types AB and BA, and queues A and B.

Lemma 3 (Kelly, 1979, Corollary 3.5)**.** *When a customer of type $\psi \in \{AB, BA\}$ reaches station $\xi \in \{A, B\}$, the probability that she finds κ customers at station ξ is equal to the steady-state probability that there are κ customers at station ξ.*

We use "type" to refer to a realized route through the network (AB or BA) and "group" to refer to a subset of the population that shares an ex-ante routing probability. Specifically, let $\mathbf{p} = (p_1, p_2, \ldots)$ be a (possibly infinite-dimensional) vector of routing probabilities corresponding to the strategies chosen by the corresponding groups, e.g., p_1 is the probability that a customer in group 1 chooses route AB. Let $\mathbf{z} = (z_1, z_2, \ldots)$ be a vector of strictly positive numbers summing to 1 representing the percentage of the population in each group. The vectors $\mathbf{p}$ and $\mathbf{z}$ together specify a strategy profile for the entire population. Then, we have that external arrivals of each type occur via independent Poisson streams with rate λp_A to route AB and λp_B to route BA, where p_A and p_B are given by

$$p_A = \langle \mathbf{p}, \mathbf{z} \rangle \text{ and } p_B = 1 - p_A.$$

The following result implies that the expected system time is the same for both routes under any strategy profile.

Corollary 6 (Expected System Time in Steady-State)**.** *For any strategy profile, i.e., for any pair of* $\mathbf{p}$ *and* $\mathbf{z}$*, the expected system time is the same for customers taking route* AB *and* BA*.*

Proof. Let $\mathbb{E}[S^A]$ be the expected system time for a customer who takes route AB, and similarly let $\mathbb{E}[S^B]$ be the expected system time for a customer who takes route BA. Also, take $L_A(\mathbf{p}, \mathbf{z})$ and $L_B(\mathbf{p}, \mathbf{z})$ to be the steady-state expected number of customers at stations A and B, respectively, when the strategy profile specified by $\mathbf{p}$ and $\mathbf{z}$ is played. By Lemma 3, we have

$$\mathbb{E}[S^A] = \frac{L_A(\mathbf{p}, \mathbf{z})}{\mu_A} + \frac{L_B(\mathbf{p}, \mathbf{z})}{\mu_B} = \mathbb{E}[S^B].$$

□

In this setting, we define our notion of equilibrium to be that customers of every group must be best-responding to the strategy profile defined by $\mathbf{p}$ and $\mathbf{z}$. As the expected system time is the same for either route, customers of every group are playing a best-response to all other groups, and there is no incentive for members of any group to change their strategy. Therefore, any $\mathbf{p}$ and $\mathbf{z}$ form an equilibrium strategy profile.

Observe that when the system is allowed to reach steady state, the herding profiles in which all customers choose the same route are indeed equilibria. However, as any other feasible routing profiles also form an equilibrium, there is no reason to give special preference to the herding profiles in this setting. Consequently, these results suggest that the herding behavior which prevails in congested service networks no longer predominates in systems which are not very busy.

2.9 Simulation Study for Systems with Stochastic Arrivals

We have shown that herding occurs in congested open-routing service networks if all customers are present at the start of service. Moreover, our results in Section 2.8 suggest that herding probably does not occur in a system which is not congested, i.e., with a service rate faster than the arrival rate. These analyses lead naturally to the following hypothesis.

Hypothesis 1. *Herding occurs when a service system is congested, that is, the arrival rate is higher than the service rates of either station until the arrival of the last customer.*

We are interested in Hypothesis 1 because the arrival rate to a service system is often not constant, and a service system may experience a high customer arrival rate especially at the start of its service availability. Note that if we assume that customers are fully rational in the open routing game with stochastic arrival times, then each customer would be required to perform prohibitively difficult analysis to solve for a Bayesian Nash Equilibrium. Even highly intelligent customers are unlikely to implement such outcomes. Accordingly, we next perform a simulation study designed to test Hypothesis 1. We assume that the simulated customers learn about the system through repeated rounds of play, and we simulate systems in which customers arrive stochastically over time and the service times at both stations are stochastic.

2.9.1 Simulation Setup

The setting for the simulation is as follows. Play proceeds for multiple rounds, and customers update their beliefs after each round of play based on the wait time that they experience. In each round, customers also observe what their wait time would have been had they chosen the opposite route, fixing the moves of the other

customers. Thus, in each round all customers get samples for both routes. Each customer's assessment of her expected system time on a particular route is equal to the empirical average of her own samples for that route across all rounds. Moves are randomly generated in the first round, and thereafter each customer chooses the route with the smaller expected system time according to her beliefs in that round.

We define the parameters γ and ϕ to control the mean and variance of the arrival times. Specifically, the arrival time of customer $i \in \{1, \ldots, N\}$ is assumed to be uniformly distributed on the interval $[i\gamma - \phi, i\gamma + \phi]$. Thus the mean arrival time for customer i is equal to $i\gamma$, and the width of the uniform distribution is 2ϕ. Arrival times of different customers are mutually independent. If ϕ is large enough relative to γ, then successive intervals may overlap and customers may not always arrive in the same order. For example, if $\gamma = .1$ and $\phi = .25$, then the arrival time of customer 1 will be uniformly distributed on $[-.15, .35]$, the arrival time of customer 2 will be uniformly distributed on $[-.05, .45]$, etc. Service times at each station are taken to be exponentially distributed, where the service rate μ_B at station B is fixed at 1 and the rate μ_A at station A varies in each experiment.

We simulated all combinations of $\gamma \in \{.001, .1, .25, .5, .75, 1\}$, $\phi \in \{0, .25, .5, .75, 1\}$, and $\mu_A \in \{.1, .25, .5, .75\}$, for a total of 120 experiments. For each parameter combination we ran 100 independent trials with 250 rounds of play in each trial. In all cases we consider $N = 50$ customers. We analyze the results of the simulation study in the next subsection.

2.9.2 Results of Simulation

In Figure 2.3, we depict the empirical frequencies for the number of simulated customers who choose route AB in the final round (round 250) for three different values of the service rate μ_A. The left panel shows the results when the arrival times are deterministic and the arrival order is fixed, i.e., when $\gamma = .001$ and $\phi = 0$. This

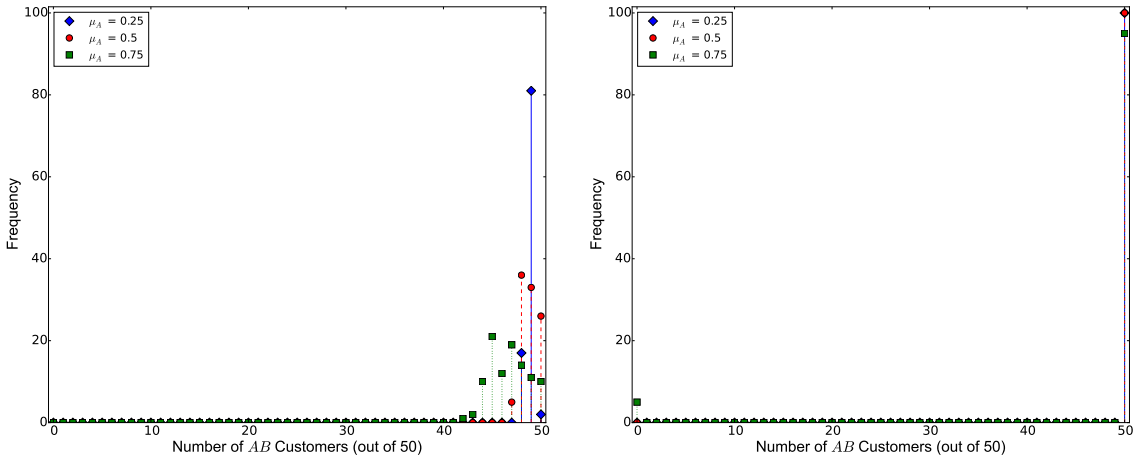

(a) $\gamma = .001$; $\phi = 0$ (b) $\gamma = .001$; $\phi = .75$

FIGURE 2.3: Frequency charts for number of AB customers in round 250 (100 trials).

example closely resembles the sequential open routing game of Section 2.5, where customers are all present at time zero and make their moves sequentially. In the simulation, unlike in Section 2.5, customers are not able to observe the moves of those before them in the order when making their routing decisions. Similarly, customers do not explicitly incorporate the strategies of later arrivals. Instead, each customer decides her route based on her experience from the historical rounds.

Recall from Proposition 7 that in the subgame perfect equilibrium of the sequential game, the first $N-1$ customers choose route AB, and the final customer chooses route BA. Referring to the plot in the left of Figure 2.3, we see that the outcome of the simulation is similar to the equilibrium; the density is mostly concentrated close to 50. As noted, in our simulated system, customers do not condition their strategies on the moves of earlier arrivals, and we therefore do not always see exactly $N-1$ customers choosing route AB. Still, experience from the previous rounds has taught them that if they take route BA, then they may be overtaken at station A by later arrivals, resulting in more total time in the system. The same observation persists for all values of μ_A.

Table 2.3: Summary statistics with $\mu_A = .75$ for number of AB customers (out of 50) in the final round.

(a) Sample Mean

γ \ ϕ	0	.25	.5	.75	1
.1	46.22	46.61	45.42	47.45	49.2
.25	46.55	45.83	46.22	46.07	45.72
.5	45.2	45.27	45	44.64	44.32
.75	41.68	42.7	41.59	41.73	42.46
1	35.9	37.38	36.47	36.23	38.17

(b) Sample First Quartile

γ \ ϕ	0	.25	.5	.75	1
.1	44	45	45	47	49
.25	45	44	44.75	44	44
.5	43	43	42	42	42
.75	39.75	41	38.75	40	40.75
1	33.75	34.75	34	34	35

In the right panel of Figure 2.3, we see results for the case in which again customers all arrive very close to the start of service, but now there is significantly more variance in the arrival order. We thus have a setting which is reminiscent of the one-shot open routing game of Section 2.4. The arrival times have the same tightly-spaced means—$\gamma = .001$—as in part (a). Now, however, we have $\phi = .75$, which is much greater than the successive difference in mean arrival times γ. Far from a fixed order, each player now could arrive in any position among the 50 customers. As she forms her beliefs over several rounds, the empirical averages will reflect this randomness in her priority.

Continuing to consult the right plot of Figure 2.3, when the service rate at station A is 75% of that at station B, we see that the outcomes in the final round are split between the two herding profiles, with more than 90% of trials ending up with herding at route AB and the rest ending with herding at route BA. When the service rate at station A is less than half of that at station B (25% and 50%), in all 100 trials we see play converging to herding at route AB. These outcomes match our analytical results from Section 2.4; that is, the herding equilibria are the only pure-strategy Nash equilibria of the open routing game; that a large class of learning rules will converge to one of the herding profiles; and that when the service rates differ by a factor greater than 2, route AB is a strictly dominant strategy for all customers.

We have thus far discussed examples with arrivals that are very close together ($\gamma = .001$), which matches closely with our theoretical models. We now discuss settings for which arrivals are more spaced out in time, i.e. $\gamma \geqslant .1$. We begin by fixing μ_A at .75, varying γ within $[.1, 1]$ and ϕ within $[0, 1]$. We note that in our simulations, herding on route BA never occurs when $\gamma \geqslant .1$. Because herding at station B does not arise, the number of AB customers in the final round—more specifically, how close this number is to 50—is a reasonable measure of the strength of herding. The mean and first quartile for this quantity can be found in Table 2.3, over ranges of γ and ϕ. We observe that all three summary statistics for the number of AB customers tend to decrease as γ increases. That is, as the arrivals occur less frequently, herding begins to dissipate. For $\gamma \leqslant .5$ we see a marked tendency toward herding. For example, when $\gamma = .25$, the mean number of AB customers in the final round for any ϕ is above 45, which is more than 90% of the total number of customers. However, once γ nears one, the herding effect is less pronounced.

Results of the simulation runs with smaller values of μ_A are qualitatively similar. First, recall that our theoretical results for service networks with all customers present at the start of service tell us that route AB is a strictly dominant strategy if $2\mu_A \leqslant \mu_B$. Our numerical analysis suggests that this extends to cases with stochastic arrivals as well. In our simulations, we find that across all trials and for all parameter combinations such that $2\mu_A \leqslant \mu_B$, the *minimum* number of AB customers in the final round is 44. We now discuss the particular case of $\mu_A = .5$, the summary statistics for which are reported in Table 2.4. We note that every entry in Table 2.4 is at least 47, and that the values in Table 2.4 are greater than the corresponding values in Table 2.3 (with $\mu_A = .75$) in all cases. This pattern of successively more prominent herding continues as μ_A decreases further. Such a pattern indicates, as also suggested by the proofs establishing the herding equilibria, that an important driver of the herding behavior is the severity of the penalty of being behind one

Table 2.4: Summary statistics with $\mu_A = .5$ for number of AB customers (out of 50) in the final round.

(a) Sample Mean

		ϕ				
		0	.25	.5	.75	1
γ	.1	48.47	48.53	49.98	50	50
	.25	48.34	48.39	47.81	48.61	49.18
	.5	48.05	48.07	48.3	48.1	47.82
	.75	47.88	47.9	47.81	47.8	47.76
	1	47.48	47.42	47.47	47.52	47.33

(b) Sample First Quartile

		ϕ				
		0	.25	.5	.75	1
γ	.1	48	48	50	50	50
	.25	48	48	47	48	49
	.5	47	47	48	47.75	47
	.75	47	47	47	47	47
	1	47	47	47	47	47

additional customer.

The outcomes of our simulation study provide strong evidence for Hypothesis 1, i.e., that herding emerges when the open routing service network is congested. Moreover, our study shows that as arrivals begin to be spaced further apart in time, herding correspondingly diminishes in prevalence, verifying our analysis in Section 2.8. Additionally, when the service rates are far apart (i.e., when $2\mu_A < \mu_B$), the pull to the slower station A is strong enough that herding is more robust to slower arrivals. These numerical results mirror qualitatively the theoretical results from previous sections, reinforcing the plausibility of herding in congested service systems in more realistic settings.

2.10 Conclusion

We model customer behavior for service networks in which self-interested customers require service at each station and are permitted to determine their routes through the network. In our base two-station model, customers are present in the system when service becomes available and make decisions about which station to visit first. In equilibrium, customers herd at one station; that is, all of the customers take the same route through the network. This behavior is motivated by the need to avoid arriving late to the congested station. If all of the other customers are visiting the

same station first, then a customer who visits the other station guarantees herself to be served last at the busy one, and she is thus better off following the crowd. However, if the service rates are far apart, then it is a dominant strategy for all players to visit the slower station first.

In our base model, we show that—under a mild restriction on the system size—the expected total time that a given customer spends in the system is a submodular function of the number of other players choosing a given route. This structural property helps to rule out alternative, non-herding equilibria, and more importantly, lead to a proof on the convergence of various learning rules to herding. The analysis of learning rules is important because in the real world, customers rarely play equilibrium strategies right away, but instead learn from their experiences. As structural properties of waiting times have been rarely applied to the analysis of queueing games, this chapter may open the door for future use of submodularity in other settings in the queueing games literature.

The herding phenomenon is quite surprising at first glance, and one may naturally ask whether such herding behavior could occur in complex, realistic environments. To verify the robustness of the herding behavior, we analyze strategic open routing in several different settings—including games with dedicated customers, stochastic arrivals, stochastic service times, and more than two stations—and find that customers continue to herd under the presence of congestion. Moreover, in establishing the herding result under stochastic service times, we employ a novel combination of existing results in queueing theory on stochastic dominance and dual queues, which together are used to establish that certain steady-state results can be applied as valid bounds on the transient system. These steps reduce what would be a difficult (and likely intractable) analysis of a transient queueing network to a straightforward manipulation of well-known bounds on stationary distributions for $GI/GI/1$ queue length.

Interestingly, by applying Kelly's classical reversibility analysis, we also demonstrate that in the steady state of an underloaded system, herding is no longer the focal equilibrium. The contrasting results suggest the insight that herding in open-routing environments is driven by congestion; we then verify this insight with extensive numerical studies. In addition, we find from an investigation of the central optimum that herding is good for social welfare. The cumulative system time—the sum of the system times for all customers—is at its lowest when all but one customer visits the faster station first. In the presence of herding, the welfare loss is usually only a single service time.

This work provides the following two main takeaways to practitioners. First, in a congested open-routing service network, experienced customers are likely to herd. This insight is important because if system planners assume that strategic customers choose routes at random or join the shortest queue, then they would incorrectly evaluate the system. Second, our analysis shows that herding achieves a near-optimal cumulative system time. Therefore, if cumulative system time is the main objective, then instead of requiring customers to take a certain route, system planners can attempt to incentivize herding, or simply help the customers to understand that herding is their best strategy.

Service networks with open routing are an exciting area of study, and this chapter generates multiple avenues for future work. For example, in some cases, customers might have complex utility functions; if customers were risk averse or if they preferred not to wait too long at any one station, then the equilibrium analysis would be substantially different from ours. Another system worth studying would be one with more than two stations wherein each customer visits only a subset of the stations, extending our results on two-station systems with dedicated customers. Intriguing also would be a system with more than two stations where customers dynamically choose their routes after each successive service completion. Finally, it would be

of interest to learn about customer routing behavior in a system when jockeying between queues is allowed.

3

Service Delivery Platforms: Pricing and Revenue Implications

3.1 Introduction

Food delivery platforms, such as GrubHub, Caviar, and UberEats, act as intermediaries and offer consumers deliveries from restaurants, while Postmates offers deliveries from other types of services in addition to food. These platforms represent a growing segment of the so-called "sharing economy," which refers to businesses whose revenue model relies on connecting customers with service providers—who may or may not be directly employed by the business—often via a mobile app (see Wallenstein and Shelat, 2017).

While substantial research interest has recently been devoted to understanding the interaction between drivers and platforms in the sharing economy (see, e.g., Cachon et al., 2017; Taylor, 2017, and references therein), a similarly important but less studied issue for food delivery platforms is the interaction between the platform and the restaurant. Ride-sharing platforms like Uber and Lyft and short-term

Research in this chapter was supervised by Pnina Feldman and Robert Swinney.

housing platforms like AirBnB have disrupted existing business models by cannibalizing demand from the established players in their respective industries (taxis and hotels, respectively). Such platforms are in direct competition with the established providers. Food delivery platforms, by contrast, maintain a symbiotic relationship with the existing providers in their industry; rather than "stealing" demand from an established player, these platforms work with restaurants to connect customers with the restaurant's product by providing an additional purchase channel.

However, while this relationship entails cooperation between the platform and the restaurant, there is substantial opportunity for misalignment between them. Specifically, the delivery platform earns revenue only from customers who order delivery, and receives nothing from dine-in customers. When solving a revenue maximization problem, the platform will set its price so as to maximize what it receives from delivery customers, which will likely not yield the system-optimal revenue. Similarly, the restaurant earns different revenue per customer from the different channels, and it may have a relative preference for one type or the other which does not match the coefficients of the aggregate revenue function.

We take a queueing approach to modeling this system, treating the restaurant's kitchen as an observable $M/M/1$ queue with customer waiting costs. On arrival, a customer observes the queue length and chooses to dine in, order delivery, or balk (i.e., not purchase), based on her reward, the prices, and her expected cost of waiting on each channel.

This chapter makes two main contributions. First, we solve the revenue maximization problem faced by a monopolist who controls both the dine-in and delivery prices and receives all revenue from the system. Our results on the optimal pricing scheme are related to the priority queueing and pricing literature (see, e.g., Adiri and Yechiali, 1974, Alperstein, 1988, and references therein) and are of independent interest. Second, we investigate means of coordinating the decentralized supply chain via

different contracts between the restaurant and the platform, where "coordinating" has the usual meaning of implementing an equilibrium outcome which achieves the maximum possible aggregate revenue. Our system shares some characteristics with the traditional manufacturer-retailer (here restaurant-platform) setup, but a crucial difference is that in addition to serving as a wholesaler to the platform, the restaurant also has an independent revenue stream from dine-in customers. This feature fundamentally changes the incentives from the systems studied, e.g., in Cachon and Lariviere (2005).

For the monopoly case, we follow an approach similar to that in Adiri and Yechiali (1974) and Alperstein (1988), and we find the optimal prices for fixed values of the "control limits," which define the customers' strategy mapping observed queue lengths to purchase decisions. In equilibrium, for short queues, customers dine in; for medium-length queues, they order delivery; and for long queues, they balk. Expressing the prices as a function of the control limits allows us to optimize over the control limits themselves, and then to substitute the optimal solution to deduce the corresponding prices.

We show that the optimal dine-in control limit does not depend on the maximum total system size. More precisely, the optimal threshold between dine-in and platform customers is a constant which does not depend on the maximum total queue length but is a function of only the model primitives (unless the maximum total queue length is less than said constant, in which case it is optimal to induce all arriving customers to dine in). In terms of comparative statics, for the special case with $\rho = 1$, the optimal dine-in control limit is increasing in the customer valuation, the service rate, and the amount of quality loss in delivery, and decreasing in the difference in unit waiting cost between waiting at home and waiting at the restaurant.

While deriving the optimal control limits, we find a surprising feature of the optimal prices: a monopolist will optimally set the delivery price *lower* than the

dine-in price. This result suggests that there may be substantial miscoordination, as in a decentralized setting the platform must charge a higher price than the restaurant in order to turn a profit. To investigate this issue further, we study a decentralized system where the restaurant and platform set prices independently, and the platform must pay the restaurant full retail price for each unit that it delivers to a customer. We show the existence of a pure-strategy subgame perfect equilibrium for the pricing game. More importantly, we observe that in the decentralized setting, the equilibrium prices can lead to a sub-optimal aggregate revenue, so there is indeed a coordination loss.

We then move to analyzing contracts which may recover the optimal aggregate revenue. We identify a contract that is related to the revenue sharing contract of Cachon and Lariviere (2005); this contract coordinates the system and can assign arbitrary fractions of the aggregate revenue to each party. Specifically, a two-way revenue sharing contract coordinates, where the restaurant shares some of its dine-in revenue with the platform, and the platform shares some of its revenue with the restaurant. Notably, the share from the platform to the restaurant may be less than the retail price that the restaurant charges dine-in customers. What is required to coordinate the supply chain is that the fraction of revenue that each party earns is the same for both channels.

Finally, we investigate numerically the extent of the coordination loss under different contract terms, specifically a "no-contract" setting and one with a traditional one-way revenue sharing contract. Overall, we find that both the restaurant and the platform tend to benefit from each other's presence, with both parties being more valuable to each other for intermediate values of the offere load ρ. In addition, in certain parameter regimes, the restaurant extracts an especially large benefit from the platform's presence because of the lack of a price-demand tradeoff. In other words, there are parameter regions where, in equilibrium, the restaurant both (a)

charges a higher price with than without the platform and (b) sees a larger throughput with than without the platform. The restaurant can accomplish this feat because of the reduced waiting-cost coefficient for platform customers; the additional channel allows the restaurant to segment customers. The segmentation is not by intrinsic characteristics, but rather by the queue length observed upon the customer's arrival. Customers who arrive to a short queue care less about their waiting-cost coefficient and more about the difference in gross reward between the channels, while to customers arriving to a long queue the difference in waiting costs is more salient.

The chapter is organized as follows. Section 3.2 reviews the relevant literature. In Section 3.3, we describe our model, and in Section 3.4 we establish the optimal customer purchasing strategies given problem parameters and restaurant and platform prices. Section 3.5 studies the pricing problem faced by a hypothetical monopolist who controls both the restaurant's and the platform's prices and seeks to maximize aggregate revenue. Section 3.6 studies the pricing game played by an independent restaurant and platform, and it also introduces a new contract type to coordinate the system. Sections 3.7-3.8 study the system numerically to uncover the key features directing customer and firm behavior, and Section 3.9 concludes.

3.2 Related Literature

This work is closely connected to the methodology-driven literature on pricing and priority classes in queueing systems. Naor (1969) was the first to rigorously study equilibrium customer queue-joining behavior, and a vast literature has grown since then on strategic customer, server, and firm behavior in queueing networks. Hassin and Haviv (2003) provides an introduction to the area and a survey of the relevant literature, and Hassin (2016) covers the more recent work in this stream.

The closest work to ours in this area is the stream of literature on the setting of prices for different priority levels in queues. Adiri and Yechiali (1974) introduces the

problem of equilibrium customer purchasing behavior and firm pricing decisions when customers arrive to an observable $M/M/1$ queue and decide what level of priority to purchase (with the additional option of not purchasing, i.e., balking). They show that when there are M priority levels from which to choose, the optimal customer strategy is to implement "control limits". That is, a set of thresholds determines intervals of the queue length, such that each interval maps to a priority level that the customer should optimally purchase if they observe a queue length within said interval. This result is a generalization of Naor's result that with a single priority class, there is a threshold queue length above which it is optimal for customers to balk. They propose a procedure for determining the optimal priority prices for the firm to offer, but stop short of providing closed-form expressions for these prices.

Alperstein (1988) extends Adiri and Yechiali (1974) in several ways. She shows that the optimal priority prices result in control limits of one for all priority classes. That is, if the firm charges the optimal prices, then each arriving customer will purchase the lowest priority level that does not yet have a customer in the queue, up until a maximum threshold at which an arriving customer balks. This process results in last-come, first-serve (LCFS) service because whoever arrives will pay to move in front of everyone in currently in the queue. In her analysis she employs a procedure which we will wield frequently: by observing that, for a given set of control limits, it is optimal for the firm to charge the highest price consistent with such limits, she is able to express the prices in terms of the control limits. Substitution then allows her to express and optimize the firm's revenue function in terms of only the control limits, eliminating the prices from the expressions and greatly simplifying the analysis.

The crucial difference between our work and that of Adiri and Yechiali (1974) and Alperstein (1988) is that we do not have priority classes, although our system operates similarly to a priority system. We do have two classes of customers, but

rather than the price determining the length of the queue that a customer will face, for us the price paid determines the waiting cost coefficient. This distinction leads to some surprising results regarding the optimal prices of the two "classes" (dine-in or delivery) in our model.

Hassin and Haviv (1997) study a two-class version of the model of Adiri and Yechiali (1974), giving a detailed treatment of equilibrium customer joining behavior. They extend the definition of a threshold strategy to allow customers to mix between choosing the priority or the regular queue when the queue length is at a certain threshold value. They uncover cases of multiple equilibria and even an example of an equilibrium customer strategy that does not conform to even their extended definition of a threshold strategy. By contrast, in our model there always exists a unique customer equilibrium, and this equilibrium involves a pure threshold strategy. The reason is that in our model, a customer's optimal strategy does not change with the strategy employed by others, since her choice of purchase channel does not change her position in the queue, but only changes her waiting-cost coefficient.

Cachon and Feldman (2011) study a system under congestion where a firm must decide whether to operate a subscription service, where customers pay a flat fee for unlimited use of the service, or a pay-per-use scheme. They find that subscription can be the preferred option for the firm even when customers strongly dislike congestion, a counterintuitive result given that under a subscription model participating customers have no incremental cost to use the service.

Debo and Veeraraghavan (2014) find nonmonotonic queue joining behavior in a different setting from Hassin and Haviv (1997). The base setting is still an observable single-server queue, but they assume that the service rate and service value are unknown to customers and positively correlated. Under these assumptions, a longer queue upon arrival gives the customer a higher expected waiting time/cost and also a higher expected reward from service. This effect can lead to so-called "sputtering

equilibria," where at a shorter queue length customers mix between joining and balking, while for a longer queue length they certainly join. Unusual queue-joining strategies have also been discovered by the same authors in Veeraraghavan and Debo (2009, 2011). In these works, customers must choose between two service providers. Some customers are informed of the service value while others are not, meaning that a longer queue can signal higher quality. Because of this signaling property, customers may "herd," i.e, choose the service provider with the longer queue.

Erlichman and Hassin (2015) study a setting related to the model of Adiri and Yechiali (1974) in which arriving customers choose upon arrival how many of the currently queued customers to overtake. This model is similar to one with many priority classes, but the price function is restricted to be linear in the number of customers overtaken, and also in this setting even by paying the highest available price a customer cannot guarantee that she will not be overtaken by a later arrival. They find some surprising equilibria in which the number of customers to overtake upon arrival is nonmonotonic in the queue length.

Gavirneni and Kulkarni (2016) investigate a problem related to two-class priority queues, focusing on an unobservable setting. They model customer waiting costs using a Burr distribution and study the welfare impact on service systems of going from a FCFS setting to a two-class priority setting. They find that adding a priority option can achieve a Pareto-improvement in welfare and that the greatest net benefit is realized for heavily-loaded systems with significant disparity in waiting costs.

Wang et al. (2018) also study a priority queueing system with two classes, regular and priority. Theirs is perhaps the most closely related recent work to ours in the priority queueing and pricing literature. Customers arrive to the system and can choose to balk, purchase regular service, or purchase priority. They consider both observable and unobservable queues and characterize customer equilibrium joining/balking decisions and optimal priority prices. In the observable case, similar to our results they

find that in equilibrium the customer equilibrium strategy involves two thresholds: with a short queue, purchase regular service; with an intermediate queue, purchase priority; and with a long queue, balk. They also conclude that the service provider is better off with an observable queue when the system is either very lightly or very heavily loaded, and better off with an unobservable queue for intermediate levels of offered load. In addition, they address the relationship between the queueing discipline (preemptive or nonpreemptive) and the potential for multiple equilibria like those found in Hassin and Haviv (1997). Other recent work in this area includes Yang and Debo (2017), Yang et al. (2016), Afeche and Sarhangian (2015), Allon et al. (2011), and Gurvich et al. (2018).

It should be noted that our work shares some features with the literature on supplier encroachment. Arya et al. (2007), for example, find that a manufacturer adding a direct channel to sell to customers can benefit both the manufacturer and the retailer in the right circumstances. More recent work in this area (see, e.g., Li et al., 2013, Ha et al., 2015) investigates conditions under which encroachment benefits and hurts each party, and David and Adida (2015) investigate contracts to coordinate a supply chain with encroachment. See Mai et al. (2018b) for related work on sourcing issues faced by innovators who may need to use a competitor as a supplier. Apart from myriad differences in setting and model characteristics (for example, the papers mentioned above do not account for congestion or queueing effects), another key distinction between our problem and those in the encroachment literature is *who* is doing the encroaching. With supplier encroachment, the status quo is for the retailer to be the only party directly selling to customers, while the manufacturer sells only through the retailer. For us it is the opposite: the status quo involves the "manufacturer" (restaurant) selling directly to customers, which is then disrupted by the "retailer" (platform) entering the market and encroaching on the manufacturer's territory.

Finally, this chapter is related to the emerging stream of literature on strategic customer and firm behavior in the sharing economy. Work on pricing and wages in ride-sharing platforms includes Banerjee et al. (2015), Cachon et al. (2017), Taylor (2017), Afèche et al. (2018), Bai et al. (2018), Benjaafar et al. (2018), Bimpikis et al. (2018), and references therein. Other work looks at issues of information asymmetry and the release of information in ride-sharing platforms and other platforms in the sharing economy. Romanyuk (2017) studies the role played by the platform in releasing (or withholding) information about buyers to (from) the sellers. Mai et al. (2018a) and Jin et al. (2018) study bilateral rating systems on ride-sharing platforms, while Bolton et al. (2013) experimentally study reciprocity effects—i.e., settings in which the probability that one party to a transaction submits a rating is influenced by whether or not the other party submits a rating—in buyer and seller ratings on an e-commerce platform. Kanoria and Saban (2017) study the matching of agents in two-sided markets with screening costs. Nikzad (2018) also studies the impact of a platform's matching algorithm, as well as the effect of the size of the labor pool and the number of platforms on the welfare of workers.

3.3 The Model

Suppose that there is one firm which operates a sit-down restaurant. Suppose also that there is an independent delivery platform which, for an extra fee, delivers food from the restaurant to customers who wish to dine at home.

The kitchen serves customers at rate μ, with exponentially distributed service times, and there is a Poisson stream of potential customers which is generated at rate Λ. Let ρ be the offered load, i.e., $\rho := \Lambda/\mu$. If $\rho \geqslant 1$, then if all arriving customers choose to join the system, the kitchen queue will grow without bound. In that case, then, the only way for the system to be stable is if some arriving customers choose the no-purchase option. If instead $\rho < 1$, then the system will be stable even if

all customers purchase; in equilibrium, however, it may still be that some customers choose not to purchase.

We assume that customers extract a reward R (homogeneous across the population) from being served their food at the restaurant. The quality of the food is degraded in the delivery process according to a constant factor $\delta \in (0, 1]$, which is deterministic and uniform across all customers; thus the reward from eating at home is δR. Waiting at the restaurant incurs a cost c_R per unit time, and waiting at home incurs a cost $c_H < c_R$ per unit time. Throughout, we will assume that $R\mu > c_R$, so that a customer arriving to an empty queue strictly prefers dining in to not purchasing if the restaurant price $p = 0$. Finally, we assume that the kitchen operates on a first-come, first-serve (FCFS) basis, irrespective of the source of the demand. The kitchen queue length is observable to customers, who make their purchasing decisions based on the queue and on the prices set by the firm and platform.

Customers are risk-neutral and maximize their expected reward. The firm charges customers a price $p \geqslant 0$. The delivery platform charges customers a price $\theta \geqslant 0$. Transfer payments between the restaurant and the platform are discussed in Section 3.6.

There is a long history of research on priority queueing, some of which considers the priority of a given customer or customer class to be a decision variable for the firm, and other work allowing customers to purchase a higher priority. In the latter case, the convention is for customers to pay a fee in order to move closer to the front of the queue. Our work is related to this stream of literature in that customers have the option to pay to change the nature of their time in the queue, in particular to Adiri and Yechiali (1974) and Alperstein (1988). However, a novel feature of our model is that, rather than paying for a shorter time in queue, here customers pay to make the queueing experience itself less painful (by reducing their waiting cost coefficient). This "payment" has two distinct elements: (i) an increase in the price

paid from p to θ (if the platform price is higher); and (ii) a loss in utility of $(1-\delta)R$ due to the quality degradation in delivery.

3.4 Customer Equilibrium

We note that our model reduces to that of Naor (1969) in the absence of a platform. In the next subsection, we study customer equilibrium joining decisions in the presence of the delivery platform.

3.4.1 Customer Equilibrium

Let $U_{\text{Rest}}(L)$ be the (random) utility that a customer gleans from dining at the restaurant, when the queue length is L. We have

$$\mathbb{E}[U_{\text{Rest}}(L)] = R - p - c_R\Big(\frac{L+1}{\mu}\Big).$$

Letting $U_{\text{Home}}(L)$ be the utility from dining at home, we similarly have

$$\mathbb{E}[U_{\text{Home}}(L)] = \delta R - \theta - c_H\Big(\frac{L+1}{\mu}\Big).$$

Note that, here and in the sequel, we make the following assumptions regarding customer behavior under indifference.

Assumption 1. *If indifferent between purchasing and not, a customer chooses to purchase.*

Assumption 2. *If indifferent between dine-in and delivery, a customer chooses to dine in.*

Following the convention in Adiri and Yechiali (1974), we use the term "control limit" in the case of a threshold strategy to refer to the queue length at which the optimal channel decision of an arriving customer switches from one channel to the other.

Lemma 4 (Control Limit Strategy)**.** *Control limits L_R and L_H exist such that:*

(i) *If the queue length $L \leqslant L_R - 1$, then an arriving customer prefers to dine at the restaurant;*

(ii) *If $L_R \leqslant L \leqslant L_R + L_H - 1$, then she prefers to use the platform;*

(iii) *If $L \geqslant L_R + L_H$, then she chooses not to purchase.*

It is possible that $L_R = 0$ or $L_H = 0$, in which case one or more of the above scenarios may never occur.

Proof. An arriving customer who observes a queue length L prefers dining at the restaurant to dining at home if and only if $\mathbb{E}[U_{\text{Home}}(L)] \leqslant \mathbb{E}[U_{\text{Rest}}(L)]$, that is,

$$\begin{aligned} R - p - \frac{c_R(L+1)}{\mu} &\geqslant \delta R - \theta - \frac{c_H(L+1)}{\mu} \\ \iff L + 1 &\leqslant \frac{\mu((1-\delta)R + \theta - p)}{c_R - c_H}. \end{aligned}$$

An arriving customer prefers dining at the restaurant to not purchasing if and only if $\mathbb{E}[U_{\text{Rest}}(L)] \geqslant 0$, i.e.,

$$\begin{aligned} R - p - \frac{c_R(L+1)}{\mu} &\geqslant 0 \\ L + 1 &\leqslant \frac{\mu(R-p)}{c_R}. \end{aligned}$$

We can thus define the control limit L_R by

$$L_R := \left\lfloor \left[\frac{\mu(R-p)}{c_R} \wedge \frac{\mu((1-\delta)R + \theta - p)}{c_R - c_H} \right]^+ \right\rfloor, \tag{3.1}$$

where $\lfloor \cdot \rfloor$ as usual is the floor function. An arriving customer prefers ordering from the platform to not purchasing if and only if $\mathbb{E}[U_{\text{Home}}(L)] \geqslant 0$, i.e., if

$$\begin{aligned} \delta R - \theta - \frac{c_H(L+1)}{\mu} &\geqslant 0 \\ L + 1 &\leqslant \frac{\mu(\delta R - \theta)}{c_H}. \end{aligned}$$

The control limit L_H is then written as

$$L_H := \left[\left\lfloor \frac{\mu(\delta R - \theta)}{c_H} \right\rfloor - L_R \right]^+. \tag{3.2}$$

□

As the kitchen operates on a first-come, first-serve basis, the state of the system is fully described by the length of the queue. The system therefore evolves as a finite-capacity $M/M/1$ queueing system with a maximum size equal to $L_R + L_H$. Hereafter, customers are assumed to follow the optimal strategies defined in this section in response to platform and restaurant prices.

3.5 Monopoly Pricing

With the goal of coordinating the supply chain, we first study the problem of a monopolist, who controls the prices of both the restaurant and the platform and seeks to maximize the total aggregate revenue across both channels. We note that the first part of our derivation follows a similar path to that in Alperstein (1988) for the case of pricing priority queues. By manipulating the prices p and θ, it is possible for the monopolist to choose the optimal control limits L_R and L_H that will be implemented by the customers.

Let π_i be the steady-state probability of an arriving customer finding i customers in the system. Similarly, let $\Pi_{j,k}$ be the steady-state probability of finding at least j and at most k customers in the system, i.e.,

$$\Pi_{j,k} := \sum_{i=j}^{k} \pi_i. \tag{3.3}$$

The aggregate revenue per unit time, $Z(\cdot)$, can then be expressed as a function of

the prices p and θ by

$$Z(p,\theta) = \Lambda\Big[p\Pi_{0,L_R-1} + \theta\Pi_{L_R,L_R+L_H-1}\Big].$$

The control limits L_R and L_H defined in equations (3.1) and (3.2) are step functions. So, for a given pair (L_R, L_H) it is optimal to set the maximum prices p and θ which will result in these control limits. We give expressions for the optimal prices for a given pair of control limits in the next lemma.

Lemma 5 (Optimal Prices for Fixed Control Limits)**.** *Consider a control-limit pair* (L_R, L_H)*. The revenue-maximizing prices consistent with such control limits are given by*

$$p = R - \frac{c_H(L_R + L_H)}{\mu} - \frac{(c_R - c_H)L_R}{\mu} \tag{3.4}$$

and

$$\theta = \delta R - \frac{c_H(L_R + L_H)}{\mu}. \tag{3.5}$$

Proof. First suppose that both channels are active, so $L_R \geqslant 1$ and $L_H \geqslant 1$. For a given pair of control limits (L_R, L_H), if L_H is nonzero then the platform price θ must satisfy

$$\delta R - \frac{c_H(L_R + L_H + 1)}{\mu} < \theta \leqslant \delta R - \frac{c_H(L_R + L_H)}{\mu}. \tag{3.6}$$

Note that the left inequality in equation (3.6) is strict because we assume that when indifferent a customer chooses to purchase. Similarly, the restaurant price must satisfy

$$R - p - \frac{c_R(L_R + 1)}{\mu} < \delta R - \theta - \frac{c_H(L_R + 1)}{\mu} \tag{3.7}$$

and

$$R - p - \frac{c_R L_R}{\mu} \geqslant \delta R - \theta - \frac{c_H L_R}{\mu}. \tag{3.8}$$

Rearranging equations (3.7) and (3.8), we get the combined relation

$$p+\frac{c_R L_R}{\mu} \leqslant (1-\delta)R+\theta+\frac{c_H L_R}{\mu} < (1-\delta)R+\theta+\frac{c_H(L_R+1)}{\mu} < p+\frac{c_R(L_R+1)}{\mu}. \quad (3.9)$$

Consider any restaurant price p that satisfies equation (3.9) with

$$p + \frac{c_R L_R}{\mu} < (1-\delta)R + \theta + \frac{c_H L_R}{\mu},$$

so p satisfies the leftmost inequality strictly in equation (3.9). Then there exists $\epsilon > 0$ such that the price $p + \epsilon$ also satisfies (3.9) and results in a strictly higher revenue because the control limits remain the same. Thus the optimal p must make the leftmost inequality tight in equation (3.9), implying that the optimal restaurant price for a given control limit pair (L_R, L_H) and platform price θ is given by

$$p = (1-\delta)R + \theta - \frac{(c_R - c_H)L_R}{\mu}. \quad (3.10)$$

Now, consider any θ that satisfies the rightmost inequality in equation (3.6) strictly. Because we know that p must satisfy equation (3.10), an increase of ϵ in the platform price results in a corresponding increase of ϵ in the restaurant price. If θ satisfies the rightmost inequality in equation (3.6) strictly, then there exists $\epsilon > 0$ such that $\theta + \epsilon$ still satisfies equation (3.6), and $p + \epsilon$ still satisfies equation (3.9) because in effect we have added ϵ to every piece of the inequality. Therefore, the prices $(p + \epsilon, \theta + \epsilon)$ result in the same control limits as (p, θ), and they yield strictly greater revenue. We conclude that the optimal platform price θ must make tight the rightmost inequality in equation (3.6), and we then have an expression for the optimal θ given by

$$\theta = \delta R - \frac{c_H(L_R + L_H)}{\mu}, \quad (3.11)$$

which matches equation (3.5). Substituting equation (3.11) into equation (3.10) yields the expression for the optimal restaurant price in equation (3.4).

Now we address the case where one of the channels is inactive. If only dine-in is active (i.e., if $L_H = 0$ and $L_R \geqslant 1$), then the optimal restaurant price makes a customer indifferent between purchasing and not if she joins at position L_R, so

$$p = R - \frac{c_R L_R}{\mu} \tag{3.12}$$

We recover this expression by setting $L_H = 0$ in equation (3.4). Similarly, if only the platform is active (so $L_R = 0$ and $L_H \geqslant 1$), then the optimal platform price makes a customer indifferent between purchasing and not if she joins at position L_H. Therefore, in this case we have

$$\theta = \delta R - \frac{c_H L_H}{\mu}. \tag{3.13}$$

We recover this expression by setting $L_R = 0$ in equation (3.5). Also, if $L_R = 0$ then the restaurant price p does not affect the revenue, and similarly for θ when $L_H = 0$. Thus, we can use the expressions for the prices in equations (3.4)-(3.5) even at the boundaries when one channel is inactive.

□

Lemma 5 allows us to rewrite the revenue as a function of only the control limits, i.e.,

$$Z(L_R, L_H) = \Lambda\Big[(1 - \pi_{L_R+L_H})\Big(\delta R - \frac{c_H(L_R + L_H)}{\mu}\Big) + \tag{3.14}$$

$$\Pi_{0,L_R-1}\Big((1-\delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big)\Big]. \tag{3.15}$$

Using equation (3.15), we arrive at our first result regarding the optimal pricing policy for the monopolist.

Proposition 13 (All-In for Platform When $\boldsymbol{\delta = 1}$.)**.** *If $\delta = 1$, then an optimal pricing policy must set the control limit $L_R = 0$, so that all purchases come via the platform.*

Proof. Consider any pricing policy resulting in control limits L_H and $L_R = k \geqslant 1$. We have

$$Z(k, L_H) = \Lambda\Big[(1 - \pi_{k+L_H})\Big(R - \frac{c_H(k + L_H)}{\mu}\Big) - \Pi_{0,k-1}\Big(\frac{(c_R - c_H)k}{\mu}\Big)\Big]$$

$$< \Lambda\Big[(1 - \pi_{k+L_H})\Big(R - \frac{c_H(k + L_H)}{\mu}\Big)\Big] = Z(0, L_H + k).$$

Thus, any policy resulting in a control limit $L_R = k \geqslant 1$ can be improved by setting $L_R = 0$ and increasing L_H to $L_H + k$. □

Thus, when $\delta = 1$ the problem reduces to setting a single control limit in a system with reward R and waiting cost c_H. This problem was studied by Naor (1969). The managerial implication of this result is simple: if a firm's product retains its quality well in the process of delivery, then offering delivery can be beneficial to the firm, and may render in-person service superfluous. A related phenomenon is observed in practice in the, where many brick-and-mortar retailers struggle to compete with online retailers like Amazon.com that deliver products directly to customer's homes.

If $\delta < 1$, then we may optimally set both $L_R > 0$ and $L_H > 0$. We first recast the revenue function. Defining $\boldsymbol{L} := L_R + L_H$, we express the revenue function $Z(L_R, \boldsymbol{L})$ by

$$Z(L_R, \boldsymbol{L}) = \Lambda\Big[(1 - \pi_{\boldsymbol{L}})\Big(\delta R - \frac{c_H \boldsymbol{L}}{\mu}\Big) + \Pi_{0,L_R-1}\Big((1 - \delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big)\Big], \quad (3.16)$$

where the domain of Z is restricted such that $L_R \leqslant \boldsymbol{L}$. In the special case with $\rho = 1$, equation (3.16) becomes

$$Z(L_R, \boldsymbol{L}) = \Lambda\Big[\frac{\boldsymbol{L}}{\boldsymbol{L}+1}\Big(\delta R - \frac{c_H \boldsymbol{L}}{\mu}\Big) + \frac{L_R}{\boldsymbol{L}+1}\Big((1 - \delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big)\Big]. \quad (3.17)$$

For $\rho \neq 1$, we compute the aggregate revenue function as follows. For an $M/M/1$ queueing system with finite capacity $\boldsymbol{L}$, the probability π_i is given by

$$\pi_i = \rho^i\Big(\frac{1 - \rho}{1 - \rho^{\boldsymbol{L}+1}}\Big), \quad 0 \leqslant i \leqslant \boldsymbol{L} \text{ and } \rho \neq 1. \quad (3.18)$$

Substituting equation (3.18) into (3.3), the probability $\Pi_{0,j}$ can now be written as

$$\Pi_{0,j} = \frac{1-\rho^{j+1}}{1-\rho^{\boldsymbol{L}+1}}, \quad 1 \leqslant j \leqslant \boldsymbol{L} \text{ and } \rho \neq 1,$$

and the aggregate revenue function from equation (3.16) is given by

$$Z(L_R, \boldsymbol{L}) = \Lambda\Big[\frac{1-\rho^{\boldsymbol{L}}}{1-\rho^{\boldsymbol{L}+1}}\Big(\delta R - \frac{c_H \boldsymbol{L}}{\mu}\Big) + \frac{1-\rho^{L_R}}{1-\rho^{\boldsymbol{L}+1}}\Big((1-\delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big)\Big]. \quad (3.19)$$

We now have the following result regarding the optimal control limit for dine-in customers, L_R^*.

Proposition 14 (Optimal Dine-In Control Limit)**.** *Let $\lfloor\cdot\rceil$ be the nearest-integer function. If $\rho = 1$, then the optimal dine-in control limit $L_R^*(\boldsymbol{L})$ is given by*

$$L_R^*(\boldsymbol{L}) = \begin{cases} \left\lfloor \dfrac{(1-\delta)R\mu}{2(c_R - c_H)} \right\rceil & \text{if } \boldsymbol{L} \geqslant \left\lfloor \dfrac{(1-\delta)R\mu}{2(c_R - c_H)} \right\rceil, \\ \boldsymbol{L} & \text{otherwise.} \end{cases} \quad (3.20)$$

Then, define

$$\widetilde{L_R} := \begin{cases} \min\Big\{ n \in \mathbb{N} \;\Big|\; \rho^n\Big[(1-\rho)(1-\delta)R + \dfrac{\rho(c_R - c_H)}{\mu} - \dfrac{n(1-\rho)(c_R - c_H)}{\mu}\Big] \leqslant \dfrac{c_R - c_H}{\mu} \Big\} & \text{if } \rho < 1, \\ \min\Big\{ n \in \mathbb{N} \;\Big|\; \rho^n\Big[(\rho-1)(1-\delta)R - \dfrac{\rho(c_R - c_H)}{\mu} - \dfrac{n(\rho-1)(c_R - c_H)}{\mu}\Big] \leqslant -\dfrac{c_R - c_H}{\mu} \Big\} & \text{if } \rho > 1. \end{cases} \quad (3.21)$$

If $\rho \neq 1$, then the optimal dine-in control limit $L_R^*(\boldsymbol{L})$ is defined as

$$L_R^*(\boldsymbol{L}) = \begin{cases} \widetilde{L_R} & \text{if } \boldsymbol{L} \geqslant \widetilde{L_R}, \\ \boldsymbol{L} & \text{otherwise.} \end{cases} \quad (3.22)$$

Proof. **Case 1: $\boldsymbol{\rho = 1}$.** First, we observe that the revenue as a function of L_R depends only on the second term inside the square brackets of equation (3.20). In fact, the optimal L_R at a given $\boldsymbol{L}$ can be found by optimizing the function $g(L_R)$, where

$$g(L_R) := L_R\Big((1-\delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big),$$

a negative quadratic function of L_R. Let us temporarily expand the domain of g to include all of $\mathbb{R}_+$. Differentiating, we have

$$g'(L_R) = (1-\delta)R - \frac{2L_R(c_R - c_H)}{\mu}.$$

As $g(L_R)$ is a negative quadratic function, the first-order condition (FOC) is sufficient for a global maximum, and solving the FOC for L_R gives

$$\overline{L_R} = \frac{(1-\delta)R\mu}{2(c_R - c_H)}. \tag{3.23}$$

Note, however, that we must have $L_R \leqslant \boldsymbol{L}$, and also L_R must be integer-valued. If $\boldsymbol{L}$ is less than the right-hand side (RHS) of equation (3.23), then the permissible values of L_R all fall into the increasing interval of the quadratic, and it is therefore optimal to set L_R to the maximum permissible value, which is $\boldsymbol{L}$. Otherwise, we round the RHS of equation (3.23) to the nearest integer, and the result is the optimal value for L_R, as in equation (3.20).

Case 2: $\boldsymbol{\rho < 1}$. If $\rho < 1$, then for a given aggregate control limit $\boldsymbol{L}$ the optimal dine-in control limit can be found by solving the single-variable optimization problem

$$\max_{L_R \in \{0,1,\ldots,\boldsymbol{L}\}} h(L_R) := \big(1-\rho^{L_R}\big)\Big((1-\delta)R - \frac{(c_R - c_H)L_R}{\mu}\Big).$$

Consider the continuous relaxation of this problem, i.e.,

$$\max_{0 \leqslant x \leqslant \boldsymbol{L}} h(x) := \big(1-\rho^{x}\big)\Big((1-\delta)R - \frac{(c_R - c_H)x}{\mu}\Big).$$

Letting $\log x$ denote the natural logarithm of x and differentiating, we have

$$h'(x) = -\big(1-\rho^x\big)\big(\frac{c_R - c_H}{\mu}\big) - \rho^x\Big((1-\delta)R - \frac{x(c_R - c_H)}{\mu}\Big)\log\rho. \tag{3.24}$$

Evaluated at zero, the function $h(\cdot)$ and its derivative $h'(\cdot)$ give respectively

$$h(0) = 0$$

and

$$h'(0) = -(1-\delta)R\log\rho > 0,$$

implying that there must exist $\epsilon > 0$ such that $h(\epsilon) > 0$. Further, it is readily verified that

$$h(x) < 0 \quad \forall x > \frac{\mu(1-\delta)R}{c_R - c_H},$$

which confirms that there exists x' with $\epsilon < x' \leqslant \mu(1-\delta)R/(c_R - c_H)$ and $g'(x') < 0$.

Setting the RHS of equation (3.24) to zero and rearranging, we arrive at the FOC given by

$$\frac{c_R - c_H}{\mu\rho^x} = \frac{c_R - c_H}{\mu} - (1-\delta)R\log\rho + x\Big(\frac{c_R - c_H}{\mu}\Big)\log\rho. \tag{3.25}$$

The fact that $g'(0) > 0$ and $g'(x') < 0$ implies that there exists at least one solution to the FOC (3.25), and moreover that this solution is a local maximizer of the function $h(\cdot)$. The LHS of the FOC (3.25) is increasing in x because the denominator is decreasing in x (since $\rho < 1$). Again because $\rho < 1$, the RHS is decreasing in x. As the LHS is increasing and the RHS is decreasing, there is at most one solution to the FOC (3.25). We thus conclude that (i) the function $h(\cdot)$ is first increasing and then decreasing, (ii) that it has a unique maximizer, and (iii) that this maximizer is the sole solution to the FOC (3.25).

Returning to the discrete setting, since $h(\cdot)$ is first increasing and then decreasing, the optimal dine-in control limit L_R^* will depend on the aggregate control limit $\boldsymbol{L}$.

Define $\Delta h(L_R)$ to be the forward difference of $h(\cdot)$, evaluated at L_R, i.e.,

$$\Delta h(L_R) = h(L_R + 1) - h(L_R).$$

Because $h(\cdot)$ is first increasing and then decreasing, the unconstrained optimal dine-in control limit will be equal to $\widetilde{L_R}$, where

$$\widetilde{L_R} = \min\{n \in \mathbb{N} | \Delta h(n) \leqslant 0\}. \tag{3.26}$$

We arrive at equation (3.21) by calculating $\Delta h(\cdot)$, substituting into equation (3.26), and rearranging.

We are, however, constrained by the aggregate control limit $\boldsymbol{L}$. Given a value of $\boldsymbol{L}$, if $\Delta h(\boldsymbol{L}) \leqslant 0$, then the optimal dine-in control limit is $\widetilde{L_R}$. Otherwise, we are restricted to the increasing portion of the function $h(\cdot)$. In that case, the optimal choice is the maximum allowable value, i.e., we have $L_R^* = \boldsymbol{L}$.

Case 3: $\boldsymbol{\rho > 1}$. The case with $\rho > 1$ is completely analogous to $\rho < 1$, and is therefore omitted. □

We next turn to finding the optimal overall control limit $\boldsymbol{L}$.

Proposition 15 (Optimal Aggregate Control Limit)**.** *If $\rho = 1$, then there is at least one local maximum of the aggregate revenue function in $\boldsymbol{L}$ if $L_R = L_R^*(\boldsymbol{L})$. Moreover, for any value of ρ the global maximum can be found via a one-dimensional search over a bounded set of nonnegative integers.*

Proof. Define the function $\chi(\boldsymbol{L})$ by

$$\chi(\boldsymbol{L}) := Z\big(L_R^*(\boldsymbol{L}), \boldsymbol{L}\big),$$

i.e., the aggregate revenue as a function of the overall control limit $\boldsymbol{L}$, evaluated at the optimal dine-in control limit $L_R^*(\boldsymbol{L})$. For the moment we allow $\boldsymbol{L}$ to take any

nonnegative real value. Note that $L_R^*(\boldsymbol{L})$ is easily extended to this domain also. Differentiating with respect to $\boldsymbol{L}$, we get

$$\chi'(\boldsymbol{L}) = \frac{-\Lambda}{(\boldsymbol{L}+1)^2}\Big[\delta R\boldsymbol{L} - \frac{c_H\boldsymbol{L}^2}{\mu} + L_R^*(\boldsymbol{L})\Big((1-\delta)R - \frac{(c_R-c_H)L_R^*(\boldsymbol{L})}{\mu}\Big)\Big] + \frac{\Lambda}{\boldsymbol{L}+1}\Big[\delta R - \frac{2c_H\boldsymbol{L}}{\mu} + \frac{\mathsf{d}L_R^*(\boldsymbol{L})}{\mathsf{d}\boldsymbol{L}}\Big((1-\delta)R - \frac{2(c_R-c_H)L_R^*(\boldsymbol{L})}{\mu}\Big)\Big].$$

The derivative $\mathsf{d}L_R^*(\boldsymbol{L})/\mathsf{d}\boldsymbol{L}$ is continuous everywhere except at $\boldsymbol{L} = \overline{L_R}$, where it is undefined. More precisely, we have

$$\frac{\mathsf{d}L_R^*(\boldsymbol{L})}{\mathsf{d}\boldsymbol{L}} = \begin{cases} 1 & \text{if } \boldsymbol{L} < \overline{L_R} \\ 0 & \text{if } \boldsymbol{L} > \overline{L_R}. \end{cases} \tag{3.27}$$

The derivative $\chi'(\boldsymbol{L})$ is thus also well-defined everywhere except at $\boldsymbol{L} = \overline{L_R}$. Substituting from equation (3.27), we have

$$\chi'(0) = \Lambda R > 0$$

and

$$\lim_{\boldsymbol{L}\to\infty} \chi'(\boldsymbol{L}) = -\frac{c_H\Lambda}{\mu} < 0.$$

Although the derivative of $\chi(\boldsymbol{L})$ has a discontinuity, the function itself is continuous on $\mathbb{R}_+$. Because we assume that $R\mu > c_R$, it is possible to achieve a strictly positive revenue with a control limit of 1. Finally, there is a finite value $\bar{\boldsymbol{L}}$ which is the maximum control limit achievable with a nonnegative price. Any higher control limits will require negative prices, which will result in strictly negative revenue. Combining the facts that $\chi(0) = 0$, $\chi(1) > 0$, and $\chi(\boldsymbol{L}) < 0$ for $\boldsymbol{L} > \bar{\boldsymbol{L}}$, we conclude that $\chi(\cdot)$ is first increasing and then decreasing. Thus, there is at least one local maximum, both of the continuous extension and of the version restricted to the integers. The global maximum (for any value of ρ) can be found from at most $\bar{\boldsymbol{L}} + 1$ evaluations of $\chi(\cdot)$.

□

From these results, we next demonstrate an important corollary regarding the ordering of the optimal prices.

Corollary 7 (Ordering of Optimal Prices)**.** *Let p^* and θ^* be the optimal monopoly prices for dine-in and delivery, respectively. Then, we must have*

$$p^* \geqslant \theta^*.$$

That is, the optimal dine-in price is higher than the optimal platform price when both are set by a monopolist.

Proof. Proof by Contradiction. Lemma 5 implies that, for a given pair of control limits $(L_R, \boldsymbol{L})$, we have

$$p - \theta = (1-\delta)R - \frac{(c_R - c_H)L_R}{\mu}. \tag{3.28}$$

Suppose by way of contradiction that the optimal prices satisfy

$$p^* < \theta^* \iff (1-\delta)R - \frac{(c_R - c_H)L_R^*}{\mu} < 0, \tag{3.29}$$

where L_R^* is the optimal dine-in control limit. We observe from equations (3.17) and (3.19) that choosing the optimal dine-in control limit L_R^* for a given aggregate control limit $\boldsymbol{L}$ entails maximizing the function $h(L_R)$, where

$$h(L_R) := \begin{cases} (1-\rho^{L_R})\Big((1-\delta)R - \dfrac{(c_R - c_H)L_R}{\mu}\Big) & \text{if } \rho < 1 \\ L_R\Big((1-\delta)R - \dfrac{(c_R - c_H)L_R}{\mu}\Big) & \text{if } \rho = 1 \\ (\rho^{L_R}-1)\Big((1-\delta)R - \dfrac{(c_R - c_H)L_R}{\mu}\Big) & \text{if } \rho > 1. \end{cases}$$

Equation (3.29) implies that at the optimal control limit L_R^* we have

$$\left((1-\delta)R - \frac{(c_R - c_H)L_R^*}{\mu}\right) < 0, \tag{3.30}$$

which ensures for each interval of ρ that $h(L_R^*) < 0$. But an immediate lower bound of zero on the optimal value of $h(L_R)$ is achieved in all three cases by setting $L_R = 0$, implying that the chosen control limit L_R^*—and by extension the prices p^* and θ^*—must be suboptimal, contradicting the initial assumption. □

This result is surprising given the prevailing business model of charging customers extra for delivery services. It also suggests that the decentralized system may operate with a substantial lack of coordination, as the platform must add a surcharge to the base price of the food to avoid losing money.

At first glance, it seems counterintuitive that it would be better to charge a customer less for delivery because delivery is considered an "add-on" service on top of preparing the food. However, in this case, the result is all about customer surplus and willingness to pay. The expression on the right-hand side of equation (3.28) is the "potential" surplus in net reward (before subtracting price) that the threshold customer—that is, one who takes position L_R in the queue—receives from dining in rather than ordering delivery.[1] Optimally, the restaurant should extract all of this surplus from the threshold customer to make him exactly indifferent between the two purchase channels (while still preferring both of them to not purchasing at all). Because $L_R \leqslant \boldsymbol{L}$, the customer in position L_R has more potential surplus than the one in position $\boldsymbol{L}$ for either channel. The platform price θ is set based on the latter customer, who joins at position $\boldsymbol{L}$ and purchases from the platform expecting a net reward of exactly zero. In order to induce this customer to join, who faces a much

[1] Note that the word "surplus" here is used as the gain in expected utility relative to the alternative of ordering delivery, not in the usual sense of its distance from zero.

longer queue, the platform must set a lower price than the restaurant, whose price is based on the customer in position L_R.

The above comparison is not perfect because the parameters for dine-in and delivery are different, but the distinctions will further clarify the result. The intrinsic differences between the platform and the restaurant are (i) the reduced gross reward when ordering delivery and (ii) the reduced waiting cost coefficient when waiting for delivery. It is possible to imagine a scenario in which the waiting costs were far apart ($c_R \approx c_H$), while the gross rewards were quite close ($\delta \approx 1$), suggesting that even though the customer in position $\boldsymbol{L}$ faces a longer queue, she could still have more potential surplus than the one in position L_R because by purchasing from the platform the waiting cost advantage could outweigh the reward reduction. But, looking back at equation (3.28), we see that this tradeoff between waiting cost and gross reward is just what the dine-in control limit is set to optimize over, and it is chosen so that at position L_R the potential surplus (which will exactly equal the restaurant price) will be larger than θ.

Following an intuitive explanation with a more technical argument, let us for the moment consider a fixed value of the aggregate control limit $\boldsymbol{L}$. We can then think of $\theta = \delta R - c_H \boldsymbol{L}/\mu$ as the "baseline" revenue per customer. Consulting the aggregate revenue function of equation (3.16), the optimal choice when setting the dine-in control limit L_R (given a value of $\boldsymbol{L}$) is to maximize the second term inside the brackets, which represents the incremental revenue that we receive for each customer who chooses to dine in. Clearly, this involves making the rightmost parenthetical term positive. It happens that this term is exactly equal to $p - \theta$, and therefore it is optimal to set $p \geqslant \theta$.

In the next section we investigate the ramifications of this result further by examining the equilibrium outcome when the platform and restaurant set their prices independently and without a contract to coordinate them.

3.6 Decentralized System

Having solved the monopoly pricing problem, we turn to a decentralized setting, where the restaurant and platform set their prices independently and in their private interest. We establish the existence of a pure-strategy subgame perfect equilibrium in Section 3.6.1, and we propose a contract to coordinate the system in Section 3.6.2.

3.6.1 Existence of Equilibrium

Throughout this subsection, we assume that there are no detailed contract terms and that the platform must pay the full retail price to the restaurant for each unit, like any other customer. We assume the following sequence of events, which together comprise the *service delivery pricing game*. First, the restaurant sets a price p which will be paid by both dine-in customers and the platform. Then, the platform sets a customer-facing price θ, inclusive of all delivery fees and surcharges; the platform's profit (or loss) on each unit is thus equal to $\theta - p$. Once chosen, these prices are made public, whereupon the customer arrival and channel-selection process ensues as discussed in Section 3.4. The restaurant and platform make their pricing decisions in order to maximize their long-run average revenue per unit time, which we define below. We note here that as defined, the pricing game played by the platform and restaurant amounts to a Stackelberg game, where the restaurant is the leader and the platform is the follower. Also, we make the following assumptions for the case of indifference between two prices.

Assumption 3. *If multiple values of θ will result in the same revenue for the platform, then the platform chooses the lowest such price.*

We observe here that, because the platform moves second, the set of possibly optimal prices for the platform in the decentralized setting is the same as the set of possibly optimal platform prices in the monopoly case. Importantly, this set is finite,

so a minimum price consistent with a given revenue always exists. This assumption can perhaps be justified by noting that customers certainly prefer lower prices, and if the same net revenue can be achieved with multiple prices, then a firm may choose the lowest such price to placate customers.

In this decentralized setting, the restaurant accrues revenue p for each unit served, whether to a dine-in customer or to a delivery customer via the platform. We can thus define the restaurant's revenue function $z^R(\cdot)$ by

$$z^R(p, \theta) := \Lambda p(1 - \pi_{\boldsymbol{L}}),$$

where customer control limits L_R and L_H are as defined in equations (3.1) and (3.2), respectively, and once again $\boldsymbol{L} = L_R + L_H$. We can similarly express the platform's revenue function $z^H(\cdot)$ by

$$z^H(p, \theta) := \Lambda(\theta - p)\Pi_{L_R, \boldsymbol{L}}.$$

Proposition 16 (Existence of Pure-Strategy Subgame Perfect Equilibrium)**.** *Suppose that the price spaces for the restaurant and platform are discrete. Then the service delivery pricing game has at least one subgame perfect equilibrium in pure strategies.*

Proof. We first show that the set of possibly optimal prices for the restaurant and platform can be reduced to a finite set. A clear upper bound on the optimal price for the restaurant is that price at which a customer who arrives and observes an empty system would be indifferent between dining in or not purchasing. Therefore, we have

$$R - p - \frac{c_R}{\mu} \geqslant 0$$

$$p \leqslant R - \frac{c_R}{\mu}.$$

We can similarly obtain an upper bound on the optimal platform price by

$$\delta R - \theta - \frac{c_H}{\mu} \geqslant 0$$

$$\theta \leqslant \delta R - \frac{c_H}{\mu}.$$

As both p and θ must be nonnegative, we have compact intervals in which the optimal prices for both parties must lie, regardless of the choice made by the other party. Assuming a discrete space of prices, we conclude that the game has finitely many feasible histories, and it therefore satisfies the definition found in Osborne and Rubinstein (1994) of a finite extensive game with perfect information. Hence, by their Proposition 99.2—attributed to Kuhn (1950, 1953)—the service delivery pricing game has at least one subgame perfect equilibrium in pure strategies.

□

3.6.2 Coordinating Contracts

We now investigate possible contract structures to identify terms which will induce both the restaurant and the platform to choose system-optimal prices. We give the name "two-way revenue sharing" to the general contract form that we study. Specifically, let α be the fraction of dine-in revenue that the restaurant keeps for itself (so the platform receives a fraction $1-\alpha$ of the dine-in revenue). Similarly, let β be the fraction of platform revenue that the platform cedes to the restaurant (so the platform keeps a fraction $1-\beta$ of the platform revenue). Under two-way revenue sharing, the restaurant's revenue $z^R(p,\theta)$ is given by

$$z^R(p,\theta) = \Lambda\Big[\alpha p \Pi_{0,L_R-1} + \beta\theta\Pi_{L_R,\boldsymbol{L}-1}\Big]. \tag{3.31}$$

The platform's revenue is similarly expressed as

$$z^H(p,\theta) = \Lambda\Big[(1-\alpha) p \Pi_{0,L_R-1} + (1-\beta)\theta\Pi_{L_R,\boldsymbol{L}-1}\Big]. \tag{3.32}$$

This general form encompasses various types of contracts between the restaurant and the platform. The traditional one-way revenue sharing contract of Cachon and Lariviere (2005) is obtained by letting $\alpha = 1$ and $0 < \beta < 1$. Importantly, while demand is indeed stochastic, our model is analogous to a pull system, so the platform does not "place an order" before the selling season but rather makes unit orders as customers arrive. As demand is always met, there is no need to distinguish wholesale payments from revenue sharing payments: both the wholesale price and the revenue share fraction are incorporated into β. Another important contract structure is the lack of a contract, i.e., a no-relationship context modeled in Section 3.6. This structure is recovered by setting $\alpha = 1$ and $\beta = p/\theta$.

We propose another contract, suggested by the "two-way" moniker, which involves payments both ways: the restaurant gets a positive fraction of the platform revenue, and the platform gets a positive fraction of the dine-in revenue.

Proposition 17 (Two-Way Revenue Sharing Coordinates the Supply Chain)**.** *Under a contract which sets $\alpha = \beta$ (i.e., each party's share of the dine-in revenue is the same as his share of the platform revenue), both the restaurant and the platform will choose prices in equilibrium which achieve the system-optimal revenue.*

Proof. From equation (3.31), if $\alpha = \beta$, then we have

$$\begin{aligned} z^R(p,\theta) &= \Lambda\Big[\alpha p \Pi_{0,L_R-1} + \alpha\theta\Pi_{L_R,\boldsymbol{L}-1}\Big] \\ &= \alpha Z(p,\theta). \end{aligned}$$

Thus, the restaurant's revenue is directly proportional to the aggregate revenue, and its optimal price will be the system-optimal price. The same reasoning can be applied to equation (3.32) for the platform, since $1-\alpha = 1-\beta$. □

We believe that our proposal of a "two-way" revenue sharing contract is unique in the OM literature. While it shares similarities with traditional one-way revenue

Table 3.1: Parameter ranges for numerical tests. In all cases $\mu = 10$ and $c_H = 10$.

Parameter	Range	# Values
ρ	[.05,3]	59
δ	[.7,.98]	15
c_R	[11,50]	40

sharing contracts, it may be more difficult to implement in practice. Specifically, restaurants may be hesitant to give up a portion of their dine-in revenue to platforms when the platform has no direct role in these sales. A possible approach to implementing this contract would be to charge a "listing" fee to the restaurant to appear on the platform. This payment structure is operated, e.g., by GrubHub (see GrubHub, 2018), although perhaps with different motivation.

Furthermore, regardless of the feasibility of its implementation, the two-way revenue sharing contract is a straightforward method to coordinate the system. It thus serves as a benchmark to which to compare other types of contracts. As we have established that a coordinating contract exists, we have a tight upper bound with which to evaluate the quality of any contract. In the next section, we will study numerically how the system performs (i) in the absence of any contract, and (ii) under the traditional one-way revenue sharing contract, where the only sharing is from the platform to the restaurant.

3.7 Numerical Tests: Coordination Loss and Mutual Gain

In order to understand the strategic interaction between the restaurant and platform, we compute equilibrium prices, control limits, and revenues over a range of parameter values. Specifically, we run 35,400 distinct parameter combinations, with the parameter ranges detailed in Table 3.1. We highlight several specific examples and summarize the overall results in Table 3.2.

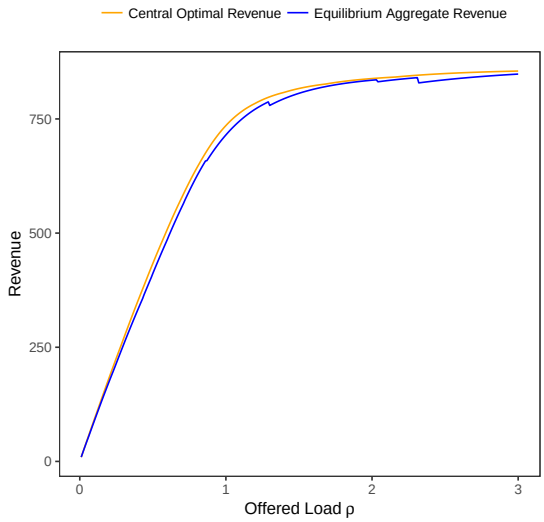

FIGURE 3.1: Coordination loss versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

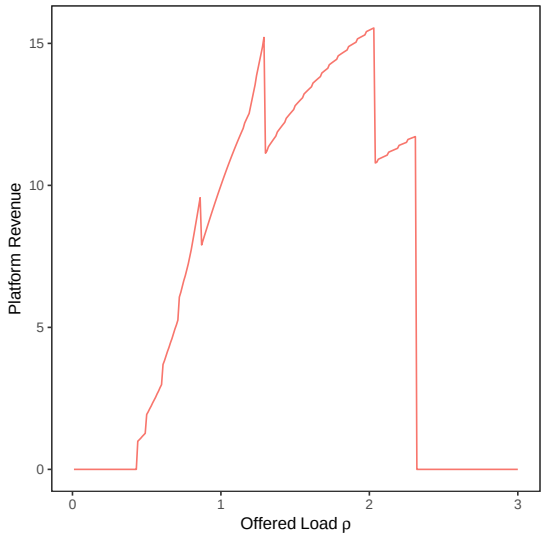

FIGURE 3.2: Equilibrium platform revenue versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

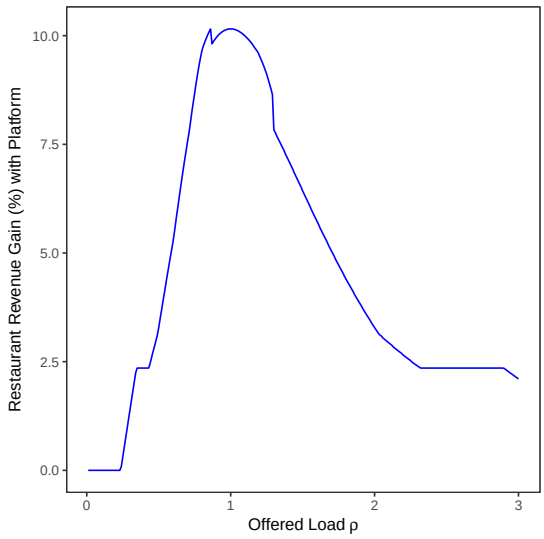

FIGURE 3.3: Restaurant's revenue gain with platform versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

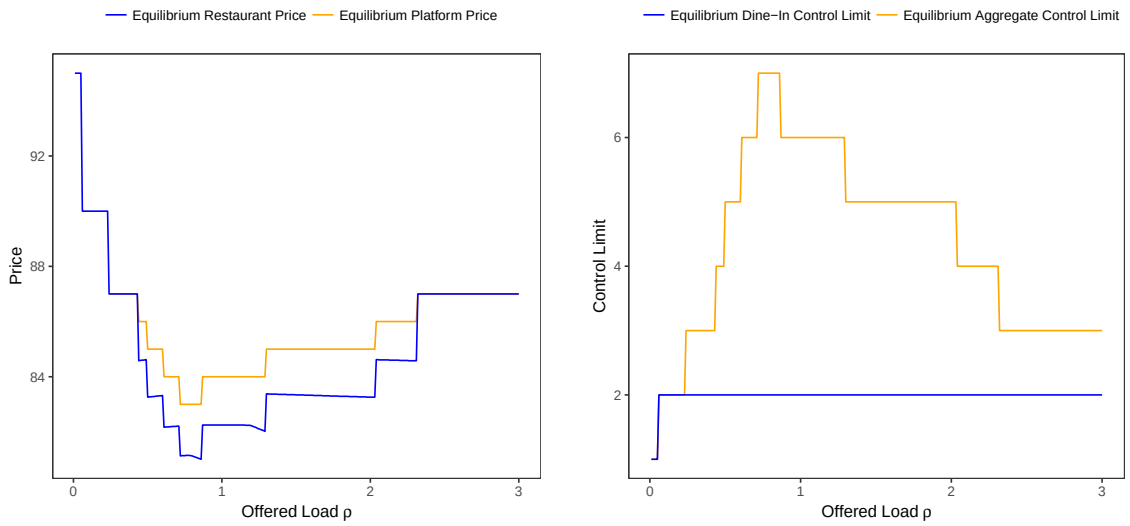

(a) Equilibrium Prices (b) Equilibrium Control Limits

FIGURE 3.4: Equilibrium prices and control limits versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

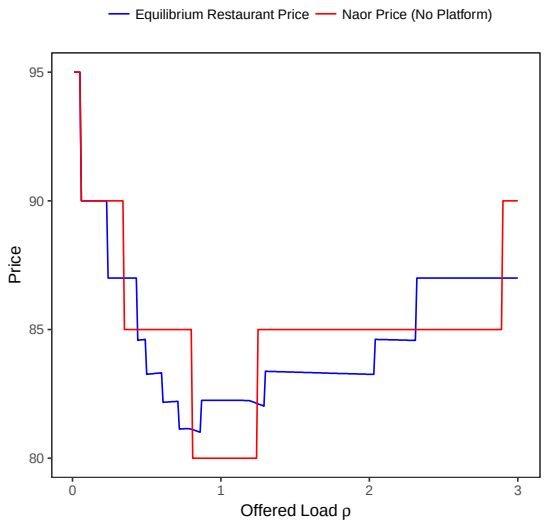

FIGURE 3.5: Restaurant's price with and without platform versus ρ ($\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

3.7.1 Coordination and Individual Revenues.

Consult first Figure 3.1. This figure plots the monopoly optimal revenue and the equilibrium aggregate revenue as a function of the offered load ρ. As implied by Corollary 7, we observe that the aggregate revenue in the subgame perfect equilibrium is indeed less than the coordinated optimal revenue, for all but very small ρ. Observationally, for these small values of ρ usually the platform is inactive, and the equilibrium and Naor-optimal revenues coincide, while the central optimal revenue makes the platform active and enjoys only a small incremental gain in revenue.

As ρ begins to increase and the platform starts earning revenue, the strategic interaction between the platform and the restaurant begins to have more of an effect, and the coordination loss increases. However, although the system operates less efficiently as ρ starts to increase, the efficiency loss does not result in a net negative impact on the two parties. Figures 3.2 and 3.3 plot the platform's revenue and the restaurant's percent gain in revenue over the case with no platform, respectively, both as a function of ρ. Both of these take their highest values for intermediate

values of ρ. While the offered load is not usually under the firms' direct control, this insight can help to identify situations where a platform and a restaurant can mutually benefit from each other's presence.

3.7.2 Prices and Control Limits.

We next consult Figure 3.4. The left panel of Figure 3.4 plots the equilibrium restaurant and platform prices, and the right panel plots the equilibrium control limits. When the offered load is very low, the restaurant essentially prices the platform out of the market. With a gross reward from delivery of $\delta R = 90$ in the figure, when the restaurant sets its price above 90, the platform cannot possibly attract any customers while making positive net revenue. Similarly, for very large values of ρ, the restaurant sets its price so that the platform is just indifferent between participating in the market and not; in fact, the restaurant and platform prices are equal, so the platform makes zero net revenue. The presence of the platform is still important to the restaurant, however. Looking back at Figure 3.3, we see that the restaurant still does better with the platform even when the platform is making no net revenue. The reason is that even though the platform is just breaking even, its presence still increases the aggregate control limit from 2 to 3, and this allows the restaurant to charge a price consistent with a lower control limit, while still getting the throughput associated with a higher control limit.

For intermediate values of ρ we see some more unusual dynamics at play. The prices in the left panel of Figure 3.4 are for the most part close to piecewise constant, but for values of ρ between roughly 0.7 and 2 the restaurant price behaves differently. It is constant for a stretch, then decreases with ρ, then jumps up, and repeats the process. We can understand the cause of this behavior by looking at the control limit chart in the right panel of the figure. If for a given offered load it is optimal for the restaurant to induce a certain pair of control limits, as ρ increases there is an

interval of ρ where it must decrease its price to prevent the platform from reducing the aggregate control limit and reducing the throughput. Eventually the offered load increases enough that it is no longer worth it to the restaurant to keep dropping its price to maintain throughput, and at that point occurs the jump.

More broadly, for values of ρ when the platform is about to change its price (and the control limits), the restaurant first makes small adjustments to its price to maintain the status quo control limits, but it eventually accepts the change. When the platform is about to reduce its price and increase its control limit, the restaurant first increases its price to keep the current control limits. Then, once the increase in throughput from higher control limits becomes more valuable than the incremental revenue from a higher price, the restaurant acquiesces to the change in control limits by reducing its price along with the platform. Conversely, as discussed, when the platform is about to increase its price and reduce its control limit, the restaurant first lowers its price to keep the control limits (and throughput) from being reduced, and then acquiesces by raising its price.

As mentioned, it is in this intermediate region of the offered load that the restaurant experiences the greatest benefit from the platform's presence. Intuitively, we would expect the entrance of the platform into the marketplace to result in reduced prices due to increased competition. However, we observe at times in equilibrium that the restaurant's price in the decentralized setting is in fact *higher* than the optimal Naor price without the platform (see Figure 3.5).

It is tempting to attribute this phenomenon to the fact that when the platform is active, customers ordering delivery have a lower waiting cost coefficient. So, the platform could profit by setting a low price while the restaurant charged a high price to the "premium" customers who arrive to a short queue and thus offer more surplus to extract. This intuition holds in the monopoly case when the platform's price can dip lower than the restaurant's price because there are no transfer payments.

However, in the decentralized setting the platform must either charge at least as much as the restaurant or exit the market.

The effect in the decentralized setting is related but distinct. The platform serves customers who arrive to a longer queue, and thus the restaurant only serves customers who arrive to a short queue. Since the price is set to make the threshold customer indifferent, if the dine-in threshold becomes lower—which often occurs when we bring in the platform—then the restaurant can afford to charge a higher price because its threshold customer has a smaller waiting cost than in the Naor setting, and the platform can still serve longer-waiting customers profitably at an even higher price because of the lower waiting-cost coefficient.

The crucial piece that allows such equilibria to hold is that the restaurant still receives full price from platform customers. So the restaurant can set its price to make the threshold customer indifferent for a short queue, but it still gets to reap full price from delivery customers due to transfer payments from the platform.

In sum, the restaurant can set a higher price in the decentralized equilibrium than in the absence of a platform because the dine-in control limit is lower than the no-platform setting. Normally, a smaller control limit would mean reduced throughput and a price-demand tradeoff. But because the platform customers have a lower waiting cost, the platform control limit is quite high, often higher than the Naor control limit. Because it still gets paid for the platform customers, the restaurant is able both to increase its price and increase its throughput, quite a desirable development!

3.7.3 Overall Outcomes.

Moving our attention to Table 3.2, we observe that all three broad outcomes are possible, i.e., in equilibrium it is possible that (i) $L_R = 0$, so only the platform is active and the restaurant serves no dine-in customers; (ii) $L_R = \boldsymbol{L}$, so only dine-in is active and the platform serves no customers; or (iii) $0 < L_R < \boldsymbol{L}$, so both channels

are active.

Looking closer at Table 3.2, we see that, for the parameter values we consider, the most common outcome is for only dine-in to be active and for the platform not to serve any customers. This is true even though we consider a range of δ concentrated close to 1. Earlier experiments with smaller values of δ show that when the quality is degraded substantially in delivery, it is difficult for the platform to compete with the restaurant because the base reward is so much less than what the restaurant can offer. We correspondingly observe that in the outcomes where the platform is indeed active, the average value of the dine-in waiting cost c_R is much higher than those when it is not (43.34 when only the platform is active and 36.73 when both channels are active, as compared to 28.46 when only dine-in is active). When waiting at the restaurant brings much more disutility than waiting at home, it is easier for the platform to compensate for the decreased base reward due to δ.

Regarding the offered load ρ, it appears to be more likely that only the platform will be active when the offered load is smaller. From a capacity perspective, this finding makes intuitive sense. If the restaurant has excess capacity because the offered load is small, then it is more likely to be beneficial for it to cooperate with the platform because it needs the additional demand that the platform can bring. When ρ is very small, it is sometimes optimal for the restaurant to allow the platform to handle all transactions. It is also important to note that the restaurant still receives full price from the platform on every transaction in this outcome. It is thus quite possible, even if the restaurant serves no dine-in customers in equilibrium, that it can be better off with than without the platform. In fact, the restaurant benefits the most from the platform in the cases where only the platform is active, although this outcome occurs in a very small number of cases. These are cases with particularly high dine-in waiting costs and particularly high delivery quality, so the platform has the most to offer. Even though the aggregate revenue isn't maximized in these cases,

Table 3.2: Output from numerical tests. Besides frequency, values listed are averages for the given outcome.

Active	Freq.	ρ	δ	c_R	p	θ	Rev. Loss	Naor Gain
Platform	0.70%	0.55	0.98	43.34	93.45	94.67	0.83%	7.87%
Dine-In	75.84%	1.54	0.81	28.46	89.82	76.98	0.43%	0.00%
Both	23.46%	1.54	0.94	36.73	87.65	88.99	1.59%	4.16%

the restaurant is still substantially better off than in the complete absence of the platform.

Interestingly, the greatest loss in aggregate revenue due to lack of coordination comes when both channels are active. Intuitively, if both channels are active in equilibrium, then the parameters are such that both parties have nonnegligible market power. This regime corresponds to a δ that is high but not too close to 1, and a dine-in waiting cost that is relatively but not too high. Thus, the strategic interactions and competition are likely to be more intense, bringing with them the opportunity for loss. We see that in this regime the restaurant's average price is lower than in either other outcome (87.65 versus 89.82 and 93.45), and the platform's average price is lower than the only other outcome when it is active (88.99 versus 94.67). So, when both channels are active it appears that the competition between them has the most negative impact on the average revenue, resulting in an average coordination loss of 1.59%.

3.8 Numerical Tests: One-Way Revenue Sharing

We now discuss the traditional one-way revenue sharing contract where the platform shares revenue with the restaurant. First, it is trivially true that a one-way revenue sharing contract can coordinate the system; if the platform shares all of its revenue with the restaurant (and also cedes its pricing decision), then the restaurant operates as a monopolist and achieves the optimal system revenue. It is not clear, however, for what values of the revenue-sharing coefficient will the system achieve coordination,

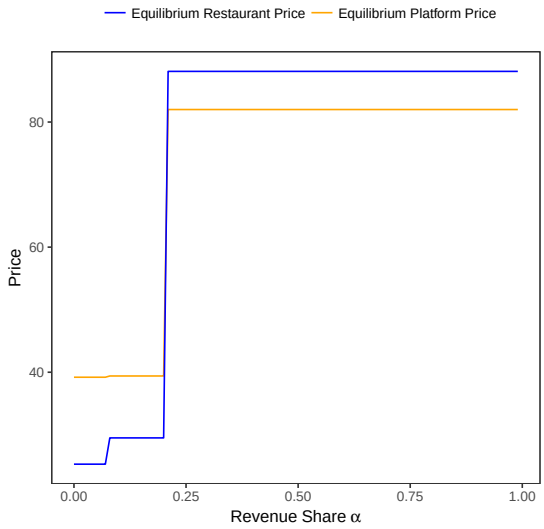

FIGURE 3.6: Equilibrium prices versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

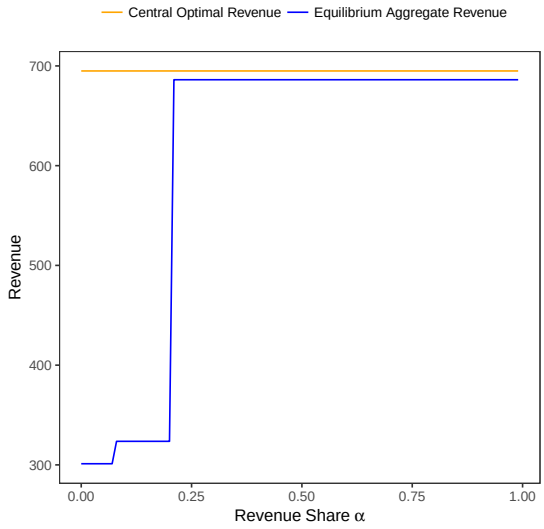

FIGURE 3.7: Coordination loss versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

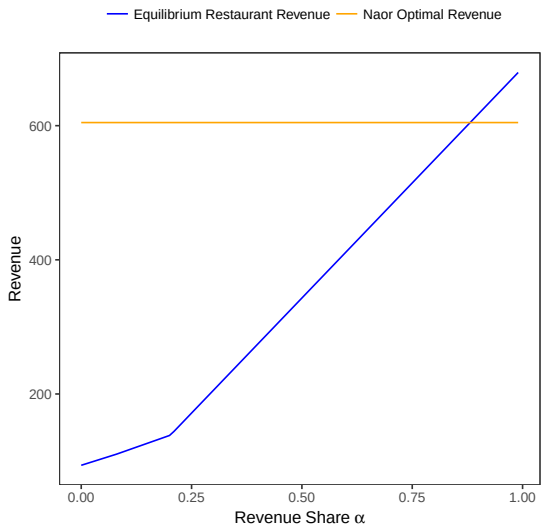

FIGURE 3.8: Restaurant revenue with and without platform versus α ($\rho = .9$, $\mu = 10$, $\delta = .9$, $R = 100$, $c_R = 50$, $c_H = 10$).

or how badly miscoordinated the system will be when coordination loss arises.

We study how the equilibrium prices and revenues change with the amount of revenue shared back to the restaurant. Figure 3.6 plots the equilibrium restaurant and platform prices against the fraction α of the platform's revenue that is given to the restaurant. The parameters for this example are the same as those for Figures 3.1-3.5, with ρ now fixed at 0.9. We see immediately that when the platform keeps most of its revenue (i.e., when α is small), the equilibrium prices are extremely low compared to the prevailing prices in the fully decentralized setting. This phenomenon reflects the fact that when the platform keeps most of its revenue, the incentive for cooperation between the parties is gone. The result is that the restaurant and the platform compete on price, driving the equilibrium prices down and greatly harming aggregate revenue, as seen in Figure 3.7. As the restaurant's share of platform revenue increases, however, a point is reached where the incentives of the restaurant and platform are more closely aligned, and the aggregate revenue approaches (but does not reach) the central optimum.

An important distinction should be drawn between the decentralized and one-way revenue sharing cases. It would be natural, but incorrect, to think of the decentralized setting—where the platform pays full price to the restaurant for each unit—as a special case of one-way revenue sharing with α set equal to p/θ. What this interpretation misses is that in the decentralized setting, α is different for different values of p and θ, while with one-way revenue sharing, α is fixed in advance and is the same regardless of the actual prices chosen. Thus, the incentives with one-way revenue sharing are wholly distinct from those in the decentralized setting, and accordingly we observe different outcomes. Even for values of α that are far from 1, the aggregate revenue under one-way revenue sharing draws closer to the coordinated optimum than the decentralized setting.

However, although the aggregate revenue performs close to the first-best outcome for a wide range of α, we must also be mindful of the revenue allocation. Figure 3.8 depicts the restaurant's equilibrium revenue as a function of α, as well as the revenue that it would achieve in the absence of the platform. In the decentralized setting, we found that the restaurant almost always preferred for the platform to exist, even though the system was not perfectly coordinated. However, with one-way revenue sharing, when the restaurant does not receive full price for each platform customer, it might prefer for the platform not to enter the market. Put another way, in a revenue-sharing contract, the platform may have to heavily compensate the restaurant to successfully entice cooperation. This insight holds whether the alternative is the complete absence of the platform or a "no-contract" setting because for a small revenue share, the restaurant prefers either other alternative.

3.9 Conclusion

Service delivery platforms occupy a special position in the sharing economy because they require a form of cooperation between established providers and disruptive

platforms offering a new service to customers. Incentives are neither perfectly aligned nor directly opposed, and we illuminate the nature and consequences of strategic interaction between an incumbent restaurant and an encroaching platform.

We first address the idealized setting where a hypothetical monopolist sets the prices of both the restaurant and the platform to maximize the aggregate revenue generated by both channels. To solve this problem, we bring tools from the priority queueing and pricing literature into the context of our model. We also note an interesting connection between our setting and that of priority queues: rather than the price paid determining a customer's position in a priority queue, in our model the price determines a customer's waiting cost coefficient. We believe that this phenomenon is new to the literature and could warrant further research in a more general setting than what we study here.

Surprisingly, we find that optimally the monopolist will set the delivery price *lower* than the price for dining in. This result ensures a coordination problem when the two channels set their prices independently, since the platform must charge higher than the dine-in price in order to turn a profit. Moving on to the decentralized setting, we find that an equilibrium exists to the pricing game played by the platform and the restaurant. We then investigate numerically the extent of the coordination loss in equilibrium for different parameter regimes.

We also study different contract structures in an attempt to coordinate the system. We propose a new contract type, which we call two-way revenue sharing, where the restaurant shares dine-in revenue with the platform, and the platform shares delivery revenue with the restaurant. We find that a two-way revenue sharing contract coordinate the system and can allocate the revenue arbitrarily among the two parties. While this contract form would likely be difficult to directly implement in practice, it is a straightforward way to achieve coordination and serves as a useful benchmark for the performance of other contract types.

In the absence of a contract, we find numerically that the restaurant usually benefits from the presence of the platform, even though the aggregate revenue achieved is sub-optimal. The platform tends to earn the most revenue for intermediate values of the offered load ρ, when the system is neither saturated with customers nor finds them so scarce that the restaurant prices the platform out of the market. Because the restaurant tends to benefit from the platform even without coordination, the platform may have to cede a relatively large share of its revenue in order to induce the restaurant to participate in a coordinated contract. The allocation of revenue to the two parties must be constrained to ensure that both parties prefer the contract to operating without one.

Our findings are important to practitioners on both sides of the restaurant-platform divide because they imply that both parties can lose out from a lack of coordination, and also that prescriptions for the optimal pricing scheme may not match conventional wisdom. In addition, there may be other spaces in the sharing economy whose stakeholders could benefit from understanding the symbiotic relationship between established and disruptive providers. We believe that our insights regarding the unintuitive nature of the optimal prices and the form of optimal contracts could potentially carry over into such settings.

This research opens up multiple veins for future work. First, it would be worthwhile to study a model which incorporates competition among restaurants and among different platforms. Such an analysis would be quite challenging technically given the large number of decision makers and different utility functions. Second, it would be of great interest to perform an empirical analysis with data from real platforms and restaurants to assess the extent of coordination loss. Such work could provide a detailed understanding of the circumstances where more and less revenue is left on the table in practice.

4

Lost-Sales Inventory Problems with Short Lifecycles and Forecast Updates

4.1 Introduction

The problem of predicting demand for parts and finished goods, and, as a function of these predictions, choosing a quantity of goods to stock in one or more warehouses, is a familiar one. Inventory planning is required in myriad industries and settings, in particular those in which unforeseen events can cause the demand or supply patterns to change dramatically (see, e.g., Simchi-Levi et al., 2015 and Simchi-Levi et al., 2018).

Recently, the pace of innovation in high tech industries has become so rapid that individual products often have extremely short shelf lives. A new iPhone is released every year, for example, and many other high tech goods undergo similarly short life cycles. In such a world, some of the prevailing assumptions in research on inventory management have become less applicable to practice. Some inventory problems are further complicated by a multi-echelon structure in which components must be

Research in this chapter was supervised by Alessandro Arlotto, Saša Pekeč, and Yehua Wei.

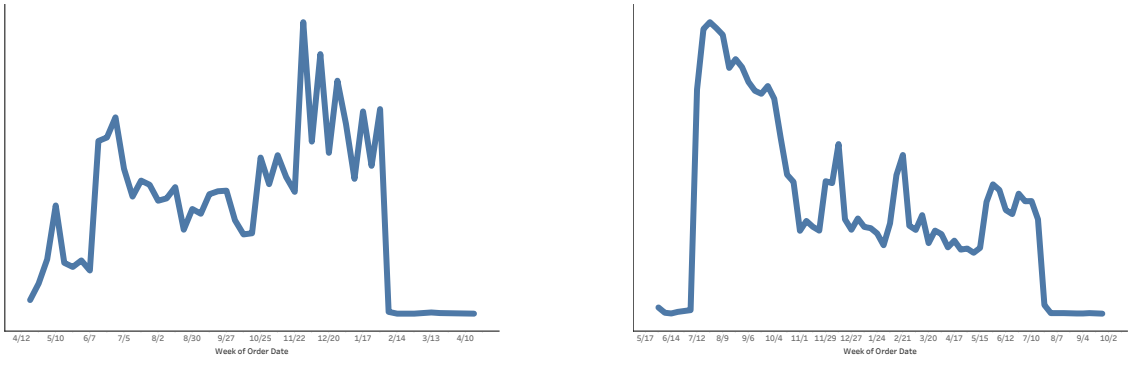

(a) Item A (b) Item B

FIGURE 4.1: Realized demand process for two sample items.

kept as work-in-process to support the production of finished-goods inventory, as in DeValve et al. (2018). For instance, it is common to assume that the demand distribution is (i) known and (ii) stationary over time (see, e.g., Graves, 1999 for a discussion of the need to effectively model nonstationary demand processes with unknown parameters). For a product with a life cycle of little more than a year, it is perhaps unlikely that managers will be able in that time to pin down the true demand distribution, and just as unlikely that the distribution can be considered stationary over the product's lifespan.

This chapter studies the inventory management problem faced by a seller of high-tech goods—we will refer to this company as Firm X throughout—whose business is growing rapidly. Specifically, the company had been using a deterministic optimization model to make stocking decisions, but unpredictability of demand necessitated the incorporation of demand uncertainty into the planning process. The demand processes for two different items are plotted in Figure 4.1, with exact quantities disguised for confidentiality. It is clear from visual inspection that the demand process is nonstationary for each of these items, and moreover that the evolution of the demand process differs greatly between the items. For instance, Item A sees the

demand grow slowly over time, not reaching its peak until the end of its lifetime, and then plummeting rapidly. By contrast, Item B sees a large initial spike in demand, which then tapers off slowly until discontinuation. This phenomenon appears to hold broadly: the demand process varies greatly over time for a given item, and the life cycle curves differ in shape across items. See Hu et al. (2018) for a detailed study of product life cycle curves and how to forecast similar items. A study of distributional properties of the cost function in an inventory problem can be found in Arlotto and Steele (2016a), and more general results on optimal sequential decisions can be found in the literature on distributional properties of Markov decision problems, e.g., Arlotto et al. (2014) and Arlotto and Steele (2016b). We will take a different approach here, not directly addressing life cycle curves but incorporating existing forecasts provided by our industry partner to account for the latest demand conditions.

Because of the rapid growth of the business, an effective solution must be scalable as both the number of items and customer demand grow. Thus, speed of computation and adaptability are paramount. In this chapter, we will therefore focus on heuristics that can be applied in a variety of environments without changing the underlying methodology, with the goal of producing a method that is easy to compute and has "good" performance.

In keeping with this philosophy, we develop a novel heuristic based on the power approximation developed in Ehrhardt (1979, 1984). The power approximation computes a near-optimal (s, S) policy for an infinite-horizon backorder inventory problem, where a policy is determined based on only the first two moments of the demand distribution. This approximation was originally developed to alleviate the computational effort required to compute the exact optimal policy for a given demand distribution. With advances in computational power, the barrier to computing the optimal policy for a given distribution no longer exists. However, as mentioned, in our context it is difficult to know the exact form of the demand distribution. It is

much more realistic to expect that a manager or demand planner could give a good estimate of the mean and variance of the demand over a given horizon, and thus an approximation using only these two moments is quite attractive.

We propose a new form of the power approximation, modified in several ways for our setting. First, while Ehrhardt's power approximation is for a stationary demand setting, Firm X operates in a nonstationary environment. We thus devise a heuristic that takes forecasts for the demand over a given horizon as inputs to determine the mean and variance of demand to feed into the power approximation. Second, Ehrhardt develops his power approximation based on a sample of demand distributions and parameters which, while it performed well for his test cases, is not as applicable to our setting. We fit a new power approximation using real demand data from Firm X for a "training set" of items, and apply this to a test set. Third, while the power approximation is designed for the backorder case, we operate in a lost-sales environment. Leveraging results from Zipkin (2008), we prove that in a lost-sales setting, for a given net inventory position the optimal order quantity is *increasing* in the proportion of the inventory that is on hand (i.e., that has already been delivered). The intuition for this result is that when more inventory is concentrated in the warehouse than in the pipeline, fewer lost sales are likely to occur. When lost sales occur, those units of demand do not deplete the stock as they would in the backorder case because they will never be filled. Thus, the higher the fraction of inventory that is on hand, the fewer lost sales are likely to occur, and thus the *lower* the inventory position is likely to be when the current order arrives. Because the inventory position is likely to be lower when the on-hand fraction is higher, all else being equal the optimal order quantity is increasing in the on-hand fraction. We thus propose a correction factor that adjusts the values of the (s, S) policy returned by the power approximation, based on the distribution of inventory in the order pipeline. In addition to these contributions, we have also provided a

GUI-based tool to Firm X which implements our heuristic.

The outline of the remainder of the chapter is as follows. We devote Section 4.2 to a detailed description of the inventory problem that we study and the data that we received from Firm X. We review the relevant literature on lost-sales inventory problems in Section 4.3. In Section 4.4, we introduce the power approximation of Ehrhardt (1979, 1984); describe the procedure for fitting a new power approximation using Firm X's demand data; and present our dynamic heuristic for combining the power approximation with demand forecasts for regular updates of the inventory policy. In Section 4.5 we present our analytical result regarding the impact of the on-hand fraction on the optimal order quantity, along with an adjustment to our heuristic to account for this feature of the optimal policy. We describe the GUI-based tool and evaluate the performance of our heuristic on real demand data from Firm X in Section 4.6, and we offer concluding remarks in Section 4.7.

4.2 The Problem and the Data

We now detail the inventory problem that we study and describe the data provided to us by Firm X. The first important aspect of our problem is the distinct nature of the demand process: namely, it has a very short life cycle and is highly volatile (i.e., nonstationary). As discussed in Section 4.1, the life cycle of products for Firm X is on the order of 1-2 years (usually much less than 2 years). It would thus be quite difficult to deduce the form of the demand distribution even if demand were stationary over time. Because the demand is decidedly *not* stationary (see Figure 4.1), predicting demand becomes even more difficult, and Firm X is forced to rely heavily on demand forecasts. It is therefore important to incorporate both the demand forecasts and knowledge about historical demand into the planning process.

The second key feature is that we operate in a lost-sales regime. Lost-sales inventory problems are notoriously difficult to solve to optimality. They are unfortunately

devoid of the structure of the backorder setting that admits optimality of base stock and (s, S) policies under a variety of assumptions—see Bijvank and Vis (2011) for a review of the status of research on lost-sales inventory problems, and Porteus (2002, Section 7.1) for a discussion of the optimality of (s, S) policies in backorder problems.

Specifically, we study a finite-horizon, periodic review (i.e, discrete time) inventory problem with:

- Linear holding (h) and shortage (p) costs;
- Positive fixed order cost K;
- Variable order cost c for each unit ordered;
- Stochastic lead times with a bounded (possibly unknown) cumulative distribution function $F^{(L)}$ having a known upper bound $\hat{L}$ and a strictly positive, integer-valued support;
- Nonstationary demand, with a (possibly unknown) cumulative distribution function Φ_t in a given period t;
- Lost sales, i.e., no backlogging.

Firm X generously provided us with more than two years worth of detailed sales, forecast, purchase order, and inventory history for a variety of products. We analyzed and combined the data to render it capable of feeding our heuristic.

4.3 Related Literature

Karlin and Scarf (1958) (hereafter KS) is often viewed as the seminal work in the literature on periodic-review, dynamic inventory models with positive (deterministic) lead times. They study both the backorder case and the lost-sales case in a quite general setting, with few assumptions on the cost structure and the demand

distribution. In the backorder case, KS prove that the optimal ordering policy can be expressed as a function of only the net inventory position (i.e., the sum of on-hand inventory and all outstanding orders), neither distinguishing on-hand inventory from on-order inventory nor distinguishing among on-order inventory that will be delivered in different periods. This work was a step towards proving the optimality of base-stock policies for proportional—i.e, linear with no fixed cost—ordering costs and (s, S) policies for systems with fixed order costs. The optimality of (s, S) policies in different inventory models, along with their computation, has been of great interest over the last half-century. See, e.g., Veinott (1966) for an early paper proving optimality of an (s, S) policy under certain conditions, and Zheng and Federgruen (1991) and references therein for algorithms to compute optimal (s, S) policies.

Intuitively, the reason that the state-space reduction is possible is because in the backorder case, all demand will eventually be filled. By the time the current order arrives and can impact the cost, all outstanding orders will have been delivered. Moreover, because any shortage is handled by shipping out units that arrive later, the inventory level at the moment when the current order arrives can be fully determined with knowledge of (a) the current net inventory position and (b) the aggregate demand between now and then. There is no need to distinguish among the different types of inventory because any unit of inventory, whether on-hand or on-order, has the same effect on the inventory level at the moment of delivery of the current order. Naturally, different sequences of orders and demands can lead to different holding and backorder cost realizations, but the current order cannot possibly affect these costs.

The lost-sales setting does not allow this reduction. Because excess demand is lost, different sequences of deliveries and demands can result in different inventory levels at the time when the current order is delivered. Thus, the net inventory position does not encode all of the relevant information for deciding on an order

quantity. Accordingly, KS prove that no optimal policy for the lost-sales problem can be a function of only the net inventory position. For the special case of proportional ordering costs and a lead time of one period, they also show that the optimal order quantity is decreasing in the current stock level. This result is intuitive, as if more stock is currently available, then it seems natural to place a smaller order because there is less risk of a stockout.

Zipkin (2008) extends this intuition to the case with longer lead times. For proportional ordering costs, Zipkin applies the structural properties of $L^\natural$-convexity and submodularity to show that the optimal expected total cost is nondecreasing in the inventory vector (the vector of order quantities for each previous period going back to the oldest outstanding order, augmented with the on-hand inventory quantity), and also that the optimal order quantity is nonincreasing in the inventory vector. Importantly, the result in Zipkin (2008) applies to a (weak) increase in *all* elements of the inventory vector. We prove a surprising consequence of his result that goes unmentioned, pertaining to what happens with an increase in one element of the inventory vector (namely, the on-hand inventory entry) and an offsetting decrease in the remaining elements. Specifically, when an increase in the on-hand inventory level is accompanied by offsetting decreases in the remaining elements, we prove that the optimal order quantity actually *increases* in the on-hand inventory level. KS's result, meanwhile, applies only to one-period lead times, when an increase in the on-hand inventory cannot be offset by a decrease in outstanding order quantities because no orders are outstanding when we place an order.

Bijvank (2009) presents an approximation called the restricted order-up-to policy (ROUTP) for lost sales problems. He attempts to account for the fact that, for a given sequence of orders and demands, the inventory level will always be weakly higher in a lost-sales model than in a backorder model because lost demand does not deplete inventory. The approximation chooses an order-up-to-level and, to correct

for the discrepancy, places a limit on the maximum order size (a smaller limit than S). He also numerically computes the exact optimal policy for an example, and this policy conforms to our result, although Bijvank (2009) does not take notice of this feature of the optimal policy.

4.4 Power Approximations Old and New

We next turn our attention to power approximations. These power approximations are intended for a backorder setting; later we will introduce a correction factor to account for the specifics of the lost-sales environment. First, we describe the original power approximation of Ehrhardt (1979, 1984). Then, we fit a new power approximation based on demand data from Firm X, and finally we devise a new heuristic that leverages our power approximation to provide an updated inventory policy in each period based on the latest demand forecasts.

4.4.1 The Power Approximation for (s, S) Policies

Ehrhardt (1979)—hereafter referred to as [E1]—studies a periodic review inventory problem for a backorder system with an infinite horizon, a deterministic lead time L, linear holding (h) and shortage (p) costs, and a fixed order cost K. Demand is assumed to be stationary, with mean μ and variance σ^2. Ehrhardt (1984) proposes modifications to the methods of [E1] for stochastic lead times, which we will also incorporate. In Section 4.4.3 and following we will discuss our adaptation of this framework to our finite-horizon, lost-sales setting with nonstationary demand.

Ehrhardt's goal is to create an approximation that takes cost and leadtime parameters, along with the first and second moments of the demand distribution, and returns an approximately optimal (s, S) policy without resorting to an iterative algorithm. His motivation for attempting this was to ensure tractability. With improved computational power and faster algorithms for computing optimal (s, S)

policies—see, e.g., Zheng and Federgruen (1991) and Federgruen and Zipkin (1984)—computing optimal (s, S) policies for given demand distributions, lead times, and cost parameters can be done relatively quickly. However, in practical contexts, rarely do managers know the true underlying demand distributions for their products. It is, however, reasonable to expect that they can give sound estimates for the mean and standard deviation of the demand distribution. Power approximations can therefore still prove useful as they provide good approximations to optimal policies without requiring the full demand distribution.

In approximating the optimal values s^* for the reorder point and $D^* := S^* - s^*$ for the difference between the order-up-to level and the reorder point, [E1] computes the optimal policy for 288 numerical examples using an algorithm given in Veinott and Wagner (1965) and fits a regression model based on results from Roberts (1962). Roberts (1962) studies the system in the asymptotic regime for very large K and p.

Before a logarithmic transformation is applied, the regression model for D^* in [E1] is given by

$$\overline{D} = C\mu^{\alpha}\left(\frac{K}{h}\right)^{\beta}(L+1)^{\gamma}(\sigma)^{\delta}\left(\frac{p}{h}\right)^{\epsilon}. \tag{4.1}$$

Letting $\sigma_L := \sigma\sqrt{L+1}$, the fitted regression equation—after removing insignificant terms and bringing together ones with close coefficients—is

$$D_p = 1.463\mu^{.364}\left(\frac{K}{h}\right)^{.498}(\sigma_L^2)^{.0691}. \tag{4.2}$$

[E1] follows a similar procedure to find, given a value of D_p, a value s_p for the reorder point; this quantity approximates not the optimal reorder point s^*, but instead the best value to use for the reorder point given that the difference D between the order-up-to level and the reorder point is set to the approximately optimal D_p. He computes the optimal "companion" reorder point, given the approximately optimal

D_p, for each of the 288 numerical examples, and he regresses for this quantity with the regression variable s_p. Defining

$$z := \sqrt{\frac{D_p}{\sigma_L(1 + p/h)}} \tag{4.3}$$

and defining $\mu_L := (L + 1)\mu$, the resulting regression equation is given by

$$s_p = \mu_L + \sigma_L^{.0832}\left(\frac{\sigma^2}{\mu}\right)^{.187}\left(\frac{.220}{z} + 1.142 - 2.866z\right). \tag{4.4}$$

The expression for z is motivated by asymptotic results in Roberts (1962) on the optimal value for the reorder point s as K and p become large.

Define S_0 as

$$S_0 := \mu_L + v\sigma_L, \tag{4.5}$$

where v is the $p/(p+h)$ standard normal quantile. For small values of D_p/μ, an order will be placed in most periods. If an order will be placed in almost every period, then the setup cost K becomes more or less a sunk cost, and it may therefore be better to use a policy consistent with a setup cost of zero. The order-up-to level of such a policy is approximated by S_0 in equation (4.5). [E1] accordingly modifies the policy set by the power approximation such that, if $D_p/\mu < 1.5$, then the reorder point is set to

$$S = \min\{S_0, s_p + D_p\}.$$

The above relation ensures that the order-up-to level will not be set too much higher than the reorder point in cases where the setup cost should not play much of a role in the decision-making process. The power approximation can then be described by:

Step 1: Use equations (4.2)-(4.4) to calculate D_p, z, and s_p.

Step 2: If D_p/μ is less than or equal to 1.5, then use $s = \min\{s_p, S_0\}$ and $S = \min\{s_p + D_p, S_0\}$. If D_p/μ is greater than 1.5, then use $s = s_p$ and $S = s_p + D_p$.

Step 3: Round s and S to the nearest whole number.

Now, recall that [E1] assumes a known, fixed lead time L; the stochastic lead time case is treated in Ehrhardt (1984) (hereafter [E2]). The power approximation of [E1] is modified in [E2] by replacing μ_L and σ_L^2 in equations (4.2)-(4.5) with μ_ℓ and σ_ℓ^2, respectively, where

$$\mu_\ell = (\mathbb{E}[L] + 1)\mu \tag{4.6}$$

and

$$\sigma_\ell^2 = (\mathbb{E}[L] + 1)\sigma^2 + \mu^2 \mathrm{Var}(L). \tag{4.7}$$

Incorporating this modification, we are left with the following expressions for the stochastic lead time approximations $\widetilde{D}_p$, $\tilde{z}$, $\widetilde{s}_p$, and $\widetilde{S}_0$, namely

$$\widetilde{D}_p = 1.463\mu^{.364}\left(\frac{K}{h}\right)^{.498}(\sigma_\ell^2)^{.0691}, \tag{4.8}$$

$$\tilde{z} := \sqrt{\frac{\widetilde{D}_p}{\sigma_\ell(1 + p/h)}}, \tag{4.9}$$

$$\widetilde{s}_p = \mu_\ell + \sigma_\ell^{.0832}\left(\frac{\sigma^2}{\mu}\right)^{.187}\left(\frac{.220}{\tilde{z}} + 1.142 - 2.866\tilde{z}\right), \tag{4.10}$$

and

$$\widetilde{S}_0 := \mu_\ell + v\sigma_\ell, \tag{4.11}$$

where v as noted is the $p/(p + h)$ standard normal fractile. The stochastic lead time version of the power approximation is obtained by respectively substituting $\widetilde{D}_p$, $\tilde{z}$, $\widetilde{s}_p$, and $\widetilde{S}_0$ for D_p, z, s_p, and S_0 in the procedure described above.

4.4.2 A New Power Approximation Using Firm X Data

The power approximation of [E1], as mentioned, was fit using 288 numerical examples. Ehrhardt uses Poisson and negative binomial distributions for demand, with a

Table 4.1: Parameter values used to fit new power approximation.

Parameter	Values	Count
Item	1-14	14
Lead Time	0,1,2,3,4,5,6	7
Setup Cost (K)	32,64	2
Shortage Cost (p)	4,9,24,99	4
Holding Cost (h)	1	1

variety of mean demands, lead times, setup costs, and shortage costs. The approximation performs remarkably well both in and out of sample for the examples that are tested.

However, the demand pattern for Firm X naturally does not match a simple distribution family like Poisson or negative binomial. Thus, rather than using Ehrhardt's power approximation "off the shelf," we take real sales data for Firm X and fit a new approximation based on its unique real-world demand profile. We should note that there may be a demand censoring phenomenon affecting the quality of the data; we do not directly address this issue, but as we will see, the approximation nonetheless performs quite well.

Similar to [E1], we compute the optimal (s, S) policies for each of the parameter/item combinations in Table 4.1, a total of 784 data points. We use the same values of K, p, and h that Ehrhardt does, but we study a larger variety of lead times and a variety of demand distributions, each one an empirical demand distribution for an item sold by Firm X. In computing the optimal policies, we use the expressions for long-run average cost of a given (s, S) policy given in Federgruen and Zipkin (1984). These expressions are exact but assume zero lead time, so we modify them for a positive lead time following equation (1) in Veinott and Wagner (1965) to compute the exact long-run average cost. We implement the algorithm for finding optimal (s, S) policies found in Zheng and Federgruen (1991).

After computing the optimal policies for the set of test examples, we first run a

regression model akin to that found in equation (6) of [E1]—reproduced as equation (4.1) above—to determine an approximation for the difference $\Delta_p \approx S - s$. The model that we settle on is fit as

$$\Delta_p = .2757\mu^{.7478}\left(\frac{K}{h}\right)^{.4861}(L+1)^{.3625}\left(\frac{p}{h}\right)^{-.1149}. \tag{4.12}$$

Comparing equation (4.12) with equation (4.1), we observe that the coefficient for the K/h term in our model is very close to that in Ehrhardt's, while the other coefficients diverge considerably. We also find that our model has more interpretive power if we exclude σ as a predictor and replace it with p/h. The adjusted R^2 for the regression model is .7188. While this figure is considerably lower than the .98 reported in [E1], we would expect a model based on real-world data to be more difficult to fit. Moreover, as [E1] is quick to stress, the measure of the strength of the approximation comes not from the goodness-of-fit on the training cases, but from how well it performs when tested. We will test the approximation—on items not used to fit the model—in Section 4.6.2.

The next step is to fit a model for the reorder point s. Motivated by equation (4.3) in Ehrhardt's procedure, we define

$$\hat{y} := \frac{\Delta_p}{\sigma_L(1 + p/h)}. \tag{4.13}$$

We fit a regression model similar to equation (9) in [E1]. An important difference is that, while [E1] attempts to estimate the s value that is best given the approximately optimal D_p, in our case we instead attempt to estimate the value of s in the exact optimal (s, S) policy. Because we will later incorporate minimum order quantities, the inventory level after an order will not necessarily be exactly equal to Δ_p. Therefore, it is more important to be as close as possible to the reorder point from the optimal (s, S) policy than for the reorder point to be optimal for the given

value of Δ_p. Hence, we regress for the reorder point s from the optimal (s, S) policy computed for each of the 784 test cases, with our regression variable denoted by ϕ_p. The form of our regression model for the approximately optimal reorder point ϕ_p, where $\hat{z} = \sqrt{\hat{y}}$, is

$$\phi_p = C_0 + C_1[(L+1)\mu] + C_2\Big[\frac{\sigma\sqrt{L+1}}{\hat{z}}\Big] + C_3[\sigma\sqrt{L+1}] + \epsilon. \tag{4.14}$$

The output of the model gives the regression equation

$$\phi_p = 148.1668 + .9284[(L+1)\mu] + .0181\Big[\frac{\sigma\sqrt{L+1}}{\hat{z}}\Big] + .7082[\sigma\sqrt{L+1}]. \tag{4.15}$$

The regression model (4.15) has an adjusted R^2 of .9611. [E1] takes several further steps to refine the regression model which we do not undertake here. When executing this new power approximation, we apply the same steps as described in Section 4.4.1, replacing equations (4.2) and (4.4) with equations (4.12) and (4.15), respectively. Note also that the modification described in Section 4.4.1 to incorporate stochastic lead times can be applied to our power approximation as well as Ehrhardt's. We will compare the performance of Ehrhardt's power approximation with the new model in Section 4.6.2.

4.4.3 Dynamic Heuristic for Computing (s_t, S_t) Policies

We proceed to adapt the power approximations described in the previous subsections for use in an inventory policy which is updated periodically based on demand forecasts. We adopt the following notation:

- d_t: random demand realized in period t;
- $\hat{d}_t^{(j)}$: forecast in period j for the demand in period t;
- ν_t: on-hand inventory in period t, before orders are delivered;

- ψ_t: on-order inventory in period t, before orders are delivered;
- $\eta_t = \nu_t + \psi_t$: net inventory position at the beginning of period t, including inventory on hand plus any outstanding orders;
- O_t: total quantity delivered in period t, from one or more outstanding orders;
- Q_t: quantity ordered in period t (zero if no order placed);
- L_t: random lead time in integer number of periods for an order placed in period t.

By convention, we assume that deliveries are received at the beginning of a period, after which the order quantity (possibly zero) is determined, and finally demand for the period is realized. Suppose that we have estimates $\overline{L}$ and φ for the mean and standard deviation of the lead time, respectively, and that the lead times of all orders are identically distributed. Denote by $\hat{\chi}_{t,k}^{(j)}$ the mean of the forecasts for the demands starting in period t and continuing for k periods, as of period j, i.e.,

$$\hat{\chi}_{t,k}^{(j)} := \frac{1}{k} \sum_{i=t}^{t+k-1} \hat{d}_i^{(j)}.$$

The last period included is thus period $t+k-1$. Similarly, denote by $\hat{\xi}_{t,k}^{(j)}$ the sample standard deviation of the forecasts for the demands in periods t through $t+k-1$, as of period j, given by

$$\hat{\xi}_{t,k}^{(j)} := \sqrt{\frac{1}{k-1} \sum_{i=t}^{t+k-1} \left(\hat{d}_i^{(j)} - \hat{\chi}_{t,k}^{(j)} \right)^2}.$$

The only parameters required by the power approximation are the first and second moments of the lead time distribution and the same for the demand distribution, along with setup cost and unit shortage and holding costs.

In a nonstationary demand environment for a product with a short life cycle, it is very difficult to ascertain the true form of the demand distribution. Recognizing this difficulty, we devise a procedure that estimates the lead time distribution using historical data and estimates the demand distribution using the current demand forecasts for a rolling horizon of fixed length. We estimate the mean demand by the sample average of the forecasts in the rolling horizon. We estimate the standard deviation of demand using the sample standard deviation of the forecasts in the rolling horizon. We acknowledge that variation in forecasts is not equivalent to variation in demand, and that this method could lead to difficulties if, for instance, the demand forecast was constant for the entirety of the rolling horizon, which would cause the demand variance to be estimated as zero. Our procedure could potentially be improved with other methods of estimating demand variance, perhaps by incorporating data on historical forecast errors.

The goal of our method is to set an inventory policy in each period that (i) is easy to compute using available information, (ii) captures as well as possible any demand volatility that is likely to impact the current order, and (iii) adapts to changing demand and supply conditions. In this spirit, and given a setup cost K, a unit shortage cost p, and a unit holding cost h, we propose the following adaptation of the power approximations described in previous subsections:

Step 0: Choose a whole number $w \geqslant 2$ to represent the number of periods to use in estimating the mean and standard deviation of the one-period demand. Setting $w = 4$, for example, would use the forecasts for a four-week period beginning with the current period. Similarly, estimate the mean and standard deviation of the lead time using the sample moments from the empirical distribution of all previous orders, yielding estimates $\overline{L}$ and φ.

Step 1: In period t, use either of the power approximations from Sections 4.4.1-4.4.2 to

compute a policy (s_t, S_t) using the lead time parameters $\overline{L}$ and φ and setting $\mu = \hat{\chi}_{t,w}^{(t)}$ and $\sigma = \hat{\xi}_{t,w}^{(t)}$.

Step 2: If the inventory position $\eta_t \leqslant s_t$, then set $Q_t = S_t - \eta_t$. Otherwise, set $Q_t = 0$. After the demand d_t is realized, update the on-hand, on-order, and net inventory positions via

$$\nu_{t+1} = \left[\nu_t + O_t - d_t\right]^+, \tag{4.16}$$

$$\psi_{t+1} = \psi_t - O_t + Q_t, \tag{4.17}$$

$$\eta_{t+1} = \nu_{t+1} + \psi_{t+1}. \tag{4.18}$$

Step 3: If period t is the end of the planning horizon, terminate. Otherwise, return to Step 1 and repeat for period $t + 1$.

Minimum Order Quantities

The above procedure can be modified to incorporate a minimum order quantity (MOQ). Suppose that the minimum order quantity is m. Then, determine a fraction κ and replace Step 2 above with

Step 2: If the inventory position $\eta_t \leqslant s_t$ and $S_t - \eta_t \geqslant \kappa m$, then set $Q_t = \max\{m, S_t - \eta_t\}$. Otherwise, set $Q_t = 0$. After the demand d_t is realized, update the on-hand, on-order, and net inventory positions via

$$\nu_{t+1} = \left[\nu_t + O_t - d_t\right]^+, \tag{4.19}$$

$$\psi_{t+1} = \psi_t - O_t + Q_t, \tag{4.20}$$

$$\eta_{t+1} = \nu_{t+1} + \psi_{t+1}. \tag{4.21}$$

Our procedure sets a new inventory policy in each period, using historical supply information to estimate the lead time distribution and evolving demand forecast information to estimate the demand distribution. It adapts to changing conditions

and uses the latest available information in each period to set the new policy. It is, however, based on approximations that were designed for backorder systems. In the next section, we derive an analytical result about the role of on-hand versus on-order inventory in determining the optimal order quantity in lost-sales inventory systems, based on results from Zipkin (2008). Using insights from this result, we then propose a modification to the procedure described in this subsection to adjust to the lost-sales environment.

4.5 The Optimal Order Quantity: On-Hand vs. On-Order Inventory

Consider the inventory model of Zipkin (2008). Specifically, we have an inventory system with proportional ordering costs; lost sales; periodic review (i.e, discrete time); a fixed lead time of L periods; linear holding and shortage costs; a finite horizon; and a discount factor r. This setting is different from that which we have so far studied, but we will obtain insights from it which will be valuable in our setting also, and will lead us to a modification of our dynamic heuristic for lost-sales systems. Our notation follows Zipkin, except that we replace his z_t for the order quantity in period t with Q_t, to avoid confusion with the meaning of z in Section 4.4.1:

- d_t = period t demand;
- y_t = on-hand inventory at period t, after the order from period $t - L$ arrives (note that we would have $y_t = \nu_t + O_t$ in Section 4.4.3's notation);
- Q_t = order quantity at period t;
- $x_{0t} = y_t$;
- $x_{kt} = Q_{t+k-L}, \quad 0 \leqslant k \leqslant L-1$ (order to arrive in period $t + k$);
- $\boldsymbol{x}_t = (x_{0t}, x_{1t}, \ldots, x_{L-1,t})$;

- $\bar{x}_t = \sum_{i=0}^{L-1} x_{kt}$.

Importantly, like Zipkin we allow the order quantity Q_t and the state variables x_{kt} to take any nonnegative value, i.e, they are continuous rather than discrete.

Corollary 5 of Zipkin shows that, holding the rest of the inventory pipeline fixed, the optimal order quantity is decreasing in the on-hand inventory level. This result matches the natural intuition that the more stock that we have on hand in the warehouse, the smaller should be the order that we place from our supplier because with more inventory we are at lower risk of a stockout. This intuition is accurate, ceteris paribus.

However, we show now a result which has a strikingly different—and perhaps surprising—implication. The ensuing Proposition 18 holds conditions constant in a different sense: it compares the optimal order quantities for two different inventory vectors, *when the sum of the inventory pipeline vectors is the same for both.* In other words, we are interested in how the optimal order quantity is affected by shifting inventory earlier or later in the pipeline, keeping the total inventory (on-hand plus on-order) constant.

Proposition 18 (Order Quantity Increasing in On-Hand Inventory for Fixed $\bar{x}$)**.** *For a fixed value of the net inventory position $\bar{x}_t$, the optimal order quantity $\hat{Q}_t$ is nondecreasing in the on-hand inventory y_t.*

Proof. By Corollary 5 and following in Zipkin, if we assume that the optimal order quantity $\hat{Q}_t(\boldsymbol{x}_t)$ is a smooth function, we have

$$\begin{aligned} \frac{\partial \hat{Q}_t}{\partial x_{0t}} &\leqslant 0, \\ \frac{\partial \hat{Q}_t}{\partial x_{kt}} &\leqslant \frac{\partial \hat{Q}_t}{\partial x_{k-1,t}}, \quad 0 < k < L. \end{aligned} \tag{4.22}$$

Now, drop the time index on $\hat{Q}(\boldsymbol{x})$. Consider the inventory vector

$$\tilde{\boldsymbol{x}} = (\tilde{x}_0, \tilde{x}_1, \ldots, \tilde{x}_{L-1}),$$

and define

$$\tilde{\boldsymbol{x}}_{k,\epsilon} := (\tilde{x}_0 + \epsilon, \ldots, \tilde{x}_{k-1}, \tilde{x}_k - \epsilon, \tilde{x}_{k+1}, \ldots, \tilde{x}_{L-1}) \text{ for } \epsilon \geqslant 0 \text{ and } k \in \{1, \ldots, L-1\}.$$

Define also the line segment $C : \boldsymbol{y}(u) = \tilde{\boldsymbol{x}}_{k,u}, u \in [0, c]$. By the fundamental theorem of line integrals (see Salas et al., 2007, Section 18.2), we have

$$\begin{aligned}
\hat{Q}(\tilde{\boldsymbol{x}}_{k,c}) - \hat{Q}(\tilde{\boldsymbol{x}}) &= \int_C \nabla \hat{Q}(\boldsymbol{y}) \mathsf{d}\boldsymbol{y} \\
&= \int_0^c \Big[\nabla \hat{Q}\big(\boldsymbol{y}(u)\big) \cdot \boldsymbol{y}'(u) \Big] \mathsf{d}u \\
&= \int_0^c \left[\frac{\partial \hat{Q}}{\partial x_0} - \frac{\partial \hat{Q}}{\partial x_k} \right]_{\boldsymbol{x} = \tilde{\boldsymbol{x}}_{k,u}} \mathsf{d}u \\
&\geqslant 0.
\end{aligned}$$

The last inequality follows because the integrand is always nonnegative by equation (4.22). This process can be repeated with different k to show that, for any two inventory pipeline vectors with the same sum, the optimal order quantity will be higher for the vector with a higher on-hand inventory level. □

The intuition about the relationship between on-hand inventory and the likelihood of a stockout can be applied here also, although with a very different implication. As mentioned, the more inventory that we have on hand in the warehouse, the lower the risk of a stockout in coming periods. This fact would seem to suggest that we should place a smaller order, but here the distinctive features of a lost-sales system play an important role. In a backorder system, the location of inventory in the pipeline is irrelevant to this period's order because any unfilled demand is backlogged and will eventually be satisfied. Thus, if we keep the sum constant, in a

backorder system our inventory position when the current order arrives will be the same for any configuration of the inventory pipeline vector. The same is not true for a lost-sales system.

To illustrate, consider two possible inventory pipeline vectors—corresponding to scenarios denoted (a) and (b)—both with the same net inventory position $\bar{x}_t$. First, let $\boldsymbol{x}_t^{\mathrm{a}}$ be such that all of the inventory is on-hand, and there are no outstanding orders (i.e., $x_{0t}^{\mathrm{a}} = \bar{x}_t$, and $x_{kt}^{\mathrm{a}} = 0$ for $k \geqslant 1$). Second, let $\boldsymbol{x}_t^{\mathrm{b}}$ be such that all of the inventory is scheduled to be delivered next (i.e., $x_{1t}^{\mathrm{b}} = \bar{x}_t$, and $x_{kt}^{\mathrm{b}} = 0$ for $k \neq 1$). Suppose that the demand in period t is realized as $d_t < \bar{x}_t$. In scenario (a), all of the demand is met, and the starting inventory next period will be $x_{0,t+1}^{\mathrm{a}} = \bar{x}_t - d_t$ (recall that no outstanding orders are to be delivered next period in this scenario). In scenario (b), we have a stockout so none of the demand is met, and the starting inventory next period will be $x_{0,t+t}^{\mathrm{b}} = \bar{x}_t$ because all of the demand is lost, and at the start of next period an order for $\bar{x}_t$ units will be delivered. In the scenario where we had more on-hand inventory, the inventory level at the start of the next period was *lower*.

So, as before, more on-hand inventory means a lower stockout probability. But in a lost-sales system, all else equal, a stockout in the current period results in a higher inventory level in future periods because unmet demand does not deplete the stock. Therefore, the more of our inventory pipeline that exists as on-hand inventory, the lower the probability of a stockout between now and the arrival of the current period's order, and thus the lower our inventory level is likely to be when the order arrives. Consequently, for a fixed net inventory position, the higher the on-hand stock, the higher the optimal order quantity.

We note that Zipkin studies systems with a finite horizon, a deterministic lead time, a discount factor, and no fixed ordering cost. As Zipkin states that its results extend to the case with an infinite horizon and stochastic lead times, we conjecture

that our Proposition 18 does as well. While we do not here address settings with a fixed ordering cost or with the long-run average cost criterion, the underlying principle behind our result applies to these settings also, and we believe that a similar result should hold.

Hence, we propose a modification to the dynamic heuristic of Section 4.4.3 to account for this feature of a lost-sales system. Suppose, in period t, after orders are delivered and before demand is realized, that a fraction α_t of the inventory pipeline exists as on-hand inventory. That is, in the notation of Section 4.4.3, we define

$$\alpha_t := \frac{\nu_t + O_t}{\eta_t + O_t}.$$

If $\alpha_t = 1$ and all of the inventory is on-hand inventory, then this lost-sales system is as close as it can get to a backorder system: for a given sequence of demand realizations, all $\eta_t + O_t$ units of inventory will be used the same way in either setting. There is still a difference between the two settings, however. If the demand over the next few periods exceeds $\eta_t + O_t$ and no orders arrive in the meantime, then the lost-sales system will end up with more inventory because unlike the lost-sales system the backorder system will eventually have to satisfy the unfilled demand. We define an additional parameter ω, and, given the policy (s_t, S_t) that is generated for the period, propose to replace S_t with $\tilde{S}_t$, where

$$\tilde{S}_t = \max\{s_t, \omega S_t + \alpha(1-\omega)S_t\}. \tag{4.23}$$

When $\alpha = 1$, we have $\tilde{S}_t = S_t$, and the policy is unchanged. The order-up-to level decreases as α decreases below 1. When $\alpha = 0$, we have $\tilde{S}_t = \omega S_t$; thus, the parameter ω can be interpreted as the degree to which the order-up-to level should be reduced when there is no on-hand inventory and everything is in the pipeline, subject to the constraint that it cannot be less than s_t. We will evaluate the performance of our heuristic with and without this modification in Section 4.6.2.

4.6 GUI-Based Tool and Performance of Heuristic

In this section, we will describe the tool that we created for Firm X to use in inventory planning, and using this tool, evaluate the performance of our dynamic heuristic on Firm X's actual data.

4.6.1 GUI-Based Tool

We have implemented our model in R, and on top of the model we have built a user interface for Firm X's inventory planning with Shiny, an R package for creating web applications.

The tool allows Firm X to upload SKU, purchase order, inventory, sales, and forecast data (see Figure 4.2). After uploading the relevant data, as seen in Figure 4.3, Firm X can view the inventory policy recommendations for a given date (e.g., the current date if setting a policy for the current period). Note that the version of the tool given to Firm X includes true SKU labels and also displays actual quantities for the inventory policy, both of which are disguised in Figure 4.3, along with detailed figures about the model and policy. The user can adjust both the date and the number of weeks to use in estimating demand parameters (the value w from Section 4.4.3), and the application will compute an updated policy. The tool also allows the user to compare the hypothetical performance of our model against realized historical inventory levels. A sample output of the tool comparing historical and model inventory levels is shown in Figure 4.4. Firm X has been evaluating our method over a trial period, and we are now training employees of Firm X to use the tool as a part of the inventory planning process.

4.6.2 Evaluation of Dynamic Heuristic

Firm X generously provided us with a long enough sales history to separate the data into training and test sets. We used the training data set to fit the new power

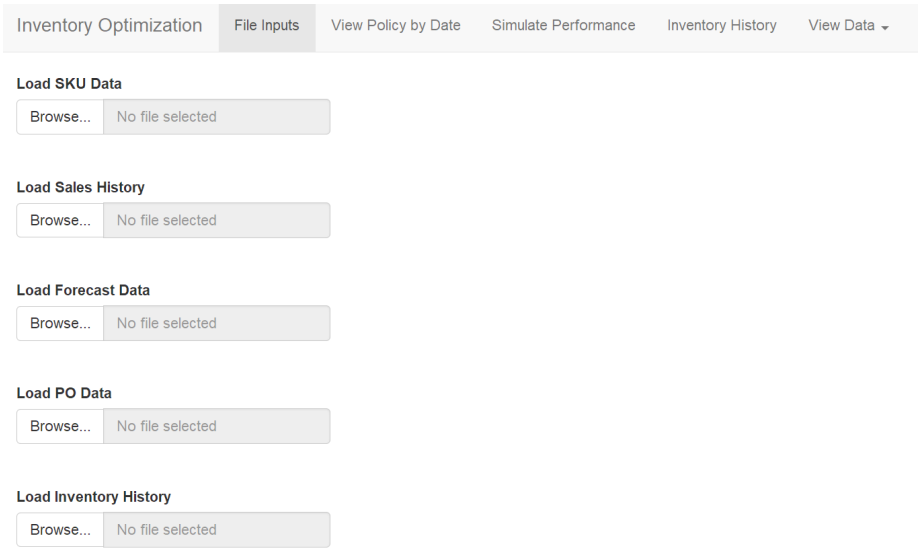

FIGURE 4.2: File input screen in GUI-based tool.

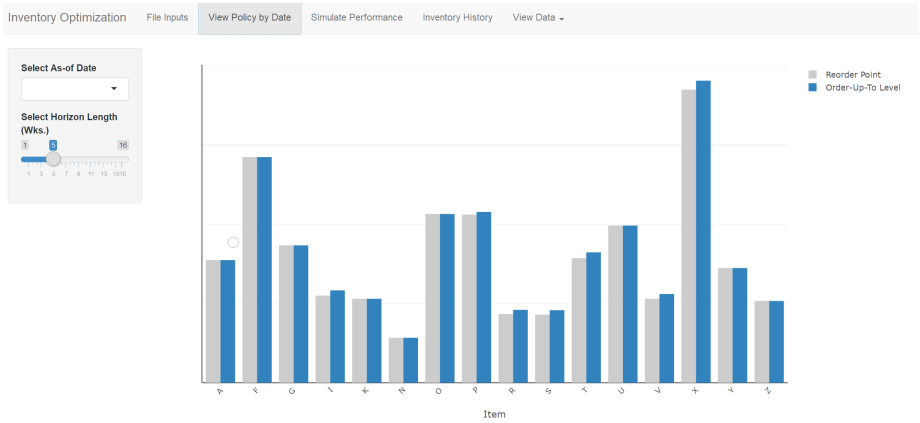

FIGURE 4.3: Inventory policy for a given date and forecast horizon.

approximation of Section 4.4.2. In this section, we test the performance of our heuristic, using both Ehrhardt's original power approximation as well as our own. We evaluate the two versions of the heuristic against each other and against realized historical inventory levels and in-stock performance for Firm X.

Our testing data contains sales, inventory, purchase order, and forecast data for 11 items over roughly a six-month period. Two items were excluded which were near

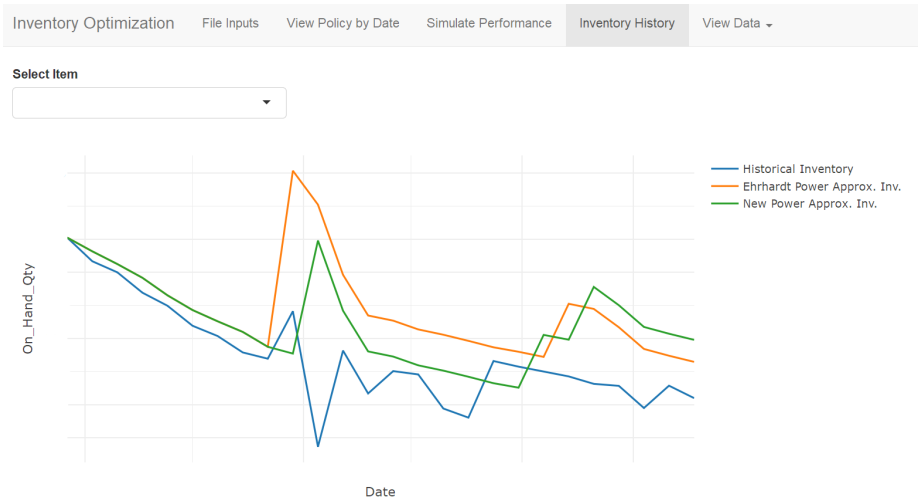

FIGURE 4.4: Historical and simulated inventory levels for an example item.

the end of their lifetime and for which Firm X received no deliveries during the time window of interest, and 11 others were excluded which were not active at the start of our time window. In a given period t and for item j, we estimate the lead time parameters $\overline{L}$ and φ^2 by the corresponding sample statistics for all purchase orders that (i) were from the same supplier that supplies item j and (ii) were received prior to period t. We estimate the demand parameters according to the procedure outlined above in our heuristic, using a horizon of $w = 5$ periods, and using actual forecasts generated by Firm X in each period. There were several periods in the data when Firm X did not update the forecasts. For a period t when Firm X did not update the forecasts, we assume that the forecast for any period $t' > t$ remained the same as the forecast for period t' generated in period $t - 1$. In addition, there were some periods for which we did not have recorded inventory levels. For these periods, we attempted to back-fill the data by looking backwards until finding a period with a valid inventory level, and reducing this value by the sums of the sales quantities during the intervening periods. Firm X does have minimum order quantities imposed by suppliers. We employ the MOQ modification to Step 2 of our heuristic proposed

at the end of Section 4.4.3, setting the fraction $\kappa = 1/3$.

We simulate our method 1000 times for each item, and the out-of-stock rate is the number of trial-weeks where we were out of stock divided by the total number of trial-weeks in the sample. Though not known by our method in advance, the realized demand in a period is the same across all trials and is set equal to the actual historical sales in that period. The reason for replication is because of stochastic lead times. Note that the lead time distribution used to estimate $\overline{L}$ and φ^2 is, as mentioned, based only on information that would have been available to Firm X at each decision epoch. However, to simulate the actual lead time for a purchase order we use the realized empirical lead time distribution for the corresponding item.

Table 4.2 summarizes the results of simulating our method with respect to the out-of-stock rate. The numbers in the table are the changes in percentage "points" from the historical values to the hypothetical values using our method. Both power approximations (Ehrhardt and our own) either improve the out-of-stock rate or increase it by less than two percentage points on 8 out of the 11 items (a negative number indicates that our method had a lower out-of-stock rate than the historical data). The magnitude of the improvement appears to be comparable for the two approximations.

Table 4.3 presents the change in average inventory level versus Firm X 's historical inventory levels from each method. The average inventory level increases for most items, but Firm X is potentially willing to accept more inventory for the sake of good in-stock performance. Our new approximation results in less of an inventory increase than Ehrhardt's power approximation for all but two items, and in some cases by a wide margin.

So, although the two methods perform similarly with respect to stockout performance, our approximation is able to achieve this performance at a lower inventory cost than Ehrhardt's power approximation. Overall, our new power approximation

Table 4.2: Percentage point change in out of stock rate.

Item	Ehrhardt	New
A	0.07	1.67
B	0.63	1.59
C	0.00	0.00
D	-3.85	-3.80
E	-10.47	-10.65
F	-4.77	-3.84
G	0.17	0.20
H	3.58	4.00
I	-3.17	-3.58
J	12.26	4.73
K	6.25	5.46

achieves on average a 0.38% reduction in out-of-stock rate, with much larger reductions on several items and an improvement on a majority of items, as compared to an average of 0.06% increase for Ehrhardt's approximation. We should stress that this performance is specific to Firm X's products, as is our approximation. Our approximation was developed with and for Firm X's particular demand profile, and it is encouraging that there is opportunity for improved performance from fitting a special approximation for a given application. The methodology behind the fitting of the approximations is agnostic to the distributional form of the demand and can thus be applied with any demand distributions; we see this as a potential avenue both for future research and future practical applications.

Lost-Sales Correction Factor

Thus far, when testing our methodology we have made no modifications to the policies with respect to our Proposition 18. We now test the modification suggested by equation (4.23) on Firm X's data, following the same procedure as before.

We again simulate our methods 1000 times for each item, with the same model parameters as before. We set the lost-sales correction parameter $\omega = .8$. This figure was chosen so that when none of the inventory is on-hand, the order-up-to level

Table 4.3: Percentage change in average inventory level.

Item	Ehrhardt	New
A	-1.62%	-16.93%
B	35.72%	22.51%
C	40.21%	18.78%
D	31.07%	2.10%
E	28.45%	25.12%
F	9.07%	-3.85%
G	26.21%	17.32%
H	100.01%	66.63%
I	72.27%	77.51%
J	40.00%	46.41%
K	158.01%	143.75%

would be slightly more than half of that in the backorder setting. With minimum order quantities, a value closer to 1 might make no difference in the order quantities, whereas a value close to zero could cause the inventory levels to fall too far. The impact of the correction on stockouts is presented in Table 4.4; note that this is again reported relative to Firm X's performance. To compare it to our methods in the absence of the lost-sales correction, we compare the numbers in Table 4.4 with the corresponding numbers in Table 4.2. Similarly, to compare the impacts on average inventory level, we compare the numbers in Table 4.3 with those in Table 4.5.

The stockout performance with and without the lost-sales correction is quite similar. With the lost-sales correction, for many items the average inventory level is roughly the same as without the correction, and for others the inventory level is somewhat reduced with the correction. This reduction in average inventory level is consistent with the purpose of the correction, as the lost-sales correction reduces the order-up-to level from that prescribed by the original heuristic.

An important part of Firm X's inventory planning process is accounting for MOQ's. Because there is a minimum order quantity, unless the difference between S_t and the inventory position is quite large, it is not advisable to place an order. Thus,

Table 4.4: Percentage point change in out of stock rate with lost-sales correction.

Item	Ehrhardt	New
A	0.03	1.69
B	0.52	1.75
C	0.00	0.00
D	-3.85	-3.79
E	-8.30	-10.43
F	-4.60	-3.45
G	0.25	0.22
H	3.71	4.17
I	-3.47	-3.53
J	9.80	2.40
K	6.79	6.71

Table 4.5: Percentage change in average inventory level with lost-sales correction.

Item	Ehrhardt	New
A	-1.10%	-17.20%
B	36.30%	20.70%
C	37.40%	13.80%
D	29.60%	-2.70%
E	24.70%	24.00%
F	8.50%	-5.10%
G	26.80%	17.20%
H	101.60%	64.20%
I	47.90%	47.70%
J	12.30%	15.90%
K	109.90%	88.80%

both with and without the correction for lost-sales, there are likely many cases where an order would be placed for fewer units than the MOQ, and if the suggested order is small relative to the MOQ, then our model chooses not to place an order. These cases can cause the performance of the methods with and without the lost-sales correction factor to look quite similar, although they might perform differently in the absence of minimum order quantities.

4.7 Conclusion

Using real demand data from an industry partner, we have fit and tested a new version of the power approximation first introduced in Ehrhardt (1979). We have proposed and implemented a fast, adaptive heuristic based on this approximation that adjusts to changing demand conditions. We have also proven a new structural result on general lost-sales inventory problems, and we use the insights from this result to inform a modification of our heuristic for lost-sales systems.

In addition, we have developed and described a user-friendly interactive tool that Firm X is planning to incorporate into the inventory planning process. We have tested our method on another data set from the same industry partner and found that it reduces stockouts in comparison to the methods currently in use. Moreover, while these current methods rely on a large deterministic optimization model that will not scale with the business, our method is efficient enough to scale quite easily as the company grows. Our method also accounts for both supply and demand uncertainty, and we believe that it has great practical potential.

5

Conclusion

This dissertation has studied operations management problems in three different settings: strategic routing in service networks; pricing and revenue implications of service delivery platforms; and lost-sales inventory problems with short life cycles and forecast updates. In each setting we have discovered insights that illuminated the nature of strategic customer and firm behavior in operations management systems.

In Chapter 2, we studied a model of strategic customer behavior in a service network with open routing, in which customers must visit multiple stations in the network but can strategically choose their routes. We uncovered a surprising herding behavior, such that in equilibrium all customers follow the same route through the network. We also identified a key structural property—submodularity of the expected system time—that greatly aided much of the analysis in the chapter. As structural properties of waiting times have been rarely applied to the analysis of queueing games, this chapter may open the door for future application of submodularity to other settings in the queueing games literature.

In Chapter 3, in an environment with an incumbent restaurant and an independent platform which offers delivery of the restaurant's food to customers, we

investigated the effect of congestion and strategic interaction on the purchasing decisions of customers and on the pricing decisions of the restaurant and platform. We discovered that when operating in a decentralized setting with no contract, the restaurant and the platform may leave revenue on the table, and that the restaurant tends to have bargaining power over the platform because of its attractive outside options. We adapted techniques from the priority queueing and pricing literature for our context, and we believe that these methods could prove useful in studying related models of congestion and strategic queue-joining behavior.

Finally, in Chapter 4, we studied a theoretically challenging lost-sales inventory problem for products with short life cycles and forecast updates. We used a training data set from an industry partner to fit a new version of the classical power approximation of Ehrhardt (1979). We evaluated our approximation, combined with a new dynamic heuristic that leverages historical supply information and the latest demand forecasts, on a separate test data set from that used to create the approximation. We found that our methods outperformed both our industry partner and Ehrhardt's approximation with respect to out-of-stock rates. We believe that there may be opportunities for future work related to inventory management to follow a similar approach: fit a specialized approximation for a given application, based on historical data, and then feed the model with updated forecasts to set a policy in each period. This approach could be used to fit other power approximations in different supply and demand environments, or, if the context called for thoroughly different methods, then the overall framework could be adapted as required.

Bibliography

Adiri, I. and Yechiali, U. (1974), "Optimal priority-purchasing and pricing decisions in nonmonopoly and monopoly queues," *Operations Research*, 22, 1051–1066.

Afèche, P. and Mendelson, H. (2004), "Pricing and priority auctions in queueing systems with a generalized delay cost structure," *Management Science*, 50, 869–882.

Afeche, P. and Sarhangian, V. (2015), "Rational abandonment from priority queues: Equilibrium strategy and pricing implications," *Working Paper, University of Toronto*.

Afèche, P., Liu, Z., and Maglaras, C. (2018), "Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance," *Working Paper, Columbia University, New York*.

Allon, G., Bassamboo, A., and Gurvich, I. (2011), ""We will be right with you": Managing customer expectations with vague promises and cheap talk," *Operations Research*, 59, 1382–1394.

Alperstein, H. (1988), "Note: Optimal pricing policy for the service facility offering a set of priority prices," *Management Science*, 34, 666–671.

Arlotto, A. and Steele, J. M. (2016a), "A central limit theorem for costs in Bulinskayas inventory management problem when deliveries face delays," *Methodology and Computing in Applied Probability*, pp. 1–16.

Arlotto, A. and Steele, J. M. (2016b), "A central limit theorem for temporally nonhomogenous Markov chains with applications to dynamic programming," *Mathematics of Operations Research*, 41, 1448–1468.

Arlotto, A., Gans, N., and Steele, J. M. (2014), "Markov decision problems where means bound variances," *Operations Research*, 62, 864–875.

Arlotto, A., Frazelle, A. E., and Wei, Y. (2018), "Strategic open routing in service networks," *Management Science*, Articles in Advance.

Arya, A., Mittendorf, B., and Sappington, D. E. (2007), "The bright side of supplier encroachment," *Marketing Science*, 26, 651–659.

Bai, J., So, K. C., Tang, C. S., Chen, X., and Wang, H. (2018), "Coordinating supply and demand on an on-demand service platform with impatient customers," *Manufacturing & Service Operations Management*, Articles in Advance.

Banerjee, S., Riquelme, C., and Johari, R. (2015), "Pricing in ride-share platforms: A queueing-theoretic approach," *Working Paper, Stanford University*.

Baron, O., Berman, O., Krass, D., and Wang, J. (2016), "Strategic idleness and dynamic scheduling in an open-shop service network: Case study and analysis," *Manufacturing & Service Operations Management*, 19, 52–71.

Benjaafar, S., Ding, J.-Y., Kong, G., and Taylor, T. (2018), "Labor welfare in on-demand service platforms," *Working Paper, University of Minnesota, Minneapolis, MN*.

Bijvank, M. (2009), "Service inventory management: Solution techniques for inventory systems without backorders," *Ph.D. Dissertation, Vrije University, Amsterdam, NL*.

Bijvank, M. and Vis, I. F. (2011), "Lost-sales inventory theory: A review," *European Journal of Operational Research*, 215, 1–13.

Bimpikis, K., Candogan, O., and Saban, D. (2018), "Spatial pricing in ride-sharing networks," *Operations Research*, forthcoming.

Bolton, G., Greiner, B., and Ockenfels, A. (2013), "Engineering trust: reciprocity in the production of reputation information," *Management Science*, 59, 265–285.

Braess, D. (1968), "Über ein Paradoxon aus der Verkehrsplanung," *Unternehmensforschung*, 12, 258–268.

Brown, G. W. (1951), "Iterative solution of games by fictitious play," *Activity Analysis of Production and Allocation*, 13, 374–376.

Cachon, G. and Feldman, P. (2011), "Pricing services subject to congestion: Charge per-use fees or sell subscriptions?" *Manufacturing & Service Operations Management*, 13, 244–260.

Cachon, G. P. and Lariviere, M. A. (2005), "Supply chain coordination with revenue-sharing contracts: strengths and limitations," *Management Science*, 51, 30–44.

Cachon, G. P. and Swinney, R. (2009), "Purchasing, pricing, and quick response in the presence of strategic consumers," *Management Science*, 55, 497–511.

Cachon, G. P. and Swinney, R. (2011), "The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior," *Management science*, 57, 778–795.

Cachon, G. P., Daniels, K. M., and Lobel, R. (2017), "The role of surge pricing on a service platform with self-scheduling capacity," *Manufacturing & Service Operations Management*, 19, 368–384.

Cohen, J. E. and Kelly, F. P. (1990), "A paradox of congestion in a queuing network," *Journal of Applied Probability*, 27, 730–734.

Cui, S. and Veeraraghavan, S. (2016), "Blind queues: The impact of consumer beliefs on revenues and congestion," *Management Science*, 62, 3656–3672.

David, A. and Adida, E. (2015), "Competition and coordination in a two-channel supply chain," *Production and Operations Management*, 24, 1358–1370.

Debo, L. and Veeraraghavan, S. (2014), "Equilibrium in queues under unknown service times and service value," *Operations Research*, 62, 38–57.

DeValve, L., Pekeč, S., and Wei, Y. (2018), "A primal-dual approach to analyzing ATO systems," *Working Paper, Duke University*.

Ehrhardt, R. (1979), "The power approximation for computing (s, S) inventory policies," *Management Science*, 25, 777–786.

Ehrhardt, R. (1984), "(s, S) policies for a dynamic inventory model with stochastic lead times," *Operations Research*, 32, 121–132.

Enders, P., Gandhi, A., Gupta, V., Debo, L., Harchol-Balter, M., and Scheller-Wolf, A. (2008), "Inducing optimal scheduling with selfish users," *Working Paper, Carnegie Mellon University, Pittsburgh, PA*.

Erlichman, J. and Hassin, R. (2015), "Strategic overtaking in a monopolistic M/M/1 queue," *IEEE Transactions on Automatic Control*, 60, 2189–2194.

Federgruen, A. and Zipkin, P. (1984), "An efficient algorithm for computing optimal (s, S) policies," *Operations Research*, 32, 1268–1285.

Feldman, M. and Tamir, T. (2012), "Conflicting congestion effects in resource allocation games," *Operations Research*, 60, 529–540.

Feldman, P., Frazelle, A. E., and Swinney, R. (2018), "Service delivery platforms: Pricing and revenue implications," *Working Paper, Duke University, Durham, NC*.

Fudenberg, D. and Levine, D. K. (1998), *The Theory of Learning in Games*, MIT Press, Cambridge, MA.

Gavirneni, S. and Kulkarni, V. G. (2016), "Self-selecting priority queues with Burr distributed waiting costs," *Production and Operations Management*, 25, 979–992.

Glazer, A. and Hassin, R. (1983), "?/M/1: On the equilibrium distribution of customer arrivals," *European Journal of Operational Research*, 13, 146–150.

Graves, S. C. (1999), "A single-item inventory model for a nonstationary demand process," *Manufacturing & Service Operations Management*, 1, 50–61.

Grimmett, G. and Stirzaker, D. (2001), *Probability and Random Processes*, Oxford University Press, Oxford, UK.

GrubHub (2018), "GrubHub for Restaurants FAQ," *Accessed May 14, 2018*, `https://get.grubhub.com/faq`.

Gurvich, I., Lariviere, M. A., and Ozkan, C. (2018), "Coverage, coarseness, and classification: Determinants of social efficiency in priority queues," *Management Science, Articles in Advance*, Forthcoming.

Ha, A., Long, X., and Nasiry, J. (2015), "Quality in supply chain encroachment," *Manufacturing & Service Operations Management*, 18, 280–298.

Hassin, R. (2016), *Rational Queueing*, CRC Press, New York.

Hassin, R. and Haviv, M. (1997), "Equilibrium threshold strategies: The case of queues with priorities," *Operations Research*, 45, 966–973.

Hassin, R. and Haviv, M. (2003), *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*, vol. 59, Springer Science & Business Media, New York.

Honnappa, H. and Jain, R. (2015), "Strategic arrivals into queueing networks: The network concert queueing game," *Operations Research*, 63, 247–259.

Hu, K., Acimovic, J., Erize, F., Thomas, D. J., and Van Mieghem, J. A. (2018), "Forecasting New Product Life Cycle Curves: Practical Approach and Empirical Analysis," *Manufacturing & Service Operations Management, Articles in Advance*.

Jin, C., Hosanagar, K., and Veeraraghavan, S. (2018), "Do ratings cut both ways? Impact of bilateral ratings on platforms," *Working Paper, University of Pennsylvania, Philadelphia, PA*.

Kanoria, Y. and Saban, D. (2017), "Facilitating the search for partners on matching platforms: restricting agents' actions," *Working Paper, Columbia University, New York*.

Karlin, S. and Scarf, H. (1958), "Inventory models of the Arrow-Harris-Marschak type with time lag," in *Studies in the Mathematical Theory of Inventory and Production*, eds. K. A. Arrow, S. Karlin, and H. E. Scarf, pp. 155–178, Stanford University Press, Stanford, CA.

Kelly, F. P. (1979), *Reversibility and Stochastic Networks*, Wiley, New York.

Kingman, J. (1962), "Some inequalities for the queue GI/G/1," *Biometrika*, 49, 315–324.

Koutsoupias, E. and Papadimitriou, C. (2009), "Worst-case equilibria," *Computer Science Review*, 3, 65–69.

Kuhn, H. W. (1950), "Extensive games," *Proceedings of the National Academy of Sciences*, 36, 570–576.

Kuhn, H. W. (1953), "Extensive games and the problem of information," in *Contributions to the Theory of Games*, eds. H. W. Kuhn and A. W. Tucker, pp. 193–216, Princeton University Press, Princeton, NJ.

Leshno, J. D. (2017), "Dynamic matching in overloaded waiting lists," *Working Paper, Columbia University, New York*.

Li, Z., Gilbert, S. M., and Lai, G. (2013), "Supplier encroachment under asymmetric information," *Management Science*, 60, 449–462.

Mai, Y., Hu, B., Hu, Y., Pekeč, S., and Zou, Z. (2018a), "Courteous or crude? Understanding and shaping user behavior in ride-sharing," *Working Paper, Duke University, Durham, NC*.

Mai, Y., Pekeč, S., and Hu, B. (2018b), "Managing product innovation spillover in sourcing," *Working Paper, Duke University, Durham, NC*.

Milgrom, P. and Roberts, J. (1990), "Rationalizability, learning, and equilibrium in games with strategic complementarities," *Econometrica: Journal of the Econometric Society*, 58, 1255–1277.

Müller, A. and Stoyan, D. (2002), *Comparison Methods for Stochastic Models and Risks*, vol. 389, John Wiley & Sons, Hoboken, NJ.

Naor, P. (1969), "The regulation of queue size by levying tolls," *Econometrica: Journal of the Econometric Society*, 37, 15–24.

Nikzad, A. (2018), "Thickness and competition in ride-sharing markets," *Working paper, Stanford University, Stanford, CA*.

Osborne, M. J. and Rubinstein, A. (1994), *A Course in Game Theory*, MIT Press, Cambridge, MA.

Parlaktürk, A. K. and Kumar, S. (2004), "Self-interested routing in queueing networks," *Management Science*, 50, 949–966.

Pinedo, M. L. (2012), *Scheduling: Theory, Algorithms, and Systems*, Springer Science & Business Media, New York.

Pinilla, J. M. and Prinz, F. B. (2003), "Lead-time reduction through flexible routing: application to Shape Deposition Manufacturing," *International Journal of Production Research*, 41, 2957–2973.

Porteus, E. L. (2002), *Foundations of Stochastic Inventory Theory*, Stanford University Press, Stanford, CA.

Roberts, D. M. (1962), "Approximations to optimal policies in a dynamic inventory model," in *Studies in Applied Probability and Management Science*, eds. S. K. K. Arrow and H. Scarf, pp. 207–229, Stanford University Press, Stanford, CA.

Romanyuk, G. (2017), "Ignorance is strength: Improving the performance of matching markets by limiting information," *Working paper, Harvard University, Cambridge, MA*.

Roughgarden, T. and Tardos, É. (2002), "How bad is selfish routing?" *Journal of the ACM (JACM)*, 49, 236–259.

Salas, S. L., Hille, E., and Etgen, G. J. (2007), *Calculus: One and Several Variables*, John Wiley & Sons, Hoboken, NJ.

Schummer, J. (2017), "Influencing waiting lists," *Working paper, Northwestern University, Evanston, IL*.

Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Gusikhin, O., Sanders, M., and Zhang, D. (2015), "Identifying risks and mitigating disruptions in the automotive supply chain," *Interfaces*, 45, 375–390.

Simchi-Levi, D., Wang, H., and Wei, Y. (2018), "Increasing supply chain robustness through process flexibility and inventory," *Production and Operations Management, Forthcoming*.

Su, X. and Zenios, S. (2004), "Patient choice in kidney allocation: The role of the queueing discipline," *Manufacturing & Service Operations Management*, 6, 280–301.

Swinney, R. (2011), "Selling to strategic consumers when product value is uncertain: The value of matching supply and demand," *Management Science*, 57, 1737–1751.

Taylor, T. A. (2017), "On-demand service platforms," *Manufacturing & Service Operations Management*, forthcoming.

Topkis, D. (1998), *Supermodularity and Complementarity*, Princeton University Press, Princeton, NJ.

Veeraraghavan, S. and Debo, L. (2009), "Joining longer queues: Information externalities in queue choice," *Manufacturing & Service Operations Management*, 11, 543–562.

Veeraraghavan, S. K. and Debo, L. G. (2011), "Herding in queues with waiting costs: Rationality and regret," *Manufacturing & Service Operations Management*, 13, 329–346.

Veinott, A. F. and Wagner, H. M. (1965), "Computing optimal (s, S) inventory policies," *Management Science*, 11, 525–552.

Veinott, Jr, A. F. (1966), "On the optimality of (s, S) inventory policies: new conditions and a new proof," *SIAM Journal on Applied Mathematics*, 14, 1067–1083.

Wallenstein, J. and Shelat, U. (2017), *Hopping aboard the sharing economy*, https://www.bcg.com/en-us/publications/2017/strategy-accelerating-growth-consumer-products-hopping-aboard-sharing-economy.aspx.

Wang, J., Cui, S., and Wang, Z. (2018), "Equilibrium strategies in $M/M/1$ priority queues with balking," *Production and Operations Management*, in press.

Yang, L. and Debo, L. (2017), "Invite your friend and you'll move up in line: Leveraging social ties via operational incentives," *Management Science*, forthcoming.

Yang, L., Debo, L., and Gupta, V. (2016), "Trading time in a congested environment," *Management Science*, 63, 2377–2395.

Zheng, Y.-S. and Federgruen, A. (1991), "Finding optimal (s, S) policies is about as simple as evaluating a single policy," *Operations Research*, 39, 654–665.

Zipkin, P. (2008), "On the structure of lost-sales inventory models," *Operations Research*, 56, 937–944.

Biography

Andrew Edward Frazelle is a Ph.D. Candidate in Decision Sciences at The Fuqua School of Business at Duke University. He was born in 1991 in Atlanta, Georgia, USA. Prior to his doctoral studies, he received a B.S. in Industrial Engineering from the Georgia Institute of Technology. He is married to Gabriela Frazelle.